A regime-switching stochastic volatility model for forecasting electricity prices

Oskar Knapik [a, *] Peter Exterkate [a, b]

[a] *CREATES, Aarhus University, 8210 Aarhus V, Denmark*
[b] *School of Economics, University of Sydney, NSW 2006, Australia*

31 January 2017

Abstract

In a recent review paper, Weron (2014) pinpoints several crucial challenges outstanding in the area of electricity price forecasting. This research attempts to address all of them by *(i)* showing the importance of considering fundamental price drivers in modeling, *(ii)* developing new techniques for probabilistic (i.e. interval or density) forecasting of electricity prices, and *(iii)* introducing an universal technique for model comparison. We propose a new regime-switching stochastic volatility model with three regimes (negative jump or "drop", normal price or "base", positive jump or "spike") where the transition matrix depends on explanatory variables. Bayesian inference is employed in order to obtain predictive densities. The main focus of the paper is on short-term density forecasting in the Nord Pool intraday market. We show that the proposed model outperforms several benchmark models at this task.

Keywords: Electricity prices, density forecasting, Markov switching, stochastic volatility, ordered probit model, Bayesian inference, fundamental price drivers, seasonality, Nord Pool power market, electricity price forecasting, probabilistic forecasting.

JEL Classification: C22, C24, Q41, Q47.

*Corresponding author. Address: CREATES, Department of Economics and Business Economics, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V, Denmark. Email address: oknapik@econ.au.dk.

Acknowledgment: The authors appreciate support from CREATES, the Center for Research in Econometric Analysis of Time Series, funded by the Danish National Research Foundation with grant number DNRF78.

1 Introduction

Electricity is a unique commodity, characterized by a high variability. It cannot be stored and requires immediate delivery. The end-user demand shows high variability and strong dependence on the weather and the business cycle. Moreover, events like power plant outages and the (un)reliability and complexity of the transmission grid reduce predictability. The resulting electricity price series are characterized by strong seasonality at different frequencies (annual, weekly, daily, and hourly). However, the most distinct feature of electricity prices is their very high volatility and abrupt, short-lived, and generally unanticipated extreme price changes known as "spikes" or "jumps" (see Serati et al., 2008; Janczura et al., 2013, among others). Electricity prices from the Nord Pool power market have also been found to be nonstationary (Lisi and Nan, 2014; Weron, 2014) and to exhibit long-memory properties (Haldrup and Nielsen, 2006a,b).

There is a large body of literature on the topic of electricity price forecasting; see Weron (2014) for a recent review. A picture that emerges from this literature is that the need for realistic models of electricity price dynamics, capturing their unique characteristics, and for adequate derivatives pricing techniques, has not been fully satisfied. It is the aim of this paper to address some of the crucial challenges pointed out by Weron (2014); specifically, the use of fundamental price drivers in modeling and developing new tools for the probabilistic forecasting of electricity prices.

When building a model for electricity prices, one of the crucial steps is to find an appropriate description of their seasonal pattern. Moreover, electricity prices present various forms of nonlinear dynamics, the crucial one being the strong dependence of the variability of the series on its own past. The resulting clustering of large shocks is well documented; see Karakatsani and Bunn (2008), among others. Furthermore, the "spiky" character of electricity prices suggests that there exists a nonlinear switching mechanism between normal and extreme states, or regimes. The requirement of stochastic jump arrival probabilities directly leads to regime-switching models. Markov regime-switching (MS) models seem to be a natural candidate for modeling such a nonlinear and complex structure; see, for example, Andreasen and Dahlgren (2006), Geman and Roncoroni (2006), Handika et al. (2014), Heydari and Siddiqui (2010), Huisman and Mahieu (2003), Kanamura and Ohashi (2008), Kosater and Mosler (2006), Mount et al. (2006). MS models have been successfully applied for modeling electricity prices by many researchers; see Janczura and Weron (2010a), Haldrup and Nielsen (2006a), among others.

This paper introduces a new regime-switching stochastic volatility model, with a time-varying transition matrix that depends on explanatory variables. The core of the model is an autoregressive process with a stochastic volatility (SV) error term. The main focus of this research is on short-time density forecasting of electricity prices. Although important, this topic is barely touched upon in the electricity prices forecasting literature (see Weron, 2014). Serinaldi (2011) forecasts the distribution of electricity prices using a generalized additive model, but computes and discusses only predictive intervals. Huurman et al. (2012) consider GARCH-type time-varying volatility models and find that models augmented with weather forecasts statistically outperform the ones without this information. They utilize the probability integral transform scores of the realization of the variables with respect to the forecast densities. Jónsson et al. (2014) develop a semi-parametric methodology for generating densities of day-ahead electricity prices in Western Denmark based on quantile regression.

The topicality and importance of electricity price forecasting is further evidenced by a special issue of the International Journal of Forecasting (Volume 32, Issue 3, 2016) dedicated entirely to that topic, and by the energy forecasting competition organized by this journal. This competition and the current state of probabilistic energy forecasting research was summarized by Hong et al. (2016). One of the presented approaches by Maciejowska et al. (2016) introduces a new methodology involving quantile regression to average large numbers of point forecasts, and principal component analysis to extract the major factors driving the individual forecasts. Their approach outperforms the benchmark of an autoregressive model with exogenous regressors (ARX), as well as quantile regression averaging without factor extraction.

Our paper explores the use of Bayesian inference in order to construct predictive densities for future electricity prices. In addition to the novel setup of our model, which nests various important special cases including ARX and several standard MS and SV models, we also introduce a natural, universal method for model comparisons via predictive Bayes factors to the electricity prices literature. Bayesian approaches have been used in the context of electricity price modelling by several authors. Panagiotelis and Smith (2008) use a vector autoregressive model with exogenous effects and a skew t distributed disturbance term for hourly Australian electricity spot prices. They use a Markov Chain Monte Carlo approach in order to construct the predictive distribution of future spot prices. Smith (2010) proposes using Bayesian inference for a Gaussian stochastic volatility model based on periodic autoregressions. Both of these studies include demand and day types as exogenous explanatory variables in the mean and log-volatility equations. They confirm that there is a nonlinear relationship between demand and mean prices, and they construct the predictive density of prices evaluated over a horizon of one week. Our work can be seen as extending this line of research in two ways, with a Markov-switching structure to flexibly accommodate such nonlinearities, and by allowing for many more predictors.

The remainder of this paper is structured as follows. In Section 2 we describe the data and the main electricity price drivers. Section 3 presents the model and our methods for Bayesian inference. Empirical results are discussed in Section 4, and Section 5 concludes.

2 Data

Our data set comes from the Nordic power exchange, Nord Pool, which is one of the leading power markets in Europe, owned by the Nordic and Baltic transmission system operators. We consider information from two different markets within Nord Pool; the day-ahead auction market Elspot and the intraday market Elbas.

About 380 companies from 20 countries trade on the Elspot market, including both producers and large consumers, for a trading volume of approximately 500 terawatt-hours in 2015. Elspot is the Nord Pool's auction market for day-ahead electricity delivery. Its web-based trading system enables participants to submit bids and offers for each individual hour of the next day. Orders can be made between 8 AM and 12 noon Central European Time (CET). The aggregated buy and sell orders form demand and supply curves for each delivery hour of the next day. The intersection of these curves constitutes the system price for each hour (quoted in euros per megawatt-hour, MWh). The resulting hourly prices are announced to the market at 12:42 PM and contracts are invoiced between buyers and sellers between 1 PM and 3 PM. Thus, all 24 prices on day $t+1$ are determined on a given day t and released simultaneously. A detailed review of the operation of the market is given in Nord Pool (2016).

The intraday market, Elbas, supplements the day-ahead market and helps secure the necessary balance between supply and demand in the power market for Northern Europe. The majority of the volume handled by Nord Pool is traded on the day-ahead market. For the most part, the balance between supply and demand is secured here. However, incidents may take place between the closing of the day-ahead market at noon and delivery the next day. On the intraday market, buyers and sellers can trade volumes almost in real time to bring the market back to balance. Elbas is a continuous market, and trading takes place every day around the clock until one hour before delivery. Prices are set based on a first-come,

Figure 1: Hourly electricity prices from the Elbas power market, y_t.

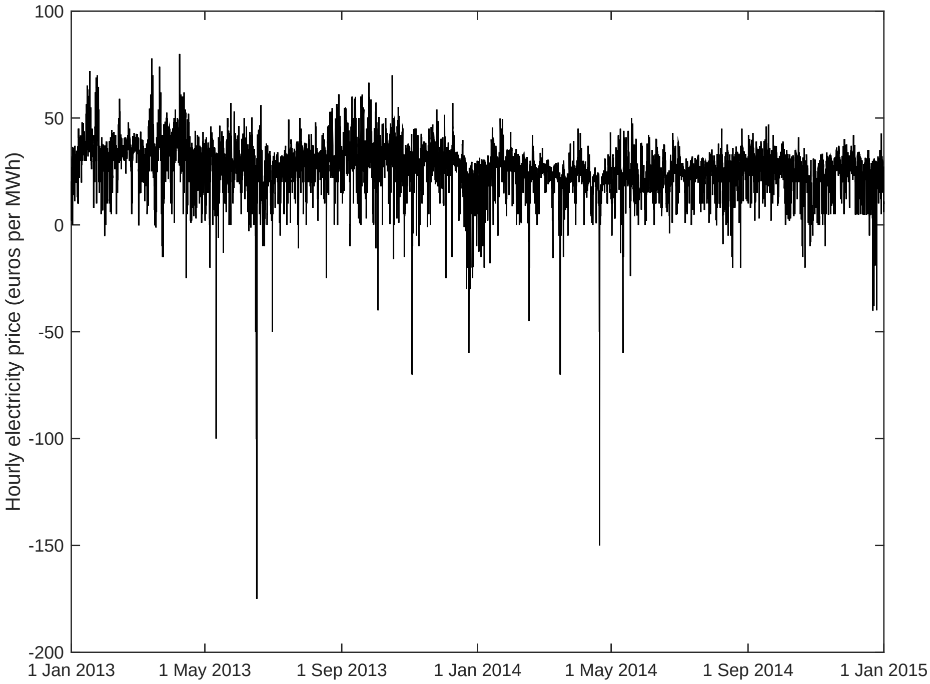

first-served principle, where best prices (highest buy and lowest sell) come first. The Elbas market is becoming increasingly important as the amount of wind power entering the grid rises. Wind power is unpredictable by nature, and imbalances between day-ahead contracts and produced volume often need to be offset.

The data for the estimation period consist of a series of hourly observations of electricity prices in these two Nord Pool markets. The study is conducted using hourly electricity prices for the whole area (system prices) from the Elspot market $spot_t$ and the corresponding hourly volume-weighted average prices from the Elbas market y_t. The focus of our study is on modeling and forecasting the Elbas prices. The in-sample data covers the period from 1 January 2013 to 31 December 2014 (17520 observations). Figures 1 and 2 present these time series. The prices exhibit many characteristics that have typically been found in this literature, including seasonality and spikes. A further sample of 12 months (8760 observations) from 1 January 2015 to 31 December 2015 is reserved for assessing the out-of-sample performance of the model.

In order to model the hourly prices from the Elbas market y_t, we use the following explanatory variables, which are considered to be the main price drivers in the Nord Pool power market:

Figure 2: Hourly electricity prices from the Elspot power market, $spot_t$.

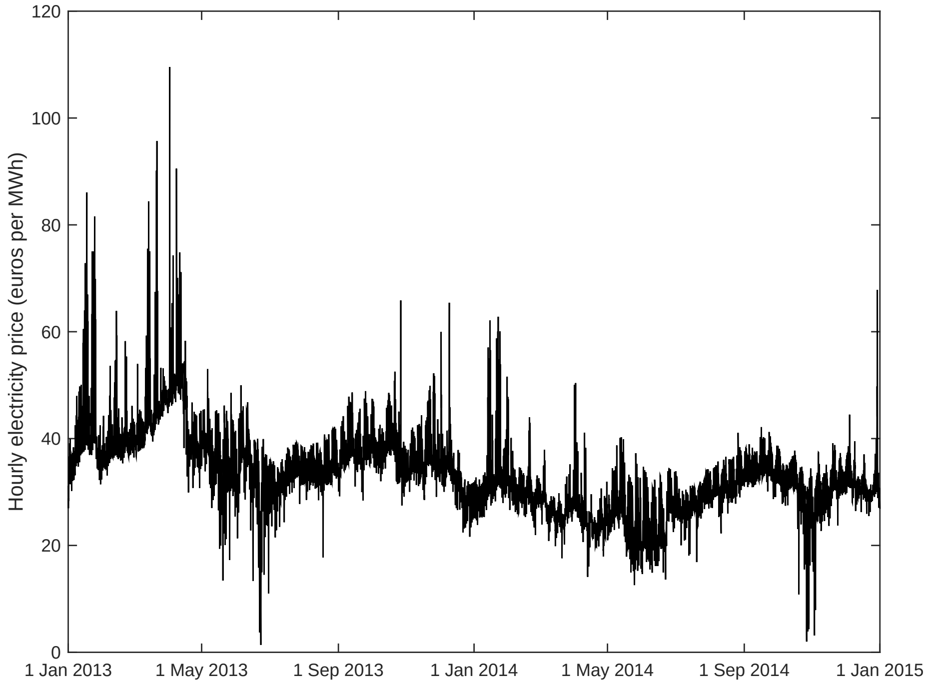

- the hourly Elspot electricity system price, $spot_t$;
- the turnover at system price from the Elspot market, $load_t$;
- the water reservoir level, res_t;
- heating degrees days, hdd_t;
- wind power production, $wind_t$; and
- a seasonal component, $seas_t$.

Obviously, the Elspot market is a significant source of information about Elbas electricity prices. The situation is analogous to that in futures markets, where the basis is the difference between futures price and the underlying spot price. We summarize this information by $spot_t$ and $load_t$. The seasonality in the electricity price series is captured by $seas_t$, which consists of sine and cosine terms taken at the daily, weekly, and annual frequencies. The rationale for including the other price drivers is as follows.

In the Nord Pool electricity market, about 53% of the power production is generated from hydropower reservoirs. The influence of water reservoir levels on electricity prices in this market has been studied by Gjolberg and Johnsen (2001), Botterud et al. (2002), Førsund and Hoel (2004) and von der Fehr et al. (2005). These researchers conclude that hydropower reservoir levels are an important factor to explain both futures and spot prices. The ability to store energy in hydro reservoirs causes less variation in the Nordic price structure, compared to that of for example Germany. Inflow during summer and in periods with low demand can be used in the winter. The reservoir levels and capacity data are from Norwegian Water Resources and Energy Directorate (NVE), Svensk Energi (Swedenergy AB), and the Finnish Environment Institute (SYKE). The data is published at the weekly frequency.[1] Our summary statistic res_t is the volume stored in reservoirs as a percentage of the total hydropower capacity available.

Temperature is the main price driver in the Nordic countries. Cold temperatures increase heating demand, since electricity is very much used for heating in these countries. Colder temperatures usually increase prices because of higher power demand. However, in special cases, the demand for the heat could trigger secondary electricity production and cause the prices to decrease. The behavior of weather variables can also produce some predictable seasonal pattern in electricity prices. The relationship between weather variables and electricity load and price has been studied by many researchers. Li and Sailor (1995) and Sailor et al. (1998) show in a few US states that temperature is the most significant weather variable explaining electricity and gas demand. The influence of air temperature has also been described by other authors, who obtained a significant explanatory power in their modeling; see, for example, Peirson and Henley (1994) and Peirson and Henley (1998). The heating degree day (hdd_t) variable that we use shows the demand for energy needed for heating. It is taken from measurements of the outside air temperature. The heating requirements for a specific structure at a specific place tend to be directly proportional to the HDD at that location. In this study we use the average temperature measured on a daily basis in 13 Nordic cities (Oslo, Bergen, Trondheim, Tromsø, Helsinki, Sodankyla, Vaasa, Tampere, Stockholm, Göteborg, Östersund, Luleå, and Copenhagen).

Finally, wind power production ($wind_t$) is also an important electricity price driver. Due to the fact that there is no fuel cost for production and a lot of unpredictability, additional wind energy can lead to price decreases. This type of energy may in some cases cause even negative prices in hours with low demand and unanticipated additional supply. On the other hand, when wind production falls short of expected values, it can trigger high prices, in both day-ahead and intraday markets.

[1]The series of water reservoir levels is very regular. Therefore, the transformation from the weekly to hourly frequency is done by simple linear interpolation.

3 Model

A Markov regime-switching (MS) model allows for various types of temporal dependence within the regimes, and 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. However, several modeling questions have to be addressed to build a proper MS model, as discussed by Janczura and Weron (2010a).

First, the number of regimes has to be chosen. The important advantage of MS models over the alternatives is that there is no need for explicitly specifying a threshold variable and level for the regimes, and as a result they are preferred for modeling risk purposes. There is no fundamental reason for considering any specific number of regimes for electricity price modeling (see Janczura and Weron, 2010a). However, almost all published papers consider only two-regime models due to their computational convenience. In addition to the base regime, a "spike" (or "excited") regime was introduced to capture extreme price behavior. Karakatsani and Bunn (2008) introduced a third regime for capturing the most extreme prices. The existence of an additional "down-spike" or "drop" regime can also be justified by occurrence of many very low prices. In this paper we consider a Markov regime-switching stochastic volatility model with three regimes, which we label drop, base, and spike. We also introduce a data-driven mechanism of switching between different regimes, in the form of an ordered probit model. This approach can easily accommodate a further extension of the number of regimes.

Secondly, the model defining the price dynamics within each of the regimes has to be selected. The base regime is usually modeled by a mean-reverting autoregressive (see Ethier and Mount, 1998; Deng, 1999, among others) or diffusion model (for reviews see Huisman, 2009; Janczura and Weron, 2010b), which is sometimes heteroscedastic (Janczura and Weron, 2009). For the spike regime(s), on the other hand, a number of different specifications have been suggested in the literature, ranging from mean-reverting diffusions (Karakatsani and Bunn, 2008), to Gaussian (Huisman and de Jong, 2003; Liebl, 2013), lognormal (Weron et al., 2004; Bierbrauer et al., 2004), exponential (Bierbrauer et al., 2007), heavy tailed (Weron, 2009) and non-parametric (Eichler and Türk, 2013) random variables, to mean-reverting diffusions with Poisson jumps (Arvesen et al., 2013; de Jong, 2006; Keles et al., 2012; Mari, 2008). One of the advantages of the regime-switching framework is that we can explicitly model the short-lived characteristics of power spikes. We consider autoregressive models with additional exogenous regressors and stochastic volatility (ARX-SV) for each of the regimes.

Finally, the dependence between the regimes has to be decided. Dependent regimes with the same random noise process in all regimes (but different parameters) are computationally less demanding than independent ones. However, an independent regimes model enables greater flexibility and seems to be a natural choice for a process which significantly changes its dynamics across regimes. We follow the latter approach.

The empirical comparison by Janczura and Weron (2010a) shows that the best structure is that of an independent spike three-regime model with time-varying transition probabilities, heteroscedastic diffusion-type base regime dynamics, and shifted spike regime distributions. We propose a model which captures each of these characteristics more accurately. Below, we provide the details on our Markov-switching stochastic volatility model for electricity prices. We first introduce the stochastic volatility model that governs the price dynamics within each of the three regimes (drop, base, spike), followed by a description of the transition dynamics between the regimes. Finally, we derive a Gibbs sampler and discuss how to use it for parameter estimation, density forecasting, and forecast evaluation.

3.1 Stochastic volatility model

Denote the (scalar) price y_t. Any exogenous observables will be dated t for convenience, but it is assumed that they are known when the forecast is being made, which can be at any time from $t-24$ to $t-1$. To avoid any potential scaling issues, all variables including the regressand are studentized over the estimation window. The latent regime is denoted $R_t \in \{1,2,3\}$, which correspond to the drop, base, and spike regimes, respectively. All parameters $(\alpha, \beta, \gamma, \delta, \tau)$, to be introduced below, are collected in θ.

We specify our stochastic volatility model as

$$\begin{aligned} y_t \,|\theta, y_{t-1}, \ldots, y_{t-p}, \sigma_t, R_t &\sim \mathcal{N}\left(x_t' \beta_{R_t}, \sigma_t^2\right), \\ \log \sigma_t \,|\theta, \sigma_{t-1}, \ldots, \sigma_{t-q}, R_t &\sim \mathcal{N}\left(z_t' \gamma_{R_t}, \tau_{R_t}^{-1}\right). \end{aligned}$$

The regressors in the mean equation are $x_t' = (1, y_{t-1}, \ldots, y_{t-p}, spot_t, seas_t')$, and in the volatility equation, $z_t' = (1, \log\sigma_{t-1}, \ldots, \log\sigma_{t-q}, spot_t, seas_t')$. We set the lag lengths to $p = q = 48$, so our price process has a two-day memory.

For notational simplicity, introduce $T \times 1$ vectors y, σ, and $\log \sigma$, the $T \times N$ matrix X, and the $T \times M$ matrix Z. It will also be convenient to collect all T_r observations that belong to regime r in separate vectors and matrices, for $r = 1, 2, 3$. The $T_r \times 1$ vectors y_r, σ_r, and $\log \sigma_r$, the $T_r \times N$ matrices X_r, and the $T_r \times M$ matrices Z_r contain only those rows of the original vectors and matrices with $R_t = r$. Finally, let $\Sigma = \text{diag}\left(\sigma_t^2\right)$ be a $T \times T$ matrix, and create diagonal $T_r \times T_r$ matrices Σ_r similarly. We may then write our stochastic volatility model as

$$y_r \,|\theta, \sigma_r, R \;\sim\; \mathcal{N}\left(X_r \beta_r, \Sigma_r\right), \qquad \log \sigma_r \,|\theta, R \;\sim\; \mathcal{N}\left(Z_r \gamma_r, \tau_r^{-1} I_{T_r}\right), \qquad \text{for } r = 1, 2, 3.$$

We use a standard set of priors, which are independent across regimes. For each regime r, we specify $p\left(\beta_r\right)$ as $\mathcal{N}\left(0, \lambda^{-1} I_N\right)$, $p\left(\gamma_r\right)$ as $\mathcal{N}\left(0, \mu^{-1} I_M\right)$, and the uninformative $p\left(\tau_r\right) \propto \tau_r^{-1}$. Preliminary experiments suggest that our results are largely insensitive to the choice of hyperparameters λ and μ; we use $\lambda = \mu = 1$ in our application. Finally, a prior needs to be specified for the pre-sample volatilities $\sigma_{1-q}, \ldots, \sigma_0$. The standard approach of using the unconditional distribution implied by τ and the autoregressive parameters in γ, as advocated by Jacquier et al. (2002), is not feasible in our setup, since we are not imposing stationarity on the volatility process. Instead, we follow de Jong and Shephard (1995) and set $\log \sigma_t \,|R_t \sim \mathcal{N}\left(0, \tau_{R_t}^{-1}\right)$ independently for $t = 1 - q, \ldots, 0$.

3.2 State dynamics

We model the regime R_t according to a hidden Markov process, where each state transition is governed by an ordered probit model. Our specification is an extension of the two-regime model in Filardo and Gordon (1998). Specifically, the transition probabilities $P\left[R_{t+1} \,|R_t = r, \theta\right]$ are given implicitly by

$$R_{t+1} = \begin{cases} 1 & \text{if} \qquad\;\; w_t' \delta_r + \varepsilon_t < 0, \\ 2 & \text{if} \quad 0 \leq w_t' \delta_r + \varepsilon_t \leq \alpha_r, \\ 3 & \text{if} \quad \alpha_r < w_t' \delta_r + \varepsilon_t \end{cases}$$

where ε_t is i.i.d. $\mathcal{N}(0, 1)$, $w_t' = \left(spot_t, res_t, load_t, hdd_t, wind_t\right)$, and the parameters δ_r and α_r may again be different for each regime r. Note that no generality is lost by fixing the mean and variance of ε_t, as well as the threshold between the first two regimes; these restrictions serve to identify the model.

To simplify notation, we write $R_{t+1}^* = w_t' \delta_{R_t} + \varepsilon_t$. We create $T_r \times K$ matrices W_r and $T_r \times 1$ vectors R_r^* as above, so that $R_r^* \,|\theta \sim \mathcal{N}\left(W_r \delta_r, I_{T_r}\right)$, for $r = 1, 2, 3$. In particular, R_r^* includes R_{t+1}^* if $R_t = r$; the r-th ordered probit model describes all transitions from regime r to any of the three regimes.

As in the mean and volatility equations, the regression coefficients in these probit models are also given independent priors $\delta_r \sim \mathcal{N}\left(0, \nu^{-1} I_K\right)$, where we set $\nu = 1$ after finding that the results are not very sensitive to this choice. As in Albert and Chib (1993), the regime thresholds α_r have uninformative priors, uniform over $(0, \infty)$. Finally, pre-sample states $R_{1-q}, \ldots, R_0$ are required in the prior specification for the volatilities. We specify $P\left[R_{1-q} = 1\right] = P\left[R_{1-q} = 3\right] = 0.05$ and $P\left[R_{1-q} = 2\right] = 0.90$, and the ordered probit model then automatically implies a prior for $R_{2-q}, \ldots, R_0$.

3.3 Gibbs sampler

Because of its modular nature, our model lends itself well to estimation using a Gibbs sampler with data augmentation. We can obtain draws from all required conditional posteriors analytically. Standard results (Koop, 2003) apply for all regression coefficients; for $r = 1, 2, 3$, we may draw

- β_r from $\mathcal{N}\left(\left(X_r' \Sigma_r^{-1} X_r + \lambda I_N\right)^{-1} X_r' \Sigma_r^{-1} y_r, \left(X_r' \Sigma_r^{-1} X_r + \lambda I_N\right)^{-1}\right)$,
- γ_r from $\mathcal{N}\left(\left(\tau_r Z_r' Z_r + \mu I_M\right)^{-1} \left(\tau_r Z_r' \log \sigma_r\right), \left(\tau_r Z_r' Z_r + \mu I_M\right)^{-1}\right)$, and
- δ_r from $\mathcal{N}\left(\left(W_r' W_r + \nu I_K\right)^{-1} W_r' R_r^*, \left(W_r' W_r + \nu I_K\right)^{-1}\right)$.

The conditional posterior for each τ_r is the usual gamma distribution, with shape parameter $T_r/2$ and scale parameter $2 \big/ \left[\left(\log \sigma_r - Z_r \gamma_r\right)' \left(\log \sigma_r - Z_r \gamma_r\right)\right]$.

The conditional posterior for the latent volatilities is nonstandard, but an auxiliary variable v_t can be introduced to obtain draws analytically. For $t = 1, \ldots, T$, define v_t to be $\left(y_t - x_t' \beta_{R_t}\right)^2 \big/ \left(2\sigma_t^2\right)$ plus a draw from the exponential distribution with mean one, and then draw $\log \sigma_t$ from $\mathcal{N}\left(z_t' \gamma_{R_t} - \tau_{R_t}^{-1}, \tau_{R_t}^{-1}\right)$, truncated to the interval $\left(\frac{1}{2} \log\left(\left(y_t - x_t' \beta_{R_t}\right)^2 / (2 v_t)\right), \infty\right)$; for further details see Damien et al. (1999).

Sampling the state dynamics R_t is done along the same lines as in Filardo and Gordon (1998). For $t = 1-q, 2-q, \ldots, T$, the conditional posterior of R_t has support $\{1, 2, 3\}$, with probability for state r proportional to $p\left(R_t = r \left| R_{t-1}, \theta\right.\right) \cdot p\left(R_{t+1} \left| R_t = r, \theta\right.\right) \cdot p\left(\log \sigma_t \left| \log \sigma_{t-1}, \ldots, \log \sigma_{t-q}, \theta, R_t = r\right.\right) \cdot p\left(y_t \left| \sigma_t, y_{t-1}, \ldots, y_{t-p}, \theta, R_t = r\right.\right)$. The factors in this expression can be explicitly computed as

$$p\left(R_t = r \left| R_{t-1}, \theta\right.\right) = \begin{cases} \Phi\left(-w_{t-1}' \delta_{R_{t-1}}\right) & \text{if } r = 1, \\ \Phi\left(\alpha_{R_{t-1}} - w_{t-1}' \delta_{R_{t-1}}\right) - \Phi\left(-w_{t-1}' \delta_{R_{t-1}}\right) & \text{if } r = 2, \\ 1 - \Phi\left(\alpha_{R_{t-1}} - w_{t-1}' \delta_{R_{t-1}}\right) & \text{if } r = 3, \end{cases}$$

$$p\left(R_{t+1}\left|R_t=r,\theta\right.\right) \;=\; \begin{cases} \Phi\left(-w_t'\delta_r\right) & \text{if } R_{t+1}=1, \\ \Phi\left(\alpha_r-w_t'\delta_r\right)-\Phi\left(-w_t'\delta_r\right) & \text{if } R_{t+1}=2, \\ 1-\Phi\left(\alpha_r-w_t'\delta_r\right) & \text{if } R_{t+1}=3, \end{cases}$$

$$p\left(\log\sigma_t\left|\log\sigma_{t-1},\ldots,\log\sigma_{t-q},\theta,R_t=r\right.\right) \;=\; \phi\left(\log\sigma_t;z_t'\gamma_r,\tau_r^{-1}\right),$$

$$p\left(y_t\left|\sigma_t,y_{t-1},\ldots,y_{t-p},\theta,R_t=r\right.\right) \;=\; \phi\left(y_t;x_t'\beta_r,\sigma_t^2\right),$$

where Φ is the standard normal CDF, and ϕ is the normal PDF with specified mean and variance.

Finally, the sampling distributions for the latent R_{t+1}^* as well as the thresholds α_r were obtained by Albert and Chib (1993). For $t=1,2,\ldots,T$, the latent R_{t+1}^* can be drawn from $\mathcal{N}\left(w_t'\delta_{R_t},1\right)$, truncated to the correct interval, which is $(-\infty,0)$ if $R_{t+1}=1$, $(0,\alpha_{R_t})$ if $R_{t+1}=2$, and (α_{R_t},∞) if $R_{t+1}=3$. The conditional posterior for α_r is uniform with lower bound $\max\left\{\max\left\{R_{t+1}^*: R_t=r \text{ and } R_{t+1}=2\right\},0\right\}$ and upper bound $\min\left\{R_{t+1}^*: R_t=r \text{ and } R_{t+1}=3\right\}$, for $r=1,2,3$.

3.4 Density forecasting

We can obtain draws from the one-step-ahead predictive density within the Gibbs sampler. At each step $d=1,2,\ldots,D$, we draw $R_{T+1}^{*(d)}\sim\mathcal{N}\left(w_T'\delta_{R_T^{(d)}}^{(d)},1\right)$ to find $R_{T+1}^{(d)}$, which is 1 if $R_{T+1}^{*(d)}<0$, 2 if $0\le R_{T+1}^{*(d)}\le\alpha_{R_T^{(d)}}^{(d)}$, and 3 otherwise. Finally, we draw $\log\sigma_{T+1}^{(d)}\sim\mathcal{N}\left(z_{T+1}'\gamma_{R_{T+1}^{(d)}}^{(d)},\tau_{R_{t+1}^{(d)}}^{(d)\,-1}\right)$, and use it to draw $y_{T+1}^{(d)}\sim\mathcal{N}\left(x_{T+1}'\beta_{R_{T+1}^{(d)}}^{(d)},\sigma_{T+1}^{(d)\,2}\right)$. The empirical distribution of these draws after discarding D_0 burn-in draws, $F(c)=\dfrac{1}{D-D_0}\sum_{d=D_0+1}^{D}\mathbf{1}\left\{y_{T+1}^{(d)}\le c\right\}$, approximates the CDF of $y_{T+1}\left|y\right.$. A kernel density estimate of the corresponding PDF is used to visualize this distribution in Section 4.

In our empirical application we will also be interested in h-step-ahead forecasts for $h=2,3,\ldots,24$; that is, density forecasts for every hour of the next day. We may recursively obtain draws of each y_{T+h}, using the same procedure as outlined for $h=1$ above. Most exogenous regressors in x_t, z_t, and w_t are available one day ahead, and highly accurate forecasts are available for the others. For the endogenous R_t^*, R_t, σ_t, and y_t that are needed for $t>T$, we may simply substitute the forecasts that were made at shorter horizons. This procedure is justified by the standard decomposition

$$p\left(y_{T+1},y_{T+2},\ldots,y_{T+24}\left|y\right.\right)=p\left(y_{T+1}\left|y\right.\right)\cdot p\left(y_{T+2}\left|y_{T+1},y\right.\right)\cdot\;\cdots\;\cdot p\left(y_{T+24}\left|y_{T+23},\ldots,y_{T+1},y\right.\right).$$

Table 1: The model specifications considered in the empirical study.

Mnemonic	Restrictions on the model introduced in Section 3
ARX	no regime switching ($R_t = 2$ for all t), no stochastic volatility ($\sigma_t^2 = \sigma^2$ for all t)
ARX-SV	no regime switching ($R_t = 2$ for all t)
MS-ARX	no stochastic volatility ($\sigma_t^2 = \sigma_{R_t}^2$ for all t)
MS-ARX-SV	no restrictions

3.5 Forecast evaluation

We evaluate the quality of our density forecasts using predictive Bayes factors, as suggested by Geweke and Amisano (2010). The predictive Bayes factor comparing two competing models is given by the ratio of the predictive densities implied by these models, evaluated at the realized prices. For one-step-ahead forecasts, we may approximate the predictive density evaluated at the realized price y_{T+1} using

$$\begin{aligned} p\left(y_{T+1}\left|y\right.\right) &= \int\int\int p\left(y_{T+1}\left|y,\theta,\sigma_{T+1},R_{T+1}\right.\right)dR_{T+1}d\sigma_{T+1}d\theta \\ &\approx \frac{1}{D-D_0}\sum_{d=D_0+1}^{D}\phi\left(y_{T+1};x_{T+1}'\beta^{(d)}_{R^{(d)}_{T+1}},\sigma^{(d)\,2}_{T+1}\right). \end{aligned}$$

A similar iterative procedure as outlined above for density forecasting can then be used to approximate $p\left(y_{T+2}\left|y_{T+1},y\right.\right),\ldots,p\left(y_{T+24}\left|y_{T+23},\ldots,y_{T+1},y\right.\right)$, except that now realized rather than simulated values of y_t need to be substituted into x_t, for $t > T$. Finally, the joint predictive density is given by

$$p\left(y_{T+1},y_{T+2},\ldots,y_{T+24}\left|y\right.\right) = p\left(y_{T+1}\left|y\right.\right)\cdot p\left(y_{T+2}\left|y_{T+1},y\right.\right)\cdot\ \cdots\ \cdot p\left(y_{T+24}\left|y_{T+23},\ldots,y_{T+1},y\right.\right).$$

4 Results

We study the hourly volume-weighted electricity prices from Elbas power market in order to understand the data generating mechanism and to examine the proposed model. In this empirical study, we consider four model specifications: the basic autoregressive process with explanatory variables (ARX), the autoregressive model with stochastic volatility error and explanatory variables (ARX-SV), the three-state Markov regime-switching model (MS-ARX) and our proposed three-state Markov regime-switching model with stochastic volatility (MS-ARX-SV). These specifications are summarized in Table 1. Formal Bayesian model comparison in terms of the predictive adequacy is measured by predictive Bayes factors.

Table 2: Predictive performance measures in the empirical study.

Model	Log predictive density	Log predictive Bayes factor against . . .		
		MS-ARX	ARX-SV	ARX
MS-ARX-SV	–4571.9895	+2242.4402	+1565.7482	–1219.1797
MS-ARX	–2329.5493		–676.6921	–3461.6199
ARX-SV	–3006.2413			–2784.9278
ARX	–5791.1691			

4.1 Full-sample results

As a preliminary check, we run the classical neural network test for neglected nonlinearity introduced by Lee et al. (1993) on the ARX model in each of our 363 estimation windows. The null hypothesis of linearity is rejected in the vast majority of cases, as was to be expected based on the literature surveyed in Section 1. Specifically, nonrejection at the 5% level occurs on only 39 days, all of which are in the last 2.5 months of the year. Moving to the 10% level, only one nonrejection remains. We conclude that the in-sample evidence is strongly in favor of nonlinear modeling.

Below we present out-of-sample forecasting results obtained for the four model specifications in Table 1, estimated for hourly electricity prices, for every hour of 2015.[2] In each case posterior analysis is based on 15000 MCMC samples from the relevant joint posterior, preceded by 5000 burnin draws. Calculations have been carried out with the authors' own codes run under Matlab. MCMC convergence is deemed satisfactory, as measured using standardized CUSUM plots (see Yu and Mykland, 1998).

Predictive density values and predictive Bayes factors, the relevant quantities for model comparison, are displayed in Table 2. It is clear that all three nonlinear models provide a better out-of-sample performance than the simple linear ARX model. To quantify the differences, we follow the interpretation suggested by Kass and Raftery (1995): a Bayes factor greater than three provides "positive" evidence of the outperformance, and the evidence is "strong" for Bayes factors greater than twenty and "very strong" beyond 150. Thus, there is very strong evidence that MS-ARX-SV outperforms the ARX benchmark, since $\exp(1219.1797) \approx 3 \times 10^{529}$. In fact, the same could be said for every pairwise model comparison in this table; the evidence in favor of both stochastic volatility (ARX-SV versus ARX) and especially regime switching (MS-ARX versus ARX) is very strong. We conclude that the forecasting performance of the MS-ARX, ARX-SV, and MS-ARX-SV models is much better than that of the ARX model.

[2] A preliminary analysis of the results did not reveal any obvious daily, weekly, or annual patterns in model performance. For this reason, only aggregate results are reported here.

Table 3: Predictive performance measures in the empirical study, Monday 12 January 2015.

Model	Log predictive density	Log predictive Bayes factor against . . . MS-ARX	ARX-SV	ARX
MS-ARX-SV	–62.7767	+33.2098	+27.9920	+38.0507
MS-ARX	–29.5669		–5.2178	+4.8409
ARX-SV	–34.7847			+10.0587
ARX	–24.7260			

However, our proposed MS-ARX-SV model is outperformed by each of the simpler nonlinear MS-ARX and ARX-SV models, which are special cases of it. Our interpretation of this result is that regime switching and stochastic volatility are both good ideas for modeling the electricity prices, since each of these features individually strongly enhances the performance of a pure ARX model. Joining both of these ideas, on the other hand, still requires some further investigation. Perhaps our highly parametrized specification is simply asking to much from the data; perhaps three regimes are not needed for this particular data set and two would be sufficient. We are currently in the process of assessing these issues.

4.2 Subsample results

In order to better understand the differences in performance between the four models, as well as to highlight the ease of obtaining predictive densities within our framework, we repeat the analysis, this time restricting the out-of-sample period to four specific days. These days were selected to be the ones where each model performed best relative to its competitors: Monday 12 January (ARX performed best), Sunday 18 January (MS-ARX-SV), Sunday 10 May (MS-ARX), and Tuesday 25 August (ARX-SV).

Tables 3–6 below are analogous to the full-sample Table 2. We observe that, except in the special case where ARX performs best (Table 3), this simple linear benchmark performs far worse than all nonlinear competitors (Tables 4–6). When our full MS-ARX-SV model performs best (Table 4), it does so by a very wide margin; note that the smallest Bayes factor is $\exp(12.3544) \approx 2 \times 10^5$ already.

Table 4: Predictive performance measures in the empirical study, Sunday 18 January 2015.

Model	Log predictive density	Log predictive Bayes factor against . . . MS-ARX	ARX-SV	ARX
MS-ARX-SV	–1.2465	–12.3544	–15.7705	–20.9411
MS-ARX	–13.6009		–3.4161	–8.5867
ARX-SV	–17.0169			–5.1706
ARX	–22.1875			

Table 5: Predictive performance measures in the empirical study, Sunday 10 May 2015.

		Log predictive Bayes factor against . . .		
Model	Log predictive density	MS-ARX	ARX-SV	ARX
MS-ARX-SV	–83.2826	+39.4483	–7.5094	–84.2646
MS-ARX	–43.8343		–46.9577	–123.7129
ARX-SV	–90.7920			–76.7552
ARX	–167.5473			

This leaves us with the intermediate cases to analyze, where either Markov switching turned out to be useful for forecasting but stochastic volatility did not, or vice versa. Table 5 presents a case in which MS-ARX strongly outperforms all other models, and we observe that the MS-ARX-SV model is still "the best of the rest". That is, the full, highly-parametric model is preferred over the ARX-SV model, which gets the nature of the nonlinearity wrong in this instance.

In the opposite case (Table 6), where ARX-SV is the preferred model, the performance of the other nonlinear models MS-ARX-SV and MS-ARX is virtually indistinguishable, with a predictive Bayes factor of $\exp(0.8272) \approx 2$ between them. These results confirm the intuition that we gained based on the full-sample results: leaving out Markov switching when we need it has a larger negative impact on the quality of our forecasts than leaving out stochastic volatility when we need it.

To further investigate what drives these results, Figures 3–6 show the predictive densities obtained from each model for two selected hours on these days, one in the afternoon ($h = 15$, referring to 2–3 PM) and one in the evening ($h = 21$, referring to 8–9 PM). The ex-post realized price is also included in each of the figures.

Firstly, Monday 12 January 2015 (illustrated in Figure 3) was a relatively uneventful day, for which the linear ARX model was "good enough" and all of its nonlinear extensions turned out to be needless complications. The full-sample results presented in Table 2 indicate that such days must be quite rare in our sample.

Table 6: Predictive performance measures in the empirical study, Tuesday 25 August 2015.

		Log predictive Bayes factor against . . .		
Model	Log predictive density	MS-ARX	ARX-SV	ARX
MS-ARX-SV	–11.0719	+0.8272	+22.3026	–2.8123
MS-ARX	–10.2447		+21.4755	–3.6395
ARX-SV	+11.2308			–25.1150
ARX	–13.8842			

Figure 3: Predictive densities for two selected hours on Monday 12 January 2015. Left: $h = 15$. Right: $h = 21$.

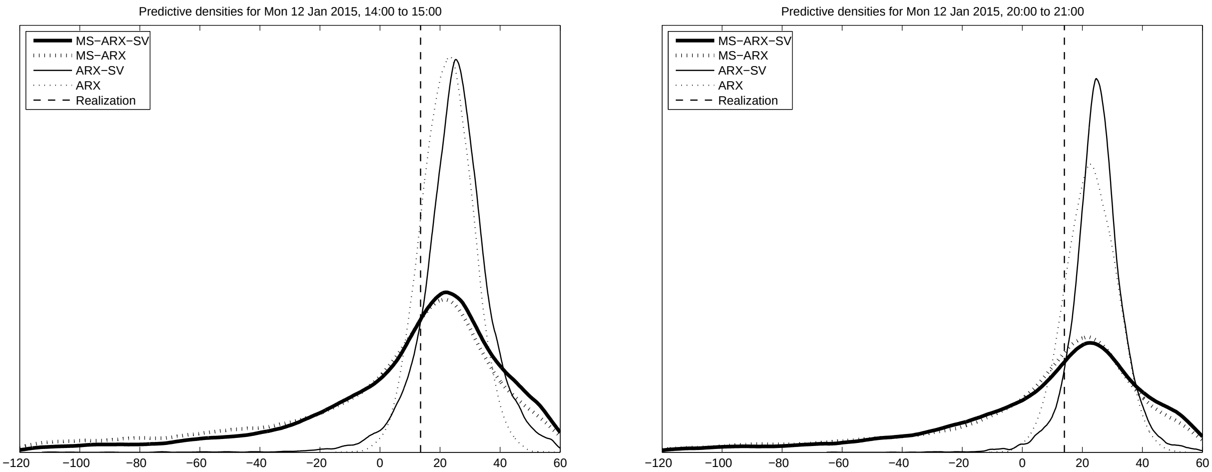

The MS-ARX-SV model performed best on Sunday 18 January (Figure 4). It appears that the models without a Markov switching component provide forecasts that are centered at the wrong location on this day. Clearly, allowing for multiple regimes provides a safeguard against such problems. The price fluctuated considerably on this day, a fact that is picked up by the relatively flat predictive densities produced by the stochastic volatility models.

A large negative jump occurred in the afternoon of Sunday 10 May (Figure 5), which explains the good performance of MS models on this date. Finally, Tuesday 25 August (Figure 6) saw large price fluctuations but no jumps, so that SV turned out to be an important model component on that day.

Figure 4: Predictive densities for two selected hours on Sunday 18 January 2015. Left: $h = 15$. Right: $h = 21$.

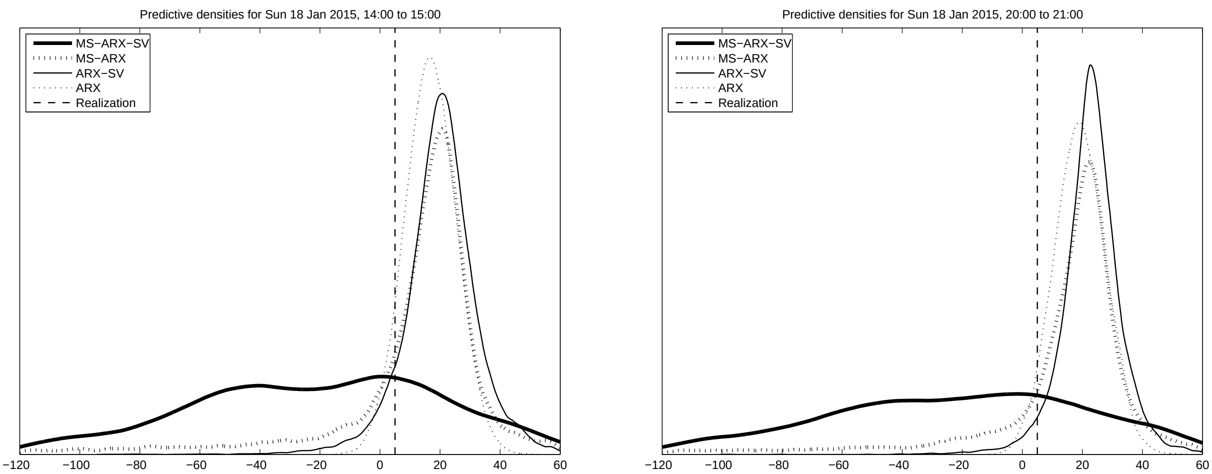

Figure 5: Predictive densities for two selected hours on Sunday 10 May 2015. Left: $h = 15$. Right: $h = 21$.

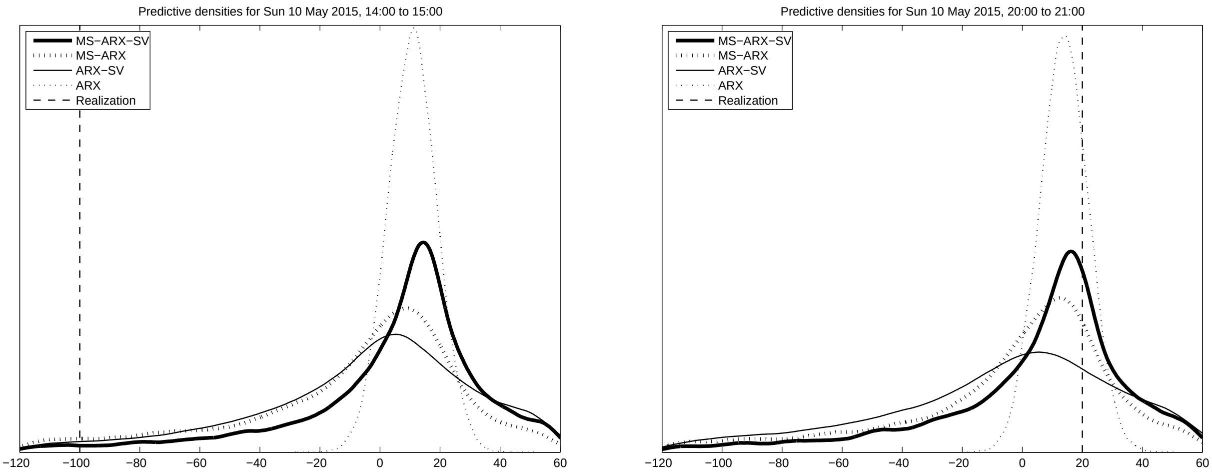

5 Conclusions

This study addresses several fundamental questions in the area of forecasting electricity prices. We propose a new regime-switching stochastic volatility model with three regimes, taking fundamental price drivers into account. We show how the predictive densities of future electricity prices can be constructed via Bayesian inference. Moreover, we introduce a universal method for model comparisons, predictive Bayes factors, to the electricity prices literature. Using this measure, we show that our model's Markov switching and stochastic volatility components both contribute to its improved performance at short-time density forecasting in the Nord Pool intraday market, relative to a model that lacks such features.

Figure 6: Predictive densities for two selected hours on Tuesday 25 August 2015. Left: $h = 15$. Right: $h = 21$.

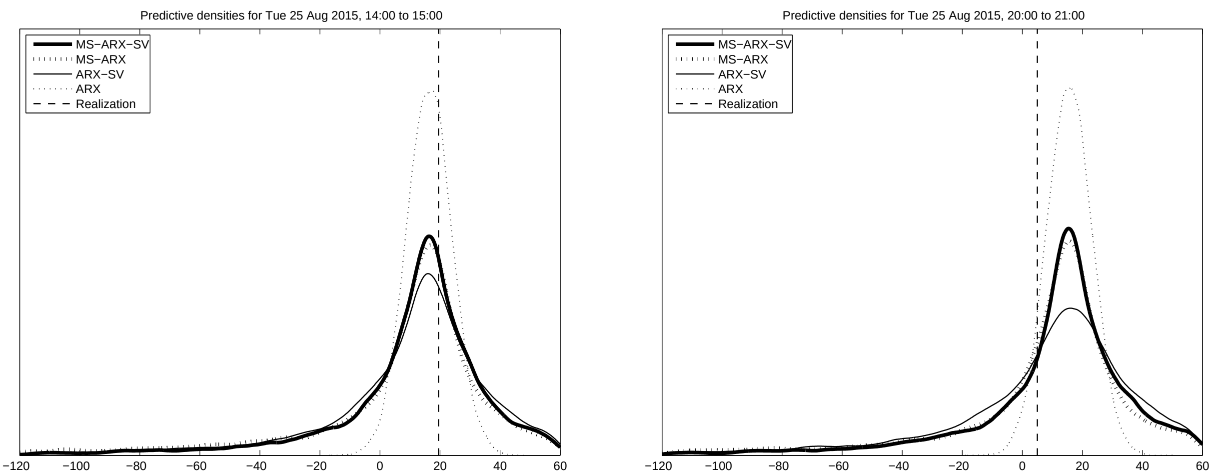

Both Markov switching models and stochastic volatility models provide very good forecasts on some occasions but poor ones on some others, and our subsample analysis suggests that there may be a complementarity between these two features. Since our MS-ARX-SV model nests both types of models, it strikes us as a useful contribution to the literature. However, more research is needed in order to obtain a desirable empirical performance from this rich model. One possible way to reduce the dimensionality of its parameter space would be to simplify the volatility dynamics, e.g. $\gamma_{R_t} = \gamma$ for each regime.

Another avenue, which appears more promising in our view, is to go back to models with two rather than three regimes. Most evidence in favor of the existence of a third regime (Karakatsani and Bunn, 2008; Janczura and Weron, 2010a) is several years old by now. As integrated energy markets have matured, a "spike" regime may no longer be necessary to describe the dynamics in electricity prices. Tentatively, visual inspection of Figure 1 confirms this intuition, but a more thorough investigation of this issue is warranted.

References

J.H. Albert and S. Chib. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88:669–679, 1993.

J. Andreasen and M. Dahlgren. At the flick of a switch. *Energy Risk*, 10:71–75, 2006.

T. Arvesen, V. Medbø, S.-E. Fleten, A. Tomasgard, and S. Westgaard. Linepack storage valuation under price uncertainty. *Energy*, 52:155–164, 2013.

M. Bierbrauer, S. Trück, and R. Weron. Modeling electricity prices with regime switching models. In M. Bubak, G.D. van Albada, P. Sloot, and J. Dongarra, editors, *Proceedings of the Fourth International Conference on Computational Science*, pages 859–867. Springer, Berlin/Heidelberg, 2004.

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

A. Botterud, A.K. Bhattacharyya, and M. Ilic. Futures and spot prices: An analysis of the Scandinavian electricity market. In *Proceedings of the North American Power Symposium*, Tempe, 2002.

P. Damien, J. Wakefield, and S. Walker. Gibbs sampling for Bayesian non-conjugate and hierarchical models by using auxiliary variables. *Journal of the Royal Statistical Society: Series B*, 61:331–344, 1999.

C. de Jong. The nature of power spikes: A regime-switch approach. *Studies in Nonlinear Dynamics and Econometrics*, 10:27–46, 2006.

P. de Jong and N. Shephard. The simulation smoother for time series models. *Biometrika*, 82:339–350, 1995.

S.J. Deng. Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes. *POWER working paper*, University of California Energy Institute, Berkeley, 1999.

M. Eichler and D. Türk. Fitting semiparametric Markov regime-switching models to electricity spot prices. *Energy Economics*, 36:614–624, 2013.

R. Ethier and T. Mount. Estimating the volatility of spot prices in restructured electricity markets and the implications for option values. *Working paper*, Cornell University, Ithaca, 1998.

A.J. Filardo and S.F. Gordon. Business cycle durations. *Journal of Econometrics*, 85:99–123, 1998.

F.R. Førsund and M. Hoel. Properties of a non-competitive electricity market dominated by hydroelectric power. *Working paper*, University of Oslo, 2004.

H. Geman and A. Roncoroni. Understanding the fine structure of electricity prices. *Journal of Business*, 79:1225–1261, 2006.

J. Geweke and G. Amisano. Comparing and evaluating Bayesian predictive distributions of asset returns. *International Journal of Forecasting*, 26:216–230, 2010.

O. Gjolberg and T. Johnsen. Electricity futures: Inventories and price relationships at Nord Pool. *Working paper*, Norwegian School of Economics and Business Administration, Bergen, 2001.

N. Haldrup and M.Ø. Nielsen. A regime switching long memory model for electricity prices. *Journal of Econometrics*, 135:349–376, 2006a.

N. Haldrup and M.Ø. Nielsen. Directional congestion and regime switching in a long memory model for electricity prices. *Studies in Nonlinear Dynamics and Econometrics*, 10:1–24, 2006b.

R. Handika, C. Truong, S. Trück, and R. Weron. Modelling price spikes in electricity markets: The impact of load, weather and capacity. *Working paper*, Hugo Steinhaus Center, Wrocław, 2014.

S. Heydari and A. Siddiqui. Valuing a gas-fired power plant: A comparison of ordinary linear models, regime-switching approaches, and models with stochastic volatility. *Energy Economics*, 32:709–725, 2010.

T. Hong, P. Pinson, S. Fan, H. Zareipour, A. Troccoli, and R.J. Hyndman. Probabilistic energy forecasting: Global Energy Forecasting Competition 2014 and beyond. *International Journal of Forecasting*, 32:896–913, 2016.

R. Huisman. *An introduction to models for the energy markets*. Risk Books, London, 2009.

R. Huisman and C. de Jong. Option pricing for power prices with spikes. *Energy Power Risk Management*, 7:12–16, 2003.

R. Huisman and R. Mahieu. Regime jumps in electricity prices. *Energy Economics*, 25:425–434, 2003.

C. Huurman, F. Ravazzolo, and C. Zhou. The power of weather. *Computational Statistics and Data Analysis*, 56:3793–3807, 2012.

E. Jacquier, N.G. Polson, and P.E. Rossi. Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics*, 20:69–87, 2002.

J. Janczura and R. Weron. Regime-switching models for electricity spot prices: Introducing heteroskedastic base regime dynamics and shifted spike distributions. In *IEEE Proceedings of the International Conference on the European Energy Market*, Leuven, 2009.

J. Janczura and R. Weron. An empirical comparison of alternate regime-switching models for electricity spot prices. *Energy Economics*, 32:1059–1073, 2010a.

J. Janczura and R. Weron. Goodness-of-fit testing for regime-switching models. *Working paper*, MPRA, München, 2010b.

J. Janczura, S. Trück, R. Weron, and R.C. Wolff. Identifying spikes and seasonal components in electricity spot price data: A guide to robust modeling. *Energy Economics*, 38:96–110, 2013.

T. Jónsson, P. Pinson, H. Madsen, and H.A. Nielsen. Predictive densities for day-ahead electricity prices using time-adaptive quantile regression. *Energies*, 7:5523–5547, 2014.

T. Kanamura and K. Ohashi. On transition probabilities of regime switching in electricity prices. *Energy Economics*, 30:1158–1172, 2008.

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

R.E. Kass and A.E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90:773–795, 1995.

D. Keles, R. Hartel, D. Möst, and W. Fichtner. 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*, 5:53–84, 2012.

G. Koop. *Bayesian Econometrics*. Wiley, Chichester, 2003.

P. Kosater and K. Mosler. Can Markov regime-switching models improve power-price forecasts? Evidence from German daily power prices. *Applied Energy*, 83:943–958, 2006.

T.-H. Lee, H. White, and C.W.J. Granger. Testing for neglected nonlinearity in time series models. *Journal of Econometrics*, 56:269–290, 1993.

X. Li and D.J. Sailor. Electricity use sensitivity to climate and climate change. *World Resource Review*, 7:334–346, 1995.

D. Liebl. Modeling and forecasting electricity spot prices: A functional data perspective. *Annals of Applied Statistics*, 7:1565–1592, 2013.

F. Lisi and F. Nan. Component estimation for electricity prices: Procedures and comparisons. *Energy Economics*, 44:143–159, 2014.

K. Maciejowska, J. Nowotarski, and R. Weron. Probabilistic forecasting of electricity spot prices using factor quantile regression averaging. *International Journal of Forecasting*, 32:957–965, 2016.

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

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

Nord Pool. `http://www.nordpoolspot.com/`, accessed 15 July, 2016.

A. Panagiotelis and M. Smith. Bayesian density forecasting of intraday electricity prices using multivariate skew t distributions. *International Journal of Forecasting*, 24:710–727, 2008.

J. Peirson and A. Henley. Electricity load and temperature: Issues in dynamic specification. *Energy Economics*, 16:235–243, 1994.

J. Peirson and A. Henley. Residential energy demand and the interaction of price and temperature: British experimental evidence. *Energy Economics*, 20:157–171, 1998.

D.J. Sailor, J.N. Rosen, and J.R. Muñoz. Natural gas consumption and climate: A comprehensive set of predictive state-level models for the United States. *Energy*, 23:91–103, 1998.

M. Serati, M. Manera, and M. Plotegher. Modeling electricity prices: From the state of the art to a draft of a new proposal. *Working paper*, Fondazione Eni Enrico Mattei, Milan, 2008.

F. Serinaldi. Distributional modeling and short-term forecasting of electricity prices by generalized additive models for location, scale and shape. *Energy Economics*, 33:1216–1226, 2011.

M.S. Smith. Bayesian inference for a periodic stochastic volatility model of intraday electricity prices. In T. Kneib and G. Tutz, editors, *Statistical Modelling and Regression Structures: Festschrift in Honour of Ludwig Fahrmeir*, pages 353–376. Physica-Verlag, Heidelberg, 2010.

N.-H.M. von der Fehr, E.S. Amundsen, and L. Bergman. The Nordic market: Signs of stress? *The Energy Journal*, 26:71–98, 2005.

R. Weron. Heavy-tails and regime-switching in electricity prices. *Mathematical Methods of Operations Research*, 69:457–473, 2009.

R. Weron. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30:1030–1081, 2014.

R. Weron, M. Bierbrauer, and S. Trück. Modeling electricity prices: Jump diffusion and regime switching. *Physica A*, 336:39–48, 2004.

B. Yu and P. Mykland. Looking at Markov samplers through CUSUM path plots: A simple diagnostic idea. *Statistics and Computing*, 8:275–286, 1998.