Identifying spikes and seasonal components in electricity spot price data: A guide to robust modeling

Janczura, Joanna and Trueck, Stefan and Weron, Rafal and Wolff, Rodney

Wrocław University of Technology, Poland, Macquarie University, Australia, The University of Queensland, Australia

6 June 2012

Identifying spikes and seasonal components in electricity spot price data: A guide to robust modeling

Joanna Janczura[a], Stefan Trück[b], Rafał Weron[c], Rodney C. Wolff[d]

[a]Hugo Steinhaus Center, Institute of Mathematics and Computer Science, Wrocław University of Technology, Wrocław, Poland
[b]Faculty of Business and Economics, Macquarie University, Sydney, Australia
[c]Institute of Organization and Management, Wrocław University of Technology, Wrocław, Poland
[d]WH Bryan Mining and Geology Research Centre, The University of Queensland, Brisbane, Australia

Abstract

An important issue in fitting stochastic models to electricity spot prices is the estimation of a component to deal with trends and seasonality in the data. Unfortunately, estimation routines for the long-term and short-term seasonal pattern are usually quite sensitive to extreme observations, known as electricity price spikes. Improved robustness of the model can be achieved by (a) filtering the data with some reasonable procedure for outlier detection, and then (b) using estimation and testing procedures on the filtered data. In this paper we examine the effects of different treatment of extreme observations on model estimation and on determining the number of spikes (outliers). In particular we compare results for the estimation of the seasonal and stochastic components of electricity spot prices using either the original or filtered data. We find significant evidence for a superior estimation of both the seasonal short-term and long-term components when the data have been treated carefully for outliers. Overall, our findings point out the substantial impact the treatment of extreme observations may have on these issues and, therefore, also on the pricing of electricity derivatives like futures and option contracts. An added value of our study is the ranking of different filtering techniques used in the energy economics literature, suggesting which methods could be and which should not be used for spike identification.

Keywords: Electricity spot price, Outlier treatment, Price spike, Robust modeling, Seasonality.

1. Introduction

Electricity as a commodity is unique. It is essentially non-storable, while end-user demand shows large variability and strong weather and business cycle dependence. Effects like power plant outages or transmission grid (un)reliability add complexity and reduce predictability. The resulting spot price series exhibit strong seasonality at the annual, weekly and daily levels, as well as mean reversion, very high volatility and abrupt, short-lived and generally unanticipated

Email addresses: joanna.janczura@pwr.wroc.pl (Joanna Janczura), stefan.trueck@mq.edu.au (Stefan Trück), rafal.weron@pwr.wroc.pl (Rafał Weron), rodney.wolff@uq.edu.au (Rodney C. Wolff)

extreme price changes known as spikes or jumps (Benth et al., 2008; Eydeland and Wolyniec, 2012; Huisman, 2009; Weron, 2006).

The first crucial step in defining a model for electricity spot price dynamics consists of finding an appropriate description of the seasonal pattern. The classical approach to seasonal decomposition of a given spot price series P_t into the trend-cycle or long-term seasonal component T_t, the periodic or short-term seasonal component s_t and remaining variability, error, or stochastic component X_t, has its origins in the *Census I method* developed in the 1950s (Makridakis et al., 1998). The functional relationship between these components can assume different forms. Two straightforward possibilities are that they combine in an additive, $P_t = T_t + s_t + X_t$, or a multiplicative fashion, $P_t = T_t \cdot s_t \cdot X_t$. While the former is more popular, if it is applied to logarithms of prices (or log-prices): $\log P_t = T_t + s_t + X_t$, it essentially boils down to the multiplicative specification for the prices themselves: $P_t = e^{T_t+s_t+X_t} = e^{T_t} \cdot e^{s_t} \cdot e^{X_t}$.

There are different suggestions in the energy economics literature for dealing with the trend-seasonal component ($T_t + s_t$ or $T_t \cdot s_t$) of electricity spot prices. For instance, some authors use piecewise constant functions (or dummies) for the months (Bhanot, 2000; Fanone et al., 2012; Fleten et al., 2011; Haldrup et al., 2010; Higgs and Worthington, 2008; Knittel and Roberts, 2005; Lucia and Schwartz, 2002). Other model the seasonal pattern by sinusoidal functions or sums of sinusoidal functions of different frequencies (Bierbrauer et al., 2007; Borovkova and Permana, 2006; Cartea and Figueroa, 2005; Erlwein et al., 2010; Geman and Roncoroni, 2006; Keles et al., 2012; Lucia and Schwartz, 2002; Pilipovic, 1998; Seifert and Uhrig-Homburg, 2007; Weron, 2008), sometimes coupled with an exponentially weighted moving average (EWMA) (De Jong, 2006). As a more robust to outliers and less strictly periodic alternative to Fourier analysis, wavelet decomposition and smoothing have been applied by Janczura and Weron (2010, 2012), Stevenson (2001), Stevenson et al. (2006), Weron (2006, 2009) and Weron et al. (2004a,b).

A critical issue in estimation of the seasonal pattern is that it might be substantially affected by price spikes. While it is clear that price spikes should be captured by an adequate stochastic model, like mean reverting jump-diffusion (Bierbrauer et al., 2007; Borovkova and Permana, 2006; Cartea and Figueroa, 2005; Clewlow and Strickland, 2000; Geman and Roncoroni, 2006; Jabłońska et al., 2011; Nomikos and Soldatos, 2010; Seifert and Uhrig-Homburg, 2007; Weron, 2008) or a regime-switching model (Becker et al., 2007; De Jong, 2006; Higgs and Worthington, 2008; Hirsch, 2009; Huisman and Mahieu, 2003; Janczura and Weron, 2010, 2012; Keles et al., 2012; Mari, 2008; Mount et al., 2006; Weron, 2009), the literature does not agree on whether these observations have to be included or excluded in the estimation of the seasonal pattern.

Even worse: despite the fact that price spikes are among the most pronounced features of electricity markets and account for a large part of the total variation of changes in spot prices, there is no commonly accepted definition of a price spike (Weron, 2006). A variety of methods for identification has been suggested. However, excluding the preliminary results of Trück et al. (2007), so far there has been no thorough empirical study on the effects of alternative treatment of the price spikes on parameter estimates for the seasonal pattern or the stochastic component of the spot price. It is exactly our goal to examine the consequences of the treatment of such extreme events in the estimation procedures. To identify the spikes we will consider a variety of different approaches. After such 'cleaning' of the observed spot prices we will then compare the resulting seasonal patterns and parameter estimates of the stochastic component.

This paper is intended as a guide to robust modeling of spot electricity prices. In Section 2 we review some stylized facts as well as possible reasons for price spikes in electricity markets. We further provide an overview of methods (or filters) that have been suggested in the literature to identify price spikes. In Section 3 we outline different procedures for deseasonalizing the data and in Section 4 we describe the applied stochastic models focusing on Markov regime-switching models. In Section 5 we present an empirical study of spot prices from the German EEX and Australian New South Wales electricity markets, where we identify outliers and estimate seasonal patterns. Then, in Section 6 we summarize the results of a Monte Carlo simulation study that is aimed at comparing the different techniques with respect to estimating the correct seasonal pattern as well as the parameters of the underlying stochastic process. Finally, in Section 7 we conclude.

2. Price spikes

2.1. Occurrence of spikes

One of the most pronounced features of electricity markets are the abrupt and generally unanticipated extreme changes in the spot prices known as spikes (or jumps). Despite their rarity, price spikes are the very motive for designing insurance and hedging protection against electricity price movements. This is one of the most serious reasons for including discontinuous components in realistic models of electricity price dynamics. Failing to do so will greatly underestimate, say, the option premium, and thus increase the risk for the writer of the option on the electricity spot price.

Within a very short period of time, the system price can increase substantially and then drop back to the previous level, see Figure 1 where spot (log-)prices from two major power markets are depicted. These temporary price escalations account for a large part of the total variation of changes in spot prices. Indeed, as the literature points out, electricity prices are far more volatile than any other commodity (Kaminski, 2004; Simonsen, 2005; Weron, 2006). Measuring volatility by the daily standard deviation of returns, electricity exhibits extreme volatility up to 50% while in general even very volatile stocks mostly do not exceed 4%.

The spike intensity is also non-homogeneous in time. Generally, the spikes are especially notorious during on-peak hours, i.e., around 9 am and 6 pm on business days, and during high consumption periods: winter in Scandinavia, summer in mid-western U.S., Australian summer in Australia, and so on. As the time horizon increases and the data are aggregated, the spikes are less and less apparent. For weekly or monthly averages, the effects of price spikes are usually neutralized in the data. However, it is not uncommon that prices from one day to the next, or even within just a few hours, increase tenfold. The ‘spiky’ nature of spot prices is the effect of non-storability of electricity. Electricity to be delivered at a specific hour cannot be substituted for electricity available shortly after or before. As currently there is no efficient technology for storing vast amounts of power, it has to be consumed at the same time as it is produced. Hence, extreme load fluctuations – caused by severe weather conditions often in combination with generation outages or transmission failures – can lead to price spikes (Huisman, 2009). The spikes are normally quite short-lived, and as soon as the weather phenomenon or outage is over, prices fall back to a normal level.

But one may wonder: why are the spikes so extreme? Simonsen (2005) and Weron (2006) argue that the main reason for spikes in electricity spot prices is the bidding strategies used by

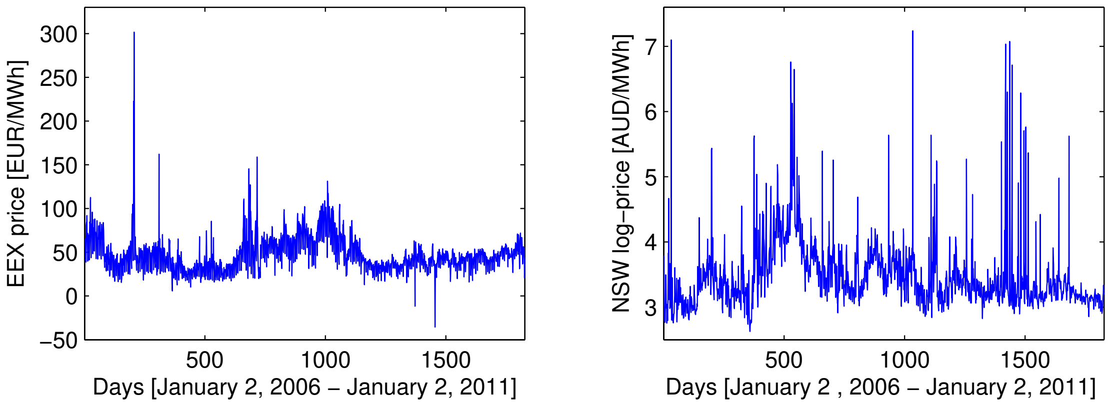

Figure 1: Mean daily electricity spot prices from the European Energy Exchange (EEX, Germany; *left*) and the New South Wales Electricity Market (NSW, Australia; *right*). Note that, in the right panel, the log-prices (and not the prices themselves) are plotted. The logarithmic scale dampens the extreme spikiness of the Australian NSW prices, which can reach up to a bid-cap of 10000 AUD/MWh during peak hours, and since 2012, 12500 AUD/MWh. The price spikes are clearly visible in both cases. The EEX dataset even includes a few downward spikes or price drops.

the market players. Since electricity is an essential commodity for many market participants, it is quite important to secure a sufficient and continuous supply of power at any time. Hence, some of the agents may be willing to pay almost any price to secure this supply. As a consequence, on a regular basis, bids are placed at the maximum allowed level (i.e., bid-cap) for the amount of electric power anticipated by the market players to be needed for that particular hour, or for the entire 24 hours of the next day (in the case of the EEX). It is important to recall that in uniform-price (or marginal) auction markets the spot price is what a buyer has to pay for each unit of power independently of the initial bid, as long as the bid was above (or equal to) the spot price. Hence, with this type of strategy, the worst case scenario is that a wholesale buyer has to pay the high prices for a maximum of 24 hours. After this period, the buyer is free to try to get cheaper power from alternative sources. With this type of bidding, there will always be some buyers that are willing to pay a considerable amount in order to cover their need of electricity.

2.2. *Identification of spikes*

The identification of spikes is a very important issue as it can influence the estimation of the deterministic and stochastic components for models of electricity spot price dynamics. However, in the literature the definition of a spike so far has been a rather subjective matter. Obviously, price spikes are defined as prices that surpass a specified threshold for a brief period of time. But it is difficult to gain any consensus on what that threshold or time interval should be. The identification methods suggested in the literature include:

- *Fixed price thresholds* where all prices exceeding some subjectively chosen price level (e.g., 100 EUR/MWh) are classified as spikes (Boogert and Dupont, 2008; Lapuerta and Moselle, 2001). Graphical techniques, like the sample mean excess function, might be used for the selection of a more optimal cutoff level (Fanone et al., 2012); however, this procedure can hardly be automated.

- *Variable price thresholds* where a certain percentage of the highest (and/or lowest) prices, e.g., the upper 1% of prices (Trück et al., 2007), is classified as outliers. A related technique is to classify as spikes all prices exceeding the mean price level by three standard deviations, with the outlying observations removed one by one in a 'recursive filter' fashion (see the RFP method below).

- *Fixed price change thresholds* where price increments or price returns exceeding some threshold, e.g., 30% as in Bierbrauer et al. (2004), are treated as outliers.

- *Variable price change thresholds*, more commonly known as the 'recursive filter' technique, where prices corresponding to the price increments or returns exceeding three standard deviations of all returns are removed one by one in an iterative procedure (Cartea and Figueroa, 2005; Clewlow and Strickland, 2000; Weron et al., 2004b; Weron, 2008).

- *Wavelet filtering* where the signal (the price series) is first decomposed using the wavelet transform, then reconstructed up to a certain detail level. Leaving out the highest one or two detail levels leads to price series with eliminated price spikes, at least for half-hourly data (Stevenson, 2001; Stevenson et al., 2006). Alternatively, more detail levels can be removed leading to a much smoother signal, which then can be subtracted from the original prices to yield deseasonalized and detrended data. In the latter, spikes can be found as observations exceeding three standard deviations of the differences (Trück et al., 2007; Weron, 2006).

- *Thresholds implied by Gaussian 90% prediction intervals*, as in Borovkova and Permana (2006) who considered as jumps those price moves that were outside 90% prediction intervals, implied by the normal distribution with the mean and variance given by the 60-days moving average and 60-days moving variance of the price moves.

- *Thresholds yielding the best model in terms of matching kurtosis*, as in Geman and Roncoroni (2006) who filtered raw price data using different thresholds and selected the one leading to the best calibrated model in view of its ability to match the kurtosis of observed daily price variations.

- *MRS model classification* as a byproduct of calibrating a Markov regime-switching model to deseasonalized and detrended data (Weron, 2009; Janczura and Weron, 2010). The Expectation-Maximization (EM) algorithm (Hamilton, 1990; Kim, 1994; Janczura and Weron, 2012) is an iterative two-step procedure which computes the conditional probabilities for the process being in a certain regime at a given time. All prices with the probabilities of being in one of the extreme regimes exceeding a certain probability threshold, say $\frac{1}{2}$, can be classified as outliers.

One more method is proposed in this article:

- *Recursive seasonal model* where in each step the number of outliers is reduced based on the marginal reduction of the mean squared error (MSE) of fitting a seasonal pattern to the new series. In each step the observations above the determined new threshold are replaced by the threshold and then a new seasonal pattern is estimated. In each recursive step the next threshold is chosen to be the mean of the three highest prices below the threshold. The procedure stops when the reduction in the MSE of the fitted seasonal pattern is less than 1%, or some similar nominated tolerance.

Obviously, the different techniques may lead to quite different identification of spikes in electricity spot prices. To our best knowledge, excluding the preliminary results of Trück et al. (2007), so far no thorough empirical study has been conducted to examine the effects the different spike identification methods may have on the determination of the seasonal pattern or parameter estimates of the stochastic components of electricity prices.

Finally, it is important to consider also seasonality and trend in the data. Otherwise the identification of certain observations as outliers will be clearly dependent on the weekday or month if there is also a yearly pattern in the data. The deseasonalization techniques will be discussed in Section 3.

2.3. Replacing spikes by 'normal' values

Once the spikes are identified they have to be replaced by 'normal', less spiky values. Again there is no consensus what these 'normal' prices should be. A non-exhaustive list of solutions proposes that spikes are replaced by:

- the chosen threshold (Shahidehpour et al., 2002);
- the mean of the two neighboring prices (Weron, 2008);
- one of the neighboring prices (Geman and Roncoroni, 2006); or
- 'similar day' values, e.g., the median of all prices having the same weekday and month (Bierbrauer et al., 2007).

Since we will be working with deseasonalized data, yet another simple method is proposed and exclusively used in this article in which the spikes are replaced by:

- the mean of the deseasonalized prices.

Finally, note that if we are not interested in the time evolution, only in distributional properties, we do not have to replace the spikes at all.

3. Deseasonalization

There are different approaches in the literature to dealing with the seasonal pattern in electricity price dynamics. 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 (i.e., a short-term seasonal component, STSC) and a long-term trend-seasonal component (LTSC) T_t, which represents the long-term non-periodic fuel price levels, the changing climate/consumption conditions throughout the years and strategic bidding practices. As mentioned above, multiplicative models become additive under logarithmic transform.

There are essentially three different suggestions in the energy economics literature for dealing with the trend-seasonal component f_t of electricity spot prices. The first is to use piecewise constant functions (Bhanot, 2000; Fanone et al., 2012; Fleten et al., 2011; Haldrup et al., 2010; Higgs and Worthington, 2008; Knittel and Roberts, 2005; Lucia and Schwartz, 2002). Since this approach yields a non-smooth trend-seasonal component and, hence, requires additional smoothing treatment, it will not be used in this study.

3.1. The 'sin-EWMA' long-term seasonal component

The second suggestion is to model the trend-seasonal pattern by sinusoidal functions (Bierbrauer et al., 2007; Borovkova and Permana, 2006; Cartea and Figueroa, 2005; Geman and Roncoroni, 2006; Lucia and Schwartz, 2002; Pilipovic, 1998; Seifert and Uhrig-Homburg, 2007; Weron, 2006). Since for many datasets sinusoidal functions are too regular in their periodicity to describe the long-term spot price evolution, they can be coupled with an exponentially weighted moving average (EWMA), as in De Jong (2006). One of two models for T_t used in this study is given by:

$$T_t = a_1 \sin\left\{2\pi\left(\frac{t}{365} + a_2\right)\right\} + a_3 + a_4 \mathrm{EWMA}_t^{0.975}. \tag{1}$$

Following De Jong (2006), the exponentially weighted moving average uses a decay factor of $\lambda = 0.975$, i.e., $\mathrm{EWMA}_t^{\lambda} = (1-\lambda)\cdot P_t + \lambda \cdot \mathrm{EWMA}_{t-1}^{\lambda}$. The parameters $a_1, ..., a_4$ can be estimated using nonlinear least squares. This LTSC model is denoted by 'sin-EWMA' later it in the text.

3.2. The wavelet long-term seasonal component

The third approach is to use wavelet decomposition and smoothing as more robust to outliers and a less periodic alternative to Fourier analysis (Janczura and Weron, 2010; Stevenson et al., 2006; Weron, 2006). Recall, that wavelets belong to families and come in pairs of a father and a mother wavelet for a given order, like the Daubechies wavelets of order 24 used here. The different families and orders of wavelets make different trade-offs between how compactly they are localized in time and their smoothness (Percival and Walden, 2000). 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. This procedure, also known as lowpass filtering, yields a traditional linear smoother. Here we use the S_6 approximation, which roughly corresponds to bimonthly ($2^6 = 64$ days) smoothing. Note that in terms of smoothing properties wavelet filters show some resemblance to kernel density estimators (KDEs). In fact, a KDE with a Gaussian kernel and standard deviation of about 90 days yields seasonal components similar to the wavelet estimated LTSC S_6. However, unless a complicated local bandwidth choice is employed, KDEs are not locally adaptive and for datasets with seasonal components of varying time frequencies and amplitudes, like electricity spot prices, the bandwidth is typically either too small or too large. Wavelet filters do not share this drawback and, in general, are superior to KDEs (Härdle et al., 1998).

3.3. The 'average week' short-term seasonal component

The price series without the LTSC is obtained by subtracting T_t from P_t. Next, the weekly periodicity s_t can simply be removed by subtracting the 'average week' calculated as the arithmetic mean of the LTSC-deseasonalized (log-)prices corresponding to each day of the week. This approach is equivalent to having dummies for each day of the week, as in De Jong (2006). Public holidays can be and are treated here as the eighth day of the week. The median or a truncated mean can be used instead of the mean as an alternative more robust to outliers, however, usually the differences are not substantial.

3.4. Forecasting of the seasonal components

Although forecasting of the seasonal components is beyond the scope of the present paper, let us briefly comment on this issue. Trigonometric or periodic functions – such as the sinusoidal LTSC or the 'average week' STSC – can be easily extrapolated into the future. Unfortunately, combining sinusoidal functions with an exponentially weighted moving average complicates things very much, because EWMA_{t+1} is dependent on the unknown future price P_{t+1}. Also predicting the wavelet LTSC beyond the next few weeks is a difficult task, since individual wavelet functions are quite localized in time or (more generally) in space. Preliminary research suggests, however, that despite this feature the wavelet LTSC can be extrapolated into the future yielding an on-average comparable prediction of the level of future spot prices to that of an extrapolation of a sinusoidal LTSC (Nowotarski et al., 2011). An alternative, potentially promising approach would be to use forward looking information, like smoothed forward curves. The information carried by forward prices provides insights as to the future evolution of spot prices. However, forward prices also include the risk premium, which should somehow be separated from the spot price forecast for it to be useful (Janczura and Weron, 2012).

4. Stochastic models for the deseasonalized spot price

The specific behavior of electricity prices, in particular extreme volatility and the described price spikes have forced producers and wholesale consumers to hedge against price movements. This in turn has significantly propelled research in electricity price modeling and forecasting. One of the first models which has been examined in the context of electricity markets is the classical mean-reversion or Vasiček-process, see, e.g., Lucia and Schwartz (2002) and Knittel and Roberts (2005). On the other hand, by making certain assumptions about the functional form of the supply and demand curve, Barlow (2002) derives a non-linear Ornstein-Uhlenbeck process as appropriate model for electricity prices. Early publications on models for electricity prices with a jump component include Deng (1998), Bhanot (2000), Clewlow and Strickland (2000), Kaminski (2004) and Knittel and Roberts (2005) and are typically based on a jump-diffusion as in Merton (1976). More recently, Geman and Roncoroni (2006) suggest to model electricity log-prices with a one-factor Markov jump-diffusion model. A particularly nice feature of their model is that the jump direction and intensity is dependent of the current level of the spot price. Hambly et al. (2009)

suggest a specification of a jump-diffusion model including both a mean-reverting diffusion and mean-reverting jump process. Jabłońska et al. (2011) use a multiple mean reversion jump-diffusion model with the mean reversion taking place at several different time and price scales.

Also GARCH-type models have been applied to describe the structure of electricity prices, see, e.g., Garcia et al. (2005) and Knittel and Roberts (2005). These studies find that the additional use of GARCH error terms outperforms the predictive power of several other models. On the other hand, Misiorek et al. (2006) find that, despite the heteroscedastic nature of the residuals in the autoregressive models, the addition of a GARCH component in the specification does not improve the accuracy of point forecasts though it slightly improves interval forecasts. In a more recent study, Koopman et al. (2007) suggest an approach using periodic dynamic long memory regression models with GARCH errors.

In parallel to jump-diffusion models, Markov regime-switching models (MRS; sometimes also called hidden Markov models – HMM) have been gaining popularity in the energy economics literature since the late 1990s. To our best knowledge, they were first applied to electricity prices by Deng (1998) and Ethier and Mount (1998), who considered 2-state models with mean-reverting AR(1) processes for the log-prices in both regimes. Independent spike (IS) MRS models, where the base regime is modeled by an AR(1) process and the spike regime by an independent normal (log-normal or Pareto) random variable, were introduced by Huisman and de Jong (2003) and Weron et al. (2004a). In the following years various MRS models have enjoyed extensive use due to their relative parsimony and the ability to capture the unique characteristics of electricity spot prices (Bierbrauer et al., 2007; De Jong, 2006; Erlwein et al., 2010; Haldrup et al., 2010; Higgs and Worthington, 2008; Hirsch, 2009; Huisman, 2009; Janczura and Weron, 2010, 2012; Kanamura and Ōhashi, 2008; Karakatsani and Bunn, 2008; Keles et al., 2012; Kholodnyi, 2005; Mari, 2008; Mount et al., 2006; Weron, 2009).

The underlying idea behind Markov regime-switching is to represent the observed stochastic behavior of a specific time series by two (or more) separate states or regimes with different underlying stochastic processes. The switching mechanism between the states is assumed to be an unobserved (latent) Markov chain. In contrast to jump-diffusion models, MRS models allow for consecutive spikes in a very natural way. Also the return of prices after a spike to the 'normal' regime is straightforward, as the regime-switching mechanism admits temporal changes of model dynamics. MRS models are also more versatile than the classical hidden Markov models (or HMM in the strict sense, see Cappe et al., 2005) popular in engineering sciences, since they allow for temporary dependence within the regimes, in particular, for mean reversion. As the latter is a characteristic feature of electricity prices it is important to have a model that captures this phenomenon. Indeed, the base regime is typically modeled by a mean-reverting diffusion (Huisman, 2009), sometimes heteroskedastic (Janczura and Weron, 2010). For the spike regime(s), on the other hand, a number of specifications has been suggested in the literature, ranging from mean-reverting diffusions to heavy-tailed random variables.

The switching mechanism between the states (or regimes) of a MRS model is assumed to be an unobserved (latent) Markov chain R_t. It is described by the transition matrix $\mathbf{P}$ containing the probabilities $p_{ij} = P(R_{t+1} = j \mid R_t = i)$ of switching from regime i at time t to regime j at time

$t+1$. For instance, for $i, j \in \{1, 2\}$ we have:

$$\mathbf{P} = (p_{ij}) = \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} = \begin{pmatrix} 1 - p_{12} & p_{12} \\ p_{21} & 1 - p_{21} \end{pmatrix}. \tag{2}$$

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}.

In this paper we focus on a specification popular in the energy economics literature where the individual regimes are driven by independent processes (for a review see, e.g., Janczura and Weron, 2010). The (deseasonalized) spot price X_t is defined as:

$$X_t = \begin{cases} X_{t,1} & \text{if } R_t = 1, \\ X_{t,2} & \text{if } R_t = 2, \\ X_{t,3} & \text{if } R_t = 3. \end{cases} \tag{3}$$

The first (base) regime describes the 'normal' price behavior and is given by the mean-reverting, heteroskedastic process. In continuous time, the process is defined by a (continuous-time) stochastic differential equation that governs the dynamics of the (deseasonalized) price, X_t:

$$dX_{t,1} = (\alpha_1 - \beta_1 X_{t,1})dt + \sigma_1 dW_t. \tag{4}$$

Note that, in this specification, the Brownian motion W_t is responsible for fluctuations around the long-term mean α_1/β_1 proportional to σ_1, while β_1 is the speed of mean-reversion. In discrete time, the process can be written as

$$X_{t,1} = \alpha_1 + (1 - \beta_1)X_{t-1,1} + \sigma_1 \epsilon_t, \tag{5}$$

where ϵ_t is standard i.i.d. Gaussian noise. The second regime represents the sudden price spikes caused by unexpected supply shortages and is given by i.i.d. random variables from the shifted log-normal distribution:

$$\log(X_{t,2} - X(q_2)) \sim \mathrm{N}(\mu_2, \sigma_2^2), \quad X_{t,2} > X(q_2). \tag{6}$$

Finally, the third regime (responsible for the sudden price drops) is governed by the shifted 'inverse log-normal' law:

$$\log(-X_{t,3} + X(q_3)) \sim \mathrm{N}(\mu_3, \sigma_3^2), \quad X_{t,3} < X(q_3). \tag{7}$$

Note that the values $X(q_i)$ in the above equations denote the q_i-th quantile of the dataset. Such specification of the spike and drop regime distributions ensures that the observations below (above) the q_2-th (q_3-th) quantile will not be classified as spikes (drops). In general, the values of q_i can be obtained by an optimization procedure. Here, for simplicity, we set q_2 to be the third quartile and q_3 to be the first quartile of the dataset.

Calibration of regime-switching models with an observable state process boils down to the problem of independently estimating parameters in each regime. In case of MRS models, though, the calibration process is not straightforward, since the state process is latent and not directly observable. We have to infer the parameters and state process values at the same time. In this paper we use the Expectation-Maximization (EM) algorithm that was first applied to MRS models by Hamilton (1990) and later refined by Kim (1994). It is a two-step iterative procedure, reaching a local maximum of the likelihood function:

- **Step 1:** For a parameter vector θ compute the conditional probabilities for the process being in regime j at time t, the so-called 'smoothed inferences'.
- **Step 2:** Calculate new and more exact maximum likelihood estimates of θ using the likelihood function, weighted with the smoothed inferences from Step 1.

Note that the introduction of independent regimes results in a significantly increased computational burden. See Janczura and Weron (2012) for an efficient modification of the algorithm to overcome this problem.

5. Results for EEX and NSW market data

In this section we compare the different approaches with respect to the estimation of the seasonal pattern and the parameters of the stochastic process. In particular we illustrate the impacts of using different outlier detection techniques (or 'filters') on the number and size of observations being classified as price spikes or price drops. Since the analysis of market data cannot provide conclusive results on the performance of the individual methods – recall that the 'actual' seasonal pattern and stochastic process are unknown – in Section 6 we will conduct an extensive simulation study to further investigate the issue. In this Monte Carlo experiment we will assume that we know the 'actual' seasonal components and parameters of the stochastic process such that we can examine the effects of the outlier filters on the estimates of the seasonal pattern and stochastic models.

5.1. The data and their deseasonalization

In order to use realistic seasonal and stochastic components, the parameters of the seasonal and stochastic components are estimated from spot prices recorded in two major power markets. Namely, we use mean daily (baseload) spot prices from the European Energy Exchange (EEX; Germany) and the mean daily (baseload) spot log-prices from the New South Wales Electricity Market (NSW; being part of the National Electricity Market in Australia). Both datasets comprise five years of data ranging from January 2, 2006, to January 2, 2011, totaling 1827 observations (261 full weeks). The datasets are plotted in Figure 1. Note that for the EEX market we provide a plot of the original price data while for the NSW electricity market, the log-prices and not the prices themselves are plotted. The logarithmic scale dampens the extreme spikiness of the Australian NSW prices, which can reach up to a bid-cap of 10000 AUD/MWh during peak hours and, since 2012, 12500 AUD/MWh. The price spikes are clearly visible in both cases while the NSW market exhibits a significantly larger number of extreme price observations. Note that the EEX dataset even includes a few price drops with a negative system price (see Fanone et al., 2012, for a discussion).

We use two methods for estimating the long-term seasonal component T_t, either (i) a wavelet smoother of level 6 (i.e., S_6, see Section 3.2) or (ii) a sine function combined with an exponentially weighted moving average (i.e., 'sin-EWMA', see Section 3.1). After detrending (subtracting T_t from the spot prices), the short-term seasonal component s_t is calculated as the 'average week' (with holidays treated as the eighth day of the week; see Section 3.3) and removed from the detrended data. The remaining stochastic component $X_t = P_t - T_t - s_t$ is modeled as a Markov regime-switching model defined by eqns. (5)-(7).

5.2. The outlier filters

In order to examine the effect of outlier treatment on the estimation of the seasonal and stochastic components of electricity prices, we apply different outlier detection methods (or 'filters') to the EEX and NSW spot (log-)prices and compare the resulting seasonal patterns and the properties of the remaining stochastic part. Since the seasonal effects would obviously influence outlier identification, some detrending technique should be performed prior to applying the outlier detection methods. Here, we proceed in the following way. First, we estimate and subtract the seasonal pattern from the original (log-)prices. Next, we apply one of the filtering methods to the deseasonalized dataset. Finally, we obtain the dataset without outliers by adding the original seasonal patterns to the filtered series. In this study, we consider the following approaches for outlier detection:

(1) No removal of spikes or drops such that the original series is used for later estimation of the seasonal pattern and the stochastic process (ORG).
(2) Fixed price threshold with deseasonalized prices exceeding the range of $(-40, 40)$ – or $(-0.5, 0.5)$ for deseasonalized log-prices – identified as outliers (FPT).
(3) Variable price threshold with 2.5% highest and 2.5% lowest deseasonalized prices treated as outliers (VPT1).
(4) Variable price threshold with 10% highest and 10% lowest deseasonalized prices treated as outliers (VPT2).
(5) Recursive filter on prices – *variable price threshold* using three standard deviations on deseasonalized prices (RFP).
(6) Recursive filter on price differences – *variable change price threshold* using three standard deviations on deseasonalized price differences (RFD).
(7) Recursive seasonal model (RM).
(8) MRS model classification (RSC).

Once the price spikes and drops are identified they have to be replaced by 'normal', less spiky values. Since we are working with deseasonalized data, for simplicity, we replace them by the mean of the deseasonalized prices.

5.3. Results of outlier detection

In a first step we consider the number of observations that are being classified as price spikes by each of the methods. We find that there are significant differences between the filters both with respect to the number of identified price spikes, as well as for the average magnitude of a positive spike or negative price drop. The results for the different methods are summarized in Table 1 and illustrated for the EEX market in Figure 2. Note that the figure presents the results only for the application of a wavelet LTSC, however, as indicated by Table 1, results with respect to the number and magnitude of spikes for the sin-EWMA long-term seasonal component are generally quite similar.

Looking at the results for the EEX market we find that the fixed price-threshold (FPT) method yields the smallest number of price spikes (0.77%), while the variable price threshold with 10% highest and 10% lowest prices treated as outliers (VPT2) clearly identifies the highest number

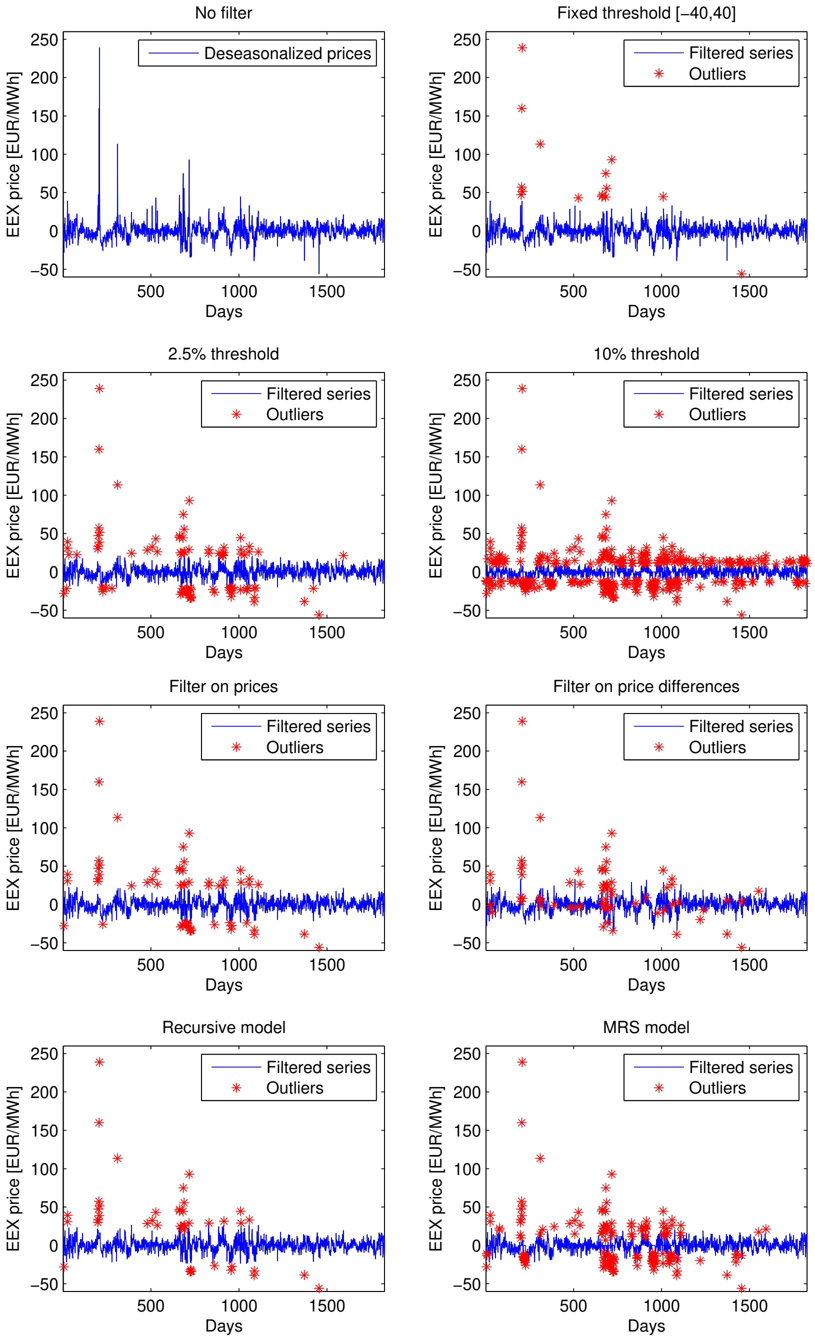

Figure 2: Price observations for the EEX market that are identified as price spikes or price drops based on different methods for outlier detection. The figures show the deseasonalized price series *(upper left)* as well as the filtered series and the observations classified as outliers based on the seven considered methods: FPT, VPT1, VPT2, RFP, RFD, RM and RSC *(left to right, top to bottom)*, see Section 5.2 for filter definitions. The deseasonalization approach used here is a wavelet approximation for the long-term seasonal pattern.

Table 1: Number of observations classified as price spikes and drops for each of the considered methods. The table also provides the average (log-)price level of observations that were identified as price spikes or price drops by the outlier detection method. Abbreviations are defined in Section 5.2.

Dataset	EEX		NSW		EEX		NSW	
LTSC	Wavelet	sin-EWMA	Wavelet	sin-EWMA	Wavelet	sin-EWMA	Wavelet	sin-EWMA
Filter	Frequency of price spikes				Average magnitude of (log-)price spike			
FPT	0.77%	0.93%	5.42%	5.91%	142.83	136.14	4.91	4.87
VPT1	2.52%	2.52%	2.52%	2.52%	107.49	106.90	5.62	5.67
VPT2	10.02%	10.02%	10.02%	10.02%	75.34	78.61	4.47	4.48
RFP	1.97%	1.97%	3.45%	3.78%	112.50	112.14	5.37	5.27
RFD	2.52%	2.57%	6.19%	6.19%	96.18	95.03	4.73	4.77
RM	1.75%	1.64%	3.17%	2.96%	115.55	116.28	5.46	5.53
RSC	4.00%	4.93%	13.36%	10.02%	97.15	88.96	4.30	4.45
Filter	Frequency of price drops				Average magnitude of (log-)price drop			
FPT	0.05%	0.16%	2.08%	2.85%	-35.57	5.67	3.25	3.44
VPT1	2.52%	2.52%	2.52%	2.52%	33.60	36.25	3.32	3.48
VPT2	10.02%	10.02%	10.02%	10.02%	38.05	36.68	3.34	3.27
RFP	1.31%	0.82%	0.05%	0.27%	28.52	27.81	2.76	3.48
RFD	1.09%	1.09%	2.08%	1.92%	36.27	36.36	3.25	3.28
RM	0.66%	0.27%	–	–	21.93	11.21	–	–
RSC	4.87%	3.61%	6.90%	0.38%	39.97	38.63	3.39	2.92

of spikes (roughly 10%). Among the other methods considered, the MRS model classifies approximately 4% to 5% of the observations as spikes, while all of the other models yield between 1.75% and 2.5% price spikes. As a result of the large differences between the models in terms of spike identification, there are also significant differences with respect to the average magnitude of a spike. Clearly, the model classifying the lowest number of observations as price spikes will yield the highest magnitude for an average spike. Therefore, we find that the average magnitude of an observation identified as a spike for the FPT is around EUR 143 while it is only EUR 75 for VPT2. For the other methods the average price spike is between EUR 96 and EUR 116. Note that these EUR values refer to the original observations and not deseasonalized price data. For price drops in the EEX market we find similar results. The FPT method yields the smallest number of observations (0.05%) being identified as price drops, while the VPT2 again gives the highest number of drops (10%), followed by the MRS model classification (3.61% for wavelet, and 4.87% for sin-EWMA LTSC). The range of observations being classified as drops by the other methods is between 0.27% and 2.52%. Note that FPT only classifies one observation with a price of EUR −35.57 as a drop, while the average price drop is EUR 39.97 for the MRS model classification and EUR 38.05 for VPT2.

Let us now consider the results for the NSW market that are reported in Table 1 and illustrated by Figure 3. Recall that, for this market, we do not use the actual spot price but log-prices, to dampen the extreme spikiness. Still, apart from the filters involving a 2.5% or 10% variable price threshold, all methods characterize a greater percentage of observations as price spikes for NSW data than for EEX data. Similar to the results for the EEX, we find that the VPT2 (10%) and RSC (13% and 10%) techniques classify the largest number of observations as price spikes. Also

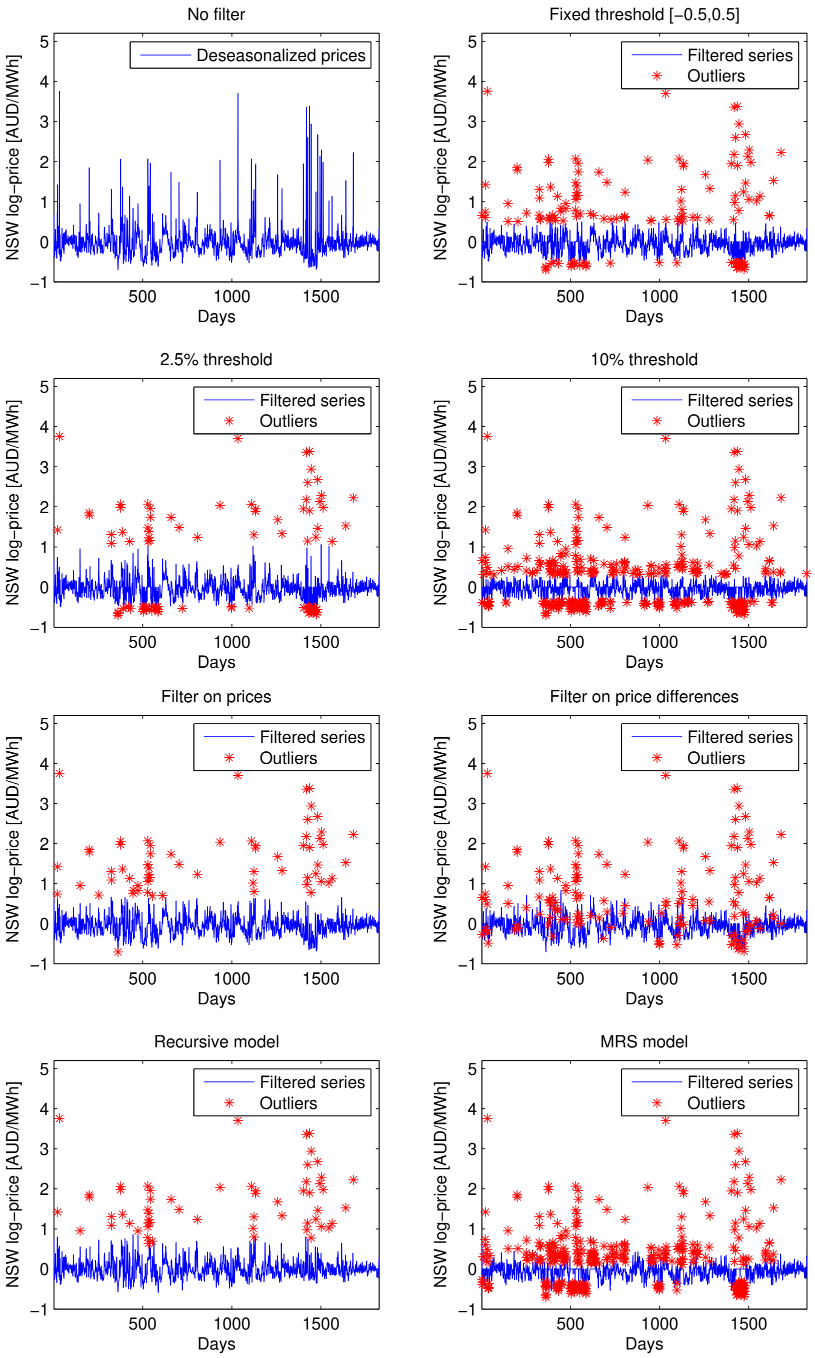

Figure 3: Price observations for the NSW market that are identified as price spikes or price drops based on different methods for outlier detection. The figures show the deseasonalized log-price series *(upper left)* as well as the filtered series and the observations classified as outliers based on the seven considered methods: FPT, VPT1, VPT2, RFP, RFD, RM and RSC *(left to right, top to bottom)*, see Section 5.2 for filter definitions. The deseasonalization approach used here is a wavelet approximation for the long-term seasonal pattern.

the recursive filter on price differences classifies approximately 6% of the observations as price spikes. Interestingly, this time also the fixed price-threshold (FPT) method yields a large number of positive spikes between 5% and 6%; however, this might be due to the rather arbitrary choice of classifying all deseasonalized prices outside the range of $(-0.5, 0.5)$ as outliers. A comparably smaller number of observations, between 2.96% and 3.78%, is identified as spikes by the RM, RFP and VPT1 methods. Again, the differences between the filters with respect to the number of observations being classified as price spikes are substantial in terms of the average magnitude of a price spike, which is between 5.27 and 5.67 for the VPT1, RM and RFP methods and around 4.4 on the log-scale for the MRS classification (RSC). Results for price drops in Table 1 indicate that for the NSW market there are clearly fewer observations classified as price drops in comparison to the number of price spikes. However, the individual methods yield quite different results with respect to the number of drops in comparison to the EEX. For the NSW market the RFP and RM techniques yield the smallest number of price drops – the recursive seasonal model does not identify even a single drop while the recursive filter on prices only classifies 0.05% and 0.27% of the observations as drops. This is also confirmed by Figure 3. On the other hand, we find that next to the VPT2 technique, also the RSC filter identifies up to 6.90% of the log-prices as price drops. Interestingly, this is only true when the wavelet LTSC is being used while for the sin-EWMA seasonal pattern the fraction of price drops is much lower. The FPT, VPT1 and RFD methods suggest around 2% to 3% of the deseasonalized log-prices as drops.

5.4. Seasonal patterns for outlier filtered data

In the next step we examine the impact of removing the outliers on the estimation of the seasonal pattern. Figures 4 to 7 show the estimated long-term (LTSC) and short-term (STSC) seasonal components for the EEX market. Results are reported for different outlier identification techniques. Figure 4 shows the original spot price series and the estimated wavelet LTSC. Results are illustrated for four different methods: using the original price series that includes all the extreme price spikes and drops (ORG); and using the VPT1, RFD or RSC filters for which the identified outliers are replaced by the mean of the deseasonalized series before reestimating the seasonal pattern. Figure 5 shows the results when instead of a wavelet approximation, sin-EWMA is used for the estimation of the long-term trend. For both figures the lower left panel in the second row zooms into the original spot prices and the estimated long-term trend for the time period June to December 2006. The lower right panel shows the estimated short-term seasonal pattern for Monday to Sunday as well as for holidays (indicated by 'Ho' on the horizontal axis).

Let us now consider the estimated long-term seasonal components. Figures 4 to 7 illustrate that the LTSC estimated for the original prices (or log-prices for NSW) can deviate quite substantially from the other long-term patterns. This is particularly true for periods where a large number of price spikes could be observed. For example, for the EEX we find that between June and August 2006 due to several price spikes in July 2006 the estimated LTSC for the original data deviates by up to EUR 15 from the estimated long-term patterns based on the outlier-corrected data. As shown in the lower left panel of Figures 4 and 5 the effect seems to be even more pronounced when instead of the wavelet approximation the sin-EWMA pattern is being used for the long-term trend. A reason for this might be that for the wavelet trend we use the S_6 approximation, which roughly corresponds to bimonthly ($2^6 = 64$ days) smoothing and is therefore less sensitive than

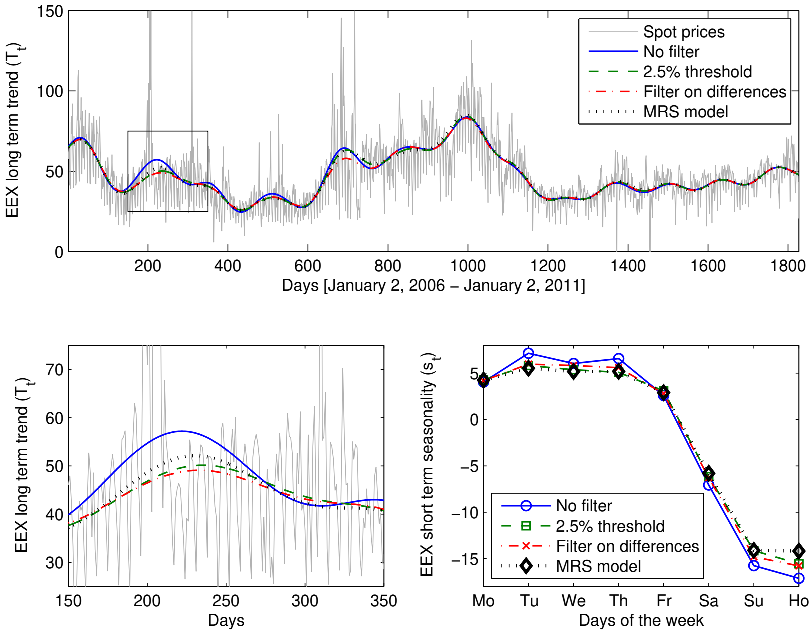

Figure 4: Price observations, estimated long-term trend T_t and short-term seasonality s_t for the EEX market. The *upper panel* shows the price series and the wavelet LTSC. Results are provided for the original price series as well as series cleaned for outliers using the VPT1, RFD and RSC filters. The lower panels provide a zoom into the LTSC for the period June to December 2006 *(lower left)* and the STSC for the same three filters *(lower right)*.

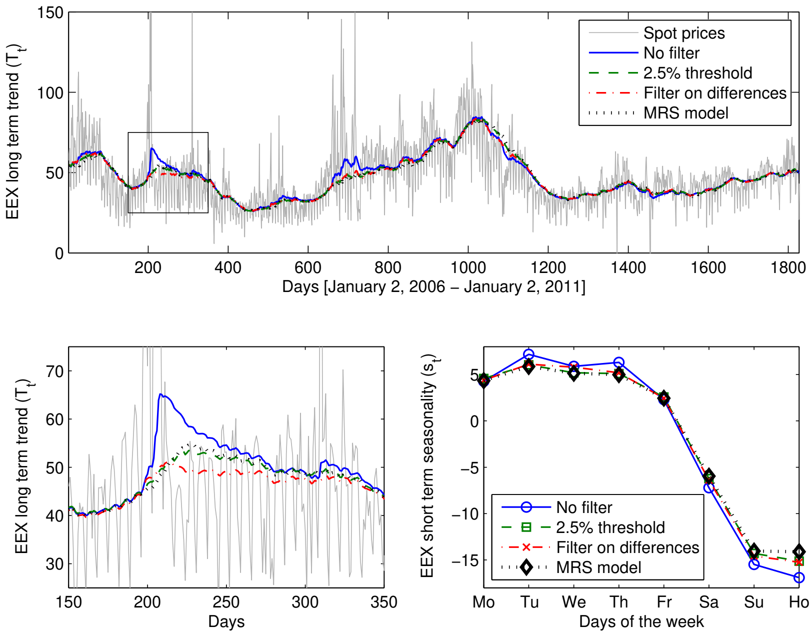

Figure 5: Price observations, estimated long-term trend T_t and short-term seasonality s_t for the EEX market. The *upper panel* shows the price series and the sin-EWMA LTSC. Results are provided for the original price series as well as series cleaned for outliers using the VPT1, RFD and RSC filters. The lower panels provide a zoom into the LTSC for the period June to December 2006 *(lower left)* and the STSC for the same three filters *(lower right)*.

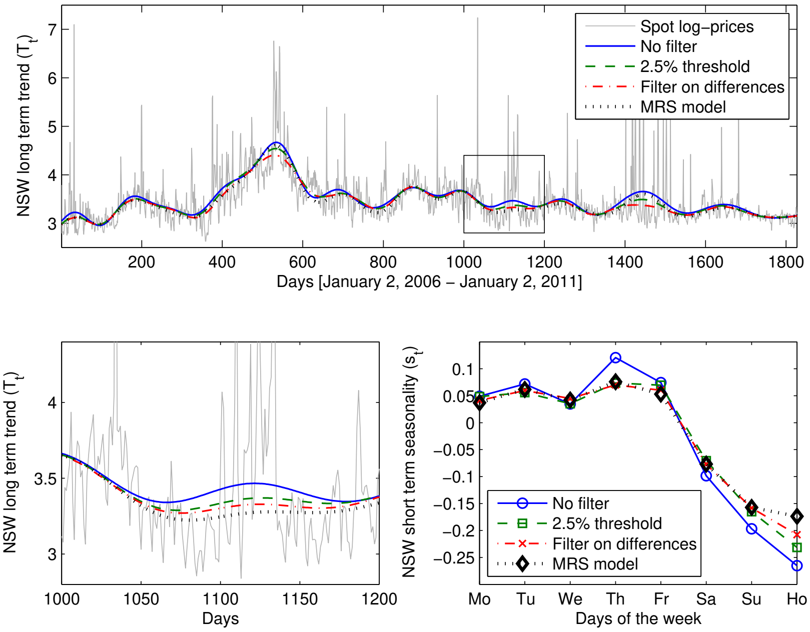

Figure 6: Price observations, estimated long-term trend T_t and short-term seasonality s_t for the NSW market. The *upper panel* shows the log-price series and the wavelet LTSC. Results are provided for the original log-price series as well as series cleaned for outliers using the VPT1, RFD and RSC filters. The lower panels provide a zoom into the LTSC for the period October 2008 to April 2009 *(lower left)* and the STSC for the same three filters *(lower right)*.

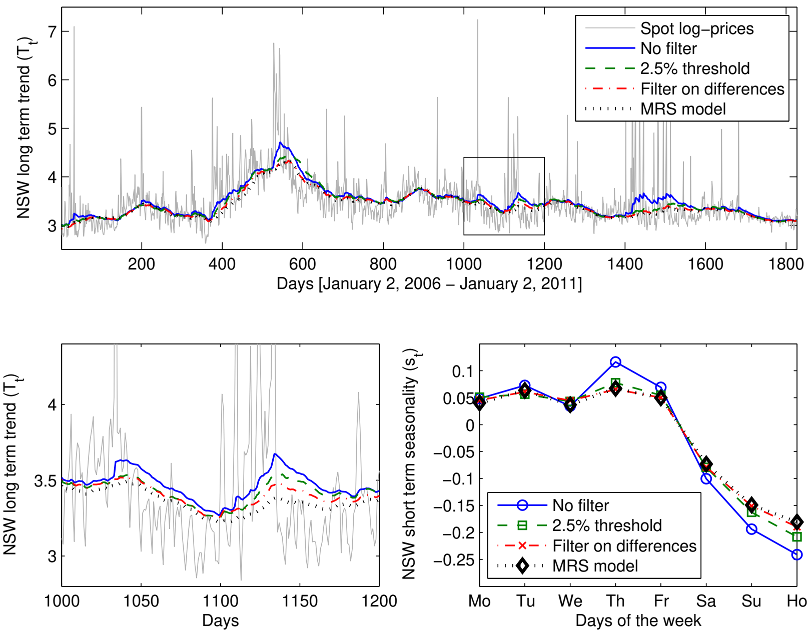

Figure 7: Price observations, estimated long-term trend T_t and short-term seasonality s_t for the NSW market. The *upper panel* shows the log-price series and the sin-EWMA LTSC. Results are provided for the original log-price series as well as series cleaned for outliers using the VPT1, RFD and RSC filters. The lower panels provide a zoom into the LTSC for the period October 2008 to April 2009 *(lower left)* and the STSC for the same three filters *(lower right)*.

EWMA to extreme observations like price spikes. Overall the wavelet approximation provides a much smoother pattern than sin-EWMA. We can also see from the zoom in the lower left panel of Figure 5 that, for the original data, the sin-EWMA estimated LTSC reacts with a certain delay to extreme price observations such that the estimated long-term trend typically increases only after a number of spikes could be observed in the market. On the other hand, using the outlier-cleaned price data, both the wavelet and sin-EWMA seasonal patterns are less volatile since the extreme observations do not enter the estimation procedure. In particular, as illustrated by Figure 5, the sin-EWMA approach yields a much smoother estimate of the LTSC for the VPT1, RFD or RSC filter in comparison to the pattern estimated for the original price data.

Results are quite similar for the NSW market, see Figures 6 and 7. Recall that, for this market, all estimates are based on log-prices instead of the original price series. We also find that for the NSW market, e.g., during the periods July to August 2007, October to November 2008, January to February 2009 or December 2009 to February 2010 the estimated long-term seasonal pattern for the original log-prices deviates quite significantly from the estimated patterns being based on data that have been cleaned for observed outliers. The lower left panels of Figures 6 and 7 zoom in on the period October 2008 to April 2009. Again, the estimated LTSC is most sensitive to extreme price observations when the original log-price data are being used; it reacts less to such observations when the estimation is based on data that have been cleaned for outliers. Like for EEX prices, the wavelet approximation provides a much smoother pattern than sin-EWMA. Also the differences between the individual methods seem to be more pronounced for the sin-EWMA technique.

Let us now consider the estimates for the short-term seasonal components that are reported in the lower right panels of Figures 4 to 7. Again we illustrate results for four of the considered techniques only: ORG, VPT1, RFD and RSC. Recall that the STSC is estimated after detrending, i.e., subtracting the LTSC from the spot or log-spot prices, and is then simply calculated as the 'average week'. Holidays are treated separately like an eight day of the week and are reported as 'Ho' in the figures.

For the EEX market the estimation technique for the long-term trend does not seem to have a strong impact on the estimated short-term pattern. Results for the STSC look quite similar, no matter whether the LTSC was estimated using the wavelet or the sin-EWMA approach. However, we find that there are visible differences between the patterns depending on the outlier detection method. Similar to the estimates of the LTSC, these differences are more substantial when the estimation is based on the original data in comparison to the methods replacing extreme observations beforehand. Applying no filter to the price spikes and drops yields clearly greater estimates for Tuesdays and Thursdays in the STSC while the estimates for Sundays and Holidays are lower. On the other hand, the estimated components look rather similar for the different filters. Among the illustrated methods, the RSC technique seems to provide slightly smaller estimates for the weekdays and slightly greater estimates for Sundays and Holidays in comparison to the VPT1 and RFD filters. This might be a result of the larger number of observations being characterized as outliers and subsequently being replaced by the RSC method. However, the differences in the estimated short-term pattern are much smaller between the individual outlier detection methods in comparison to the calculated pattern for the original data.

Also for the NSW market, the estimation technique for the long-term trend has only limited impact on the estimated short-term seasonal pattern, see Figures 6 and 7. However, the estimated STSC for the original log-prices has a slightly different shape than the component for the data that has been cleaned for outliers. We find that the pattern based on original log-prices yields a larger estimate for Thursdays than the other techniques while it gives smaller estimates in particular for Sundays and Holidays. Again the results are very similar for the VPT1, RFD and RSC methods, suggesting that it matters whether outliers are being replaced before estimation of the STSC, but not so much which filtering techniques are used.

We conclude that replacing extreme observations, such as significant price spikes or price drops, seems to have a significant impact on the estimated seasonal component, as illustrated by Figures 4 to 7. This is true for both the long-term and short-term seasonal patterns. In particular the estimated LTSC seems to be significantly affected by price spikes when the original price or log-price data are used. On the other hand, deviations between the estimated long-term and short-term seasonal pattern for different outlier detection methods are less pronounced. Overall, differences in the estimated seasonal pattern will subsequently lead to different parameter estimates for the stochastic price process. A potential mis-specification of the seasonal component might even introduce an additional bias or artificial price fluctuations that will lead to biased estimates of the parameters for the stochastic price process. However, since the actual seasonal pattern is unobservable, it is hard to argue at this point which method provides the 'best' estimate of the pattern. We will further investigate the impact of the different approaches on the estimation of the seasonal pattern and, subsequently, the parameters of the stochastic process in a simulation study in the next section.

6. Monte Carlo simulation study

In order to investigate the effect of outlier treatment on the estimation of the seasonal and stochastic components of electricity prices, we conduct the following simulation study: we apply different outlier detection methods (filters) to simulated trajectories and compare the resulting estimated seasonal patterns and stochastic model parameters with the true values. As illustrated in the previous section, outliers might have severe impact on the form of the LTSC and STSC estimated from market data, which may then also have impact on the estimated parameters for the stochastic process.

6.1. The simulation setup

In our simulation study we proceed in the following way. First, we estimate the LTSC from the electricity spot prices using either wavelets or the sin-EWMA approach (see Sections 3.1 and 3.2) and then the STSC using the 'average week' method (see Section 3.3). The EEX and NSW price series used for this purpose provide us with two distinct but realistic simulation environments. Next, the upper 2.5% and lower 2.5% of the deseasonalized prices are substituted with the mean of all remaining observations, i.e., we use the VPT1 method; we are motivated to do so by the results of the previous section, where VPT1 has been found to perform reasonably well and robust. The dataset without outliers is then recovered by adding the seasonal components (LTSC and STSC) to the filtered series. Finally, the long-term T_t^* and short-term s_t^* seasonal components are estimated from the dataset without outliers. For our simulation study, T_t^* and s_t^* then represent the 'correct' (or 'actual') long-term and short-term seasonal patterns, respectively, that we would like to retrieve or estimate from the simulated data.

Our hypothesis is that using the discussed methods for outlier detection should lead to better results with respect to specifying the LTSC and STSC correctly, and subsequently for estimating the parameters of the stochastic process. For our simulation study we also require simulated stochastic trajectories; they will be based on the Markov regime switching (MRS) model defined by eqns. (5)-(7) with parameters estimated from the T_t^* and s_t^* deseasonalized EEX and NSW (log-)prices. In total we simulate 1000 trajectories (for each of the two power markets and for each of the two LTSC patterns) and then estimate the seasonal pattern and the parameters of the stochastic process for each trajectory. The following simulation and estimation procedure is conducted:

1. Simulate a sample trajectory of $P_t = X_t + T_t^* + s_t^*$, where X_t is the stochastic component (a sample realization of the MRS model with 'correct' parameters estimated from EEX or NSW market data), T_t^* is the 'correct' LTSC (wavelets or sin-EWMA) and s_t^* is the 'correct' STSC ('average week'. This trajectory can be treated as a stochastic 'copy' of the EEX or NSW spot price series, but with known seasonal patterns and parameters of the stochastic component.
2. Fit a seasonal pattern to P_t using wavelets or sin-EWMA for the LTSC (obtaining T_t^0) and 'average week' for the STSC (obtaining s_t^0). Deseasonalize the simulated trajectory $X_t^0 = P_t - T_t^0 - s_t^0$. This step is equivalent to fitting a seasonal pattern and a stochastic model to raw electricity spot prices.
3. Filter the deseasonalized series X_t^0 using one of the outlier detection methods, yielding the deseasonalized series without outliers X_t^{0F}. Construct an outlier filtered trajectory by adding back the seasonal pattern estimated in step 2: $P_t^F = X_t^{0F} + T_t^0 + s_t^0$.
4. Fit a seasonal pattern to the filtered trajectory P_t^F using wavelets or sin-EWMA for the LTSC (obtaining T_t^F) and 'average week' for the STSC (obtaining s_t^F). Then deseasonalize the original trajectory P_t by removing the seasonal pattern estimated from the filtered series P_t^F: $X_t^F = P_t - T_t^F - s_t^F$. This step corresponds to the approach we are testing in this paper, i.e., fitting a seasonal pattern and a stochastic model to outlier filtered electricity spot prices. Note that we cannot bypass steps 2 and 3 and filter out the outliers from raw price data as the seasonal patterns substantially complicate the identification of price spikes and drops in raw data.
5. Calculate mean squared (MSE) and mean absolute errors (MAE) between the original seasonal components (T_t^*, s_t^*) and seasonal components obtained without filtering (T_t^0, s_t^0) or using one of the outlier detection methods (T_t^F, s_t^F).
6. Estimate the stochastic component parameters from the deseasonalized series $X_t^0 = P_t - T_t^0 - s_t^0$ (see step 2) or the deseasonalized outlier filtered series $X_t^F = P_t - T_t^F - s_t^F$ (see step 4) and calculate the MSE and MAE for these parameters by comparing with the 'correct' parameters of X_t.

Overall, for each trajectory, we either use the original price series, or apply any of the suggested outlier detection methods to determine the filtered price series, before estimating the seasonal pattern. Therefore, in our simulation study, in total we compare the performance of eight techniques:

(1) ORG, (2) FPT, (3) VPT1, (4) VPT2, (5) RFP, (6) RFD, (7) RM and (8) RSC; see Section 5.2 for filter definitions.

6.2. Seasonal patterns for outlier filtered trajectories

We first examine the results on estimation of the seasonal pattern for the simulation runs. Table 2 provides the results for the mean squared error (MSE), comparing the actual and estimated long-term (T_t), short-term (s_t) and overall ($f_t = T_t + s_t$) seasonal patterns based on the different outlier detection methods. Note that the results for the two markets are provided for both the wavelet and the sin-EWMA long-term seasonal components. Next to using the MSE as loss function, we also calculated the mean absolute error (MAE). However, since MSE and MAE yielded very similar results, these results are not reported here. They are, however, available upon request to the authors.

We also performed a Kruskal-Wallis test to examine significant differences between the samples of MSE for the considered outlier detection methods. The Kruskal-Wallis test is a non-parametric version of the classical one-way analysis of variance (ANOVA), and tests the null hypothesis that all samples are drawn from the same population, or equivalently, from different populations with the same distribution, see, e.g., Hollander and Wolfe (1999). Unlike a standard one-way ANOVA, the test does not require the assumption that all samples come from a population with a normal distribution. Examining the distribution of the simulation-based MSE and MAE for the deviation from the actual seasonal pattern, the assumption of normality for the population would not be justified. Therefore, a Kruskal-Wallis test is more appropriate to test for significant differences between the populations. As mentioned above the null hypothesis for such a test is that all samples are drawn from the same population. Therefore, rejecting the null only provides statistical evidence for at least one of the samples being from a population with a different distribution. However, the test does not provide detailed information on which of the samples are significantly different. A test that can provide such information is called a multiple comparison procedure (Hochberg and Tamhane, 1987). The test uses Tukey's honestly significant difference criterion, that is optimal for the comparison of groups with equal sample sizes, to test for significant differences with respect to the performance of the methods. The test is conducted with a significance level of $\alpha = 0.05$. For each market and deseasonalization method, Table 2 indicates for all of the considered filtering approaches (1)-(8) which of the other methods perform significantly worse or significantly better for the considered criterion. Note that when all other methods were significantly better (or worse) than a particular approach this is indicated by 'All', while '–' indicates that none of the other models provided significantly better (or worse) results. Being based on a nonparametric test statistic, it can be expected that the multiple comparison procedure will not be able to distinguish significant differences between all of the considered methods. However, generally most of the considered methods for outlier detection perform significantly better than the 'no filter' technique (i.e., ORG) with respect to estimating the seasonal pattern.

From a first glance at Table 2, we find that for most of the considered criteria, using the original price data yields the worst results with respect to estimating the seasonal pattern. In comparison to the 'no filter' approach, most of the considered outlier detection methods lead to a significant reduction in the MSE for the estimated long-term, short-term and overall seasonal pattern. It also becomes clear that results for the LTSC T_t usually dominate those for the STSC s_t in terms of

Table 2: Mean Squared Error (MSE) for estimation of long-term (T_t), short-term (s_t) and overall (f_t) seasonal pattern for the conducted simulation study of the EEX and NSW markets. The LTSC T_t is modeled by wavelets (*two upper panels*) or sin-EWMA (*two lower panels*), while the stochastic price behavior is simulated using the MRS model defined by eqns. (5)-(7). The best result for each of the considered measures is indicated in bold. The table also provides results for the multiple comparison procedure using Tukey's honestly significant difference criterion (Hochberg and Tamhane, 1987), indicating for each of the approaches (1)-(8) which of the other methods performs significantly worse / significantly better for the considered criterion.

			EEX (Wavelet)			
Measure	MSE (T_t)	Worse / Better	MSE (s_t)	Worse / Better	MSE (f_t)	Worse / Better
(1) ORG	7.7268	– / All	0.4297	– / All	8.1589	– / All
(2) FPT	5.7674	{1,4} / {3,5,6,7,8}	0.2396	{1,4} / {3,5,6,7,8}	5.9804	{1,4} / {3,5,6,7,8}
(3) VPT1	**4.6753**	{1,2,4,6} / –	0.1989	{1,2,4} / {8}	**4.8485**	{1,2,4,6} / –
(4) VPT2	6.3848	{1} / {2,3,5,6,7,8}	0.2889	{1} / {2,3,5,6,7,8}	6.6337	{1} / {2,3,5,6,7,8}
(5) RFP	5.6470	{1,2,4,6} / –	0.2082	{1,2,4} / {8}	5.8160	{1,2,4,6} / –
(6) RFD	5.2237	{1,2,4} / {3,5,7,8}	0.1985	{1,2,4} / {8}	5.3864	{1,2,4} / {3,5,7,8}
(7) RM	4.8250	{1,2,4,6} / –	0.2020	{1,2,4} / {8}	4.9899	{1,2,4,6} / –
(8) RSC	4.9798	{1,2,4,6} / –	**0.1778**	{1,2,3,4,5,6,7} / –	5.1229	{1,2,4,6} / –
			NSW (Wavelet)			
Measure	MSE (T_t)	Worse / Better	MSE (s_t)	Worse / Better	MSE (f_t)	Worse / Better
(1) ORG	0.0106	– / {3,4,5,6,7,8}	0.0066	– / {3,4,5,6,7,8}	0.0113	– / {3,4,5,6,7,8}
(2) FPT	0.0105	– / {3,4,5,6,7,8}	0.0060	– / {3,4,5,6,7,8}	0.0112	– / {3,4,5,6,7,8}
(3) VPT1	0.0065	{1,2,4,8} / {5,6,7}	0.0025	{1,2,4} / {5,6,7,8}	0.0068	{1,2,4,8} / {5,6,7}
(4) VPT2	0.0074	{1,2} / {3,5,6,7}	0.0033	{1,2} / {3,5,6,7,8}	0.0077	{1,2} / {3,5,6,7}
(5) RFP	0.0055	{1,2,3,4,8} / –	**0.0016**	{1,2,3,4,7,8} / –	0.0056	{1,2,3,4,8} / –
(6) RFD	0.0055	{1,2,3,4,8} / –	0.0016	{1,2,3,4,8} / –	0.0056	{1,2,3,4,8} / –
(7) RM	**0.0054**	{1,2,3,4,8} / –	0.0017	{1,2,3,4} / {5}	**0.0055**	{1,2,3,4,8} / –
(8) RSC	0.0078	{1,2} / {3,5,6,7}	0.0023	{1,2,3,4} / {5,6}	0.0080	{1,2} / {3,5,6,7}
			EEX (sin-EWMA)			
Measure	MSE (T_t)	Worse / Better	MSE (s_t)	Worse / Better	MSE (f_t)	Worse / Better
(1) ORG	26.8461	{4,8} / {2,3,5,6,7}	0.5292	– / All	27.3719	{4,8} / {2,3,5,6,7}
(2) FPT	23.2980	{1,3,4,8} / –	0.2511	{1,4} / {5,6,8}	23.5400	{1,3,4,8} / –
(3) VPT1	25.1094	{1,4,8} / {2,5,6,7}	0.2456	{1,4} / {6,8}	25.3448	{1,4,8} / {2,5,6,7}
(4) VPT2	37.0827	– / All	0.3923	{1} / {2,3,5,6,7,8}	37.4618	– / All
(5) RFP	23.4552	{1,3,4,8} / –	0.2294	{1,2,4} / –	23.6769	{1,3,4,8} / –
(6) RFD	23.0855	{1,3,4,8} / –	0.2193	{1,2,3,4} / –	23.2968	{1,3,4,8} / –
(7) RM	**22.8728**	{1,3,4,8} / –	0.2314	{1,4} / –	**23.0968**	{1,3,4,8} / –
(8) RSC	29.8795	{4} / {1,2,3,5,6,7}	**0.2150**	{1,2,3,4} / –	30.0838	{4} / {1,2,3,5,6,7}
			NSW (sin-EWMA)			
Measure	MSE (T_t)	Worse / Better	MSE (s_t)	Worse / Better	MSE (f_t)	Worse / Better
(1) ORG	0.0241	– {3,5,6,7,8}	0.0091	– / {3,4,5,6,7,8}	0.0250	– / {3,5,6,7,8}
(2) FPT	0.0242	– / {3,5,6,7,8}	0.0090	– / {3,4,5,6,7,8}	0.0251	– / {3,5,6,7,8}
(3) VPT1	0.0179	{1,2,4,8} / {5,6,7}	0.0029	{1,2,4} / {5,6,7,8}	0.0182	{1,2,4,8} / {5,6,7}
(4) VPT2	0.0235	– / {3,5,6,7,8}	0.0051	{1,2} / {3,5,6,7,8}	0.0239	– / {3,5,6,7,8}
(5) RFP	0.0167	{1,2,3,4,8} / –	0.0028	{1,2,3,4,7,8} / –	0.0170	{1,2,3,4,8} / –
(6) RFD	0.0167	{1,2,3,4,8} / –	**0.0016**	{1,2,3,4,7,8} / –	0.0169	{1,2,3,4,8} / –
(7) RM	**0.0164**	{1,2,3,4,8} / –	0.0019	{1,2,3,4}/ {5,6}	**0.0166**	{1,2,3,4,8} / –
(8) RSC	0.0212	{1,2,4} / {3,5,6,7}	0.0036	{1,2,3,4} / {5,6}	0.0216	{1,2,4} / {3,5,6,7}

contribution to the MSE for the overall seasonal pattern. Therefore, the method that performs best with respect to T_t, in most cases also yields the best results for the total seasonal pattern. In the following we report the results in more detail.

Let us now consider the results for the EEX market when a wavelet LTSC is used, see the top panel of Table 2. We find that for both the long-term T_t, short-term s_t and overall seasonal pattern the ORG or 'no filter' method performs the worst. The variable price threshold with 10% highest and 10% lowest prices treated as outliers (VPT2) does not seem to perform that well, but still reduces MSE for the estimation of the total seasonal pattern by more than 18%. For all of the other outlier detection methods the MSE is reduced by more than 25% up to 40% in comparison to the 'no filter' approach. Since T_t dominates s_t with respect to contribution to the error in the estimation of the total seasonal pattern, the VPT1 method that performs best with respect to T_t also yields the overall best results. The Kruskal-Wallis test rejects the hypothesis that the samples are drawn from the same population at all levels of significance ($p = 0.000$). With respect to multiple comparison of the groups, we find that all methods yield significantly smaller absolute errors for the long term, short-term and overall seasonal pattern than the 'no filter' approach. For the estimation of the long-term pattern T_t, the VPT1 method yields the best results, significantly outperforming four of the other methods according to Tukey's honestly significant difference criterion. The recursive filter on prices (RFP) and regime switching model classification (RSC) for detection of the outliers are only slightly worse and also outperform four of the other methods. Note that the test cannot detect significant differences between VPT1, RFP, RM and RSC with respect to the performance in estimating T_t. For the short-term pattern s_t, we find that the RSC method significantly outperforms all the other methods, while the test finds no significant differences between the VPT1, RFP, RFD, RM technique: they all outperform three of the other methods (ORG, FPT, VPT2), but are significantly worse than the RSC technique. Note that for s_t the 'no filter' method performs significantly worse than all methods using data that have been cleaned for outliers. Finally, for the total seasonal pattern, the VPT1, RFP, RM and RSC methods yield the best results, each of them significantly outperforming four of the other techniques. The test does not find significant differences between the performance of these methods in terms of estimating the overall pattern. Again, the ORG method is significantly outperformed by all of the other techniques.

Results for the NSW market with a wavelet LTSC also lead to similar conclusions with respect to the results for the ORG method, see the upper-middle panel of Table 2. Again, the 'no filter' approach yields the greatest value for the MSE for the estimation of T_t, s_t and the total seasonal pattern. Only the results for the fixed price threshold method (FPT) are of the same magnitude, while all other outlier detection methods yield much smaller values for the MSE. For the estimation of the overall seasonal pattern, the best results are obtained for the RFP, RFD and RM technique where the MSE is reduced by more than 50% in comparison to the 'no filter' method. Also the other outlier detection methods, except FPT, reduce the MSE by roughly between 30% and 40%. Again the results for the long-term component T_t dominate s_t in terms of contribution to the MSE for the overall seasonal pattern. Therefore, the methods yielding the best results for T_t also perform best for the total seasonal pattern. For estimation of s_t, ORG and FPT lead to the worst results with the largest MSE. On the other hand, the outlier detection methods RFP, RFD and RM yield the best results: these methods significantly outperform five of the other methods, while MSE in comparison to the ORG method is reduced by approximately 75%. Note that the test for multi comparison cannot significantly distinguish between the performance of RFP, RFD and RM for estimation of the short-term pattern. Overall, the recursive model estimation yields the best results for the seasonal pattern in the NSW market when a wavelet LTSC is used.

Results are not that clear-cut for the sin-EWMA long-term seasonal pattern, see the two bottom panels of Table 2. Generally the 'no filter' method is still outperformed by most of the outlier detection methods. The Kruskall-Wallis test rejects for both markets that the MSE samples are drawn from the same population at any significance level ($p = 0.000$). However, for the EEX market the variable price threshold with 10% largest and 10% smallest prices treated as outliers (VPT2) and MRS model classification (RSC) perform worse for T_t and the overall seasonal pattern estimation. Recall from Section 5 that these are the methods that generally classify the largest percentage of observations as outliers and replace them before estimation of the seasonal pattern. So in this case, it seems like removing a substantial fraction of the extreme observations does not provide an accurate estimation of the long-term seasonal sin-EWMA component. Generally, the estimation results for the sin-EWMA pattern are clearly worse for all of the considered outlier filters. In comparison to using a wavelet LTSC, the MSE for the sin-EWMA simulation study is approximately four times larger in magnitude. Note further that the reduction in MSE for the total seasonal pattern for the four best methods, FPT, RFP, RFD and RM, in comparison to the 'no filter' technique is only around 14%. Therefore, it seems as though the estimation technique for the long-term trend impacts on the overall quality of estimation of the seasonal pattern, but also on the ranking of the methods. Still for estimation of the short-term pattern s_t, ORG is significantly outperformed by all outlier detection techniques. Overall, for the EEX market the ORG method ranks sixth out of the eight groups and is significantly outperformed by five of the other techniques.

For the NSW market, using the original data for estimation of the seasonal pattern likewise yields the worst results together with the fixed price threshold (FPT) method. Five of the other techniques significantly outperform the ORG method for T_t and the total seasonal pattern. For the three best methods, RFP, RFD and RM, the MSE is reduced by more than 30%. Only the FPT and VPT2 method are not able to significantly outperform the 'no filter' approach. Similarly to the EEX market, the errors for estimation of the sin-EWMA long-term seasonal pattern are clearly larger than for the wavelet LTSC. For the NSW market they increase on average by a magnitude of 100%. Again, the estimation technique for the long-term trend dominates the overall quality of estimating the seasonal pattern and the ranking of the methods.

Overall, applying a filtering procedure for the outliers before estimating the seasonal pattern seems to provide a more appropriate specification of the long term, short-term and overall seasonal patterns. This is true for the considered EEX and NSW simulation environments and also for both of the applied estimation techniques for the long-term seasonal pattern (wavelet and sin-EWMA). The ORG method usually ranks among the three worst methods and is significantly outperformed in 23 out of 28 possible cases with respect to MSE for the estimated overall seasonal pattern.

6.3. Parameter estimates for the stochastic component of the outlier filtered trajectories

In this section we examine the impact of the misspecification of the seasonal pattern on parameter estimates for the stochastic process. Table 3 provides results on the parameter estimates for the base regime, while Table 4 summarizes the results for the parameter estimates of the spike

Table 3: Mean Squared Error (MSE) for estimation of base regime parameters for the MRS model defined by eqns. (5)-(7) for the conducted simulation study of the EEX market. The LTSC T_t is modeled using wavelets (*upper panel*) or sin-EWMA (*lower panel*), while the stochastic price behavior is simulated using the MRS model. The best result for each of the considered measures is indicated in bold. The table also provides results for the multiple comparison procedure using Tukey's honestly significant difference criterion (Hochberg and Tamhane, 1987), indicating for each of the approaches (1)-(8) which of the other methods perform significantly worse / significantly better for the considered criterion.

EEX (Wavelet, Base Regime)						
Measure	MSE (α_1)	Worse / Better	MSE (α_1/β_1)	Worse / Better	MSE (σ_1)	Worse / Better
(1) ORG	0.0616	– / All	0.2530	– / {2,3,4,5,7,8}	4.5006	– / All
(2) FPT	0.0357	{1,6} / {4,8}	0.1441	{1,6} / {4}	2.7430	{1} / {3}
(3) VPT1	0.0291	{1,6,7} / –	0.1087	{1,6,7} / –	**2.1131**	{1,2} / –
(4) VPT2	0.0236	{1,2,5,6,7} / –	0.0917	{1,2,5,6,7,8} / –	3.2302	{1} / –
(5) RFP	0.0311	{1,6,7} / {4}	0.1268	{1,6,7} / {4}	2.5537	{1} / –
(6) RFD	0.0450	{1} / {2,3,4,5,8}	0.1847	– / {2,3,4,5,7,8}	2.4833	{1} / –
(7) RM	0.0368	{1} / {3,4,5,8}	0.1498	{1,6} / {3,4,5,8}	3.1372	{1} / –
(8) RSC	**0.0185**	{1,2,6,7} / –	**0.0822**	{1,6,7} / {4}	2.3781	{1} / –
EEX (sin-EWMA, Base Regime)						
Measure	MSE (α_1)	Worse / Better	MSE (α_1/β_1)	Worse / Better	MSE (σ_1)	Worse / Better
(1) ORG	0.0505	– / All	0.6192	– / All	19.8489	{4} / {2,3,5,6,7,8}
(2) FPT	0.0195	{1} / {8}	0.2146	{1,4} / {8}	14.4930	{1,4} / –
(3) VPT1	0.0201	{1} / {8}	0.2380	{1} / {5,8}	14.7368	{1,4} / –
(4) VPT2	0.0190	{1} / {8}	0.2943	{1} / {2,5,6,7,8}	36.0074	– / All
(5) RFP	0.0165	{1} / {8}	0.1752	{1,3,4} / {8}	**14.2021**	{1,4,8} / –
(6) RFD	0.0192	{1} / {8}	0.2018	{1,4} / {8}	12.8745	{1,4,8} / –
(7) RM	0.0189	{1} / {8}	0.2028	{1,4} / {8}	13.6774	{1,4,8} / –
(8) RSC	**0.0072**	All / –	**0.0854**	All / –	15.0862	{1} / {5,6,7}

and drop regime for the EEX market; results for the NSW market lead to similar conclusions and are reported in the appendix only. Recall that these results are based on data that contain extreme price observations such as price spikes and drops. After estimation of the seasonal component, the pattern is subtracted from the spot (log-)price and the three-state regime switching model is fitted to the remaining stochastic component of the data.

Let us first consider the results for the base regime in Table 3 when either a wavelet or sin-EWMA long-term seasonal component is used. The table reports the MSE of parameter estimates for α_1, the long-term mean price level α_1/β_1 and the instantaneous volatility σ_1 in equations (4) and (5). We find that when a wavelet LTSC is used, for all parameters the 'no filter' method usually yields the worst results. For the parameter α_1 all techniques yield significantly better results for the MSE in comparison to the ORG method. The MSE is usually reduced by more than 40% for six of the seven considered outlier detection methods. Also for the long-term mean price level α_1/β_1, six out of seven filtering techniques significantly outperform the 'no filter' method. Interestingly, for these two parameters the best results are obtained for the VPT2 and RSC methods that usually classify the greatest fraction of observations as outliers. Also for the volatility parameter σ_1 of the base regime all filters yield significantly better results than the ORG method. Note that for σ_1 the test for multiple comparison of the groups does not indicate significant differences between any of

Table 4: Mean Squared Error (MSE) for estimation of the spike and drop regime parameters for the MRS model defined by eqns. (5)-(7) for the conducted simulation study of the EEX market. The LTSC T_t is modeled using wavelets (*upper panel*) or sin-EWMA (*lower panel*), while the stochastic price behavior is simulated using the MRS model. The best result for each of the considered measures is indicated in bold. The table also provides results for the multiple comparison procedure using Tukey's honestly significant difference criterion (Hochberg and Tamhane, 1987), indicating for each of the filtering approaches (1)-(8) which of the other methods perform significantly worse / significantly better for the considered criterion.

	EEX (Wavelet, Spike and Drop Regime)							
	MSE(μ_2)	Worse / Better	MSE(σ_2)	Worse / Better	MSE(μ_3)	Worse / Better	MSE(σ_3)	Worse / Better
(1)	0.0463	– / {8}	0.0417	{7} / –	0.0872	– / {3,4,5,6,7,8}	0.0306	– / {3,4,5,8}
(2)	0.0432	– / {8}	0.0481	– / {8}	0.0750	– / {3,4,5,6,7,8}	0.0283	– / {3,4,5,8}
(3)	0.0461	– / {8}	0.0482	– / {4,8}	0.0474	{1,2,5,6} / {8}	0.0192	{1,2,6,7} / {8}
(4)	0.0427	– / {8}	0.0376	{3,7} / –	0.0551	{1,2,6} / {8}	0.0197	{1,2,6,7} / {8}
(5)	0.0432	– / {8}	0.0452	– / {8}	0.0592	{1,2,6} / {3,8}	0.0264	{1,2} / {8}
(6)	0.0438	– / {8}	0.0469	– / {8}	0.0679	{1} / {3,4,5,7,8}	0.0256	– / {3,4,8}
(7)	0.0446	– / {8}	0.0502	– / {1,4,8}	0.0538	{1,2,6} / {8}	0.0240	– / {3,4,8}
(8)	**0.0331**	All / –	**0.0341**	{2,3,5,6,7} / –	**0.0372**	All / –	**0.0134**	All / –
	EEX (sin-EWMA, Spike and Drop Regime)							
	MSE(μ_2)	Worse / Better	MSE(σ_2)	Worse / Better	MSE(μ_3)	Worse / Better	MSE(σ_3)	Worse / Better
(1)	0.0487	– / {3,4,6,8}	0.0478	– / –	0.0501	{3,4,8} / –	0.0475	{4,8} / –
(2)	0.0462	– / {4,6,8}	0.0530	– / –	0.0508	{3,4,8} / –	0.0479	{4,8} / –
(3)	0.0394	{1} / –	0.0505	– / –	0.0668	{4,8} / {1,2,6,7}	0.0478	{4,8} / –
(4)	0.0357	{1,2,5,7} / –	0.0506	– / –	0.1193	– / {1,2,3,5,6,7}	0.0602	– / {1,2,3,5,6,7}
(5)	0.0417	– / {4,8}	0.0503	– / –	0.0555	{4,8} / –	0.0469	{4,8} / –
(6)	0.0387	{1,2} / –	0.0469	– / –	0.0503	{3,4,8} / –	**0.0396**	{4,8} / –
(7)	0.0422	– / {4,8}	0.0574	– / –	**0.0466**	{3,4,8} / –	0.0450	{4,8} / –
(8)	**0.0326**	{1,2,5,7} / –	**0.0456**	– / –	0.1012	– / {1,2,3,5,6,7}	0.0561	– / {1,2,3,5,6,7}

the outlier detection methods VPT1, VPT2, RFP, RFD, RM and RSC at the 5% level.

Results for the sin-EWMA technique are illustrated in the lower panel of Table 3. Similar to the results for the wavelet LTSC, we find that the 'no filter' approach yields the worst results with respect to the parameter estimates of the base regime. For the parameter α_1 all techniques yield significantly better results for the MSE in comparison to the ORG method. The MSE is reduced by at least 60% for any of the considered outlier filters. Also for the long-term mean price level α_1/β_1, the MSE is reduced by more than 50% when the seasonal pattern is estimated based on the filtered data prior to fitting the stochastic model. Interestingly, for both parameters by far the best results are obtained when outlier filtering is conducted using the RSC filter. This method outperforms all other outlier detection techniques and reduces the MSE by more than 85% in comparison to using original data. For the base regime volatility σ_1, five out of seven filtering techniques significantly outperform the ORG method: the MSE is reduced by approximately 25% to 30% for these methods. Note that in comparison to the wavelet LTSC, estimation errors for the long-term mean price level α_1/β_1 and the volatility parameter σ_1 are much larger. In particular for the latter the MSE is usually about five times larger. Recall that the wavelet LTSC also provided smaller errors for the estimation of the seasonal pattern. Therefore our results for σ_1 indicate that a bad estimator of the seasonal pattern will also have significant impacts on the estimation of the stochastic parameters and bears a much higher risk of misspecifying the stochastic process.

Results for parameter estimates of the other two regimes, the spike and drop regime, are reported in Table 4. We find that for the parameters of the second and third regime results are not so clear-cut. When a wavelet LTSC is used (upper panel), the RSC filtering technique for the outliers provides the best results and usually outperforms all other methods. For μ_2, the location parameter of the shifted log-normal distribution for the spike regime, the RSC filter yields significantly lower estimation errors than all other techniques, while the multiple comparison procedure cannot identify significant differences between any of the other methods. Also for the scale parameter σ_2 of the spike regime, as well as μ_3 and σ_3, the location and scale parameter of the drop regime, the RSC filter gives the best results. Note that for μ_3 and σ_3, the 'no filter' approach yields the largest MSE and is significantly outperformed by six, respectively, four of the filtering techniques in terms of this error measure. However, for σ_2 none of the filtering techniques provides significantly smaller MSE in comparison to using the original observations for the estimation of the seasonal pattern.

Results for parameter estimates of the second and third regime in the sin-EWMA simulation scenario are provided in the lower panel of Table 4. We find that for the parameters of the spike regime, μ_2 and σ_2, again the RSC filter leads to the smallest MSE. For μ_2, the 'no filter' approach yields the largest MSE and is significantly outperformed by four of the filtering techniques. On the other hand, for σ_2, the multiple comparison test cannot detect significant differences between any of the applied methods. For the parameters of the drop regime, μ_3 and σ_3, in particular VPT2 and RSC perform poorly and lead to greater MSEs in the estimation than most of the other methods. Interestingly, for these parameters, the ORG does not give significantly greater estimation errors than any of the filtering techniques.

Overall, we find that the estimation of the seasonal pattern based on prices that have been filtered for extreme price spikes and drops seems to not only provide a better specification of the seasonal pattern, but, subsequently, also better estimates for parameters of the stochastic process. This is true in particular for the parameters of the base regime, where the 'no filter' approach performs significantly worse for 40 out of 42 possible cases. These results are valid for both the wavelet and the sin-EWMA long-term seasonal patterns. On the other hand, we find that for parameter estimates of the second and third regime – modeling price spikes and price drops – it is more difficult to distinguish between the tested methods.

Finally, Table 5 provides a summary of the overall performance of the considered outlier detection methods. Each method is assigned a number (a 'rank') which is computed as one plus the number of methods performing significantly better in terms of MSE and according to the multiple comparison procedure (Hochberg and Tamhane, 1987) across all considered markets (EEX, NSW) and LTSC patterns (wavelet-based, sin-EWMA). We then calculate two geometric averages – one for the 'ranks' of the seasonal patterns (T_t, s_t) and one for the 'ranks' of the parameter estimates of the stochastic model ($\alpha_1, \alpha_1/\beta_1, \sigma_1, \mu_2, \sigma_2, \mu_3, \sigma_3$). We find that for seasonal pattern estimation the 'no filter' approach is performing significantly worse than the other methods, being ranked on average at 7.08 out of 8 considered cases. On the other hand, the ranking illustrates that in particular the recursive filter (RFP and RFD) and the recursive model (RM) approach yield the overall best results for seasonal pattern estimation. For estimation of the parameters of the stochastic process, simulation results suggest that the 'no filter' method is outperformed on average by more competitors (ranked at 3.27) than any of the other techniques. On the other hand, for parameter estimation it is clearly the Markov regime-switching model classification (RSC) of outliers that yields the best results. Note, however, that for most estimated parameters, it is more difficult to

Table 5: Summary of overall performance for the considered outlier detection methods based on MSE results for both markets and LTSC estimation techniques. The best results are emphasized in bold, the worst are underlined. See text for details on the derivation of these rankings.

	Seasonal pattern	Stochastic component
Measure	Average rank	
(1) ORG	<u>7.08</u>	<u>3.27</u>
(2) FPT	4.83	2.51
(3) VPT1	3.24	1.70
(4) VPT2	6.44	1.75
(5) RFP	**1.09**	1.98
(6) RFD	1.33	2.39
(7) RM	1.36	2.04
(8) RSC	2.51	**1.63**

significantly distinguish between the performance of the applied methods.

7. Conclusions

An important issue in the estimation of stochastic models for electricity spot prices is the estimation of a component to deal with trends and seasonality in the data. Unfortunately, estimation routines for the seasonal long-term and short patterns are very sensitive to the abrupt and generally unanticipated extreme changes in electricity spot prices known as spikes (or jumps). In this study we investigate the impact on model estimation for different treatment of extreme observations on determining the number of spikes and descriptive statistics of the adjusted series. Despite their rarity, price spikes are the very motive for designing insurance protection against electricity price movements, i.e., the buying or selling of derivative contracts such as options and futures. The existence of price spikes is also the reason for including discontinuous components in realistic models of electricity price dynamics.

We have shown that improved robustness of the electricity spot price model can be achieved by filtering the data with some reasonable procedure for outlier detection, and then using classical estimation and testing techniques for the seasonal pattern on the filtered data. Using data from the European Energy Exchange (EEX) and New South Wales (NSW) electricity market, we examine the effects of different filtering techniques on the estimation of the seasonal pattern and, subsequently, the parameters of the stochastic process. We find that using filtered data that have been cleaned for outliers instead of the original spot price yields quite different estimates of the seasonal long-term and short-term pattern, in particular for periods where actual price spikes could be observed in the markets. In a Monte Carlo simulation study we then show that estimates of the seasonal pattern being based on filtered data are significantly better in comparison to estimates based on the original price data. While we cannot identify a single best method for outlier detection, for the considered EEX and NSW markets the best results with respect to estimating the seasonal pattern are obtained when either a recursive filter on prices (RFP) or price differences (RFD) or a recursive seasonal model (RM) estimation is applied. More importantly, usually all of the considered approaches significantly outperform the 'no filter' approach that uses the original spot price. We further illustrate that a misspecification of the seasonal pattern generally also leads to larger errors for parameter estimates of the stochastic process, and therefore, to a potentially misspecified model. Also here the 'no filter' approach ranks last among all considered methods while a Markov regime-switching model classification (RSC) and variable price thresholds (VPT1, VPT2) tend to provide the best results.

Our findings point out the substantial impact the treatment of extreme observations may have on the estimation of the seasonal pattern and stochastic model. Based on our results we recommend that before the estimation of the deterministic seasonal pattern a filtering procedure for detection and replacement of extreme observations in electricity spot prices should be implemented. Such an approach will yield more appropriate estimates for both seasonal pattern and the parameters of the stochastic process.

Acknowledgements

This paper has benefited from conversations with the participants of the Trondheim Summer 2011 Energy Workshop, the 2011 WPI Conference in Energy Finance, the Energy Finance Christmas Workshop (EFC11) in Wrocław and the seminars at Macquarie University, University of Verona and Wrocław University of Technology. This work was supported by funds from the Australian Research Council through grant no. DP1096326 and the National Science Centre (NCN, Poland) through grant no. 2011/01/B/HS4/01077.

References

Barlow, M. (2002) A diffusion model for electricity prices. Mathematical Finance 12, 287-298.

Becker, R., Hurn, S., Pavlov, V. (2007) Modelling Spikes in Electricity Prices. The Economic Record 83(263), 371-382.

Benth, F.E., Benth, J.S., Koekebakker, S. (2008) Stochastic Modeling of Electricity and Related Markets. World Scientific, Singapore.

Bhanot, K. (2000) Behavior of power prices: Implications for the valuation and hedging of financial contracts. The Journal of Risk 2, 43-62.

Bierbrauer, M., Menn, C., Rachev, S.T., Trück, S. (2007) Spot and derivative pricing in the EEX power market. Journal of Banking and Finance 31, 3462-3485.

Bierbrauer, M., Trück, S., Weron, R. (2004) Modeling electricity prices with regime switching models. Lecture Notes in Computer Science 3039, 859-867.

Boogert, A., Dupont, D. (2008) When supply meets demand: The case of hourly spot electricity prices. IEEE Transactions on Power Systems 23(2), 389-398.

Borovkova, S., Permana, F.J. (2006) Modelling electricity prices by the potential jump-diffusion. In: A.N. Shiryaev et al. (eds.), Stochastic Finance – Proceedings of StochFin2004, Springer, 239-264.

Cappe, O., Moulines E., Ryden T. (2005). Inference in Hidden Markov Models. Springer.

Cartea, A., Figueroa, M. (2005) Pricing in electricity markets: A mean reverting jump diffusion model with seasonality. Applied Mathematical Finance 12(4), 313-335.

Clewlow, L., Strickland, C. (2000). Energy Derivatives – Pricing and Risk Management. Lacima Publications.

De Jong, C. (2006) The nature of power spikes: A regime-switch approach. Studies in Nonlinear Dynamics & Econometrics 10(3), Article 3.

Deng, S.-J. (1998) Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes. PSerc Working Paper 98-28.

Erlwein, C., Benth, F.E., Mamon, R. (2010) HMM filtering and parameter estimation of an electricity spot price model. Energy Economics 32, 1034-1043.

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

Eydeland, A., Wolyniec, K. (2012) Energy and Power Risk Management (2nd ed.). Wiley, Hoboken, NJ.

Fanone, E., Gamba, A., Prokopczuk, M. (2012) The case of negative day-ahead electricity prices. Energy Economics, In press, doi:10.1016/j.eneco.2011.12.006.

Fleten, S.-E., Heggedal, A.M., Siddiqui, A. (2011) Transmission capacity between Norway and Germany: a real options analysis. Journal of Energy Markets 4(1), 121-147.

Garcia, R.C., Contreras, J., van Akkeren, M., Garcia, J.B. (2005) A GARCH forecasting model to predict day-ahead electricity prices. IEEE Transactions on Power Systems 20(2), 867-874.

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

Haldrup, N., Nielsen, F.S., Nielsen, M.Ø. (2010) A vector autoregressive model for electricity prices subject to long memory and regime switching. Energy Economics 32, 1044-1058.

Hambly, B., Howison, S., Kluge, T. (2009) Modelling spikes and pricing swing options in electricity markets. Quantitative Finance 9(8), 937-949.

Hamilton, J. (1990) Analysis of time series subject to changes in regime. Journal of Econometrics 45, 39-70.

Härdle, W., Kerkyacharian, G., Picard, D., Tsybakov, A. (1998) Wavelets, Approximation and Statistical Applications. Lecture Notes in Statistics 129. Springer-Verlag, New York.

Higgs, H., Worthington, A. (2008) Stochastic price modeling of high volatility, mean-reverting, spike-prone commodities: The Australian wholesale spot electricity market. Energy Economics 30, 3172-3185.

Hirsch, G. (2009) Pricing of hourly exercisable electricity swing options using different price processes. Journal of Energy Markets 2(2), 3-46.

Hochberg, Y., Tamhane, A.C. (1987) Multiple Comparison Procedures. Wiley, Hoboken, NJ.

Hollander, M., Wolfe, D.A. (1999) Nonparametric Statistical Methods. Wiley, Hoboken, NJ.

Huisman, R. (2009) An Introduction to Models for the Energy Markets. Risk Books.

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

Huisman, R., Mahieu, R. (2003) Regime jumps in electricity prices. Energy Economics 25, 425-434.

Jabłońska, M., Nampala, H., Kauranne, T. (2011) The multiple-mean-reversion jump-diffusion model for Nordic electricity spot prices. Journal of Energy Markets 4(2), 3-25.

Janczura, J., Weron, R. (2010). An empirical comparison of alternate regime-switching models for electricity spot prices. Energy Economics 32, 1059-1073.

Janczura, J., Weron, R. (2012) Efficient estimation of Markov regime-switching models: An application to electricity spot prices, AStA - Advances in Statistical Analysis, Online First DOI:10.1007/s10182-011-0181-2.

Kaminski, V. (2004) Managing Energy Price Risk: The New Challenges and Solutions, 3rd ed. Risk Books, London.

Kanamura, T., Ōhashi, K. (2008) On transition probabilities of regime switching in electricity prices. Energy Economics 30, 1158-1172.

Karakatsani, N.V., Bunn, D.W. (2008) Intra-day and regime-switching dynamics in electricity price formation. Energy Economics 30, 1776-1797.

Keles, D., Hartel, R., Möst, D., Fichtner, W. (2012) Compressed-air energy storage power plant investments under uncertain electricity prices: An evaluation of compressed-air energy storage plants in liberalized energy markets. Journal of Energy Markets 5(1), 53-84.

Kholodnyi, V.A. (2005). Modeling power forward prices for power with spikes: A non-Markovian approach. Nonlinear Analysis 63, 958-965.

Kim, C.-J. (1994) Dynamic linear models with Markov-switching. Journal of Econometrics 60, 1-22.

Knittel, C.R., Roberts, M.R. (2005) An empirical examination of restructured electricity prices. Energy Economics 27, 791-817.

Koopman, S.J., Ooms, M., Carnero, M.A. (2007) Periodic seasonal reg-ARFIMA-GARCH models for daily electricity spot prices. Journal of the American Statistical Association 102(477), 16-27.

Lapuerta, C., Moselle, B. (2001) Recommendations for the Dutch electricity market. The Brattle Group Report, London.

Lucia, J.J., Schwartz, E.S. (2002) Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. Review of Derivatives Research 5, 5-50.

Makridakis, S., Wheelwright, S.C., Hyndman, R.J. (1998) Forecasting – Methods and Applications, 3rd ed.. Wiley.

Mari, C. (2008) Random movements of power prices in competitive markets: A hybrid model approach. Journal of Energy Markets 1(2), 87-103.

Merton, R.C. (1976) Option pricing when underlying stock returns are discontinuous. Journal of Financial Economics 3, 125-144.

Misiorek, A., Trück, S., Weron, R. (2006) Point and interval forecasting of spot electricity prices: Linear vs. non-linear time series models. Studies in Nonlinear Dynamics and Econometrics 10(3), Article 2.

Mount, T.D., Ning, Y., Cai, X. (2006) Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters. Energy Economics 28: 62-80.

Nomikos, N.K., Soldatos, O.A. (2010) Analysis of model implied volatility for jump diffusion models: Empirical evidence from the Nordpool market. Energy Economics 32, 302-312.

Nowotarski, J., Tomczyk, J., Weron, R. (2011) Wavelet-based modeling and forecasting of the seasonal component of spot electricity prices. The Energy Finance Christmas Workshop (EFC11), Wrocław, Dec. 19-20, 2011.

Pilipovic, D. (1998) Energy Risk: Valuing and Managing Energy Derivatives. McGraw-Hill, New York.

Percival, D.B., Walden, A.T. (2000) Wavelet Methods for Time Series Analysis. Cambridge University Press.

Seifert, J., Uhrig-Homburg, M. (2007) Modelling jumps in electricity prices: theory and empirical evidence. Review of Derivatives Research 10, 59-85.

Shahidehpour, M., Yamin, H., Li, Z. (2002) Market Operations in Electric Power Systems: Forecasting, Scheduling, and Risk Management. Wiley.

Simonsen, I. (2005) Volatility of power markets. Physica A 355, 10-20.

Stevenson, M. (2001) Filtering and forecasting spot electricity prices in the increasingly deregulated Australian electricity market. Research Paper No 63, Quantitative Finance Research Centre, University of Technology, Sydney.

Stevenson, M.J., Amaral, J.F.M., Peat, M. (2006) Risk management and the role of spot price predictions in the Australian retail electricity market. Studies in Nonlinear Dynamics and Econometrics 10(3), Article 4.

Trück, S., Weron, R., Wolff, R. (2007) Outlier treatment and robust approaches for modeling electricity spot prices. Proceedings of the 56th Session of the ISI. Available at MPRA: http://mpra.ub.uni-muenchen.de/4711/.

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 Economics 30, 1098-1115.

Weron, R. (2009) Heavy-tails and regime-switching in electricity prices. Mathematical Methods of Operations Research 69(3), 457-473.

Weron, R., Bierbrauer, M., Trück, S. (2004a) Modeling electricity prices: jump diffusion and regime switching. Physica A 336, 39-48.

Weron, R., Simonsen, I., Wilman, P. (2004b) Modeling highly volatile and seasonal markets: Evidence from the Nord Pool electricity market. In: The Application of Econophysics, H. Takayasu (ed.), Springer, Tokyo, 182-191.

Appendix

Table 6: Mean Squared Error (MSE) for estimation of base regime parameters for the MRS model defined by eqns. (5)-(7) for the conducted simulation study of the NSW market. The LTSC T_t is modeled using wavelets (*upper panel*) or sin-EWMA (*lower panel*), while the stochastic price behavior is simulated using the MRS model. Note, that, the MSE values for α_1 and σ_1 are rescaled by 10^3 or 10^4. The best result for each of the considered measures is indicated in bold. The table also provides results for the multiple comparison procedure using Tukey's honestly significant difference criterion (Hochberg and Tamhane, 1987), indicating for each of the filtering approaches (1)-(8) which of the other methods perform significantly worse / significantly better for the considered criterion.

NSW (Wavelet, Base Regime)						
Measure	MSE (α_1) $\cdot 10^3$	Worse / Better	MSE (α_1/β_1)	Worse / Better	MSE (σ_1) $\cdot 10^4$	Worse / Better
(1) ORG	0.2229	– / {3,4,5,6,7,8}	0.0019	– / {3,4,5,6,7,8}	0.1497	– / {3,4,5,6,7,8}
(2) FPT	0.2285	– / {3,4,5,6,7,8}	0.0019	– / {3,4,5,6,7,8}	0.1520	– / {3,4,5,6,7,8}
(3) VPT1	0.0397	{1,2,5,6,8} / {4}	0.0003	{1,2,5,6,8} / {4}	0.0707	{1,2} / {5,6,7,8}
(4) VPT2	**0.0224**	All / –	**0.0002**	All / –	0.0827	{1,2} / {5,6,7,8}
(5) RFP	0.0500	{1,2} / {3,4,7}	0.0004	{1,2} / {3,4,7}	0.0563	{1,2,3,4} / –
(6) RFD	0.0526	{1,2} / {3,4,7}	0.0004	{1,2} / {3,4,7}	0.0680	{1,2,3,4} / {8}
(7) RM	0.0400	{1,2,5,6,8} / {4}	0.0003	{1,2,5,6,8} / {4}	0.0549	{1,2,3,4} / {8}
(8) RSC	0.0491	{1,2} / {3,4,7}	0.0004	{1,2} / {3,4,7}	**0.0426**	{1,2,3,4,6,7} / –
NSW (sin-EWMA, Base Regime)						
Measure	MSE (α_1) $\cdot 10^3$	Worse / Better	MSE (α_1/β_1)	Worse / Better	MSE (σ_1) $\cdot 10^4$	Worse / Better
(1) ORG	0.0442	{5,6,8} / {3}	0.0008	{6,8} / {3,4,7}	0.1384	{4} / {3,5,6,7,8}
(2) FPT	0.0527	{5,6,8} / {3}	0.0010	{6,8} / {3,4,7}	0.1319	{4} / {3,5,6,7,8}
(3) VPT1	**0.0214**	All / –	**0.0002**	All / –	0.0736	{1,2,4} / {5,6,7,8}
(4) VPT2	0.0404	{5,6,8} / {3}	0.0003	{1,2,5,6,7,8} / {3}	0.1431	{1,2,3,4} / –
(5) RFP	0.0594	{6,8} / {1,2,3,4,7}	0.0007	{6,8} / {3,4,7}	0.0510	{1,2,3,4} / –
(6) RFD	0.0726	{8} / {1,2,3,4,5,7}	0.0009	{8} / {1,2,3,4,5,7}	**0.0504**	{1,2,3,4} / –
(7) RM	0.0398	{5,6,8} / {3}	0.0005	{1,2,5,6,8} / {3,4}	0.0546	{1,2,3,4} / –
(8) RSC	0.1141	– / All	0.0014	– /All	0.0577	{1,2,3,4} / –

Table 7: Mean Squared Error (MSE) for estimation of the spike and drop regime parameters for the MRS model studied in Table 6.

NSW (Wavelet, Spike and Drop Regime)								
	MSE(μ_2)	Worse / Better	MSE(σ_2)	Worse / Better	MSE(μ_3)	Worse / Better	MSE(σ_3)	Worse / Better
(1)	0.0447	– / {8}	0.0567	– / {8}	0.1324	{6,7} / –	0.0236	{6,7} / {4,8}
(2)	0.0479	– / {8}	0.0650	– / {8}	0.1282	{6} / {8}	0.0271	– / {4,8}
(3)	0.0449	– / {8}	0.0597	– / {8}	0.1265	{6,7} / –	0.0260	– / {4,8}
(4)	0.0429	– / –	0.0567	– / {8}	0.1084	{5,6,7} / –	0.0185	{1,2,3,5,6,7} / {8}
(5)	0.0427	– / {8}	0.0615	– / {8}	0.1509	– / {4,8}	0.0237	– / {4,8}
(6)	0.0422	– / {8}	0.0553	– / {8}	0.1481	– / {1,2,3,4,8}	0.0261	– / {1,4,8}
(7)	0.0409	– / –	0.0595	– / {8}	0.1430	– / {1,3,4,8}	0.0249	– / {1,4,8}
(8)	**0.0369**	{1,2,3,5,6} / –	**0.0424**	All / –	**0.1020**	{2,5,6,7} / –	**0.0137**	All / –
NSW (sin-EWMA, Spike and Drop Regime)								
	MSE(μ_2)	Worse / Better	MSE(σ_2)	Worse / Better	MSE(μ_3)	Worse / Better	MSE(σ_3)	Worse / Better
(1)	0.0192	– / –	0.0401	– / –	0.3222	– / {3,4,5,6,,8}	0.0507	{6,7} / {4,8}
(2)	0.0191	– / –	0.0384	– / –	0.3159	– / {3,4,8}	0.0506	– / {4,8}
(3)	0.0184	– / –	0.0415	– / –	0.2400	{1,2,5,6,7} / –	0.0386	– / {4,8}
(4)	0.0201	– / –	0.0372	– / –	0.2107	{1,2,5,6,7} / –	**0.0311**	{1,2,3,5,6,7} / {8}
(5)	0.0193	– / –	0.0417	– / –	0.3016	{1} / {3,4,8}	0.0411	– / {4,8}
(6)	0.0185	– / –	**0.0361**	– / –	0.2938	{1} / {3,4,8}	0.0451	– / {1,4,8}
(7)	**0.0178**	– / –	0.0369	– / –	0.3213	– / {3,4,8}	0.0475	– / {4,8}
(8)	0.0185	– / –	0.0392	– / –	**0.2085**	{1,2,5,6,7} / –	0.0321	All / –