2017-4
Oskar Knapik
PhD Dissertation

Essays on econometric modelling and forecasting of electricity prices

DEPARTMENT OF ECONOMICS AND BUSINESS ECONOMICS
AARHUS BSS • AARHUS UNIVERSITY • DENMARK

AARHUS BSS
SCHOOL OF BUSINESS AND SOCIAL SCIENCES
AARHUS UNIVERSITY

Essays on econometric modeling and forecasting of electricity prices

PhD Dissertation

Oskar Knapik

SOLIDUM PETIT IN PROFUNDIS · UNIVERSITAS ARHUSIENSIS ·

CREATES
Center for Research in Econometric Analysis of Time Series

AARHUS BSS
SCHOOL OF BUSINESS AND SOCIAL SCIENCES
AARHUS UNIVERSITY

Aarhus BSS, Aarhus University
Department of Economics and Business Economics

2016

PREFACE

This PhD dissertation was written between September 2013 and August 2016, during my PhD studies in the Department of Economics and Business Economics at Aarhus University. During my doctoral studies, I was affiliated with the Center for Research in Econometric Analysis of Time Series (CREATES), funded by the Danish National Research Foundation. I am grateful the Department of Economics and Business Economics, Aarhus University and CREATES for providing me with inspiring and excellent research environments and financial support throughout my studies. This enabled me to attend numerous courses, workshops and conferences both nationally and internationally.

I would like to take this opportunity to thank a number of people. First and foremost, I want to thank my main supervisor and co-author of one of the chapters in this dissertation, Professor Niels Haldrup, for his generous support and great advice, and for always having his door open and welcoming my questions and discussions. I sincerely appreciate his interesting ideas and comments during the research which is included in this dissertation. Working with him has been an inspiring and great learning experience. I would also like to thank to my co-authors Professor Tommaso Proietti and Peter Exterkate for their many valuable discussions and insightful comments. And I also want to thank Professor Jacek Leskow for being my friend and scientific mentor. I am very grateful for him introducing me to his friends and colleagues in the international scientific environment. I truly appreciate his friendship and support.

I am thankful for all my fellow PhD students and other colleagues in the Department of Economics and Business Economics, and for the friendly and enjoyable atmosphere we enjoyed together. One of them in particular, Carlos Vladimir Rodríguez Caballero, deserves special thanks, for the support and friendship he has given me over the last few years.

I am grateful to Solveig Nygaard Sørensen, CREATES' administrator, for all the practical help and conversations in the last three years. Administrative assistant to the head of the PhD unit at Aarhus University, Susanne Christensen, also deserves a special thank you for always being there for practical help in navigating my PhD program.

Finally, I would like to thank my family for all their love and encouragement.

Oskar Knapik
Aarhus, August 2016

UPDATED PREFACE

The predefence took place on October 24, 2016. I am very grateful to the members of the assessment committee Alessandra Luati (Università di Bologna), Luigi Grossi (Università degli Studi di Verona) and Svend Hylleberg (Aarhus University and CREATES) for their careful reading of my dissertation and their constructive comments and suggestions. Most of the suggestions have been incorporated in the current version of the dissertation while others remain for future research.

O. Knapik
Aarhus, January 2017

Contents

SUMMARY

This thesis comprises three self-contained chapters on the econometric modeling and forecasting of electricity prices in the Nord Pool power market. They present investigations of the hot topics in electricity market research today, with a focus on day-ahead forecasting of system electricity prices in the Elspot power market, main price drivers, and forecasting price jumps, as well as forecasting electricity price densities in the intraday power market of Elbas. These chapters also attempt to address three crucial and outstanding challenges in the area of electricity price forecasting, as pinpointed in the recent review paper by Weron (2014), 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) proposing a universal technique for model comparison.

Electricity prices are very complex data. They exhibit strong seasonality at the annual, weekly and daily levels, high volatility, and abrupt, short-lived and generally unanticipated extreme price changes known as spikes or jumps (see Janczura, Trueck, Weron, and Wolff, 2013; Weron, 2008; Serati, Manera, and Plotegher, 2008, among others). Characteristics of spot electricity prices also include mean-reversion and, for some markets, (seasonal) long memory (see e.g. Haldrup and Nielsen (2006a,b)). The unique characteristics of electricity prices and related derivatives contracts have boosted the demand for econometric models that can precisely capture their dynamics. In terms of model specifications, there are three broad strands in the electricity price modeling literature, namely traditional autoregressive time series models, nonlinear time series models (with a particular emphasis on Markov-switching models), and continuous-time diffusion or jump-diffusion models (see Weron (2014) for a recent review). This thesis contributes to this literature by proposing new models and addressing fundamental research questions.

The first chapter of this thesis, entitled 'A generalized exponential time series regression model for electricity prices, focuses on modeling and forecasting of daily electricity spot prices on the Nord Pool Elspot power market', proposes a new model that allows for seasonal and non-seasonal persistence, in the form of a generalized exponential (GEXP) process for electricity prices. As the presence of spikes can distort estimation of the dynamic structure of the series, we propose an iterative estimation strategy which, conditionally on a set of parameter estimates, clears the spikes using a

data cleaning algorithm, and re-estimates the parameters using the cleaned data, making the estimates more robust. Using the estimated model, the best linear predictor is then constructed. The introduced modeling approach provides a good fit within sample and outperforms competing benchmark predictors in terms of forecasting accuracy. We also find that building separate models for each hour of the day and averaging the forecasts is not a better strategy than forecasting the daily average directly.

Chapter two of the thesis, 'Modeling and forecasting electricity price jumps in the Nord Pool power market', investigates the influence of fundamental price drivers on the forecasting of price jumps in the Nord Pool intraday market. In this chapter, we propose three categorical time series models which take into account i) price drivers, ii) persistence, and iii) seasonality of electricity prices. We show that the models outperform commonly-used benchmark. The chapter shows how crucial forecasting is for price jumps, in order to incorporate additional factors that influence price drivers like loads, temperature and water reservoir level, as well as to take into account the persistence.

The last chapter of the thesis, 'A regime-switching stochastic volatility model for forecasting electricity prices', focuses on crucial challenges in the area of electricity price forecasting. This chapter addresses these challenges 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 a universal technique for model comparison. We propose a new regime-switching stochastic volatility model with three regimes (negative jump, normal price, and positive jump) where the transition matrix depends on explanatory variables controlled by ordered probit. Bayesian inference is explored in order to obtain predictive densities. The main focus of the paper is on short time density forecasting in the Nord Pool intraday market. We propose a universal technique for model comparison in terms of predictability, namely predictive Bayes factors. Based on this method, we determined that our proposed model outperforms competitors.

Danish summary

Denne afhandling består af tre selvstændige artikler, der omhandler økonometrisk modellering og forecasting af elektricitetspriser på NordPool elektricitetsmarkedet. Artiklerne beskæftiger sig med forskning inden for elektricitetsmarkedet med fokus på day-ahead forecasting af system elektricitetspriser på Elspot elektricitetsmarkedet, de vigtigste faktorer der påvirker priserne samt forecasting af pris-jumps og forecasting af elektricitetspris-densiteter på intra-day elektricitetsmarkedet Elbas. Artiklerne adresserer også tre vigtige udfordringer i forhold til forecasting af elektricitetspriser, som også belyses i en nylig artikel af Weron (2014), idet vi i) viser betydningen af at tage de grundlæggende faktorer, der påvirker priserne i forbindelse med modellering, i betragtning, ii) udvikler nye teknikker for probalistisk (i.e. interval eller densitet) forecasting af elektricitetspriser, iii) fremsætter en universel teknik for modelsammenligning.

Elektricitetspriser må karakteriseres som meget kompleks data. De er således i høj grad sæsonbestemte på års-, uge- og dags-niveau, og er kendetegnet ved meget høj volatilitet og pludselige, kortvarige og generelt uventede ekstreme prisændringer, også benævnt spikes (eller jumps) (se blandt andre Janczura et al. (2013); Weron (2008), Serati et al. (2008)). Spotpriser på elektricitet er også karakteriseret ved mean-reversion og for nogle markeder for (sæsonbestemt) lang hukommelse (se f.eks. Haldrup and Nielsen (2006a,b)). De unikke karakteristika ved elektricitetspriser og relaterede derivative kontrakter har skabt en efterspørgsel efter økonometriske modeller, der præcist kan opfange deres dynamik. I relation til modelspecifikation er der tre brede strenge inden for litteraturen, der behandler modellering af elektricitetspriser, nemlig traditionelle autoregressive tidsseriemodeller, non-lineære tidsseriemodeller (med vægt på Markov-switching modeller), og continuous-time diffusion eller jump-diffusion modeller (se Weron (2014) for en nylig gennemgang). Artiklerne i denne afhandling bidrager til litteraturen på forskellig vis ved at fremsætte nye modeller og adressere fundamentale forskningsspørgsmål.

Afhandlingens første artikel, ”A generalized exponential time series regression model for electricity prices”, fokuserer på modellering og forecasting af de daglige elektricitetsspotpriser på Nord Pool Elspot elektricitetsmarkedet. En ny model, der muliggør sæsonbestemt og ikke-sæsonbestemt persistens, fremsættes i form af en generalized exponential (GEXP) proces for elektricitetspriser. Idet tilstedeværelsen af spikes kan forvrænge estimationen af seriens dynamiske struktur, fremsætter vi en

iterativ estimationsstrategi, der, betinget af et sæt parameter-estimater, renser spikes'ene ved hjælp af en data rensningsalgoritme, og re-estimerer parametrene ved hjælp af de rensede data for derved at robustgøre estimaterne. Afhængig af den estimerede model konstrueres den bedste lineære predictor. Den introducerede modelleringstilgang er bedre end konkurrerende benchmark predictors i relation til forecastingens nøjagtighed. Vi viser, at det at lave separate modeller for hver time og beregne gennemsnittet af forecast'ene er ikke en bedre strategi end at forecaste det daglige gennemsnit direkte. Afhandlingens anden artikel, "Modeling and forecasting electricity price jumps in the Nord Pool power market", undersøger, hvilken indflydelse fundamentale faktorer, der påvirker priserne, har på forecastingen af pris jumps på NordPool intra-day markedet. Vi fremsætter tre kategoriske tidsseriemodeller, der tager højde for i) faktorer, der påvirker priserne, ii) persistens, iii) sæsonbestemte elektricitetspriser. Vi viser, at disse modeller er bedre end normalt-anvendte benchmarks. Artiklen viser, hvor vigtig forecasting er for pris-jumps i relation til at inkorporere mere viden om faktorer, der påvirker priserne, såsom belastning, temperatur og vandreservoir niveau og i relation til at tage persistensen i jumps-processen i betragtning.

Afhandlingens sidste kapitel, "A regime-switching stochastic volatility model for forecasting electricity prices", omhandler de udfordringer, der opstår i forbindelse forecasting af elektricitetspriser. Artiklen forsøger at komme omkring dem alle ved at i) vise vigtigheden af at tage grundlæggende faktorer, der påvirker priserne i forbindelse med modellering, i betragtning, ii) udvikle nye teknikker for probalistisk (interval eller densitet) forecasting af elektricitetspriser, iii) introducere en universel teknik for modelsammenligning. Vi fremsætter en ny regime-skiftende stokastisk volatilitetsmodel med tre regimer (negativ jump, normal pris og positiv jump (spike), hvor transitionsmatrixen afhænger af forklarende variabler kontrolleret af "ordered probit"). Der anvendes Bayesiansk inferens for at opnå prædiktive densiteter. Artiklen fokuserer hovedsageligt på korttids-densitets-forecasting på NordPool intra-day markedet. Vi anvender en universel teknik for sammenligning af modeller i relation til forudsigelighed, nemlig "predictive Bayes factors". På baggrund af denne metode viser vi, at den fremsatte model er bedre end konkurrerende modeller.

Bibliography

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

Haldrup, N., Nielsen, M., 2006b. A regime switching long memory model for electricity prices. Journal of Econometrics 135, 349–376.

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

Serati, M., Manera, M., Plotegher, M., 2008. Modeling electricity prices: from the state of the art to a draft of a new proposal. Tech. rep., Liuc papers in Economics, Cattaneo University (LIUC).

Weron, R., 2008. Market price of risk implied by asian-style electricity options and futures. Energy Economics 30, 1098–1115.

Weron, R., 2014. Electricity price forecasting: A review of the state-of-the-art with a look into the future. International Journal of Forecasting 30 (4), 1030–1081.

CHAPTER 1

A GENERALIZED EXPONENTIAL TIME SERIES REGRESSION MODEL FOR ELECTRICITY PRICES

Niels Haldrup
Aarhus University and CREATES
Oskar Knapik
Aarhus University and CREATES
Tommaso Proietti
Università di Roma "Tor Vergata" and CREATES

Abstract

We consider the modeling and forecasting of daily electricity spot prices on the Nord Pool Elspot power market. A method that can handle seasonal and non-seasonal persistence is proposed in the form of a generalized exponential (GEXP) process for the power prices. As the presence of spikes can distort estimation of the dynamic structure of the series, we consider an iterative estimation strategy which, conditional on a set of parameter estimates, clears the spikes using a data cleaning algorithm, and reestimates the parameters using the cleaned data so as to robustify the estimates. Conditional on the estimated model, the best linear predictor is constructed. Our modeling approach provides good fit within sample and outperforms many competing benchmark predictors in terms of forecasting accuracy. We also find that building separate models for prices of each hour of the day and averaging the forecasts is not always a better strategy than forecasting the daily average price directly.

1.1 Introduction

The daily spot prices from the Nord Pool power market (NordPool, 2016) exhibit persistent (long memory) features combined with periodic behavior related to weekly, monthly and yearly periodicity, which also appears to be rather persistent. The fractional noise model (see Granger and Joyeux (1980) and Hosking (1981)) is suitable for capturing the persistence of the series at the long-run frequency. Periodic patterns reverting slowly to deterministic cycles can be modeled by the class of generalized fractional, or Gegenbauer, processes, introduced by Hosking (1981) in his seminal paper and analyzed by Gray, Zhang, and Woodward (1989). After applying a cascade of (generalized) fractional filters, a popular approach is to assume that the filtered series is a short-memory autoregressive moving average (ARMA) process.

These processes have been used widely in many empirical applications. Ferrara and Guegan (2001) used Gegenbauer processes to analyze cointegration in the Nikkei spot index. Smallwood and Beaumont (2003) analyzed IBM trading volumes. Woodward, Cheng, and Gray (1998) considered atmospheric CO_2 data. There is also evidence of fractional seasonal integration in macroeconomic time series, see, for example Arteche and Robinson (2000) and Gil-Alaña and Robinson (2001).

Gegenbauer processes have also been used to study energy market data. Soares and Souza (2006) used Gegenbauer ARMA (GEGARMA) processes with explanatory variables to forecast electricity demand. Diongue, Guégan, and Vignal (2004) introduced the k-factor Gegenbauer integrated generalized autoregressive conditional heteroskedasticity (GEGARCH) model for modeling energy prices and, subsequently, Diongue, Guégan, and Vignal (2009) considered the model for forecasting EEX electricity spot market prices in Germany.

An alternative semiparametric approach is based on the generalized exponential model for the spectrum, as in Hsu and Tsai (2009). According to this approach, the logarithm of the spectrum of the short-memory filtered series is represented by a finite trigonometric polynomial. Bloomfield (1973) introduced the exponential (EXP) model and discussed its maximum likelihood estimation, relying on the distributional properties of the periodogram of a short-memory process, based on Walker (1964) and Whittle (1953). The model was then generalized to processes featuring long-range dependence by Robinson (1991) and J.Beran (1993), originating the fractional EXP model (FEXP). Maximum likelihood estimation of the FEXP model has been dealt with by Narukawa and Matsuda (2011) and Hurvich (2002) addressed the issue of forecasting with this model.

Electricity spot prices are very complex data. They exhibit strong seasonality at the annual, weekly and daily frequency and very high volatility and abrupt, short-lived and generally unanticipated extreme price changes, known as spikes (or jumps); (see Weron, 2006; Serati, Manera, and Plotegher, 2008; Janczura, Trück, Weron, and Wolf, 2013; de Jong, 2006, among others). Characteristics of spot electricity prices also include mean reversion and (seasonal) long memory; (see Haldrup and Nielsen, 2006a,b; Weron, 2008; Diongue et al., 2004, 2009; Koopman, Ooms, and Carnero, 2007, among others).

The unique characteristics of electricity prices and related derivatives contracts have boosted the demand for econometric models that can appropriately describe their dynamics (see (Benth and Koekebakker, 2008; Möst and Keles, 2010; Raviv, Bouwman, and Dijk, 2015)). Electricity price models are of great importance for forecasting, derivative pricing, and risk management. The recent reviews on forecasting electricity prices by Zareipour (2012) and Weron (2014) document the huge interest in the topic.

In electricity markets such as the Nord Pool Spot system, hourly prices are determined in a day-ahead market in which the daily average can be calculated. Most models are built for daily average prices, which play a key role in the electricity market. For instance, the average daily price is widely used to approximate other spot electricity prices and is used as a reference price for derivatives contracts, e.g. futures and forwards.

The paper contributes to the above literature by proposing a time series regression model relating prices to a set of explanatory variables with errors that follow a generalized exponential model. The model accounts for the most prominent features, such as long-range dependence at "the zero frequency" and the weekly cycle and its harmonics. Differently from related research on Nord Pool daily electricity spot price modeling (see Weron, Simonsen, and Wilman (2004)), we build our model directly for the levels, rather than for the first differences.

Estimation is carried out in the frequency domain by maximizing the Whittle likelihood, which is an asymptotic approximation to the likelihood function. The paper also introduces a new robust estimation procedure for dealing with price spikes, based on a robustification of the Kalman filter, which takes into account possible long memory in the data. Our proposal is validated by conducting a rolling forecasting exercise which compares the daily average electricity spot price forecasts arising from our model with several competitors.

The article is organized as follows. Section 1.2 introduces the generalized exponential model and presents the concept of generalized long memory and the relevant theory on Gegenbauer processes. Section 1.3 discusses robust inference and forecasting with this model. Section 1.4 presents empirical results for modeling and forecasting daily average electricity spot prices from the Nord Pool power market and, finally, Section 1.5 concludes.

1.2 The generalized exponential model

The crucial step in constructing a model for electricity price dynamics consists of finding an appropriate description of the seasonal pattern. There are different suggestions for that task in the literature. In order to account for seasonality we build our model on the Gegenbauer processes as in Hsu and Tsai (2009). We extend their approach by proposing a time series regression model with generalized FEXP disturbances, which is particularly well suited to accounting for most features characterizing electricity prices and allows to incorporate additional explanatory variables as covariates in the analysis. Moreover, we consider four Gegenbauer factors. Our model is similar to the time series regression model based on Gegenbauer ARMA (GEGARMA) processes proposed by Soares and Souza (2006) for electricity demand prediction, but it differs for the representation of the short memory components and we take a flexible semiparametric approach. Contrary to the work of Soares and Souza (2006), our approach allows for both nonstationarity and long-memory of the modeled time series.

Model specification

Let y_t denote a daily time series referring to the electricity prices of a particular hour of the day, or the daily average. We consider the following time series regression model:

$$y_t = x_t'\beta + u_t, \tag{1.1}$$

where x_t is a $k \times 1$ vector of explanatory variables and u_t is a zero mean random process. Denoting by B the lag operator, u_t is generated by the following fractionally integrated process

$$(1-B)^{d_0} \prod_{j=1}^{3} (1 - 2\cos\varpi_j B + B^2)^{d_j} u_t = z_t, \quad \varpi_j = \frac{2\pi}{7} j, j = 1, 2, 3, \tag{1.2}$$

where $d_j, j = 1, 2, 3$ are the fractional integration parameters at the Gegenbauer frequencies and z_t is a short-memory stationary process defined in the frequency domain characterized by the spectral density $f_z(\omega)$, which will be defined below. The factor $(1-B)^{d_0}$ accounts for the long-range dependence at the long-run frequency, whereas the Gegenbauer polynomials $(1 - 2\cos\varpi_j B + B^2), j = 1, 2, 3$, account for the persistent seasonal behavior of the process at frequencies $\varpi_j = \frac{2\pi}{7} j$, which correspond to cycles with a period of 7 days ($j = 1$), 3.5 days ($j = 2$, two cycles in a week), and 2.3 days ($j = 3$, three cycles per week). The process is stationary if $d_0 \in (0, 0.5)$ and $d_j \in (0, 0.25), j = 1, 2, 3$ (see Woodward, Gray, and Elliott (2011), pg. 418, Theorem 11.5a). The Gegenbauer process was introduced by Hosking (1981) and formalized by Gray et al. (1989).

The spectrum of short-memory component z_t follows Bloomfield's exponential model (see (Bloomfield, 1973)) of order q

$$f_z(\omega) = \frac{1}{2\pi} \exp\left(c_{z0} + 2 \sum_{k=1}^{q} c_{zk} \cos(\omega k) \right), \tag{1.3}$$

where $c_{z0}, c_{zk}, k = 1, \ldots, q$ are real-valued parameters, known as the cepstral coefficients of z_t (Bogert, Healy, and Tukey, 1963).

An important implication of the model specification is that the logarithm of the spectral generating function of u_t, denoted as $2\pi f(\omega)$, is linear in the memory coefficients and in the short-run coefficients c_{zk}

$$\ln[2\pi f(\omega)] = c_{z0}+2\sum_{k=1}^{q} c_{zk}\cos(\omega k)-2d_0\ln\left|2\sin(\omega/2)\right|-2\sum_{j=1}^{3} d_j\ln\left|4\sin\left(\frac{\omega+\varpi_j}{2}\right)\sin\left(\frac{\omega-\varpi_j}{2}\right)\right|. \tag{1.4}$$

The inverse Fourier transform of the logarithmic spectrum in (1.4) provides the cepstral coefficients of the process u_t:

$$c_k = \frac{1}{2\pi}\int_{-\pi}^{\pi}\ln[2\pi f(\omega)]\exp(\imath\omega k)\mathrm{d}\omega, k = 0,1,\ldots,q \tag{1.5}$$

By trigonometric identities (see Gradshteyn and Ryzhik (2007), formula 1.441.2 and 1.448.2) for $k \geq 1$:

$$c_k = I(k\leq q)c_{zk} + \frac{1}{k}\left(d_0 + 2\sum_{j=1}^{3} d_j\cos(\varpi_j k)\right), \quad k = 0,1,2,\ldots. \tag{1.6}$$

The sequence $\{c_k\}_{k=0,1,...}$ is referred to as the *cepstrum* of u_t (Bogert et al., 1963) and carries all of the relevant information that is needed for the linear prediction of process u_t, and y_t thereof.

In the sequel we will refer to the specification consisting of (1.2) and (1.4) as the generalized exponential (GEXP) model.

Wold representation and linear prediction

Let $\mathscr{F}_t$ denote the information available up to time t, consisting of the past values of y_t and the current and past values of x_t. If the model is correctly specified, the one-step-ahead prediction error variance (p.e.v.) of y_t, $\sigma^2 = Var(y_t|\mathscr{F}_{t-1})$, is obtained by the Szegö-Kolmogorov formula as

$$\sigma^2 = \exp\left[\frac{1}{2\pi}\int_{-\pi}^{\pi}\ln[2\pi f(\omega)]\mathrm{d}\omega\right],$$

from which it follows that $c_{z0} = \ln\sigma^2$. Moreover, if the stationarity condition is satisfied, we can obtain the Wold and autoregressive representations of the system:

$$y_t = x_t'\beta + \psi(B)\xi_t, \quad \phi(B)(y_t - x_t'\beta) = \xi_t, \quad \xi_t \sim WN(0,\sigma^2) \tag{1.7}$$

where $\psi(B) = 1+\psi_1 B+\psi_2 B^2+\ldots,\sum_j\psi_j^2 < \infty$, and $\phi(B) = \sum_{j=0}^{\infty}\phi_j B^j = \psi(B)^{-1}$, $\sum_j\phi_j^2 < \infty$, $\phi_0 = 1$. The moving average coefficients of the Wold representation are obtained recursively from the cepstral coefficients by the formula

$$\psi_j = j^{-1}\sum_{r=1}^{j} r c_r\psi_{j-r}, \quad j = 1,2,\ldots, \tag{1.8}$$

with $\psi_0 = 1$. See Jancek (1982), Pourahmadi (1983) and Hurvich (2002). The autoregressive coefficients are obtained according to the recursive formula

$$\phi_j = -j^{-1} \sum_{r=1}^{j} r c_r \phi_{j-r}, \quad j = 1, 2, \ldots, \tag{1.9}$$

with starting value $\phi_0 = 1$.

1.3 Robust estimation

Approximate (Whittle) likelihood estimation

Given a time series realization of length n, $\{(y_t, x_t), t = 1, 2, \ldots, n\}$, and letting $\omega_j = \frac{2\pi j}{n}$ denote the Fourier frequencies, for $j = 1, \ldots, \lfloor \frac{n-1}{2} \rfloor$,, where $\lfloor \cdot \rfloor$ is the largest integer not greater than the argument, estimation of regression parameters β, memory parameters $d_j, j = 0, 1, 2, 3$, and cepstral coefficients $c_k, k = 0, 1, \ldots, q$, is carried out in the frequency domain.

Let us denote the periodogram of x, y by

$$I_{(z)}(\omega) = \frac{1}{2\pi n} \left| \sum_{t=1}^{n} z_t e^{-\imath \omega t} \right|^2,$$

where (z) is either x or y. The cross-periodogram of x and y is defined as

$$I_{xy}(\omega) = \frac{1}{2\pi n} \left(\sum_{t=1}^{n} x_t e^{-\imath \omega t} \right) \left(\sum_{t=1}^{n} y_t e^{-\imath \omega t} \right).$$

The cross-periodogram and as well as the matrix periodogram of x_t are complex valued and we will denote their real part by I^*_{xy} and $I_x(\omega)^*$, respectively. The periodogram of $u_t = y_t - x'_t \beta$ can be written as

$$I(\omega) = I_y(\omega) - 2\beta' I^*_{xy}(\omega) + \beta' I^*_x(\omega) \beta.$$

If we denote by $f(\omega)$ the spectral density of u_t, as implied by (1.4), the Whittle approximation to the likelihood is

$$\mathscr{L}(\theta) = -\sum_{j=1}^{T-1} \left[\ln f(\omega_j) + \frac{I(\omega_j)}{f(\omega_j)} \right]. \tag{1.10}$$

The maximizer of (1.10) is the Whittle pseudo-maximum likelihood estimator of $\theta = (\beta, d_0, d_1, d_2, d_3, c_{z0}, c_{z1}, \ldots, c_{zq})$. The parameter β can be concentrated out of the likelihood function, yielding the frequency domain generalized least squares estimate

$$\hat{\beta} = \left[\sum_{j=1}^{T-1} \frac{1}{f(\omega_j)} I^*_x(\omega_j) \right]^{-1} \sum_{j=1}^{T-1} \frac{1}{f(\omega_j)} I^*_{xy}(\omega_j).$$

After replacing (1.10), the profile likelihood can be maximized with respect to the memory and cepstral parameters.

We refer to Dahlhaus (1989), Velasco and Robinson (2000), Giraitis, Koul, and Surgailis (2012) and Beran, Feng, Ghosh, and Kulik (2013) for the properties of the estimator in the long memory case.

Tapering

For estimation we use the tapered periodogram. Tapering aims at reducing the estimation bias that characterizes the periodogram ordinates in the nonstationary case. Velasco (1995) and Velasco and Robinson (2000) show that with adequate data tapers, the Whittle estimator of the parameters of classes of fractional integrated models, encompassing the FEXP, is consistent and asymptotically normal when the true memory parameter is in the nonstationary region. The adoption of a data taper makes the estimates invariant to the presence of certain deterministic trends.

The tapered discrete Fourier transform of u_t is defined as the squared modulus of

$$w\left(\omega_j\right) = \left(2\pi \sum_{t=1}^{n} h_t^2\right)^{-1/2} \sum_{t=1}^{n} h_t u_t e^{i\omega_j t} \tag{1.11}$$

where $\{h_t\}_{t=1}^n$ is a taper sequence, i.e. a sequence of nonnegative weights that down-weight the extreme values of the sequence on both sides, leaving the central part almost unchanged. Note that the raw periodogram arises in the case of $h_t = 1$. As in Velasco and Robinson (2000), the sequence $\{h_t\}_{t=1}^n$ is obtained from the coefficients of the polynomial

$$\left(\frac{1 - z^{[n/p]}}{1 - z}\right)^p.$$

The typical choices are $p = 2,3$. The tapered periodogram is then $I(\omega_j) = \left|w(\omega_j)\right|^2$.

Robust filtering and forecasting

As documented in section 1.4, electricity prices are characterized by abnormally high or low values, called price spikes. Their effect on the periodogram and, thus, on the Whittle estimates depends on their size, pattern and recurrence. Assume that there is no periodicity in spike occurrence and recalling that the frequency response function of a pulse dummy, $I_t(t = \tau)$, is constant across the frequency range. The results of the presence of single outlier at time $t = \tau$ are the downward biased estimates of the memory parameters and inflation of the estimates of the conditional and unconditional variances of the series. There are several strategies for robustifying the estimate of the parameters and, thus, of the spectrum. McCloskey and Hill (2013) propose to replace the sample spectrum $I(\omega)$ in (1.10) with a robust periodogram, constructed from the Fourier transform of a robust autocovariance estimate. Our alternative strategy is based on an iterative data-cleaning method extended by the robust Kalman filter introduced by Martin and Thomson (1982).

The proposed procedure entails iterating the following steps:

1. Estimate the parameters of the GEXP time series regression model of Section 3.1 through approximate Whittle likelihood estimation, maximizing (1.10).

2. Obtain the AR or MA approximation of the GEXP model as described in Section 1.2.

3. Cast the approximating AR or MA model in a state space form and apply the robust Kalman filter outlined in 1.6. In order to eliminate the influence of outliers (price spikes), the filter shrinks y_t towards its one-step-ahead prediction, depending on the size of the innovation.

4. Replace series y_t with its cleaned version (using the robust Kalman filter) and go to step 1.

5. Repeat steps 1-4 until the convergence is achieved.

1.4 Empirical results

Data from Nord Pool

Our dataset refers to the Nordic power exchange, Nord Pool Spot, owned by Nordic and Baltic transmission system operators, one of the leading power markets in Europe. About 380 companies from 20 countries trade in the Nord Pool Spot markets, with participants including both producers and large consumers with a trading volume of approximately 500 terawatt hours in 2014. The market includes Norway, Sweden, Finland, Denmark (since 2000), Estonia (since 2010), Lithuania (since 2012), and Latvia (since 2013). A detailed review of the operation of the market can be found in NordPool (2016).

Within the Nord Pool Spot, Elspot is the auction market for day-ahead electricity delivery. The Nord Pool Spot web-based trading system enables participants to submit bids and offers for each individual hour of the next day. Orders can be made between 08:00 and 12:00 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 the curves constitutes the system price for each hour (quoted per megawatt hour, MWh) and is the equilibrium price that would exist in the absence of transmission constraints within the grid. We will use the words *system* and *spot* prices interchangeably to indicate this price. The hourly prices are announced to the market at 12:42 CET and contracts are invoiced between buyers and sellers between 13:00 and 15:00. All 24 prices on day $t+1$ are determined on a given day t and released simultaneously. The system price is not necessarily the price that is paid in the single areas of the power grid, since the areas can be subject to transmission congestion. Thus, different area spot prices are likely to be determined as well as part of the day-ahead auction. The reason why the (daily average) system price is important is that it serves as the reference price for the clearing of most financial contracts; therefore, it is crucial for derivatives pricing and risk management.

The dataset consists of the 24 hourly spot electricity system prices from the Nord Pool Elspot power exchange for each day from January 1, 2000 to December 31, 2011, covering 4383 observation days. In 2006 the trading currency in Nord Pool changed from local (Nordic) currencies into the euro. Therefore, all prices are exchanged into the same currency.

Figure 1.1 plots the daily electricity log system price, y_t, where y_t is the average of 24 hourly prices, which is the main object of interest in our empirical modeling and forecasting exercises. While the daily average (based on hourly prices) represents our target variable, we will also consider building separate models for the hourly prices, with the purpose of forecasting the daily average by aggregating the forecasts of the individual hours. Looking at the data, we observe potentially nonstationary and persistent behavior which might be caused by long-memory features of the electricity spot prices. System prices vary over the week and over the year and are characterized by persistent level changes as well as spikes and drops.

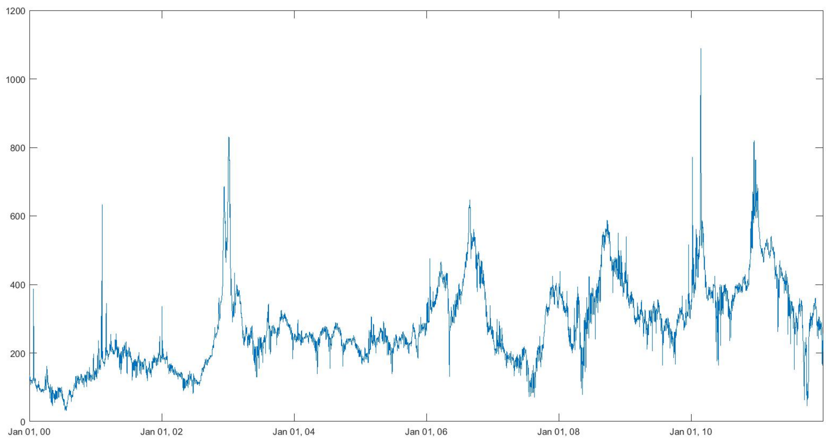

Figure 1.1. Daily average spot system price for Nord Pool power market (Norwegian kroners (NOK) per MWh), y_t. Sample: January 1, 2000 - December 31, 2011

A time series model for Nord Pool electricity spot prices

In this section we present the empirical results for the daily Nord Pool data described in the previous section, with particular reference to the logarithm of the daily spot average price time series. A generalized exponential model-based time series regression model, like (1.2), may be adequate to capture the dynamics in the conditional mean of the series. The explanatory variables that we consider are the water reservoir level and dummy variables for holiday effects, day of week effects as well as three deterministic trigonometric cycles that account for the annual cycle and two harmonics. We set off by presenting and discussing the maximum likelihood estimates of the model parameters and by assessing their empirical adequacy. We also examine the influence of outliers (price spikes) on the parameter estimates, with special attention to the memory parameter estimates.

Figure 1.2 presents the original and the transformed electricity daily average spot prices with the use of tapering of the demeaned observations. As it can be seen from the plot, tapering brings the series closer to stationarity and removes potential trends from the data.

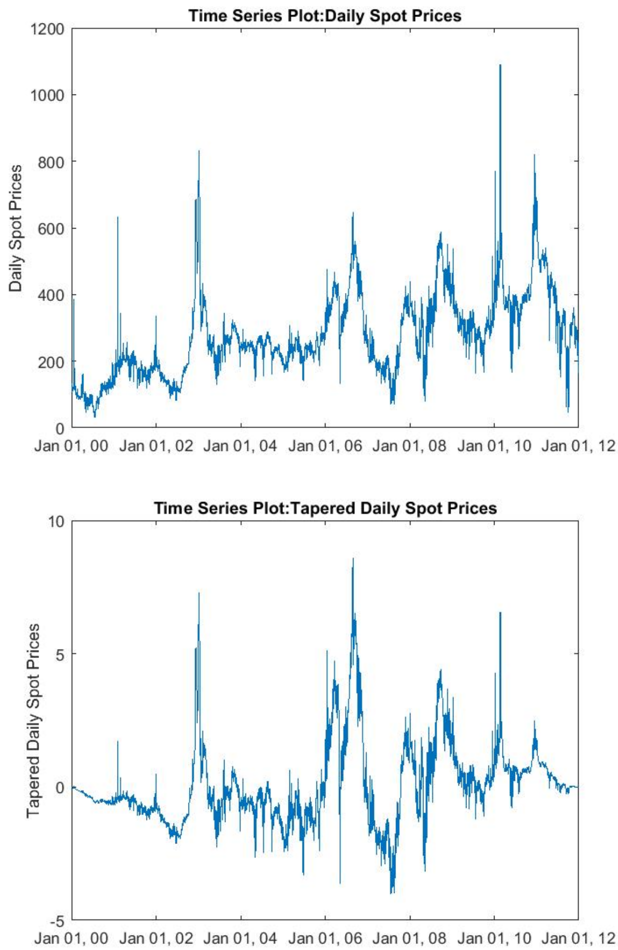

Figure 1.2. Original and tapered daily average electricity spot prices

The results of fitting the GEXP model are presented in Tables 1.1-1.3. In each table we group together classical and robust estimates of the parameters in order to evaluate the influence of robustification on the parameter estimates. Table 1.1 presents the estimates of the memory parameters at the zero and seasonal frequencies: 0, $2\pi/7, 4\pi/7, 6\pi/7$. Electricity spot prices seem to be non-stationary, although still mean-reverting, at the zero frequency, as the estimated d is significantly above 0.5. There is also clear evidence of seasonal long memory ($d_i > 0, i = 1,2,3$). For the daily average the parameter estimates are unaffected. It is seen, however, that robustification only affects the memory estimates to a minor extent.

Table 1.1. The estimates of memory parameters for the zero (d) and harmonic frequencies (d_1, d_2, d_3)

	d	d_1	d_2	d_3
Classical	**0.8806**	**0.3288**	**0.3205**	**0.2356**
	(0.0446)	(0.0311)	(0.0313)	(0.0303)
Robust	**0.8801**	**0.3276**	**0.3194**	**0.2342**
	(0.0446)	(0.0311)	(0.0313)	(0.0303)

Note: Standard errors in brackets. The statistically significant parameters at the 5% significance level appear in bold.

Table 1.2 contains the estimated short-run cepstral coefficients describing the short-run dynamics of the series. The number of short-run cepstral coefficients is chosen according to the Bayesian Information Criterion (BIC). For both estimation methods the optimal number according to the BIC of the coefficients is 7. All of the short-run cepstral coefficients are statistically significant. The estimated intercept, $\hat{c}_0$, is the estimate of the logarithm of the prediction error variance.

Table 1.2. The estimates of the short-run cepstral coefficients

	c_0	c_1	c_2	c_3	c_4	c_5	c_6	c_7
Classical	**-2.6286**	**0.1417**	**0.1371**	**0.1104**	**0.0883**	**0.0829**	**0.0645**	**-0.1488**
	(0.0214)	(0.0764)	(0.0417)	(0.0304)	(0.0249)	(0.0216)	(0.0196)	(0.0244)
Robut	**-2.6291**	**0.1409**	**0.1368**	**0.1102**	**0.0881**	**0.0828**	**0.0644**	**-0.1478**
	(0.0214)	(0.0764)	(0.0417)	(0.0304)	(0.0249)	(0.0216)	(0.0196)	(0.0244)

Note: Standard errors in brackets. The statistically significant parameter estimates at the 5% significance level appear in bold.

Table 1.3 shows the estimation results concerning the explanatory variables[1]. Most of the considered explanatory variables are statistically insignificant. Holidays and Sunday appear to be statistically significant when the classical estimator is used for estimation.

[1] The explanatory variables are holidays and day-of-the-week dummies, water reservoir levels and cosine functions that extract the annual and semiannual cycles of the data.

Table 1.3. The estimates of explanatory variable coefficients

	Holidays	**Tuesday**	**Wednesday**	**Thursday**	**Friday**	**Saturday**	**Sunday**	**Water reservoir**	$cos(2t\pi/365.25)$	$cos(4t\pi/365.25)$	$cos(3t\pi/365.25)$
Classical	**-0.2120**	-0.0334	-0.0016	-0.0160	-0.0641	0.0006	-0.0897	-0.0144	0.3896	0.0398	0.1997
	(0.0571)	(0.1452)	(0.1944)	(0.2119)	(0.2111)	(0.2005)	(0.1549)	(0.0249)	(0.7059)	(0.4080)	(0.4940)
Robust	**-0.2093**	-0.0326	-0.0015	-0.0157	-0.0634	-0.0005	-0.0900	-0.0142	0.3847	0.0393	0.1970
	(0.0571)	(0.1444)	(0.1935)	(0.2109)	(0.2101)	(0.1996)	(0.1542)	(0.0249)	(0.7055)	(0.4079)	(0.4938)

Note: The brackets contain standard errors. The statistically significant parameters at 5% significance level appear in bold.

Figure 1.3 displays the log-periodogram of the series and the estimated logarithmic spectrum. As seen, the model effectively captures the spectral peaks at the seasonal frequency and at the zero frequency.

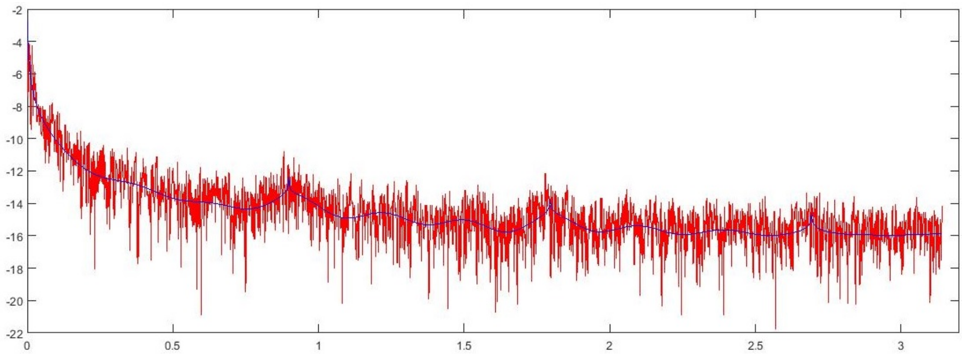

Figure 1.3. Goodness of fit in spectral domain for the cleaned log electricity prices

Figure 1.4 shows the sample autocorrelation function of the cleaned residuals. We may conclude that the proposed modeling strategy involving tapering and robust estimation provides a reasonable fit of the series by the model.

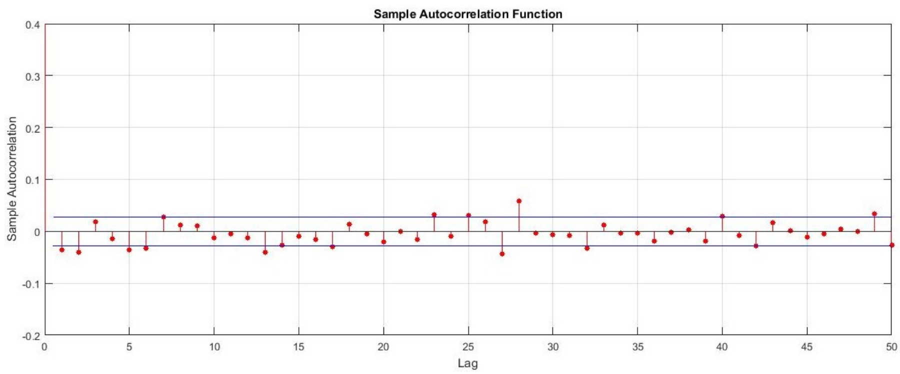

Figure 1.4. Sample autocorrelation for cleaned residuals

Forecasting

This section deals with predicting average daily electricity prices. We compare several forecasting strategies using a range of univariate and multivariate time series models. The forecasts using the generalized exponential model are obtained with the finite autoregressive approximation of the model and the use of the Kalman filter, as is described in 1.6.

We also address the question of whether the prices for the individual hours contain useful predictive information compared to the daily average price. We compare the forecasts obtained from separate models for each 24 hours of the day, which are then averaged and compared with the forecasts from the aggregate model, which is modeled directly for the daily average prices. Raviv et al. (2015), using a different model class, demonstrate that forecast averaging may be superior to direct modeling and forecasting of the (average) daily observations.

The comparative assessment of the predictive accuracy is based on a rolling forecasting experiment: starting from 01.01.2012, we estimate the model and forecast the following day's average price; we then proceed by updating the sample by adding one observation, and deleting the one at the beginning of the sample, re-estimating the parameters and forecasting the following day's prices, until the end of the sample is reached. The evaluation sample is based on 1460 (4 years) one-step-ahead forecasts and 1453 seven-steps-ahead forecasts.

Following Raviv et al. (2015), we consider the ARX model to be the benchmark model[2]. The ARX model of order p is specified as follows:

$$y_t = \sum_{j=1}^{p} \phi_j y_{t-j} + \sum_{k=1}^{K} \psi_k d_{kt} + \varepsilon_t, \;\; \varepsilon_t \sim \mathrm{WN}(0, \sigma^2),$$

where d_{kt} are dummies for Saturdays, Sundays, and for each month of the year. The order p is chosen at every point in time according to the Akaike Information Criterion.

Additionally to the ARX benchmark model from Raviv et al. (2015) we consider also the ARX model with the full set of the explanatory variables[3] (denoted by ARX all).

The next model used for comparative purposes is the seasonal random walk with a seasonal period of 7 days, $y_t = y_{t-7} + \varepsilon_t$, according to which the next day prediction is $\hat{y}_{t+1|t} = y_{t-6}$.

Another relevant predictor is the seasonal Holt-Winters method (see Holt (1957), Winters (1960), and Hyndman, Koehler, Ord, and Snyder (2008)); we have considered

[2]The comparison has also been made with the best seasonal ARIMA benchmark, as it is reported in Soares and Souza (2006). However, this benchmark was outperformed by the ARX model.

[3]The same as defined in the footnote (1).

its additive formulation:

$$\begin{aligned}
\hat{y}_{t+1|t} &= l_t + b_t + s_{t-6} \\
l_t &= \alpha\left(y_t - s_{t-7}\right) + (1-\alpha)\left(l_{t-1} + b_{t-1}\right) \\
b_t &= \beta^*\left(l_t - l_{t-1}\right) + \left(1-\beta^*\right) b_{t-1} \\
s_t &= \gamma\left(y_t - l_{t-1} - b_{t-1}\right) + \left(1-\gamma\right) s_{t-7},
\end{aligned}$$

where α, β^* and γ denote smoothing parameters, taking values in [0,1], estimated by minimizing the one-step-ahead mean square prediction error within the training sample.

Additionally, we consider two models defined in Livera, Hyndman, and Snyder (2011), namely

- exponential smoothing state space model (denoted ETS),
- trigonometric BATS model[4] (denoted TBATS).

We consider also two multivariate time series models constructed based on 24 daily series for each hour of a day

- Bayesian VAR as defined in Banbura, Giannone, and Reichlin (2010),
- factor VAR model as defined in Raviv et al. (2015).

We refer to Raviv et al. (2015), Livera et al. (2011) and Banbura et al. (2010) for additional information on the competitive models.

Forecasting accuracy measures

The statistics that are used for assessing the forecasting accuracy of the different methods and models are the mean absolute error (MAE) and the root mean squared error (RMSE),

$$RMSE = \sqrt{\frac{1}{N}\sum_{t=T}^{H-1}\left(\hat{y}_{t+1|t} - y_{t+1}\right)^2}, \tag{1.12}$$

$$MAE = \frac{1}{N}\sum_{t=T}^{H-1}\left|\hat{y}_{t+1|t} - y_{t+1}\right|, \tag{1.13}$$

see Hyndman and Athanasopoulos (2012), where H is the total number of observations, T is the length of the estimation window, and $N = H - T$ is the number of out-of-sample forecasts.

Thus, the rolling forecast experiment compares different predictors, including the one arising from the generalized exponential model, applied to the daily average price series. The predictors are also applied to the 24 individual time series referring to the

[4] BATS is an acronym for key features of the model: Box–Cox transform, ARMA errors, Trend, and Seasonal components.

single hours of the day, and the one-step-ahead and seven-steps-ahead forecasts are later aggregated into a daily average forecast. This yields a total of thirteen modeling strategies to be compared.

Table 1.4 reports the measures of one day ahead forecasting accuracy for the considered models. Analogously to the work of Raviv et al. (2015), we provide the nominal values of the forecasting accuracy measures for only the benchmark ARX model. The performance of the other models is presented in relative terms, that is, as the ratio of the accuracy measure for a specific model to that of the benchmark.

The results show that the GEXP model has similar performance to the other models. The prices for individual hours contain useful predictive information for the daily average price for one-day-ahead forecasting. The forecasts based on sample averaging of the forecasts based on factor VAR model with two factors are the most accurate. Moreover, the GEXP model provides more accurate results than the univariate models considered in Raviv et al. (2015)(2015, page 236, Table 2). Some further improvements of daily average price forecasting might be achieved by playing with the weights when computing the average; see Raviv et al. (2015) for further details.

Table 1.4. Forecasting accuracy measures for 1 day ahead forecasting

	RMSE	MAE
ARX*	124.0612	94.5606
ARX all	0.5951	0.5186
Seasonal HW	0.1381	0.1404
Seasonal RW	0.3520	0.3010
ETS	0.1361	0.1276
TBATS	0.1226	0.1177
GEXP	0.1530	0.1364
ARX: averaging	0.7793	0.7761
ARX all: averaging	0.1748	0.1746
Seasonal HW: averaging	NA	NA
Seasonal RW: averaging	0.3521	0.3012
ETS: averaging	NA	NA
TBATS: averaging	NA	NA
GEXP: averaging	0.1674	0.1512
BVAR	0.1471	0.1254
factor VAR	0.1338	0.0998

Note: The first row contains the values of forecasting accuracy measures for the ARX model (benchmark), and the remaining rows contain the ratio of the accuracy measure for a specific model and for the benchmark. The smaller the value, the better the forecasting performance of a given model. For some of the models the algorithm could not converge.

Table 1.5 contains different measures of forecasting accuracy for the models considered for the forecasting horizon of seven days.

The results clearly show that the GEXP model outperforms other model performance when RMSE is used for the comparison. MAE shows that the forecasts based on the BVAR and factor VAR models are more accurate. Moreover, the GEXP model provides more accurate results than the univariate and multivariate models considered in Raviv et al. (2015, page 236, Table 2).

Table 1.5. Forecasting accuracy measures for 7 days ahead forecasting

	RMSE	MAE
ARX*	124.2914	94.6002
ARX all	0.6219	0.5977
Seasonal HW	0.3486	0.3006
Seasonal RW	0.3517	0.3011
ETS	0.3464	0.3001
TBATS	0.3453	0.3056
GEXP	0.3374	0.2972
ARX: averaging	0.8117	0.8044
ARX all: averaging	0.3616	0.3383
Seasonal HW: averaging	NA	NA
Seasonal RW: averaging	0.3518	0.3012
ETS: averaging	NA	NA
TBATS: averaging	NA	NA
GEXP: averaging	0.3380	0.3003
BVAR	0.3408	0.2855
factor VAR	0.3383	0.2916

Note: The first row contains the values of forecasting accuracy measures for the ARX model (benchmark), and the remaining rows contain the ratio of the accuracy measure for a specific model and for the benchmark. The smaller the value, the better the forecasting performance of a given model. For some of the models the algorithm does not converge.

Testing the difference in the forecasting performance

We also test for the significance of the difference in the forecasting performance of the models with the use of the Giacomini-White test (see Giacomini and White (2006), for details).

Table 1.6 includes the test results based on the absolute error loss, which corresponds to testing the null hypothesis of equal predictive ability in terms of MAE. The absolute error loss function was chosen for comparison with the work of Raviv et al. (2015) and also to decrease the influence of the outliers on the final results. Similar results can be obtained with the square error loss function. A positive value means that the MAE of the model in the column is larger than that of the row model. The results confirm that the GEXP model performs well relative to the competitors. In addition, it can be seen that averaging of forecasts for each hour of the day is better in terms of the forecasting performance than forecasting the average electricity price directly.

Table 1.6. Giacomini-White test statistics for 1 day ahead forecasting

	ARX*	ARX all	Seasonal HW	Seasonal RW	ETS	TBATS	GEXP	ARX: averaging	ARX all: averaging	Seasonal RW: averaging	GEXP: averaging	BVAR
ARX all	20.6525											
Seasonal HW	40.7950	24.9785										
Seasonal RW	30.6906	12.9949	-17.9766									
ETS	42.2080	25.9602	7.9800	19.7508								
TBATS	41.7640	26.7085	10.9457	21.3665	4.4697							
GEXP	40.8339	25.1838	1.2388	18.6669	-2.6285	-5.8621						
ARX: averaging	41.3008	-12.6500	-38.6738	-25.2454	-40.3440	-39.9533	-38.7164					
ARX all: averaging	38.0411	25.2985	-7.1361	13.4837	-9.5139	-11.8246	-7.5862	35.1449				
Seasonal RW: averaging	30.6797	12.9850	-17.9973	-1.3877	-19.7704	-21.3879	-18.6865	25.2332	-13.5015			
GEXP: averaging	40.2977	24.3907	-3.0833	18.5599	-6.8694	-10.1430	-9.7464	37.8318	4.8298	18.5807		
BVAR	39.4080	25.7622	3.2393	20.7414	20.7414	-1.7026	2.1288	37.0007	10.1235	20.7617	5.2904	
factor VAR	42.0610	28.2709	10.1807	29.2794	29.2794	5.0489	8.2424	39.9547	16.0334	29.3038	13.0628	6.1979

Note: The test statistic is computed such that a positive value means that the MAE of the row model is smaller than the MAE of the column model.

Table 1.7 includes the test results based on the absolute error loss, which corresponds to testing the null hypothesis of equal predictive ability in terms of MAE for the case of 7 days ahead forecasting. The results show that the GEXP is one of the best model for 7 days ahead forecasting and outperforms the GEXP averaging.

Table 1.7. Giacomini-White test statistics for 7 days ahead forecasting

	ARX*	ARX all	Seasonal HW	Seasonal RW	ETS	TBATS	GEXP	ARX: averaging	ARX all: averaging	Seasonal RW: averaging	GEXP: averaging	BVAR
ARX all	7.9809											
Seasonal HW	19.2137	8.3192										
Seasonal RW	19.4058	8.3334	-0.3235									
ETS	19.2518	8.3532	0.2485	0.4306								
TBATS	30.1891	8.3656	-1.2256	-1.0328	-1.3160							
GEXP	19.7585	8.6764	0.9409	1.0301	0.8340	1.7495						
ARX: averaging	24.4350	-4.0466	-15.7834	-15.9927	-15.7662	-15.4797	-16.2149					
ARX all: averaging	20.5935	10.0299	0.0081	-2.2343	-2.3092	-1.8873	-2.6361	16.3432				
Seasonal RW: averaging	19.4049	8.3315	-0.4374	-1.3873	-0.5065	0.9980	-1.0663	15.9904	2.2247			
GEXP: averaging	19.6099	8.5125	0.2237	0.2918	0.1391	1.1270	-1.3692	15.9165	2.3565	0.3198		
BVAR	20.2272	9.2118	2.1208	2.0518	2.1662	2.9662	1.4952	16.3574	3.3921	2.0738	1.9758	
factor VAR	19.5425	8.4993	1.5659	1.6066	1.8083	2.2386	0.7692	15.7880	2.8056	1.6357	1.2565	-1.2377

Note: The test statistic is computed such that a positive value means that the MAE of the row model is smaller than the MAE of the column model.

Model Confidence Set

In addition to the previous analysis, we also estimated the model confidence set (MCS) according to the methodology proposed by Hansen, Lunde, and Nason (2011). An MCS is a set of models that contains the best model with a given level of confidence. The procedure consists of a sequence of tests which permit constructing a Superior Set of Models (SSM), where the null hypothesis of Equal Predictive Ability is not rejected at a given confidence level.

Given the initial set of models, M^0, comprising the thirteen predictors, denote the prediction error arising from method k as $v_{kt} = y_t - \hat{y}_{k,t|t-1}$. Let $d_{kl,t} = |v_{kt}| - |v_{lt}|$ be the loss differential at time t between predictors k and l for all $k, l \in M^0$, and define

$$\bar{d}_{kl} = \frac{1}{n-n_0} \sum_{t=n_0+1}^{n} d_{kl,t}, \quad \bar{d}_k = \frac{1}{K} \sum_{l=1}^{K} \bar{d}_{kl}.$$

The t-statistics associated with null $H_0 : \mathrm{E}(\bar{d}_{kl}) = 0$ (equal forecast accuracy) versus alternative $H_0 : \mathrm{E}(\bar{d}_{kl}) > 0$, are $t_{kl} = \frac{\bar{d}_{kl}}{\hat{\mathrm{Var}}(\bar{d}_{kl})}$, where $\hat{\mathrm{Var}}(\bar{d}_{kl})$ is an estimate of $\mathrm{Var}(\bar{d}_{kl})$. Also, the t-statistics associated with null $H_0 : \mathrm{E}(\bar{d}_{k.}) = 0$ (equal forecast accuracy) versus alternative $H_0 : \mathrm{E}(\bar{d}_{k.}) > 0$, are $t_k = \frac{\bar{d}_{k.}}{\hat{\mathrm{Var}}(\bar{d}_{k.})}$, where $\hat{\mathrm{Var}}(\bar{d}_{k.})$ is an estimate of $\mathrm{Var}(\bar{d}_{k.})$. If M is the current set of models under assessment, to test the hypotheses $H_{0,M} : \mathrm{E}(\bar{d}_{kl}) = 0, \forall k, l \in M$ or $H_{0,M} : \mathrm{E}(\bar{d}_{k.}) = 0, \forall k \in M$, i.e. all of the models have the same predictive accuracy, we use test statistics $T_{R,M} = \max_{kl \in M} |t_{kl}|$ and $T_{\max,M} = \max_{k \in M} t_k$. We initially set $\mathcal{M} = \mathcal{M}^0$, and test $H_{0,M}$ using the above statistics at significance level $\alpha = 0.10$. The critical values are obtained by the block-bootstrap method. If $H_{0,\mathcal{M}}$ is accepted, then the MCS at level $1-\alpha$ is $\widehat{M}^*_{1-\alpha} = M$; otherwise, we proceed to eliminate the predictor for which the t_k statistic is a maximum from the set, $k \in M$, and iterate the procedure with the surviving predictors.

The MCS p-values of the $T_{R,M}$ and $T_{max,M}$ statistics are reported in Table 1.8 for one-step-ahead forecasting. The estimated SSMs differ in the number of eliminated models as well as their compositions. We can observe that nine forecasting strategies were eliminated by the MCS procedure; and thus the superior set of models contains five models.

This empirical finding highlights the statistical equivalence of forecasting future daily electricity spot prices with the GEXP model, ETS, TBATS and factor VAR model for one-day-ahead. The GEXP model is placed among the top models for both considered test statistics. The results confirm that the best forecasting strategy is the one involving averaging of forecasts based on the multivariate time series model - the factor VAR with two factors. The TBATS model is placed on the second position. The third place is taken by the approach based on direct forecasting of the average electricity price with the GEXP model.

Table 1.8. Superior Set of Models on the 99% confidence level for a day ahead forecasting (nine models eliminated)

	$Rank_{R,M}$	$T_{R,M}$	$p-value_{R,M}$	$Rank_{max,R}$	$T_{max,R}$	$p-value_{max,R}$	$Loss$
HW	5	0.8531	0.9924	5	1.6017	0.0426	293.6601
ETS	4	0.5012	1.0000	4	1.4132	0.1440	284.8241
TBATS	1	-1.6538	1.0000	1	-0.9501	1.0000	232.3224
GEXP	3	0.3732	1.0000	3	1.3434	0.2388	282.1861
factor VAR	2	0.0599	1.0000	2	0.9566	0.9932	275.6022

Note: Comparison of the superior set of models. The p–values of the $T_{R,M}$ and $T_{max,M}$ statistics are reported in the fourth and seventh columns, respectively. The $p\text{–}value$ of the test statistic is equal to the minimum of the overall p–values. Columns $Rank_{R,M}$ and $Rank_{max,M}$ report the ranking of the models belonging to the SSMs. Last column $Loss$ is the average loss across the considered period.

The MCS p-values of the $T_{R,M}$ and $T_{max,M}$ statistics are reported in Table 1.9 for seven days ahead forecasting. The estimated SSMs differ in the number of eliminated models as well as their compositions. We can observe that twelve forecasting strategies were eliminated by the MCS procedure; and thus the superior set of models contains only one model, namely the GEXP model. This empirical finding highlights the statistical dominance of forecasting electricity spot prices with the GEXP model for a seven-days-ahead forecast horizon.

Table 1.9. Superior Set of Models on the 99% confidence level for 7 days ahead forecasting (twelve models eliminated)

	$Rank_{R,M}$	$T_{R,M}$	$p-value_{R,M}$	$Rank_{max,R}$	$T_{max,R}$	$p-value_{max,R}$	$Loss$
GEXP	1	-0.65903	1	1	-0.65903	1	1732.079

Note: Comparison of the superior set of models. The p–values of the $T_{R,M}$ and $T_{max,M}$ statistics are reported in the fourth and seventh columns, respectively. The p–$value$ of the test statistic is equal to the minimum of the overall p–values. Columns $Rank_{R,M}$ and $Rank_{max,M}$ report the ranking of the models belonging to the SSMs. Last column $Loss$ is the average loss across the considered period.

1.5 Conclusions

This paper has proposed a model for daily electricity spot prices from the Nord Pool power exchange. It can be applied also to other utilities presenting similar seasonal and long-memory patterns, e.g. electricity demand or loads. The model is formulated in the frequency domain and captures the long-range dependence and the persistent seasonal pattern of prices as well as the effect of explanatory variables (e.g. calendar effects component and water reservoir levels). One of the most challenging tasks when modeling electricity prices is to deal with extreme observations such as price spikes. To handle this problem we have proposed a novel estimation strategy based on a robust Kalman filter. The GEXP model provides a statistically coherent representation of the price dynamics. Whether using classical or robust estimates, the effect on the estimated memory appears to be minor, however.

A forecasting exercise was conducted to put the predictive ability of our model to the test. The empirical evidence suggests that the best strategy for daily average spot electricity price forecasting for one-day-ahead is to construct separate models for each hour and average the forecasts, which is in line with findings of Raviv et al. (2015). However, for the forecast horizon of seven days, it is better to build the model and forecast directly the average electricity price. The forecasting exercise shows very good performance of the proposed model for seven days ahead forecasts.

1.6 Appendix

Robust filtering and forecasting

The robust spectrum estimation and robust forecasting are based on the AR(m) or MA(m) approximation of the GEXP process. In general, an ARMA(p, q) time series model for $u_t = y_t - x_t'\beta$,

$$u_t + \phi_1 u_{t-1} + \cdots + \phi_p u_{t-p} + \xi_t + \theta_1 \xi_{t-1} + \cdots + \theta_q \xi_{t-q}, \xi_t \sim \mathrm{WN}(0, \sigma^2),$$

can be written in state space form with a measurement equation:

$$u_t = Z\alpha_t + G\xi_t, \quad t = 1, \ldots, n, \tag{1.14}$$

where α_t is a random vector with $m = \max(p, q)$ elements, $Z = [1, 0, \ldots, 0]$, $G = 1$. The evolution of the states is governed by the transition equation

$$\alpha_{t+1} = T\alpha_t + H\xi_t, \qquad t = 1, 2, \ldots, n, \tag{1.15}$$

where

$$T = \begin{bmatrix} -\phi_1 & 1 & 0 & \cdots & 0 \\ -\phi_2 & 0 & 1 & \ddots & 0 \\ \vdots & \vdots & \ddots & \ddots & 0 \\ \vdots & \cdots & \cdots & 0 & 1 \\ -\phi_m & 0 & \cdots & \cdots & 0 \end{bmatrix}, H = \begin{bmatrix} \theta_1 - \phi_1 \\ \theta_2 - \phi_2 \\ \vdots \\ \vdots \\ \theta_m - \phi_m \end{bmatrix}.$$

Initial state vector α_1, assuming stationarity (the eigenvalues of T are inside the unit circle), has a distribution with mean $\mathrm{E}(\alpha_1) = 0$ and variance $\mathrm{Var}(\alpha_1) = \sigma^2 P_{1|0}$, satisfying matrix equation $P_{1|0} = TP_{1|0}T' + HH'$.

The above state space representation is due to Pearlman (1980), and encompasses both the pure AR and MA cases.

Assuming that the process is Gaussian, define $\mathscr{F}_t = \{u_1, u_2, \ldots, u_t\}$, with the information set up to and including time t, $\tilde{\alpha}_{t|t-1} = \mathrm{E}(\alpha_t | \mathscr{F}_{t-1})$, and $\mathrm{Var}(\alpha_t | \mathscr{F}_{t-1}) = \sigma^2 P_{t|t-1}$.

The Kalman filter (KF) is the following recursive algorithm: for $t = 1, \ldots, n$,

$$\begin{aligned} \nu_t &= u_t - Z\tilde{\alpha}_{t|t-1}, & f_t &= ZP_{t|t-1}Z' + GG', \\ & & C_t &= P_{t|t-1}Z'f_t^{-1}, \\ \tilde{\alpha}_{t|t} &= \tilde{\alpha}_{t|t-1} + C_t\nu_t, & P_{t|t} &= P_{t|t-1} - C_t f_t C_t', \\ & & Q_t &= HG'f_t^{-1}, \\ \tilde{\alpha}_{t+1|t} &= T\tilde{\alpha}_{t|t} + Q_t\nu_t, & P_{t+1|t} &= TP_{t|t}T' + HH' - (Q_t f_t Q_t' + Q_t f_t C_t' T' + TC_t f_t Q_t'). \end{aligned} \tag{1.16}$$

The above equations compute the innovations $\nu_t = u_t - \mathrm{E}(u_t | \mathscr{F}_{t-1})$, and $\sigma^2 f_t$ is the prediction error variance at time t, which is $\mathrm{Var}(u_t | \mathscr{F}_{t-1})$; $\tilde{\alpha}_{t|t}$ are the updated, or real-time, estimates of the state vector, and C_t is the gain, $C_t = \mathrm{Cov}(\alpha_t, y_t | Y_{t-1})[\mathrm{Var}(y_t | Y_{t-1})]^{-1}$. It can be shown that $\alpha_t | \mathscr{F}_t \sim \mathrm{N}(\tilde{\alpha}_{t|t}, \sigma^2 P_{t|t})$. The vector Q_t is interpreted as $\mathrm{Cov}(H\xi_t, u_t | \mathscr{F}_{t-1})[\mathrm{Var}(u_t | \mathscr{F}_{t-1})]^{-1}$, so that $\tilde{\alpha}_{t+1|t}$ is the one-step-ahead state prediction and we can write $\alpha_{t+1} | \mathscr{F}_t \sim \mathrm{N}(\tilde{\alpha}_{t+1|t}, \sigma^2 P_{t+1|t})$.

Masreliez and Martin (1977) and Martin and Thomson (1982) proposed to obtain a robust filter to modify the Kalman filter updating and prediction equations (1.16) by using a bounded and continuous function of the standardized innovations so as to control the effects of outliers on the conditional mean estimates.

Let us denote $\tilde{v}_t = \frac{v_t}{f_t^{1/2}}$. The above KF is modified as follows: for $t = 1, \dots, n$,

$$
\begin{aligned}
&v_t = u_t - Z\tilde{\alpha}_{t|t-1}, && f_t = ZP_{t|t-1}Z' + GG',\\
& && C_t = P_{t|t-1}Z'f_t^{-1},\\
&\tilde{\alpha}_{t|t} = \tilde{\alpha}_{t|t-1} + C_t f_t^{1/2}\psi(\tilde{v}_t), && P_{t|t} = P_{t|t-1} - w(\tilde{v}_t)C_t f_t C_t',\\
& && Q_t = HG'f_t^{-1},\\
&\tilde{\alpha}_{t+1|t} = T\tilde{\alpha}_{t|t} + Q_t f_t^{1/2}\psi(\tilde{v}_t), && P_{t+1|t} = TP_{t|t}T' + HH' - w(\tilde{v}_t)(Q_t f_t Q_t' + Q_t f_t C_t' T' + TC_t f_t Q_t').
\end{aligned}
\tag{1.17}
$$

Here, $\psi(u)$ is Hampel's two-part redescending function:

$$
\psi(u) = \begin{cases} u, & |u| \le a,\\ \text{sign}(u)\frac{a}{b-a}(b - |u|), & a < |u| \le b,\\ 0, & |u| > b, \end{cases}
$$

for $a < b$. The weight function is $w(u) = \psi(u)/u$. Note that if $\psi(u) = u$, the identity function, $w(u) = 1$ and the above recursions yield the Kalman filter (1.16).

A "clean" estimate of y_t is then

$$
\tilde{u}_{t|t} = Z\tilde{\alpha}_{t|t} + GG'f_t^{-1/2}\psi(\tilde{v}_t). \tag{1.18}
$$

Noticing that $u_t = Z\tilde{\alpha}_{t|t} + GG'f_t^{-1}v_t$, we can write $\tilde{u}_{t|t} = u_t - GG'f_t^{-1}v_t[1 - w(\tilde{v}_t)]$, which shows that $\tilde{u}_{t|t} = u_t$ when $w(u) = 1$, which occurs for $|u| < a$, i.e. when the standardized innovations are small. On the contrary, when $w(u) = 0$, which takes place for $|u| > b$, simple manipulations show that in the presence of a large residual, the cleaned observations are shrunk towards the one-step ahead prediction:

$$
\tilde{u}_{t|t} = [1 - GG'f_t^{-1}]u_t + GGf_t^{-1}\tilde{u}_{t|t-1}, \quad \tilde{u}_{t|t-1} = Z\tilde{\alpha}_{t|t-1}.
$$

It can be shown that the steady-state Kalman filter for the AR or MA model considered has $\lim_{t\to\infty} f_t = GG'$; thus, after processing a large number of observations, the occurrence of a large outlier causes $\tilde{u}_{t|t} \to \tilde{u}_{t|t-1}$. Finally, when $a < |\tilde{v}_t| \le b$, $\tilde{u}_{t|t}$ is a weighted linear combination of u_t and $\tilde{u}_{t|t-1}$.

The theoretical underpinnings of the robust KF are provided in Masreliez and Martin (1977).

1.7 References

Arteche, J., Robinson, M., 2000. Semiparametric inference in seasonal and cyclical long memory processes. Journal of Time Series Analysis 21, 1–25.

Banbura, M., Giannone, D., Reichlin, L., 2010. Large Bayesian vector auto regressions. Journal of Applied Econometrics 25 (1), 71–92.

Benth, F., Koekebakker, S., 2008. Stochastic modeling of financial electricity contracts. Energy Economics 30, 1116–1157.

Beran, J., Feng, Y., Ghosh, S., Kulik, R., 2013. Long-memory processes: Probabilistic Properties and Statistical Methods. Berlin, Springer.

Bloomfield, P., 1973. An exponential model for the spectrum of a scalar time series. Biometrika 60, 217–226.

Bogert, B., Healy, M., Tukey, J., 1963. The frequency analysis of time series for echoes: Spectrum, pseudo-autocovariance, cross-spectrum and shape cracking. In Proc. Symp. on Time Series Analysis 1, 209–243.

Dahlhaus, R., 1989. Efficient parameter estimation for self-similar processes. The Annals of Statistics 17, 1749–1766.

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

Diongue, K., Guégan, D., Vignal, B., 2004. A k-factor gigarch process: estimation and application on electricity market spot prices. In Probabilistic Methods Applied to Power Systems, 1–7.

Diongue, K., Guégan, D., Vignal, B., 2009. Forecasting electricity spot market prices with a k-factor gigarch process. Applied Energy 86, 505–510.

Ferrara, L., Guegan, D., 2001. Forecasting with k-factor gegenbauer processes: Theory and applications. Journal of Forecasting 20, 581–601.

Giacomini, R., White, H., 2006. Tests of conditional predictive ability. Econometrica 74, 1545–1578.

Gil-Alaña, L., Robinson, P., 2001. Testing of seasonal fractional integration in uk and japanese consumption and income. Journal of Applied Econometrics 16, 95–114.

Giraitis, L., Koul, H., Surgailis, D., 2012. Large sample inference for long memory processes. Tech. rep., Imperial College Press.

Gradshteyn, I., Ryzhik, I., 2007. Table of integrals, series, and products. Elsevier/Academic Press, Amsterdam, Seventh edition.

Granger, C., Joyeux, R., 1980. An introduction to long-memory time series models and fractional differencing. Journal of Time Series Analysis 1, 15–29.

Gray, L., Zhang, N., Woodward, A., 1989. On generalized fractional processes. Journal of Time Series Analysis 10, 233–257.

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

Haldrup, N., Nielsen, M., 2006b. A regime switching long memory model for electricity prices. Journal of Econometrics 135, 349–376.

Hansen, P., Lunde, A., Nason, J., 2011. The model confidence set. Econometrica 79, 453–497.

Holt, C., 1957. Forecasting seasonals and trends by exponentially weighted moving averages. Technical report, onr research memorandum, Carnigie Institute.

Hosking, J., 1981. Fractional differencing. Biometrika 68, 165–176.

Hsu, N., Tsai, H., 2009. Semiparametric estimation for seasonal long-memory time series using generalized exponential models. Journal of Statistical Planning and Inference 139, 1992–2009.

Hurvich, C. M., 2002. Multistep forecasting of long memory series using fractional exponential models. International Journal of Forecasting 18, 167–179.

Hyndman, R., Athanasopoulos, G., 2012. Forecasting: principles and practice. OTexts.

Hyndman, R., Koehler, A., Ord, J., Snyder, R., 2008. Forecasting with Exponential Smoothing. The State Space Approach, Springer.

Jancek, G., 1982. Determining the degree of differencing for time series via the log spectrum. Journal of Time Series Analysis 3, 177–183.

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

J.Beran, 1993. Fitting long-memory models by generalized linear regression. Biometrika 80, 817–822.

Koopman, S., Ooms, M., Carnero, M., 2007. Periodic seasonal reg-arfima-garch models for daily electricity spot prices. Journal of the American Statistical Association 102, 16–27.

Livera, A. M. D., Hyndman, R. J., Snyder, R. D., 2011. Forecasting time series with complex seasonal patterns using exponential smoothing. Journal of the American Statistical Association 106 (496), 1513–1527.

Martin, R., Thomson, D., 1982. Robust-resistant spectrum estimation. Proceedings of the IEEE 70, 1097–1115.

Masreliez, C., Martin, R., 1977. Robust bayesian estimation for the linear model and robusting the kalman filter. IEEE Transactions of Automatic Control, 361–371.

McCloskey, A., Hill, J., 2013. Parameter estimation robust to low-frequency contamination. Journal of Business & Economic Statistics.

Möst, D., Keles, D., 2010. A survey of stochastic modeling approaches for liberalised electricity markets. European Journal of Operational Research 207, 543–556.

Narukawa, M., Matsuda, Y., 2011. Broadband semi-parametric estimation of long-memory time series by fractional exponential models. Journal of Time Series Analysis 32, 175–193.

NordPool, 2016. Nord pool spot.
URL `http://www.nordpoolspot.com/`

Pearlman, J. G., 1980. An algorithm for the exact likelihood of a high-order autoregressive-moving average process. Biometrika 67, 232–233.

Pourahmadi, M., 1983. Exact factorization of the spectral density and its application to forecasting and time series analysis. Communications in Statistics - Theory and Methods 12, 2085–2094.

Raviv, E., Bouwman, K., Dijk, D. V., 2015. Forecasting day-ahead electricity prices: Utilizing hourly prices. Energy Economics 50, 227–239.

Robinson, P. M., 1991. Nonparametric function estimation for long memory time series. Tech. rep., In Nonparametric and Semiparametric Methods in Econometrics and Statistics: Proc. of the 5th Int. Symposium in Economic Theory and Econometrics.

Serati, M., Manera, M., Plotegher, M., 2008. Modeling electricity prices: from the state of the art to a draft of a new proposal. Tech. rep., Liuc papers in Economics, Cattaneo University (LIUC).

Smallwood, A., Beaumont, P., 2003. Multiple frequency long memory models. Tech. rep., Department of Economics University of Oklahoma.

Soares, L., Souza, L., 2006. Forecasting electricity demand using generalized long memory. International Journal of Forecasting 22, 17–28.

Velasco, C., 1995. Non-stationary log-periodogram regression. Journal of Econometrics 91, 325–371.

Velasco, C., Robinson, P., 2000. Whittle pseudo-maximum likelihood estimation for nonstationary time series. Journal of the American Statistical Association 95, 1229–1243.

Walker, A., 1964. Asymptotic properties of least-squares estimates of parameters of the spectrum of a stationary non-deterministic time-series. Journal of the Australian Mathematical Society 4, 363–384.

Weron, R., 2006. Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach. Wiley.

Weron, R., 2008. Market price of risk implied by asian-style electricity options and futures. Energy Economics 30, 1098–1115.

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

Weron, R., Simonsen, I., Wilman, P., 2004. Modeling highly volatile and seasonal markets: evidence from the nord pool electricity market. The Application of Econophysics, 182–191.

Whittle, P., 1953. Estimation and information in stationary time series. Arkiv för Matematik 2, 423–434.

Winters, P., 1960. Forecasting sales by exponentially weighted moving averages. Management Science, 324–342.

Woodward, A., Cheng, C., Gray, L., 1998. A k-factor garma long-memory model. Journal of Time Series Analysis 19, 485–504.

Woodward, W., Gray, H., Elliott, A., 2011. Applied time series analysis, 1st Edition. CRC Press.

Zareipour, H., 2012. Short-term electricity market prices: A review of characteristics and forecasting methods. Springer Berlin Heidelberg.

CHAPTER 2

MODELING AND FORECASTING ELECTRICITY PRICE JUMPS IN THE NORD POOL POWER MARKET

Oskar Knapik
Aarhus University and CREATES

Abstract

For risk management, traders in the electricity market are mainly interested in the risk of negative (drops) or of positive (spikes) price jumps, i.e. the sellers face the risk of negative price jumps, while the buyers face the risk of positive price jumps. Understanding the mechanism that drives extreme prices and forecasting of the price jumps is crucial for risk management and market design. In this paper, we consider the problem of the impact of fundamental price drivers on forecasting of price jumps in Nord Pool intraday market. We develop categorical time series models which take into account i) price drivers, ii) persistence, and iii) seasonality of electricity prices. The models are shown to outperform commonly used benchmarks. The paper shows how crucial for price jumps forecasting is to incorporate additional knowledge in price drivers such as loads, temperature and water reservoir levels, as well as take into account the persistence in the jump occurrence process.

2.1 Introduction

Electricity prices often tend to temporarily jump to extreme levels, a phenomenon which is usually associated with with the non-storability of electricity, unexpected increases in demand, unexpected supply shortages or over production from wind turbines or the failures of the transmission infrastructure (see Geman and Roncoroni, 2006; Christensen, Hurn, and Lindsay, 2012; Eydeland and Wolyniec, 2003; Harris, 2006, among others). The spikes may occur due to the fact that the central dispatch process sometimes needs to rely on the bids of the high marginal cost of production generators in order to satisfy demand. It is important to remember though, that the severity of spikes cannot be explained only by the dispatch of units with a high marginal cost of production but it is mainly driven by gaming or speculation (see Weron, 2006). Electricity price jumps are visible for intra-day or day-ahead prices on a half-hourly, hourly or daily time grid, but not for forward prices. Extreme price events are particularly hazardous for power market participants. Many traders on the electricity market are mainly concerned with either the risk of negative or positive price jumps, i.e., buyers face the risk of positive, while sellers face the risk of negative price jumps (see Hellstrom, Lundgren, and Yu, 2012). Consequently, improving our understanding of the factors which contribute to the occurrence of extreme prices as well as accurate forecasting of these events, is crucial for effective risk management in the energy sector. It is the forecasting problem of price jumps which is the central concern of this paper.

Models for electricity prices typically fall into three categories: traditional autoregressive time series models, nonlinear time series models (with a particular emphasis on Markov-switching models) and continuous-time diffusion or jump-diffusion models (see Christensen et al., 2012). All of these models aim to characterize the trajectory of the spot price or return across time. A review of the models with special attention to forecasting of electricity prices can be found in (Weron, 2014). Taken at face value, these models appear to be different in their identification and treatment of price jumps (see Janczura, Trueck, Weron, and Wolff, 2013, for the review).

There is a scant research focusing directly on price jumps modeling and forecasting. Lu, Dong, and Li (2005) explore the reasons for price spikes using Bayesian classification and similarity searching. Their approach uses the measurement of the proposed composite indexes reflecting the relationship among electricity demand, supply and reserve capacity using Queensland half hourly prices. An important research in this area is the work of Jong (2006) representing the class of regime-switching models for electricity price models called Independent Spikes models (see Lindstrom, Noren, and Madsen, 2015). Mount, Ning, and Cai (2006) show that a stochastic regime-switching model with two regimes and the two transition probabilities being functions of the load and/or the implicit reserve margin can reflect the volatile behavior of wholesale electricity prices associated with price spikes. Becker, Hurn, and Pavlov (2007) claim that the use of a time-varying probability regime-switching model with transition probabilities modeled with logit transformation of variables capturing demand (daily average and maximum load) and weather (daily average and maximum temperature, dew-point) can help to predict price spikes for Queensland data. Amjady and Keynia (2011) propose a data mining selection technique for prediction of both likelihood and severity of price spikes. Cartea, Figueroa, and Geman (2012) implement "tight market conditions" in capacity constraints in the form of a threshold variable to a regime-switching model for England and Wales power markets. They reveal that 85% of spikes occur when the demand-to-capacity ratio is in the interval [0.908; 0.960]. Christensen et al. (2012) consider the time series of price spikes for half-hourly data from the Australian market and introduce a nonlinear variant of the autoregressive conditional hazard (ACH) model. Clements, Fuller, and Hurn (2013) propose a semi-parametric approach based on a framework developed for forecasting realized volatility for forecasting spikes in the Australian market. Maryniak and Weron (2014) analyze forward looking data that is available to all participants in the UK power market and showed that the reserve margin has a huge potential for explaining the spike probability.

The majority of models for electricity prices treat price spikes as a memoryless process. However, evidence suggests there is a significant persistence component and helps to explain the intensity of the jumping process (see Christensen, Hurn, and Lindsay, 2009; Christensen et al., 2012; Clements et al., 2013, among others). Most of the models consider also the case of two states: normal price and spike price. Here, we consider categorical time series model of three electricity price states (normal price and positive/negative spike) taking into account the persistence.

The focus of majority of the papers in the literature is on Australia, UK and US power markets. There are only few papers that refer to the Nord Pool power market, when the trends are on extreme price events. Hellstrom et al. (2012) explore the possible reasons behind electricity price jumps in the Nordic electricity market by the use of a mixed GARCH-EARJI[1] jump model. Voronin and Partanen (2013) propose data mining and time series techniques for prediction of both normal prices and price spikes in the Finish Nord Pool Spot day-ahead power market. (Voronin, Partanen, and Kauranne, 2014) propose a hybrid forecasting model for the Finnish electricity spot market and show that hybridization of the normal range price and price spikes forecasts may provide comprehensive and valuable information for electricity market participants. Lindstrom et al. (2015) extend the Independent Spike Model used to model the electricity price and find out that consumption can be used to forecast extreme events in the Nord Pool power market.

This paper studies electricity prices from the Nord Pool power market. In the Nordic countries, more than 80% of the hourly consumed electricity is traded on the Elspot market, the day-ahead electricity market. Since security of supply is very important and forecasts for demand and/or supply a day ahead are not very accurate, several other spot markets have been put in place. One of them is the Elbas market where up to an hour before delivery producers and retail suppliers can update the quantity of power traded. The Nord Pool market plays a key role in the development of intraday power trading in Europe and it is becoming increasingly important due to a visions share of wind power production. Future prospects indicate exponential growth, reaching 1.900 GW installed

[1] $EARJI(r,s)$ is an exponential autoregressive jump intensity model

wind capacity worldwide in 2020 (Source: World Wind Energy Association). This type of market can be crucial in increasing the share of renewable energy in the energy mix. On the other hand, Elbas is also almost not explored. Therefore the main focus of this research is on Elbas market.

This paper makes several contributions to the existing state of knowledge in the modeling of extreme price event occurrences in Nord Pool power markets. We extensively study the fundamental price drivers for the Nord Pool power market. We propose new categorical time series models to properly model the persistence in the electricity price jumps process. One important characteristic of the proposed econometric models which are considered is that they embed the information content of previous jumps and relate it to explanatory variables modeled in three ways with the use of a first-order Markov chain model with a time-varying transition matrix, an autoregressive ordered probit model, and an autoregressive conditional multinomial model. The paper is organized as follows. Section 2.2 describes the Nord Pool power market structure, focusing on hypothetical causes of price jumps. Section 2.3 describes the data. Section 2.4 introduces the main price drivers. Section 2.5 presents the models and statistical inference. Sections 2.6 and 2.7 provide the empirical and forecasting results. Section 2.8 concludes.

2.2 The Nord Pool power market and hypothesised causes of price jumps

The Nordic power market was fully established in 2000 when the regional electricity markets of Sweden, Norway, Finland and Denmark merged. Nowadays, the Nord Pool power market also includes the Baltic countries of Estonia, Latvia and Lithuania. Nord Pool operates as one market in which supply to a region is aggregated and generators are dispatched in order to satisfy the demand as cost-effectively as possible. If in one of the regions the local demand is higher than the local supply or the electricity in a neighboring region is cheap enough to guarantee transmission, the electricity is imported and exported in between, subject to the physical constraints of the transmission infrastructure. Hydroelectric production (which supplies around 57% of the Nord Pool capacity and nuclear generators (which supply around 18% of the Nord Pool capacity) have relatively high start-up costs and low marginal costs of production. Gas turbines and oil-fired plants together supply around 10% of the market, and only take about 20 min to start power generation, but have a comparatively high marginal cost of production, used typically for peak periods only. Wind power production supplies around 9% of electricity demand with an increasing share. This type of renewable energy is less predictable than more traditional sources of energy and may lead to price drops.

The majority of the volume handled by Nord Pool Spot is traded on the day-ahead market called Elspot. To a large extent, the balance between supply and demand is secured there. However, incidents may happen between the closing of Elspot at noon CET and delivery on the following day. A nuclear power plant can suddenly stop operating or strong winds may cause higher wind power generation than expected. Therefore, Nord Pool Spot's intraday trading system Elbas has been introduced. The importance of the intraday market is growing as more wind power enters the grid. Wind power production is unpredictable by nature and fluctuates in relation to day-ahead contracts. Therefore produced volume often needs to be offset. Elbas will play a key role in the development of intraday power trading. Covering the Nordic and Baltic regions as well as Germany and, recently, also the UK, Elbas supplements Nord Pool Spot's day-ahead market and helps to secure the necessary balance in real time between supply and demand in the power market for Northern Europe.

Figure 2.1 displays a scatter plot of the dependence between price and load (approximated as turnover at system price) in the Nord Pool Elspot power market. We might observe that price spikes appear when the load is high. On the other hand, negative jumps/drops can be observed for low load. Therefore, load should be considered an important explanatory variable.[2]

Figure 2.1. Price and demand in Nord Pool Elspot power market

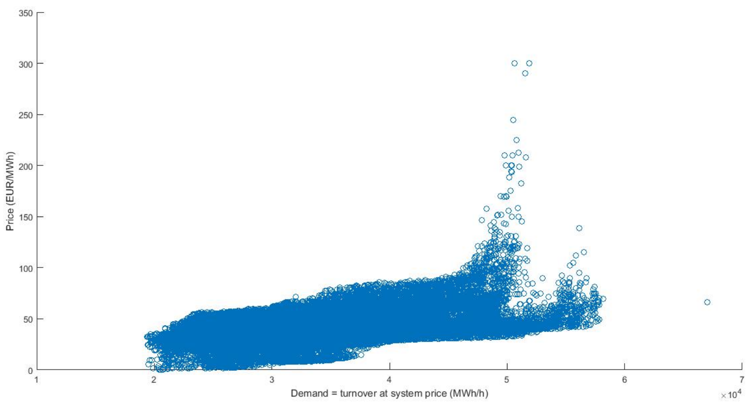

The reasons behind the occurrence of price jumps relate to the interaction between system demand and supply (see Barlow, 2002; Geman and Roncoroni, 2006, among others). The demand for electricity is very inelastic, as the demand side is protected from pool price fluctuations by retailers who buy electricity at the spot price and sell the electricity at fixed rates. Electricity during normal conditions is provided by traditional low-cost generators (coal-fired and hydroelectric generators). If the system is affected by increases in demand and/or reductions in supply, the spot price jumps above the threshold, where it is cost-effective for generators with higher costs of production (gas-

[2]In this paper we do not consider transmission capacity constraints as we work with system prices and volume-weighted prices for the entire Nord Pool. However, those constraints might be important explanatory variables for the electricity prices for single areas within the grid.

fired and diesel generators) to compete with low-cost generators. On the other side, if the system is affected by decreases in demand and/or increases in supply, the spot price jumps below the threshold, where paying the buyers for taking the additional energy (rather than shutting down the power plant) happens in the most extreme cases.

Figure 2.2. Power production price curve in the Nordic electricity market

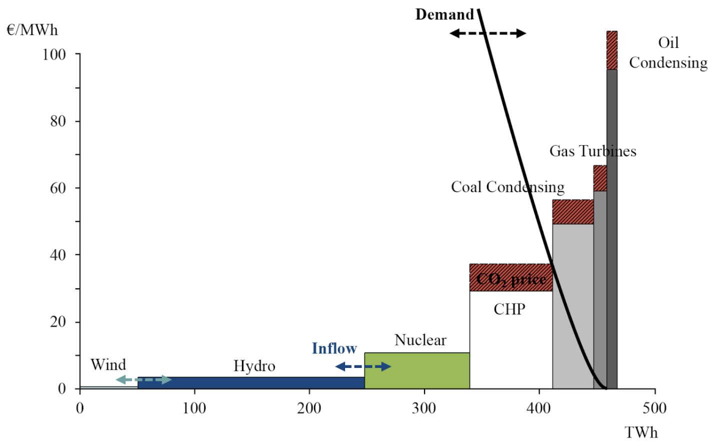

Source: `http://www.nordpoolspot.com`

Figure 2.2 shows the power production price chart in the Nordic electricity market, where the production cost is on the y-axis and the annual total production is on the x-axis. The blocks in the figure represent different means of power production. The width of the blocks reflects the generation capacity and the height represents the marginal production costs. The red-striped areas illustrate the price increases caused by the EU ETS CO2 emissions allowances. The annual power demand in the region is illustrated by the black demand curve in the picture. The chart doesn't include wind- and biomass-based power generation, as they represent a relatively small share of the total production capacity. Furthermore, as hydropower represents around half of the total power generation in the market and its marginal production cost is almost zero, fluctuations in the hydropower supply (marked with the blue-dashed arrows) shift the other means of production along the x-axis. If the nuclear power generation remains stable, and does not balance the hydropower generation fluctuations so that power demand could be satisfied with hydropower and nuclear only, the next means of generation along the x-axis are CHP (combined heat and power) and coal condensing, with both using coal as raw material input. Thus, if demand for electricity in the Nordic market exceeds the combined production capacity of hydropower and nuclear, the marginal production methods are coal-fired methods of production and, therefore, the marginal price in the market should be linked to coal prices. Another important factor affecting the cost of coal-fired power is the price of the emission allowances set by the European Union that have a huge impact on coal-based power production costs. Kara et al. (2008) estimated that the average electricity spot price would increase by 0.74 EUR/MWh for every 1 EUR per tonne CO2 allowance price in the Nordic area between 2008 and 2012. If electricity demand in the market cannot be satisfied by the above-mentioned means of power production, other power production means are gas turbines and oil condensing. As natural gas prices are highly correlated with oil price (see, for example Krichene (2002) and Villar and Joutz (2006)), the cost of these means of production is highly dependent on oil price.

2.3 Data

We study hourly volume-weighted average prices from the Elbas power market. The data covers the period from 14 September 2009 to 31 December 2013. A further sample length of 24 months from 1 January 2014 to 31 December 2015 is reserved for assessing the out-of-sample performance of the models.

The estimation period is chosen to cover a few very interesting price peak events in the Nord Pool power market. Firstly, prices peaked during the winter of 2009-2010. For Sweden, Finland, Eastern Denmark as well as Mid and Northern Norway there were three very high price peaks that occurred on 17 December 2009 at 17-18 in the evening and on 8 January 2010 at 8-9 in the morning. For these areas, the prices during the three peaks were 1400, 1000 and 1400 EUR/MWh, respectively. The system price remained at 300 EUR/MWh or below. At the same time the price in Southern Norway and Western Denmark was relatively low, at about 65 EUR/MWh. Secondly, the estimation period covers the price spike event on 24 August 2010, when spot prices in pan-Nordic power exchange Nord Pool Spot's Estonia market area reached an exceptional 2000 EUR/MWh following the loss of 160MW of generation capacity in the Baltic state. Prices for Tuesday delivery between 8-11 and 12-14 reached the maximum technical level[3] that can be recorded under Elspot day-ahead market rules.

For the purpose of this paper, a price jump will be formally defined as a situation where the spot electricity price exceeds a particular threshold value that is chosen to lie outside the normal range of daily fluctuations. To define price jump events, some method to identify jumps must be chosen. A variety of such methods can be found in the work of Janczura et al. (2013). Following the works of Becker et al. (2007) and Christensen et al. (2009), Christensen et al. (2012), Clements et al. (2013) we use static thresholds for spike/drop identification. Whilst the actual threshold used is market-specific, the reason for using a threshold to define extreme events is generic (see Kanamura and Ohashi (2007), Mount et al. (2006)). This logic comes from the fact that applications such as demand-side management do not require precise values of future spot prices but use specific price thresholds for making scheduling decisions Weron (2014). Additional

[3] At this time. Currently, the upper price limit is set to be 3000 EUR/MWh.

justification for using static thresholds is that the sample size is small enough that market conditions should not change. Using a static threshold is also justified by the fact that average marginal costs of producing power with different sources do not change. It is also more informative than dynamic thresholds for market participants.

Figure 2.3 presents the quantiles of Elbas hourly average electricity prices and Elspot hourly spot system electricity prices. It can be seen that the electricity spot price fluctuates between 20 EUR/MWh and 60 EUR/MWh under "normal" conditions. The threshold chosen for defining price spike events in Nord Pool is that of 80 EUR/MWh, which lies above the 90% spot prices for each hour of the day. This threshold is also slightly larger than the average marginal cost of gas turbine generation capacity bought online during periods of market stress (around 77 EUR/MWh). The threshold chosen for defining price drop events is chosen that of 10 EUR/MWh, which lies below the 10% spot prices for each hour of the day and is around the average marginal cost of production based on nuclear power. It should be noted that the price limits imposed by the market rules in force during the sample period are -500 EUR/MWh for the lower limit and 3000 EUR/MWh for the upper limit. The sensitivity analysis shows that the results are relatively robust for the choice of the exact values. Based on the analysis we can define the same thresholds for both considered power markets. The choice of the threshold is market specific and requires additional analysis in order to establish it for different power market. However, one can consider application of some outlier detection methods as considered in the literature (see Janczura et al., 2013; Chen and Liu, 1993a,b; Johansen and Nielsen, 2016, among others).

Figure 2.3. Quantiles of Elbas average prices by hour

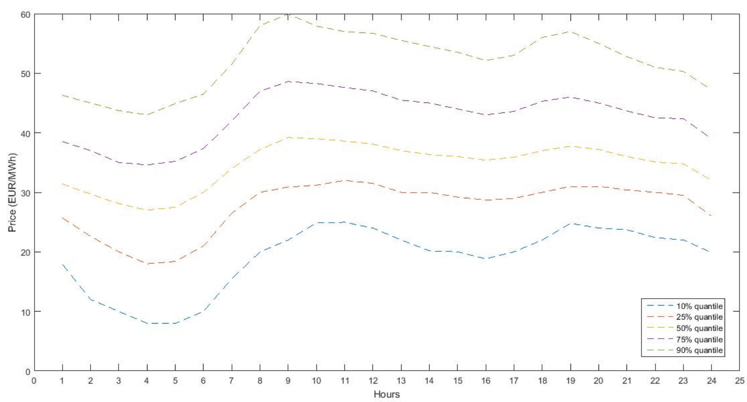

Figure 2.4. Quantiles of Elspot spot prices by hour

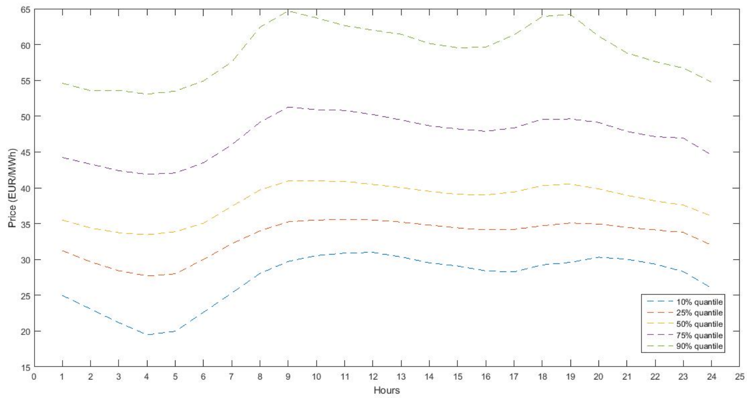

Figure 2.5 and 2.6 present hourly Elspot system electricity prices and Elbas volume-weighted average prices by hours and days in the considered time period. We can observe that there is a periodic dependency in the occurrence of price spikes and jumps.

Figure 2.5. Elbas average electricity prices by hours and days

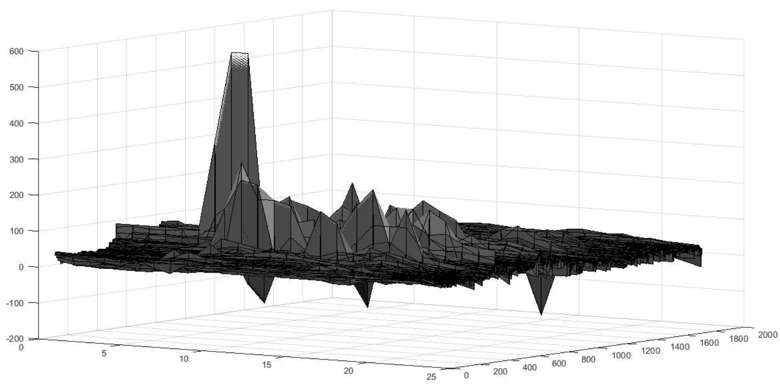

Figure 2.6. Elspot spot electricity prices by hours and days

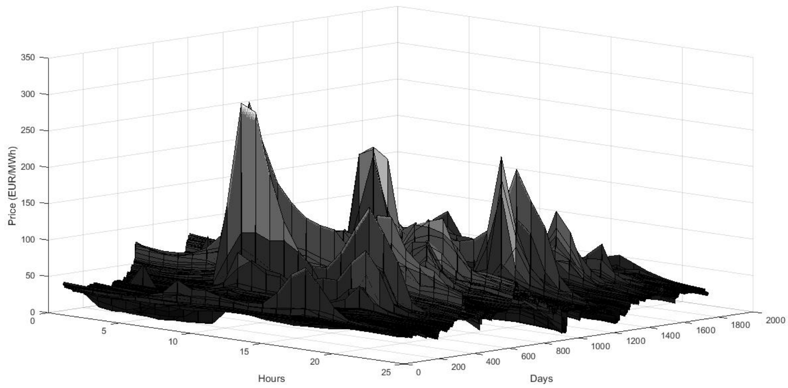

Table 2.1 shows the structure of electricity prices in the Elbas power market. 92.5% of all prices are within the range of 10 EUR/MWh to 80 EUR/MWh. Negative jumps and positive jumps are respectively 6.06% and 1.44% of the total number of observations.

Table 2.1. Elbas electricity prices: negative jumps, normal prices and positive jumps

	Negative jumps	Normal prices	Positive jumps
Number	2284	34855	541
Frequency	6.06%	92.5%	1.44%

Figure 2.7a-2.9b display the respective number of price spikes (exceedance above 80 EUR/MWh) and price drops (exceedance below 10 EUR/MWh) on hourly, daily and monthly bases, and also provide casual empirical evidence to support temporal dependency and seasonality in the Nord Pool Elspot and Elbas power markets. The percentage of spikes and drops is computed with respect to the total number of spikes and drops, respectively. Price spikes are more likely to appear early in the morning (8:00-11:00) and late afternoon (17:00-19:00), on working days (Monday-Friday) and during the winter. The figures illustrate also that price drops are more often during the night and early morning (24:00-7:00), from Wednesday to Sunday, and in July and October. This suggests that jumps are very much driven by the demand.

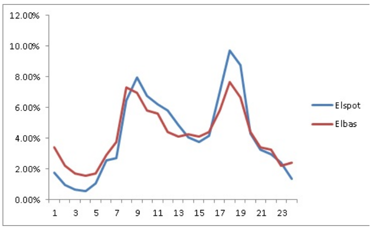

(a) 2.7a Hourly seasonality of spikes

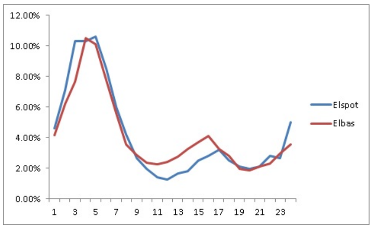

(b) 2.7b Hourly seasonality of drops

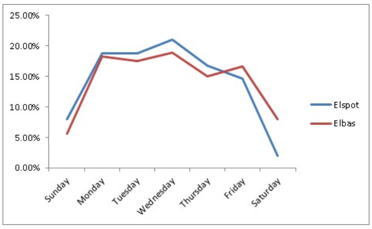

(a) 2.8a Daily seasonality of spikes

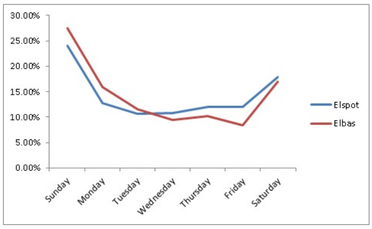

(b) 2.8b Daily seasonality of drops

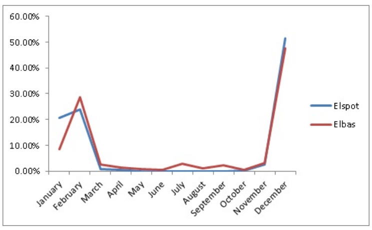

(a) 2.9a Monthly seasonality of spikes

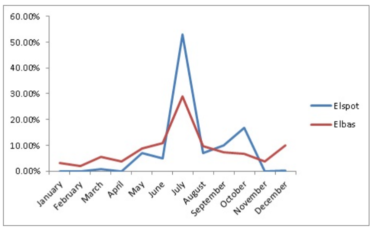

(b) 2.9b Monthly seasonality of drops

Figures 2.10a, 2.10b, 2.11a and 2.11b present proportions of jumps (spikes and drops separately) that appeared in next 10 hours, assuming that a jump occurred at a given hour in the total number of jumps (spikes and drops)[4]. The counting is done as follows: in the loop we check if for a given hour the price is a spike or drop; if it is, we check whether the prices in the next 10 hours are also in the same price state, and increase the numbers corresponding to each of the next 10 hours if needed. The whole number for each following hour is divided by the total number of jumps (spikes or drops) observed in the system. We might observe strong persistence in each case, i.e. jump clustering may be observed. The potential reasons behind abnormal price events,

[4]Technically speaking the method is based on computations of conditional probabilities of the price being in given state given price(s) in the past were in the same state, e.g. for spikes

$$P\left(Y_t = 2|Y_{t-1} = 2,\ldots,Y_{t-10} = 2\right), P\left(Y_t = 2|Y_{t-1} = 2,...,Y_{t-9} = 2\right),..., P\left(Y_t = 2|Y_{t-1} = 2\right)$$

consistent with the clustering of price jumps (shown in Figures 2.8a to 2.9b), might lie in the persistence of system stress (see Becker et al., 2007; Christensen et al., 2009, 2012, among others). This is supported by Figures 2.8a to 2.9b , which show that a lot of jumps occur in successive hours, meaning that the probability that a jump occurs in the next hour, given that one has occurred in the previous hour (hours), is higher than otherwise, excluding other temporal effects. For example, the proportion of spikes that occur one hour after another spike has occurred is around 45%. The proportion of prices being in a spike regime for the next five hours is 5%. This finding will also be supported by the estimated transition matrix of a first-order Markov chain. This reasonably mild finding has important implications for model choice, as it renders many of the current methods of capturing jumps inadequate.

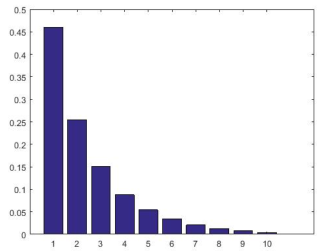

(a) 2.10a: Persistence of spikes: Elbas

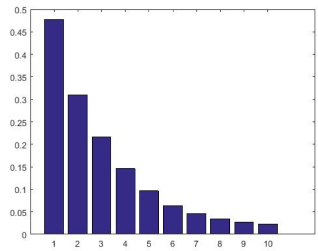

(b) 2.10b: Persistence of drops: Elbas

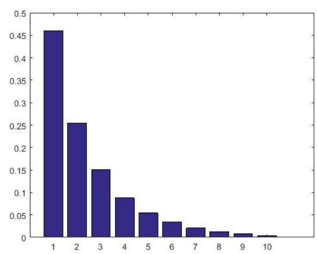

(a) 2.11a: Persistence of spikes: Elspot

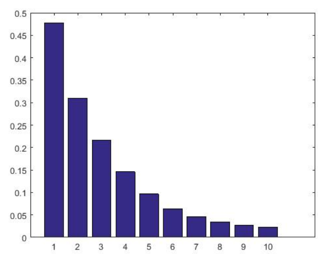

(b) 2.11b: Persistence of drops: Elspot

Section 2.5 provides a more analytical approach supporting the observation that there is persistence of exceedance in electricity prices, which shows temporal dependence. The formal framework of categorical time series can be used to determine the statistical significance of possible factors driving the occurrence of price jumps. The total number of abnormal price events arising per day may be treated as a counting measure, which measures the stress acting on the system. Here we will look at one of the measures of persistence, which is a first-order Markov chain, as the estimated transition matrix parameters.

Table 2.2 and 2.3 contain the maximum likelihood estimates of the transition matrix of the first-order Markov chain for Elbas and Elspot power markets and their standard errors. The transition matrix describes the probabilities of moving between different price states in time t, and conditionally on being in a given state in time $t-1$. For instance, if the price in the Elbas market is in a negative drop state in time $t-1$, we have around a 48% chance that it will remain in the same state in time t. On the other hand, if the price in time $t-1$ is in a normal price state, then it will be in the same state with around a 98% chance. However, if the price in time $t-1$ is in a spike state, then we have around a 50% chance that it will be in a normal state and a 45% chance that it will remain in the spike state in time t.

Table 2.2. First-order Markov chain: MLE estimates transition matrix for Elbas

			t	
		Drop	Normal price	Spike
t-1	Drop	0.4778	0.4966	0.0256
		(0.0202)	(0.0206)	(0.0047)
	Normal price	0.0140	0.9798	0.0062
		(0.0006)	(0.0048)	(0.0004)
	Spike	0.0393	0.5009	0.4560
		(0.0086)	(0.0306)	(0.0293)

Table 2.3. First-order Markov chain: MLE estimates transition matrix for Elspot

			t	
		Drop	Normal price	Spike
t-1	Drop	0.8551	0.1448	0.0000
		(0.0345)	(0.0142)	(0.0000)
	Normal price	0.0025	0.9947	0.0028
		(0.0002)	(0.0048)	(0.0002)
	Spike	0.0000	0.1604	0.8396
		(0.0000)	(0.0147)	(0.0336)

However, looking at the seasonal character of price jump occurrence, it is reasonable to suspect that those transition matrices might vary over time and can be explained with some covariates. Kanamura and Ohashi (2007) show that in the case of the U.S. electricity market PJM the transition probabilities of electricity prices cannot be constant, and depend on both the current demand level relative to the supply capacity and the trends of demand fluctuation. Therefore, it is important to have a closer look at the main price drivers in the Nord Pool power market.

2.4 The main price drivers

In order to understand the reasons behind price jumps, Geman and Roncoroni (2006) and Mount et al. (2006) analyze the generator bid curves, particularly focusing on the transition in bids from low-cost high-supply generators to high-cost low-supply generators. Systematic changes in demand due to weather or business demands, reductions in supply due to scheduled infrastructure maintenance, and non-systematic reductions in supply due to generator or network failure are some of the factors that can shift the demand and supply curves (see Christensen et al. (2012)). The empirical literature has found that the occurrence of price jumps varies across time and supports the potential reasoning behind abnormal price events. For example, Escribano, Pena, and Villaplana (2002) and Knittel and Roberts (2005) show that for some markets, the intensity parameter of the spiking process in electricity prices demonstrates seasonal dependence. Kanamura and Ohashi (2007) also observe a seasonal dependency in the transition probabilities between jump and non-jump regimes due to systematic changes in demand.

The ability of the Nordic power system to store energy in hydro-reservoirs causes less variation in the Nordic price structure than that of, for example, Germany. Inflow during summer and in periods with low demand can be used in the winter. Furthermore, hydrologic forecasts have an impact on prices, since more or less energy than expected will cause the prices to go up or down. Hydropower reservoir levels are collected on a weekly basis from the beginning to the end of the tested period. Reservoirs are taken as a percentage of the total hydropower capacity available in the Nord Pool area. 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). Reservoirs taken into account from Sweden and Finland are those after their integration in the Nord Pool market. Denmark is not included because its power production resources do not include any hydropower reservoir plant. In the Nord Pool electricity market, about 53% of power production is generated by hydropower reservoirs. The influence of reservoir levels on electricity futures prices at Nord Pool has been studied by Gjolberg and Johnsen (2001), Botterud, Bhattacharyya, and Ilic (2002), Forsund and Hoel (2004) and Fehr, Amundsen, and Bergman (2005). The researchers conclude that hydropower reservoir levels are an important factor that explains futures and spot prices. The seasonality of reservoir levels has a highly important influence on electricity spot prices. From a storage theory perspective, inventory seasonals generate seasonals in the marginal convenience yield – and in the basis (see Fama and French (1987), p. 56). Taking reservoir levels as inventories of electricity, the effect of demand and supply shocks on spot electricity prices will depend on reservoir levels and how they are managed. Thus, demand or supply shock is easily offset when reservoirs are high. Reservoirs being low, a demand or supply change, are more difficult to balance and will be persistent, allowing spot prices to rise. In order to better understand the impact of reservoir levels on prices, two extreme cases can be studied in a hydropower generation market: very high reservoir levels and very low reservoir levels. In the case of an overflow the potential gains of producers will be reduced. This causes a negative convenience yield, that is, producers would prefer to sell power at a lower price than to allow overflows. As the main idea of hydropower management is to distribute water during the periods when reservoirs are nearly full, spot prices will be lower than usual and futures prices will be above spot prices. On the other hand, when reservoirs are very low, the convenience yield will be positive and might include big values. In this case, spot prices will be higher than short-term futures prices. If reservoir levels are not enough to satisfy demand, electricity prices will probably increase together with power imports. Furthermore, precipitation can be an important factor in order to obtain an estimation of the water inflow to hydroelectric reservoirs, and the expectation of a dry or rainy period will be clearly influenced by its values.

The behavior of weather variables can also produce some predictable seasonal patterns in spot prices. The relationship between weather variables and electricity load and price has been studied by many researchers. Temperature is the main price driver in Nordic countries. Cold temperatures increase heat demand, since electricity is very much used for heating in Nordic countries. Colder temperatures usually increase prices because of higher power demand. However, in special cases, e.g. combined heat and power plants where heat is the primary product, the demand for the heat could trigger secondary electricity production and cause the prices to decrease. Li and Sailor (1995), and Munoz, Rosen, and Sailor (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 explicative power in their modeling; see, for example, Peirson and Henley (1994), Peirson and Henley (1998) and Pardo, Ridal, Murtagh, and Cernicharo (2002). Heating degree day (HDD) is a variable that shows the demand for energy needed for heating. It is taken from measurements of outside air temperature. The heating requirements for a specific structure in a specific place tend to be directly proportional to the number of HDDs in that location. In this study we will consider 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). The data on heating degree days comes from `DegreeDays.net`.

Due to the fact that there is no fuel cost for production and unpredictability, additional wind energy can lead to a price decrease. This type of energy may, in some cases, cause even negative prices in hours with low demand and additional supply. On the other hand, when wind production falls short of expected values, it can trigger high prices in both day-ahead and intra-day markets.

Finally, electricity system prices from Elspot might contain important information about Elbas electricity prices. The same situation can be seen in futures markets, where the basis is the difference between the futures price and the underlying spot price. Figure 2.12 and 2.13 present the dependence between Elspot and Elbas electricity prices. Within a range of price fluctuation (20-60 EUR/MWh) we can see a very strong correlation between Elspot and Elbas prices. In general, Elspot electricity prices are higher than Elbas average electricity prices. The reason for that could be the additional supply caused by wind power in the Elbas market. As the Elspot electricity prices are formulated in a day-ahead market, they might be treated as an important explanatory variable for the Elbas average electricity prices and, consequently, as price jump occurrence factor.

Figure 2.12. Dependence between Elspot and Elbas electricity prices

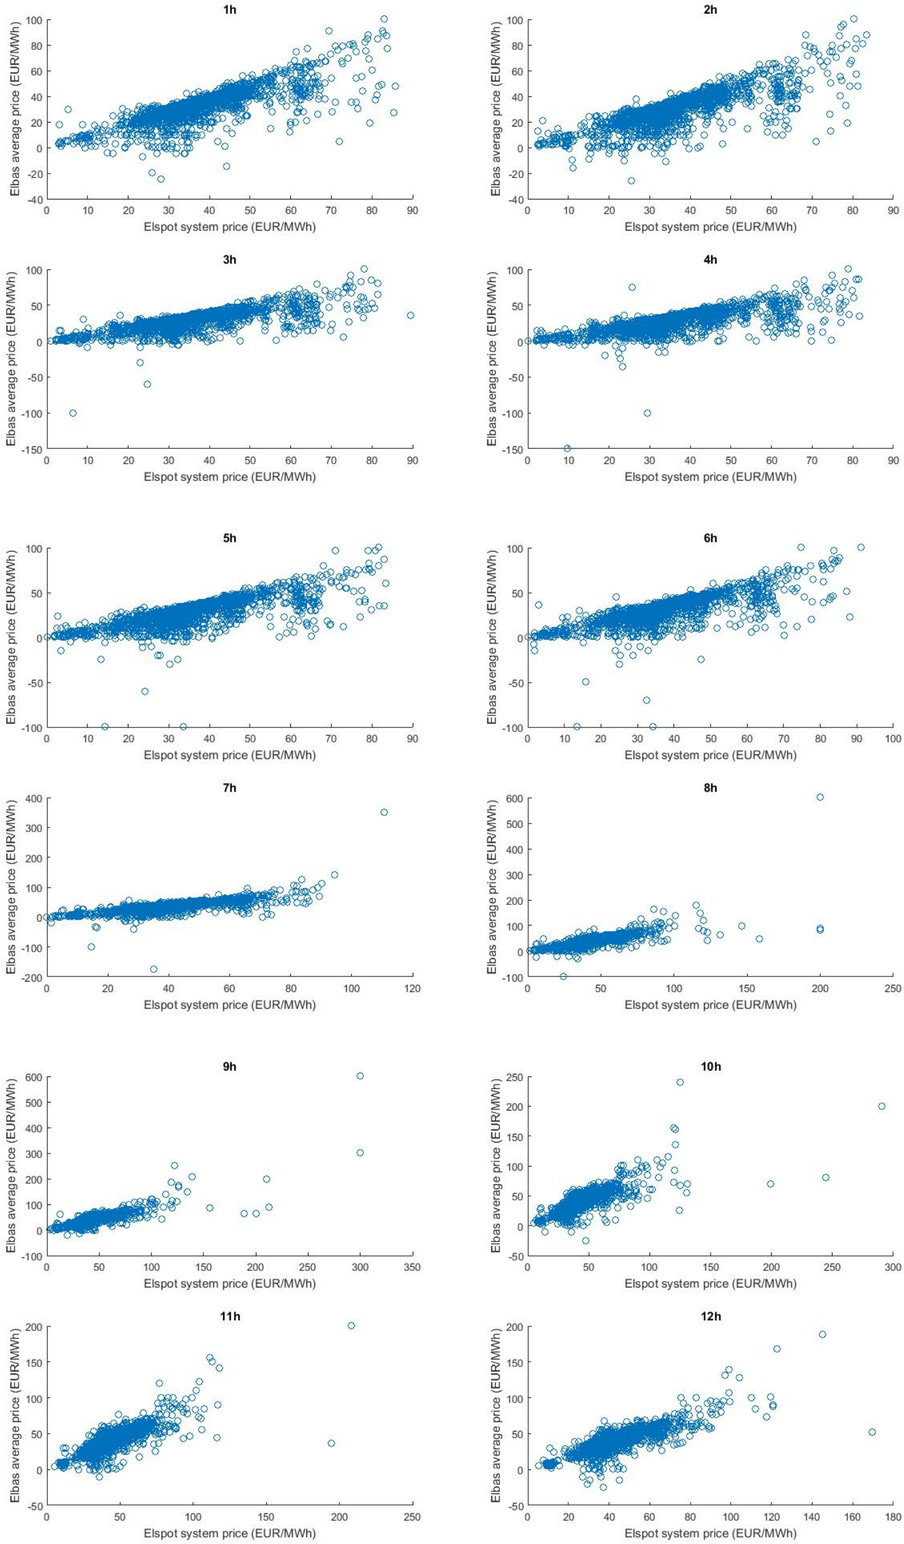

Figure 2.13. Dependence between Elspot and Elbas electricity prices

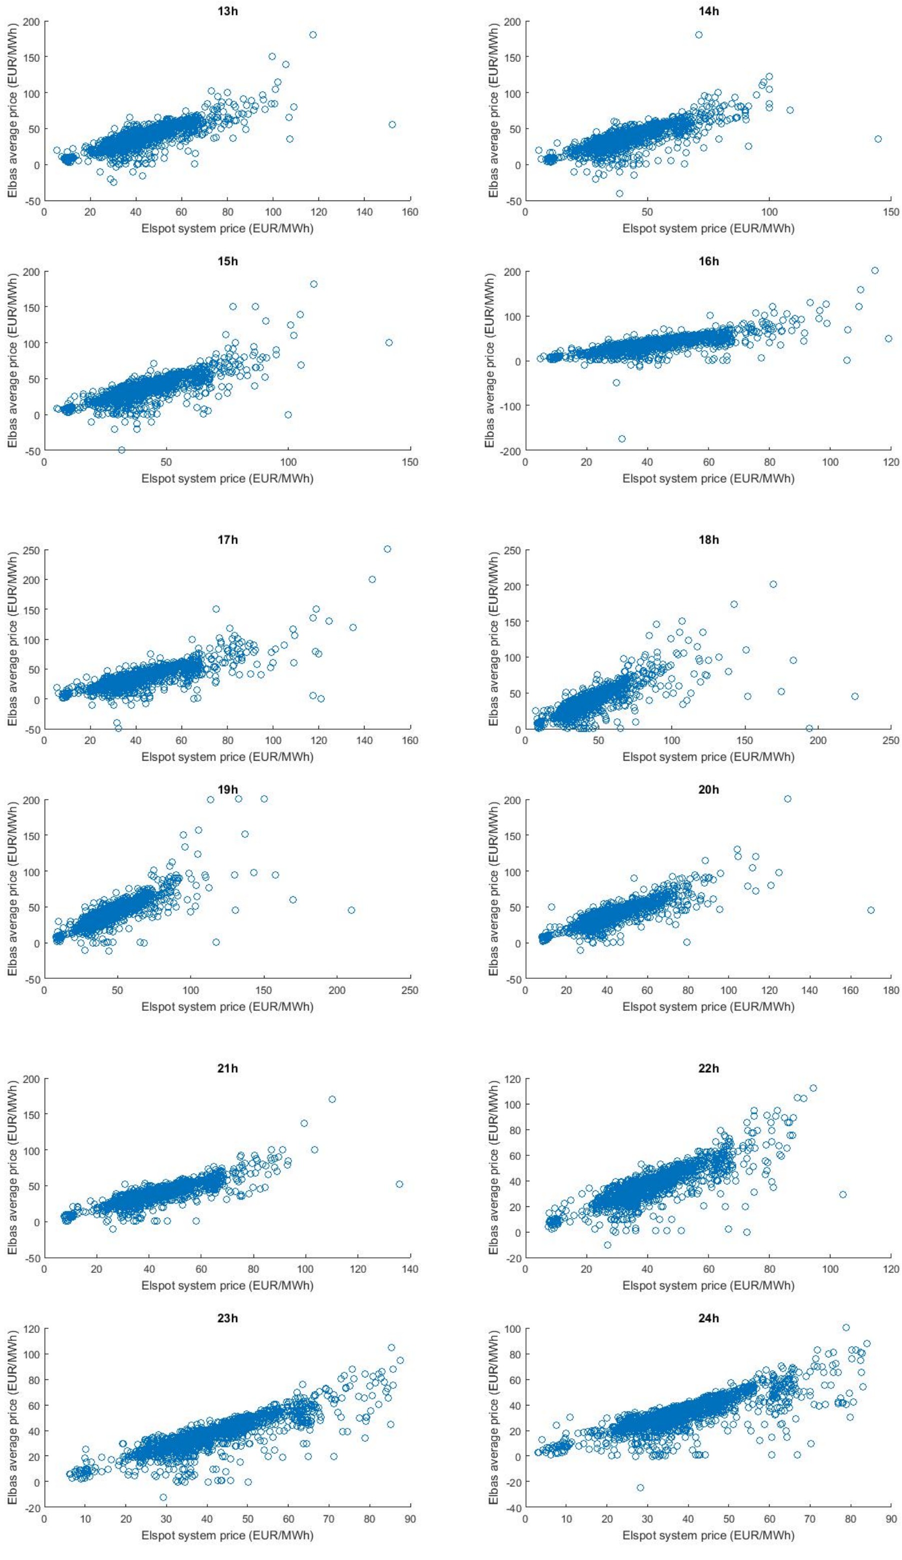

2.5 Models for electricity price jumps

Let P_t be the electricity spot price and consider the three states

$$Y_t = \begin{cases} 0, \ if \ -500 \leq P_t \leq 10 \text{ (negative jump, drop)} \\ 1, \ if \ 10 < P_t < 80 \text{ (normal price)} \\ 2, \ if \ 80 \leq P_t \leq 3000 \text{ (positive jump, spike)} \end{cases} \tag{2.1}$$

Y_t is then a categorical time series of electricity prices. There are several approaches to modeling that type of processes. Here, we will consider three models: an autoregressive ordered probit (AOP) model, a first-order Markov model with a time-varying transition matrix, and an autoregressive conditional multinomial model. Each model provides different insights into the dynamics of the data-generating mechanism.

Autoregressive ordered probit model

Let Y_t^* denote a latent (unobserved) variable. The autoregressive ordered probit model assumes the autoregressive effect in the response variable and it is built around an auxiliary/latent regression model in the following way:

$$Y_t^* = X_t\beta + \rho Y_{t-1} + \varepsilon_t, \tag{2.2}$$

Based on Y_t^* the autoregressive ordered probit model defines the mechanism leading to the price being in a specific state as follows:

$$Y_t = \begin{cases} 0 \ if \ Y_t^* \leq \mu_0 \\ 1 \ if \ \mu_0 < Y_t^* \leq \mu_1 \\ 2 \ if \ \mu_1 < Y_t^*, \end{cases} \tag{2.3}$$

We assume that ε_t is normally distributed with zero mean and unit variance. The probability that electricity price Y_t at time t is in a specific state, conditional on explanatory variables X_t and previous state Y_{t-1}, is then given by

$$\begin{aligned} P(Y_t = 0|X_t, Y_{t-1}) &= \Phi\left(-X_t\beta - \rho Y_{t-1}\right) \\ P(Y_t = 1|X_t, Y_{t-1}) &= \Phi\left(\mu_1 - X_t\beta - \rho Y_{t-1}\right) - \Phi\left(-X_t\beta - \rho Y_{t-1}\right) \\ P(Y_t = 2|X_t, Y_{t-1}) &= 1 - \Phi\left(\mu_1 - X_t\beta - \rho Y_{t-1}\right) \end{aligned} \tag{2.4}$$

where $\Phi()$ denotes the cumulative distribution function of the standard normal distribution, and $\theta = (\beta, \rho, \mu_0, \mu_1)$ is the vector of unknown parameters, including β vector, ρ scalar and $\mu = (\mu_0, \mu_1)$ parameters called cutpoints or threshold parameters.

The cutpoint parameters are estimated by the data and help to match the probabilities associated with each discrete outcome. These are subject to the obvious restriction $\mu_0 < \mu_1$. We assume that explanatory variables X_t and lagged response variable Y_{t-1} are observed, but that latent selection variable Y_t^* is not.

For the three probabilities, the partial (marginal) effects of changes in the explanatory variables and lagged response variable are

$$\begin{aligned}
\frac{\partial P(Y_t=0|X_t)}{\partial X_t} &= -\phi\left(-X_t\beta\right)\beta \\
\frac{\partial P(Y_t=1|X_t)}{\partial X_t} &= \left[\phi\left(\mu_1 - X_t\beta\right) - \phi\left(-X_t\beta\right)\right]\beta \\
\frac{\partial P(Y_t=2|X_t)}{\partial X_t} &= \phi\left(\mu_1 - X_t\beta\right)\beta \\
\frac{\partial P(Y_t=0|Y_{t-1})}{\partial Y_{t-1}} &= -\phi\left(-Y_{t-1}\rho\right)\rho \\
\frac{\partial P(Y_t=1|Y_{t-1})}{\partial Y_{t-1}} &= \left[\phi\left(\mu_1 - Y_{t-1}\rho\right) - \phi\left(-Y_{t-1}\rho\right)\right]\rho \\
\frac{\partial P(Y_t=2|Y_{t-1})}{\partial Y_{t-1}} &= \phi\left(\mu_1 - Y_{t-1}\rho\right)\rho
\end{aligned} \tag{2.5}$$

where $\phi()$ is the density function of the standard normal distribution.

We can evaluate these as sample means, or take a sample average of the marginal effects. Unlike the probit, the signs of the "interior" marginal effects are unknown and not completely determined by the signs of β and ρ. We can, however, sign the effects of the lowest and highest categories based on β, ρ. The others, however, cannot be known by the reader simply by looking at a table of point estimates.

Model parameters $\theta = (\beta, \rho, \mu)$ are estimated by partial maximum likelihood. In order to describe the estimation procedure we should note that the t'th observation of categorical time series Y_t expressed by vector $\mathbf{Y_t} = (Y_{t1}, Y_{t2})'$ of length $q = v - 1 = 2$ with elements

$$Y_{tj} = \begin{cases} 1, \text{ if the } j\text{th category is observed at time } t \\ 0, \text{ otherwise} \end{cases}$$

for $t = 2, \ldots, N$ and $j = 1, 2$. Denote by $\pi_t = (\pi_{t1}, \pi_{t2})'$ the vector of conditional probabilities given F_{t-1}

$$\pi_{tj} = E\left[Y_{tj}|F_{t-1}\right] = P\left(Y_{tj} = 1|F_{t-1}\right),\ j = 1, 2$$

for every $t = 2, \ldots, N$. At times we refer to the π_{tj} as "transition probabilities". As before, σ-field F_{t-1} stands for the whole information up to and including time t. Define

$$Y_{tm} = 1 - \sum_{j=1}^{q} Y_{tj}$$

and

$$\pi_{tm} = 1 - \sum_{j=1}^{q} \pi_{tj}.$$

In addition, put $\{W_{t-1} = (X_t, Y_{t-1}), t = 1, \ldots, N-1\}$ for the $p \times q$ matrix that represents a covariate process. Each response Y_{tj} corresponds to a vector of length p of random time-dependent covariates which forms the j'th column of W_{t-1}. The covariate matrix may contain lagged values of the response process and of any other auxiliary process as discussed earlier (Fokianos and Kedem, 2003, see). In estimation, X_t contains Elspot electricity system prices and turnover at a system price, wind power production, heating degree days, and water reservoir level.

We assume that the vector of transition probabilities - that is, conditional expectation of response vector Y_t - is linked to the covariate process through the equation

$$\pi_t(\theta) = h\left(W_{t-1}^{'}\theta\right), \tag{2.6}$$

with θ being a p-dimensional vector of time-invariant parameters. The equation (2.6) gives the general form. In our case the $\pi_t(\theta)$ is given by

$$\pi_t(\theta) = \begin{pmatrix} \pi_{t1}(\theta) \\ \pi_{t2}(\theta) \end{pmatrix} = \begin{pmatrix} P\left(Y_t = 0|X_t, Y_{t-1}\right) \\ P\left(Y_t = 1|X_t, Y_{t-1}\right) \end{pmatrix} \tag{2.7}$$

where the probabilities $P()$ are defined as in equation (2.4).

Furthermore, we can obtain

$$Var[Y_t] = \Sigma_t(\theta) =$$
$$= \begin{pmatrix} \pi_{t1}(\theta)\left(1-\pi_{t1}(\theta)\right) & -\pi_{t1}(\theta)\pi_{t2}(\theta) \\ -\pi_{t1}(\theta)\pi_{t2}(\theta) & \pi_{t2}(\theta)\left(1-\pi_{t2}(\theta)\right) \end{pmatrix}$$

The partial likelihood is a product of the multinomial probabilities

$$\prod_{j=1}^{3} \pi_{tj}^{Y_{tj}}(\theta)$$

that is,

$$PL = \prod_{t=2}^{N} \prod_{j=1}^{3} \pi_{tj}^{Y_{tj}}(\theta).$$

Thus, the partial log-likelihood is given by

$$l(\theta) = \log PL(\theta) = \sum_{t=2}^{N} \sum_{j=1}^{3} Y_{tj} \log \pi_{tj}(\theta).$$

The maximization of the log-likelihood function has to be done numerically. More details can be found in Fokianos and Kedem (2002).

Markov model

The second model is the first-order non-homogeneous Markov model for electricity price Y_t with $v = 3$ regimes whose transition matrix depends on some covariates. The transition matrix is parametrized by two $v \times 1$ vectors: π_t and ϕ_t. The π_t parameters control the probability of moving into each of the states unconditionally on the state from which it is moving, conditional on a move occurring. This single vector is not sufficient for modeling all of the dependence among the data. In particular, the probability of staying in a state given that the previous state is the same, is likely to be understated. The ϕ_t vector of probabilities is added to increase the chance of remaining in the same state. Each element of ϕ_t corresponds to the probability of moving to another state, although the new state may be the same as the old state. Thus, there are two mechanisms by which the new state can be the same as the old state: there can be no jump, or there can be a jump but to the same state. One appealing feature of this model is that it is flexible enough to accommodate states that form frequent small patches (i.e. a consecutive sequence of a state) as well as those that form occasional large patches with the same overall proportions. The parametrization based on the work of Foster, Bravington, Williams, Althaus, Laslett, and Kloser (2009) leads to transition probabilities $p_{t,ji}$ from state y_i to state y_j from observation number t to observation number $t+1$.

$$p_{t,ij} = P\left(Y_{t+1} = y_j | Y_t = y_i\right) = \begin{cases} \left(1-\phi_{ti}\right) + \phi_{ti}\pi_{tj}, & \text{if } i = j \\ \phi_{ti}\pi_{tj}, & \text{if } i \neq j \end{cases} \tag{2.8}$$

This can be represented in the form of the following time-varying transition matrix:

$$P_t = diag\left(1-\phi_t\right) + \pi_t \phi_t^T \tag{2.9}$$

The explanatory variables enter the model via logistic and additive logistic transformations of linear combination of the covariates (Aitchison (1982), Billheimer et al. (2002)). That is to say,

$$\phi_{ti} = \frac{\exp\left(x_{ti}^T \gamma_i\right)}{1+\exp\left(x_{ti}^T \gamma_i\right)} \tag{2.10}$$

$$\pi_{ti} = \begin{cases} \dfrac{\exp\left(u_t^T \beta_i\right)}{1+\sum_{j=1}^{v-1} \exp\left(u_t^T \beta_j\right)} & \text{if } 1 \leq i \leq v-1 \\ 1-\sum_{j=1}^{v-1} \pi_{tj} & \text{if } i = v \end{cases} \tag{2.11}$$

where x_{ti} $(1 \leq i \leq v)$ and u_t are the vectors of covariate values associated with the γ_i and β_j $(1 \leq j \leq v-1)$ parameter vectors respectively. We denote the entire set of parameters through two sets of vectors, namely γ and β with vector elements $|\gamma| = v$ and $|\beta| = v-1$, respectively. It is possible that the individual ϕ_{ti} are affected by different explanatory variables by specifying different x_{ti} in the logistic transformation (2.10). This is not sensible for the individual π_{ti} as they are all dependent on all parameters β_i $(1 \leq i \leq v-1)$ in the additive logistic transformation (2.11).

Estimation of the sets of parameters γ and β is carried out via direct maximization of the likelihood. Let us assume that a categorical time series or sample from chain $y_1, \ldots, y_N$ is observed. The joint probability of this realization is given by

$$P\left(Y_1 = y_1, \ldots, Y_N = y_N\right) = P\left(Y_1 = y_1\right) \prod_{t=2}^{N} P\left(Y_{t+1} = y_{t+1} | Y_t = y_t\right)$$

The log-likelihood is defined as

$$l(\gamma,\beta) = \log\Big(P(Y_1 = y_1)\Big) + \sum_{t=2}^{N} \log\Big(P(Y_{t+1} = y_{t+1} | Y_t = y_t)\Big)$$

The log-likelihood is maximized using a Quasi-Newton algorithm.

Autoregressive conditional multinomial

The autoregressive conditional multinomial model was introduced by Russell and Engle (2005) and Russel (1996). The presentation of the model follows the work of Russell and Engle. Let us consider the categorical electricity price series where each Y_t is a random variable which may take one of v states at time t, and consider a first-order Markov chain for the price change between different states.

The conditional distribution of Y_t is characterized by

$$\mathbf{\Pi_t} = \mathbf{PY_{t-1}}$$

where $\mathbf{\Pi_t}$ denote a vector of conditional probabilities that the jth element of $\mathbf{Y_t}$ takes the value of 1, $\mathbf{P}$ is the $v \times v$ transition matrix, and $\mathbf{Y_t}$ is the jth column of the identity matrix if the jth state of Y_t occurred.

Transition matrix $\mathbf{P}$ must satisfy that all elements are nonnegative and all columns must sum to unity. In general settings, $\mathbf{P}$ might vary over time. However, the restrictions on $\mathbf{P}$ are directly satisfied by simple estimators when the transition matrix is constant, and imposing those restrictions for the time-varying case is not an easy task. Therefore, Russell and Engle suggested a different approach.

Let one state of the Y_t variable, state v, be chosen as the base state. The logs of odds ratios of variable Y_t taking the mth state against the vth state are defined as

$$\log\left(\frac{\pi_{tm}}{\pi_{tv}}\right) = \log\left(\sum_{j=1}^{v} P_{mj}\mathbf{Y}_{(t-1)j}\right) - \log\left(\sum_{j=1}^{v} P_{vj}\mathbf{Y}_{(t-1)j}\right) = \sum_{j=1}^{v-1} P^*_{mj}\tilde{\mathbf{Y}}_{(t-1)j} + c_m$$

where π_{ij} is the j-th element of Π_i and P_{ij} is the j-th element of the i-th row of the $\mathbf{P}$ matrix. Here, c_m is a scalar constant and P^*_{mj} denotes ratio $\log\left(\frac{P_{mj}}{P_{vj}}\right)$; $\tilde{\mathbf{Y}}$ is now the $v-1$ dimensional vector.

$(v-1)\times 1$ elements π_{mj}/π_{vj} are collected in matrix P^* of dimension $(v-1)\times(v-1)$. The probabilities of Y_t variable taking states $1,\ldots,v$ may be computed directly from the matrix. The probability the variable Y_t taking state v follows the restriction that the probabilities sum to 1. Moving from direct modeling of the transition matrix towards modeling logs of odds ratios allows for the following dynamic specification of the odds ratio equation

$$h(\pi_t) = \sum_{j=1}^{p} A_j\Big(\mathbf{Y}_{t-j} - \Pi_{t-j}\Big) + \sum_{j=1}^{q} B_j h(\pi_{t-j}) + \gamma z_t$$

where $h()$ is the inverse logistic function, A_j and B_j denote the jth $(k-1)\times(k-1)$ parameter matrices, z_t contains exogenous variables and the constant, $\Pi_\mathbf{t}$ is the vector of transition probabilities at time t, and $\mathbf{Y_t}$ is the jth column of the identity matrix if the jth state of Y_t occurred.

The conditional probabilities of variable Y_t taking the kth state can be computed as follows:

$$\pi_{tk} = \frac{\exp\left(\sum_{j=1}^{p} A_{jk}\left(\mathbf{Y}_{t-j} - \Pi_{t-j}\right) + \sum_{j=1}^{q} B_{jk}\log\left(\frac{\pi_{t-j,k}}{\pi_{t-j,v}}\right) + \gamma z_t\right)}{1 + \sum_{m=1}^{v-1}\exp\left(\sum_{j=1}^{p} A_{jm}\left(\mathbf{Y}_{t-j} - \Pi_{t-j}\right) + \sum_{j=1}^{q} B_{jm}\log\left(\frac{\pi_{t-j,m}}{\pi_{t-j,v}}\right) + \gamma z_t\right)}$$

The model is relatively easy to interpret when $p = q = 1$, that is, for the $ACM(1,1)$ case. Matrix A determines the impact of previous periods on current transition probabilities and the eigenvalues of matrix B indicate how quickly the impact is weakening. If all transition probabilities are to be non-zero, the condition that all solutions of the equation in z, $\left|I - B_1 z - B_2 z^2 - \ldots - B_q z^q\right| = 0$ are smaller than 1 in absolute value must be satisfied.

The model is estimated with the maximum likelihood method. The likelihood can be represented as the product of the conditional densities. Letting π_{tj} denote the jth element of $\pi_\mathbf{t}$, the log-likelihood can be established as follows:

$$L = \sum_{t=1}^{N}\sum_{j=1}^{K}(Y_{tj}\log(\pi_{tj})) = \sum_{t=1}^{N} Y_t'\log(\pi_t) \tag{2.12}$$

The optimization of the likelihood function has to be done numerically with a proper optimization algorithm, e.g. the BHHH algorithm (see Berndt, Hall, Hall, and Hausman (1974) for the details). The ML estimates of the parameters are consistent and asymptotically normal under common regularity conditions.

2.6 Estimation and forecasting results

This section contains estimation results based on the models described in Section 2.5 and the data considered in Section 2.3. We study the influence of the explanatory variables on extreme price event occurrence in Elbas electricity prices. As is discussed in Section 2.4, it is anticipated that the variables relating to the load (the turnover at system price), temperature, water reservoir and wind power production ought to influence the occurrence of extreme price events. Therefore, we use the load, Elspot prices, heating degree days, water reservoir level and wind power production as explanatory variables.

Estimation results

Model 1

The results of the maximum likelihood estimation of the autoregressive ordered probit model parameters for Elbas electricity prices are exhibited in Table 2.4, where standard errors are again calculated using the typical sandwich form. All of the explanatory variables are statistically significant. In accordance with the results of other studies for different power markets (see Kanamura and Ohashi, 2007; Mount et al., 2006; Christensen et al., 2009, 2012, among others), the coefficients of the load are positive and significant, indicating that higher values of load are associated with higher states of electricity prices (spike regime) and lower values might lead to negative electricity jumps. The heating degree days also have a statistically significant effect on Elbas prices. Wind power, as it was expected, has statistically significant negative effect on Elbas electricity price. The Elbas price is more likely to be in a higher state (regime) with a higher load, lower water reservoir level, lower wind production and higher heating degree day. The Elspot electricity price also has a statistically significant positive effect on Elbas electricity prices. Statistically significant persistence has been confirmed: Elbas price from the previous hour has a statistically significant positive effect on the current Elbas price. The thresholds/intercept parameters are significantly different from each other, so the three price states should not be combined into one and suggested thresholds for price drops and spikes seem to be reasonable.

Each of the price drivers is statistically significant and might be useful when forecasting the occurrence of price jumps. Concerning the measure of goodness of fit, the categorical R^2 has been computed as well as the corrected categorical R^2. Their values are 0.9899 and 0.8648 respectively. They confirm that the model fits to the data very well.

Table 2.4. Maximum likelihood estimates of the autoregressive ordered probit model parameters

Coefficients	Estimate	Std. Error	z value	$Pr(>\lvert z \rvert)$
$y_t = 1$ (Elspot)	1.5108	0.0875	17.281	2.00E-16
$y_t = 2$ (Elspot)	2.9003	0.1127	25.726	2.00E-16
Load	0.3360	0.0290	11.606	2.00E-16
HDD	0.0095	0.0026	3.682	0.0002
water reservoir	-0.0020	0.0001	-2.906	0.0036
wind power	-0.0026	0.0001	-17.983	2.00E-16
$y_{t-1} = 1$ (Elbas)	2.2000	0.0353	62.329	2.00E-16
$y_{t-1} = 2$ (Elbas)	4.0941	0.0830	49.283	2.00E-16
0|1	2.4358	0.1187	20.52	2.00E-16
1|2	7.5855	0.1321	57.4	2.00E-16

Table 2.5 contains marginal effects of the autoregressive ordered probit model for Elbas electricity prices. Marginal effects are calculated as the mean of the independent variables. The results are in accordance with the expectations. The highest impact on the changes in probabilities associated with Elbas prices being in different states has a previous state of Elbas prices. If the past Elbas price state was in a normal state, then we will have, on average, a 46.7% decrease in the probability of the current Elbas price being in drop regime, a 46.5% increase in the probability of being in a normal price state, and a 0.2% increase in the probability of being in a spike state. If the past Elbas price was in a spike state, then we will have, on average, a 0.21% decrease in the probability of the current Elbas price being in a drop state, a 81.3% decrease in the probability of being in a normal price regime, and a 83.4% increase in the probability of being in a spike price state. If load (the turnover at system prices) increases by 10,000 MWh, then we will have, on average, a 1.5% decrease in the probability of the Elbas price being in a drop state, a 1.4% increase in the probability of being in a normal price regime, and a 0.1% increase in the probability of being in a spike regime. If the Elspot price is in a normal regime then we will have on average 24.7% decrease in the probability of Elbas price being in a drop regime, a 24.6% increase in the probability of being in a normal price state, and a 0.1% increase in the probability of being in a spike state. If the Elspot price state is in a spike state, then we will have, on average, a 2.1% decrease in the probability of Elbas price being in a drop state, a 39.1% increase in the probability of being in a normal price regime and a 41.3% increase in the probability of being in a spike price state. Changes in three of the explanatory variables - heating degree days, water reservoir level and wind power production - (although statistically significant) turn out to have a close-to-zero average effect on the probability of the price being in any of the considered states.

Table 2.5. Marginal effects

	$y_t = 0$	$y_t = 1$	$y_t = 2$
$y_t = 1$ (Elspot)	-0.247	0.246	0.001
$y_t = 2$ (Elspot)	-0.021	-0.391	0.413
Load	-0.015	0.014	0.001
HDD	0.000	0.000	0.000
water reservoir	0.000	0.000	0.000
wind power	0.000	0.000	0.000
$y_{t-1} = 1$ (Elbas)	-0.467	0.465	0.002
$y_{t-1} = 2$ (Elbas)	-0.021	-0.813	0.834

Model 2

Table 2.6 contains maximum likelihood estimates of parameters of the π_t and ϕ_t matrices. Three explanatory variables are used: Elspot prices, load, and the heating degree days. All explanatory variables turn out to be statistically significant, but their influence on the transition probabilities is different. π_t parameters control the probability of moving into each of the states unconditionally on the state from which it is moving, conditional on a move occurring. ϕ_t corresponds to the probability of moving to another state, although the new state may be the same as the old state.

Table 2.6. Maximum likelihood estimates of the parameters of ϕ_t and π_t matrices

Probability	Parameters	Explanatory variable	Estimate	St Error	z value
ϕ_1	γ_{11}	y_t (Elspot)	0.7237	0.3201	2.2605
	γ_{12}	Load	-0.4885	0.2704	-1.8063
	γ_{13}	HDD	0.0152	0.0233	0.652
ϕ_2	γ_{21}	y_t (Elspot)	0.6454	0.2685	2.4038
	γ_{22}	Load	-1.7767	0.1963	-9.0522
	γ_{23}	HDD	0.2409	0.0147	16.3883
ϕ_3	γ_{31}	y_t (Elspot)	-1.0336	0.222	-4.6554
	γ_{32}	Load	0.5409	0.1954	2.7686
	γ_{33}	HDD	-0.0195	0.0252	-0.7736
π_1	β_{11}	y_t (Elspot)	1.2018	0.3185	3.7733
	β_{12}	Load	-1.0118	0.2577	-3.9262
	β_{13}	HDD	0.2161	0.0191	11.3294
π_2	β_{21}	y_t (Elspot)	1.1486	0.2969	3.8683
	β_{22}	Load	-2.7392	0.242	-11.3211
	β_{23}	HDD	0.3695	0.0209	17.6475

Model 3

Table 2.7 contains the maximum likelihood estimates of the parameters of $ACM(1,1)$ with explanatory variables. The eigenvalues of matrix B indicate how quickly the impact fades. The B parameter matrix has the important property that all of its eigenvalues are less than one in absolute value. This means that the model is correctly specified in terms of the properties of the state change process, as the estimated process should assume every possible state infinitely many times in the future. Moreover, parameter matrix B implies that the price regime change process is of relatively high persistence, i.e. if some event increased the probability of a price regime change at time t, then at time $t+1$ the effect will still be present. The difference between the estimates of b_1 and b_2 is statistically significant. Matrix A determines the impact of previous periods on current transition probabilities. The estimates of the parameters on the diagonal of A, i.e. a_{11} and a_{22}, are both positive, which means that after each price regime increases the probability of another price regime increase. Positive value a_{12} implies that after the price regime changes, the probability of an opposite change increases. The parameters of explanatory variables are all statistically significant. Moreover, the positive values of

g_{12} and g_{22} imply that the probability of being in a normal or spike regime increases with the increase of turnover at system price. On the other hand, the negative values of g_{15} and g_{25} suggest that the probability of the price being in a normal or spike regime decreases when wind power production increases.

Table 2.7. Maximum likelihood estimates of $ACM(1,1)$ model parameters

	Estimate	Std Error	z value	$P(>\lvert z\rvert)$	Explanatory variable
c_1	-1.43	0.13	-11.00	2.1188E-27	
a_{11}	3.19	0.08	40.91	0	
a_{12}	1.23	0.27	4.51	1.4959E-05	
b_1	0.77	0.01	59.46	0	
g_{111}	1.48	0.11	13.12	1.5857E-38	$y_t = 1$ (Elspot)
g_{112}	1.20	0.16	7.44	3.7828E-13	$y_t = 2$ (Elspot)
g_{12}	0.35	0.04	10.11	2.4379E-23	Load
g_{13}	-0.01	0.00	-4.00	0.0001	HDD
g_{14}	-0.01	0.01	-1.50	0.1295	water reservoir
g_{15}	-0.02	0.00	-10.00	7.6946E-23	wind power
c_2	-12.03	18.56	-0.65	0.3234	
a_{21}	3.93	0.60	6.60	1.3563E-10	
a_{22}	4.49	0.64	7.07	5.7368E-12	
b_2	0.32	0.07	4.30	3.8985E-05	
g_{211}	4.53	18.58	0.24	0.3873	$y_t = 1$ (Elspot)
g_{212}	6.43	18.58	0.35	0.3757	$y_t = 2$ (Elspot)
g_{22}	1.30	0.24	5.39	1.9770E-07	Load
g_{23}	0.16	0.02	7.09	4.8143E-12	HDD
g_{24}	-0.21	0.06	-3.25	0.0020	water reservoir
g_{25}	-0.08	0.02	-5.47	1.2930E-07	wind power

Note:
$$h(\pi_t) = c + A\cdot(Y_{t-1} - \pi_{t-1}) + B\cdot h(\pi_{t-1}) + \gamma z_t$$
$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix},\ B = \begin{bmatrix} b_{11} & 0 \\ 0 & b_{22} \end{bmatrix}$$
$$\gamma = \begin{bmatrix} \gamma_{11} & \gamma_{12} \\ \ldots & \ldots \\ \gamma_{71} & \gamma_{72} \end{bmatrix}$$

2.7 Forecasting results

To assess the forecasting performance of the considered models, the following procedure is adopted. The parameters of the models are estimated using the sample period data (14 September 2009 to 31 December 2013). The observations from 1 January 2014 to 31

December 2015 are used for the forecast performance evaluation. For each model, one-step-ahead forecasts from an extended window are computed. The model parameters are reestimated once a new observation is included and then one-step-ahead forecasts are made. We do not attempt to forecast the explanatory variables. As the explanatory variables forecast, the actual values are taken. The prediction in the Markov model and an autoregressive conditional multinomial model is based on the mode of the conditional distribution of Y_{t+1}, given that $Y_t = y_j$ is y_j, the last realization of the categorical time series. The choice of hour-ahead forecasts is driven by the fact that the Elbas operates as a continuous-trading market for each hour interval. In the event of a price jump being forecasted for the next hour interval, effective risk management requires an immediate action on the part of retailers in order to mitigate the effects of the potential price jump. Retailers can reduce their reliance on the pool to meet their demand by activating the standby capacity that they may have available. Both hydroelectric and gasfired peaking plants can be brought from an idle state to full capacity in an hour. If this physical response is not available, an alternative strategy is to use the futures market, which trades in real time.

Table 2.8 presents the results of the forecasting exercise based on the data from 1 January 2014 to 31 December 2015. The out-of-sample time series does not contain any observations above 80 EUR/MWh. Therefore, we will focus only on the correct predicting of price drops and the full 3×3 confusion matrix can not be constructed. Table 2.8 contains information on the number of correct predictions of price jump occurrence based on four models. As a benchmark, two simple models are considered: the Markov model, in which transition matrix is constant in time, and the ordered probit model with seasonal dummies as explanatory variables. The Markov model was not able to correctly predict any of the price jump events. The ordered probit model with seasonal dummies was able to predict 2037 jump occurrences. This shows that it is necessary to incorporate some additional information in explanatory variables in order to be able to predict jump occurrences more accurately. The best forecasts of price jump occurrences are based on the ACM model. The model was able to correctly predict 2856 jump events from a total of 3527. The second model is the autoregressive ordered probit model with 2797 correct predictions. Slightly worse is the non-homogeneous Markov model with 2755 correct predictions of price jump occurrences. There is also a difference between the models with and without explanatory variables. We show that it is crucial for the price jumps prediction to include the explanatory variables in the model.

Table 2.8. Forecasting of price jump occurrences

Jump events forecasting		
autoregressive ordered probit (with explanatory variables)	Correct	2797
	False	730
autoregressive ordered probit (without explanatory variables)	Correct	1200
	False	2327
non-homogeneous Markov model (with explanatory variables)	Correct	2755
	False	772
non-homogeneous Markov model (without explanatory variables)	Correct	1158
	False	2369
autoregressive conditional multinomial (with explanatory variables)	Correct	2856
	False	671
autoregressive conditional multinomial (without explanatory variables)	Correct	1353
	False	2174
homogeneous Markov model (benchmark)	Correct	0
	False	3527
ordered probit model with seasonal dummies (benchmark)	Correct	2037
	False	1490

2.8 Conclusions

The accurate forecast of extreme price events is of high importance for risk management in the electricity sector. However, most of the electricity-pricing models are an adaptation of models for prices or returns from financial econometrics with a different degree of success. Opposed to the majority of contributions to this field, the current paper focuses on the forecasting of extreme price events, the occurrence of which is treated as a realization of categorical time series. An autoregressive ordered probit, a Markov model, and an autoregressive conditional multinomial model were used to analyze the drivers of the process and to forecast extreme price events. We show that it is crucial for correct jump occurrence prediction to build a statistical model which takes into account both persistence in the series and external information on price drivers. High loads were found to have a significant impact on the probability of a price spike and low loads were found to increase the probability of a price drop. Essentially, the persistence of the categorical price process was also confirmed to be significant in explaining the occurrence of extreme price events. The considered models provide hour-ahead forecasts of price jumps that are superior to the forecasts made by a memoryless model and an ordered probit model with seasonal dummies as explanatory variables. The best forecasts of the extreme price events are obtained based on the $ACM(1,1)$ model.

2.9 References

Aitchison, J., 1982. The Statistical Analysis of Compositional Data. Vol. 44. Journal of the Royal Statistical Society.

Amjady, N., Keynia, E., 2011. A new prediction strategy for price spike forecasting of day-ahead electricity markets. Applied Soft Computing 11, 4246–4256.

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, 371–382.

Berndt, E., Hall, B., Hall, R., Hausman, J., 1974. Estimation and Inference in Nonlinear Structural Models. Annals of Economic and Social Measurement 3, 653–665.

Botterud, A., Bhattacharyya, A., Ilic, M., 2002. Futures and spot prices–an analysis of the scandinavian electricity market. In: Proceedings of North American Power Symposium. pp. 1–8.

Cartea, A., Figueroa, M., Geman, H., 2012. Modelling Electricity Prices with Forward Looking Capacity Constraints. Applied Mathematical Finance 16, 103–122.

Chen, C., Liu, L., 1993a. Forecasting time series with outliers. Journal of Forecasting 12 (1), 13–35.

Chen, C., Liu, L., 1993b. Joint Estimation of Model Parameters and Outlier Effects in Time Series. Journal of the American Statistical Association 88 (421), 284–297.

Christensen, T., Hurn, A., Lindsay, K., 2012. Forecasting spikes in electricity prices. International Journal of Forecasting 28, 400–411.

Christensen, T., Hurn, S., Lindsay, K., 2009. It never rains but it pours: modeling the persistence of spikes in electricity prices. The Energy Journal 207, 543–556.

Clements, A., Fuller, J., Hurn, S., 2013. Semi-parametric forecasting of spikes in electricity prices. Economic Record 89, 508–521.

Escribano, A., Pena, J., Villaplana, P., 2002. Modeling electricity prices: international evidence. Ph.D. thesis, Oxford Bulletin of Economics and Statistics.

Eydeland, A., Wolyniec, K., 2003. Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging. Wiley Finance, Hoboken, N.J.

Fama, E., French, K., 1987. Dividend yields and expected stock returns. Journal of Financial Economics 22, 3–25.

Fehr, N. V. D., Amundsen, E., Bergman, L., 2005. The nordic market: signs of stress? The Energy Journal Issue on European Electricity Liberalization, 71–98.

Fokianos, K., Kedem, B., 2002. Regression Models for Time Series Analysis. Wiley Finance.

Fokianos, K., Kedem, B., 2003. Regression Theory for Categorical Time Series. Statistical Science 18, 357–376.

Forsund, F., Hoel, M., 2004. Properties of non-competitive electricity market dominated by hydroelectric power. SSRN Electronic Journal 30, 1116–1157.

Foster, S., Bravington, M., Williams, A., Althaus, F., Laslett, G., Kloser, R., 2009. Analysis and prediction of faunal distributions from video and multi-beam sonar data using markov models. Environmetrics 20, 541–560.

Geman, H., Roncoroni, A., 8 2006. Understanding the fine structure od electricity prices. The Journal of Business 79, 1225–1262.

Gjolberg, O., Johnsen, T., Sept 2001. Electricity futures: Inventories and price relationships at nord pool. Tech. rep., Norwegian School of Economics and Business Administration, working paper, Norwegian School of Economics and Business Administration.

Harris, C., 2006. Electricity Markets: Pricing, Structures and Economics. Wiley.

Hellstrom, J., Lundgren, J., Yu, H., 2012. Why do electricity prices jump? empirical evidence from the nordic electricity market. Energy Economics 34, 1774–1781.

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

Johansen, S., Nielsen, B., 2016. Asymptotic Theory of Outlier Detection Algorithms for Linear Time Series Regression Models. Scandinavian Journal of Statistics 43 (2), 321–348.

Jong, C. D., 2006. The nature of power spikes: A regime-switch approach. Nonlinear Analysis of Electricity Prices 10, 27–46.

Kanamura, T., Ohashi, K., 2007. A structural model for electricity prices with spikes: Measurement of spike risk and optimal policies for hydropower plant operation. Energy Economics.

Knittel, C., Roberts, M., 2005. An empirical examination of restructured electricity prices. Energy Economics 27, 791–817.

Krichene, N., 2002. World crude oil and natural gas: a demand and supply model. Energy Economics 24, 557–576.

Li, X., Sailor, D., 1995. Electricity use sensitivity to climate and climate change. World Resource Review.

Lindstrom, E., Noren, V., Madsen, H., 2015. Consumption management in the nord pool region: A stability analysis. Applied Energy 146, 239–246.

Lu, X., Dong, Z., Li, X., 2005. Electricity market price spike forecast with data mining techniques. Electric Power Systems Research 73 (1), 19–29.

Maryniak, P., Weron, R., 2014. Forecasting the occurence of electricity price spikes in the UK power market. HSC Research Reports 14.

Mount, T., 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.

Munoz, J. R., Rosen, J., Sailor, D., 1998. Natural gas consumption and climate: a comprehensive set of predictive state-level models for the united states. Energy the International Journal 23, 91–103.

Pardo, J., Ridal, M., Murtagh, D., Cernicharo, J., 2002. Microwave temperature and pressure measurements with the odin satellite: I. observational method. Can. J. Phys. 80, 443–454.

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

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

Russel, J., 1996. Econometric Analysis of Irregularly-Spaced TransactionData Using a New Class of Accelerated Failrue Time Models with Applicationsto Financial Transaction Data. Ph.D. thesis, University of California, San Diego.

Russell, J. R., Engle, R., 2005. A discrete-state continuous-time model of financial transactions prices and times: The autoregressive conditional multinomial: Autoregressive conditional duration model. Journal of Business & Economic Statistics 23 (2), 166–180.

Villar, J., Joutz, F., 2006. The relationship between crude oil and natural gas prices. Eia manuscript, Energy Information Administration.

Voronin, S., Partanen, J., 2013. Price Forecasting in the Day-Ahead Energy Market by an Iterative Method with Separate Normal Price and Price Spike Frameworks. Energies 6, 5897–5920.

Voronin, S., Partanen, J., Kauranne, T., 2014. A hybrid electricity price forecasting model for the nordic electricity spot market. International Transactions on Electrical Energy Systems 24, 736–760.

Weron, R., 2006. Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach. Wiley, Chichester.

Weron, R., 2014. Electricity price forecasting: A review of the state-of-the-art with a look into the future. International Journal of Forecasting 30 (4), 1030–1081.

CHAPTER 3

A REGIME-SWITCHING STOCHASTIC VOLATILITY MODEL FOR FORECASTING ELECTRICITY PRICES

Oskar Knapik
Aarhus University and CREATES
Peter Exterkate
University of Sydney and CREATES

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, iii) introducing an universal technique for model comparison. We propose new regime-switching stochastic volatility model with three regimes (negative jump, normal price, positive jump (spike)) where the transition matrix depends on explanatory variables. Bayesian inference is explored in order to obtain predictive densities. The main focus of the paper is on short-time density forecasting in Nord Pool intraday market. We show that the proposed model outperforms several benchmark models at this task.

3.1 Introduction

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

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

When building a model for electricity prices, one of the crucial steps is to find an appropriate description of 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. Some nonlinearities of these series are a nonconstant variance, and generally they are characterized by the clustering of large shocks, or heteroskedasticity. It is well documented that electricity prices exhibit volatility clustering (see Karakatsani and Bunn (2008) among the others). The "spiky" character of electricity prices suggests that there exists a nonlinear switching mechanism between normal and low/high 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 nonlinear and complex structure (see Andreasen and Dahlgren, 2006; Geman and Roncoroni, 2006; Handika, Truong, Trück, and Weron, 2014; Heydari and Siddiqui, 2010; Huisman and Mahieu, 2003; Kanamura and Ohashi, 2008; Kosater and Mosler, 2006; Mount, Ning, and Cai, 2006, among others). MS models are successfully applied by many researchers for electricity prices modeling (see Janczura and Weron, 2010a; Haldrup and Nielsen, 2006b, 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 stochastic volatility error term. The main focus of this research is on short-time density forecasting of electricity prices. Although important, this topic is however barely touched in the electricity prices forecasting literature (see Weron, 2014). Serinaldi (2011) forecasts the distribution of electricity prices using a GAMLSS approach, but computes and discusses only predictive intervals. Huurman, Ravazzolo, and Zhou (2012) consider GARCH-type time-varying volatility models and find that models that 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, Pinson, Madsen, and Nielsen (2014) develop a semi-parametric methodology for generating densities of day-ahead electricity prices in Western Denmark (Nord Pool).

The actuality and importance of electricity price forecasting is further evidenced by a special issue of the International Journal of Forecasting (Volume 32, Issue 3, Pages 585-1102, July–September 2016) dedicated to that topic, and by the energy forecasting competition organized by this journal. The competition and the current status of the probabilistic energy forecasting research is summarized in the paper of Hong, Pinson, Fan, Zareipour, Troccoli, and Hyndman (2016). One of the presented approaches by Maciejowska, Nowotarski, and Weron (2016) introduces new methodology involving quantile regression to average large numbers of point forecasts and principal component analysis (PCA) to extract the major factors driving the individual forecasts. Their approach outperforms both of benchmarks autoregressive exogenous (ARX) model and the quantile regression averaging (QRA) without PCA based on comprehensive evaluation with the unconditional coverage, the conditional coverage and the Winkler score.

The paper explores Bayesian inference in order to construct predictive densities for future electricity prices. We introduce also the electricity price modeling and forecasting literature a natural, universal method for model comparisons via predictive Bayes factors. Bayesian approaches have been used in the context of electricity prices modelling by several authors. Panagiotelis and Smith (2008) use a first order vector autoregressive model with exogenous effects and a skew t distributed disturbance for hourly Australian electricity spot prices. They use a Bayesian 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 with periodic autoregressions (PAR) in both the level and log-volatility process. They include demand and day types as exogenous explanatory variables in both the mean and log-volatility equations. They confirm that there is a nonlinear relationship between demand and mean prices and construct the predictive density of prices evaluated over a horizon of one week. Our work can be seen as extending this study in two ways, with a Markov-switching structure to flexibly accommodate such nonlinearities, and by allowing for many more predictors.

The paper is divided into five sections. Section 3.1 introduces. Section 3.2 describes the data and the main electricity price drivers. Section 3.3 presents the model and Bayesian inference. Section 3.4 shows the empirical and forecasting results. Section 3.5 concludes.

3.2 Data

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

There is about 380 companies from 20 countries that trade on the Nord Pool Spot (Elspot) market including both producers and large consumers, for a trading volume of approximately 500 terawatt hours in 2015. Within the Nord Pool Spot, Elspot is the auction market for day-ahead electricity delivery. The Nord Pool Spot web-based trading system enables participants to submit bids and offers for each individual hour of the next day. Orders can be made between 08:00 and 12:00 a.m. 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 the curves constitutes the system price for each hour (quoted per megawatt hour, MWh). The hourly prices are announced to the market at 12:42 CET and contracts are invoiced between buyers and sellers between 13:00 and 15:00. 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 CET and delivery the next day. On the intraday market, buyers and sellers can trade volumes close to real time to bring the market back in 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, first-served principle, where best prices come first – highest buy price and lowest sell price. 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 two Nord Pool markets: Elspot and Elbas. The study is conducted using hourly electricity prices for the whole area (system prices) from Elspot market $spot_t$ and corresponding hourly volume-weighted average prices from Elbas market y_t. The data covers the period from 1 January 2013 to 31 December 2014 (17520 observations). The focus of this research is on modeling and forecasting of electricity prices from the Elbas market. We will present and evaluate out-of-sample forecasts for 2 January 2015 (24 observations), the first working day after the in-sample period.

Figure 3.1 presents the hourly time series of volume-weighted electricity prices from Elbas market in the period 1/1/2013 till 31/12/2014. The prices exhibit typical characteristics, including seasonality and spikes.

Figure 3.1. Hourly electricity prices from Elbas power market, y_t

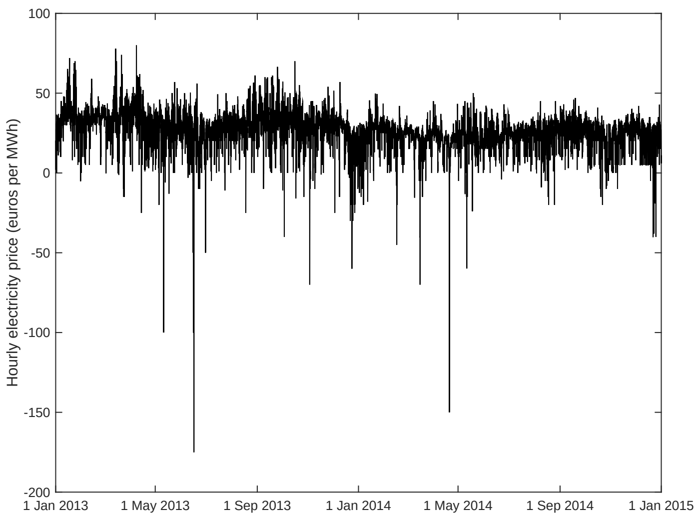

In order to model the hourly volume-weighted average prices from Elbas market y_t, we consider the following explanatory variables considered to be the main price drivers in the Nord Pool power market:

- hourly Elspot electricity system price, $spot_t$,
- turnover at system price from Elspot market, $load_t$,
- water reservoir level, res_t,
- heating degrees days, hdd_t,
- wind power production, $wind_t$,
- seasonal component, $seas_t$[1].

In the Nord Pool electricity market, about 53% of power production is generated from hydropower reservoirs. The influence of water reservoir levels in electricity prices at Nord Pool has been studied by (Gjolberg and Johnsen, 2001), (Botterud, Bhattacharyya, and Ilic, 2002), (Førsund and Hoel, 2004) and (von der Fehr, Amundsen, and Bergman, 2005). The researchers conclude that hydropower reservoir levels are an important factor that explains futures and spot prices. The ability of the Nordic power system to

[1] The seasonality component in the electricity price series is captured by sine and cosine terms taken with hourly, daily and weekly frequencies.

store energy in hydro reservoirs causes less variation in the Nordic price structure, than that of for example Germany. Inflow during summer and in periods with low demand can be used in the winter. The data on hydropower reservoir levels are collected from the first week of 2013 to the end of the last week of 2014. Reservoirs are taken as a percentage of the total hydropower capacity available in the Nord Pool area. 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). Reservoirs taken into account from Sweden and Finland are those after their integration in the Nord Pool market. The data is published on the weekly frequency[2]. The seasonality of reservoir levels has a highly important influence on electricity spot prices.

Temperature is the main price driver in the Nordic countries. Cold temperatures increase heat demand, since electricity is very much used for heating in the Nordic countries. Colder temperatures usually increase prices because of higher power demand. However, in special cases, combined heat and power plants where heat is the primary product, the demand for the heat could trigger secondary electricity production and causes 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, Rosen, and Muñoz (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 explicative power in their modeling; see, for example, (Peirson and Henley, 1994), (Peirson and Henley, 1998).Heating degree day (HDD) is a variable that shows the demand for energy needed for heating. It is taken from measurements of outside air temperature. The heating requirements for a specific structure at a specific place tend to be directly proportional to the number of HDD at that location. In this study we will consider 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).

Wind power production is also an important electricity price driver. Due to the fact that there is no fuel cost for production and unpredictability, additional wind energy can lead to price decrease. This type of energy may in some cases cause even negative prices in hours with low demand and additional supply. On the other hand, when wind production falls short of expected values, it can trigger high prices, both in the Day-Ahead and Intra-day markets.

Finally, Elspot electricity market is another significant source of information about Elbas electricity prices. The same situation is in futures markets, where the basis is the difference between futures price and the underlying spot price. Figure 3.2 presents

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

hourly system prices from Elspot power market.

Figure 3.2. Hourly electricity prices from Elspot market, $spot_t$

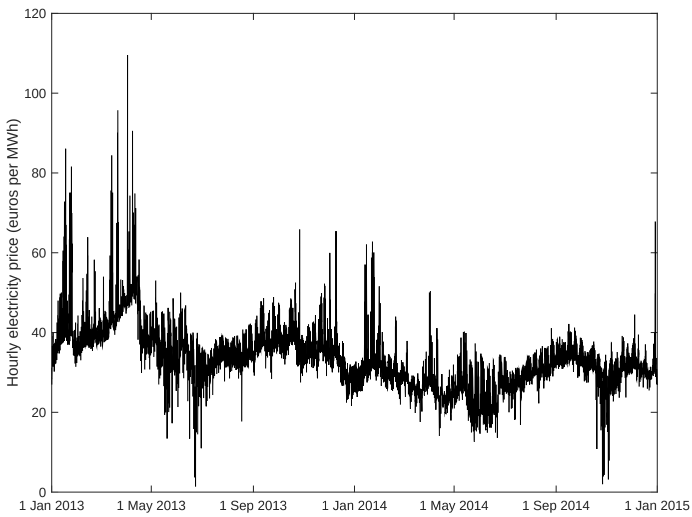

3.3 Model

The Markov regime-switching (MS) model allows for 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 (see Janczura and Weron, 2010a).

First, the number of regimes has to be chosen. The important advantage of the MS models over the alternatives is no need for explicitly specifying threshold variable and level for the regimes, and therefore they are preferred for modeling risk purposes. There is no fundamental reason for considering specific number of regimes for electricity prices modeling (see Janczura and Weron (2010a)). However, almost all published papers consider only 2-regime models due to the computational convenience. Additional to the base regime, a spike (or excited) regime was introduced to capture extreme price behavior. Karakatsani and Bunn (2008) introduced third regime for capturing the most extreme prices. Also, the existence of an additional "down-spike" or "drop" regime can by justified for many of the very low prices. This research considers three regimes (down-spike (drop), normal/base and spike) Markov regime-switching stochastic-volatility model for electricity prices. We also introduce a data-driven mechanism of switching between different regimes in the form of an ordered probit model. Our approach enables to extend the number of regimes easily.

Secondly, the model defining the price dynamics in each of the regimes has to be selected. The base regime is usually modeled by a mean-reverting AR (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, Bierbrauer, and Trück, 2004), (Bierbrauer, Trück, and Weron, 2004), exponential (Bierbrauer, Menn, Rachev, and Trück, 2007), heavy tailed (Weron, 2009) and non-parametric (Eichler and Türk, 2013) random variables, to mean-reverting diffusions with Poisson jumps (Arvesen, Medbø, Fleten, Tomasgard, and Westgaard, 2013; de Jong, 2006; Keles, Hartel, Möst, and Fichtner, 2012; Mari, 2008). One of the advantages of the regime-switching framework is that we can explicitly model the short-lived characteristics of power spike. We consider ARX-SV models 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 that independent ones. However, independent regimes model enables greater flexibility and seems to be a natural choice for a process which significantly changes its dynamics. We follow the second mentioned approach. An empirical comparison in the paper (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 capture more accurately each of the characteristics of the best structure model considered in Janczura and Weron (2010a) . In this section we provide the details on Markov-switching stochastic volatility model for electricity prices. We first introduce the stochastic volatility model that exists within each of the three regimes (negative jump, regular, positive jump), followed by a description of the transition dynamics between the regimes. Finally, we derive a Gibbs sampler that we use for parameter estimation and forecasting.

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 at time $t-1$, when the forecast is being made. 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 negative jump, regular, and positive jump regimes, respectively. All parameters $(\alpha,\beta,\gamma,\delta,\tau)$, which are introduced below, are collected in θ.

We specify our stochastic volatility model as

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

The regressors in the mean equation are $x_t' = \left(1, y_{t-1}, \ldots, y_{t-p}, spot_t, seas_t'\right)$, and in the volatility equation, $z_t' = \left(1, \log\sigma_{t-1}, \ldots, \log\sigma_{t-q}, spot_t, seas_t'\right)$. 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

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

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, Polson, and Rossi (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 $\log\sigma_t \left| R_t \right. \sim \mathcal{N}\left(0, \tau_{R_t}^{-1}\right)$ independently for $t = 1 - q, \ldots, 0$.

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 $P\left[R_{t+1} \left| R_t = r, \theta \right.\right]$ are given implicitly by

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

where ε_t is i.i.d. $\mathcal{N}\left(0, 1\right)$, $w_t' = \left(spot_t, res_t, load_t, hdd_t, wind_t\right)$, 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 regime-specific $T_r \times K$ 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 $R_{2-q}, \ldots, R_0$.

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 $\mathcal{N}\left(z'_t \gamma_{R_t} - \tau^{-1}_{R_t}, \tau^{-1}_{R_t}\right)$, truncated to the interval $\left(\frac{1}{2} \log\left(\left(y_t - x'_t \beta_{R_t}\right)^2 \Big/ \left(2 v_t\right)\right), \infty\right)$; for further details see Damien, Wakefield, and Walker (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 distribution of R_t has support $\{1, 2, 3\}$, with probability for state r proportional to the product $p\left(R_t = r \,\middle|\, R_{t-1}, \theta\right) \cdot$ $p\left(R_{t+1} \,\middle|\, R_t = r, \theta\right) \cdot p\left(\log \sigma_t \,\middle|\, \log \sigma_{t-1}, \ldots, \log \sigma_{t-q}, \theta, R_t = r\right) \cdot p\left(y_t \,\middle|\, \sigma_t, y_{t-1}, \ldots, y_{t-p}, \theta, R_t = r\right)$.

The factors in this expression can be explicitly computed as

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

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

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

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

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 $R_{t+1} = 1$, $\left(0, \alpha_{R_t}\right)$ if $R_{t+1} = 2$, and $\left(\alpha_{R_t}, \infty\right)$ 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 $r = 1, 2, 3$.

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 $R^{*(d)}_{T+1} \sim \mathcal{N}\left(w'_T\delta^{(d)}_{R^{(d)}_T}, 1\right)$ to find $R^{(d)}_{T+1}$, which is 1 if $R^{*(d)}_{T+1} < 0$, 2 if $0 \le R^{*(d)}_{T+1} \le \alpha^{(d)}_{R^{(d)}_T}$, and 3 otherwise. Finally, we draw $\log\sigma^{(d)}_{T+1} \sim \mathcal{N}\left(z'_{T+1}\gamma^{(d)}_{R^{(d)}_{T+1}}, \tau^{(d)-1}_{R^{(d)}_{t+1}}\right)$, and use it to draw $y^{(d)}_{T+1} \sim \mathcal{N}\left(x'_{T+1}\beta^{(d)}_{R^{(d)}_{T+1}}, \sigma^{(d)2}_{T+1}\right)$. The empirical distribution formed by these draws after discarding D_0 initial burn-in draws, $F(c) = \dfrac{1}{D - D_0}\sum_{d=D_0+1}^{D} \mathbf{1}\left\{y^{(d)}_{T+1} \le c\right\}$, approximates the CDF of $y_{T+1} \,|\, y$. A kernel density estimate of the corresponding PDF is used to visualize this distribution in the Results section below.

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

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

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. A number greater than one indicates a preference for the model in the numerator.

For one-step-ahead forecasts, we may approximate the predictive density evaluated at the realized price y_{T+1} using

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

A similar iterative procedure as outlined above for density forecasting can then be used to approximate $p\left(y_{T+2} \,|\, y_{T+1}, y\right), \dots, p\left(y_{T+24} \,|\, y_{T+23}, \dots, y_{T+1}, y\right)$, 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 again given by

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

3.4 Results

We study the hourly volume-weighted electricity prices from Elbas power market in order to understand the data generating mechanism and examine the proposed model. For that reason 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). Formal Bayesian model comparison in terms of the predictive adequacy is measured by predictive Bayes factors.

Table 3.1. The model specifications considered in the empirical study.

Mnemonic	Restrictions on the model introduced in Section 3.3
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

Full-sample results

As a preliminary check, we run the classical neural network test for neglected nonlinearity introduced by Lee, White, and Granger (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 expected based on the literature surveyed in Section 3.1. 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 results obtained within four model specifications: ARX, ARX-SV, MS-ARX and MS-ARX-SV (see Table 3.1), estimated for hourly electricity prices (see Figure 3.1) for every hour of 2015.[3] 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), which are not reported here but may be obtained from the authors upon request.

We report on the model comparison first. Relevant quantities, including the predictive density values and predictive Bayes factors, are displayed in Table 3.2. It is clear that all three nonlinear models provide a better out-of-sample forecasting performance than the simple linear ARX model. To quantify the performance differences, we follow the interpretation of Bayes factors 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 simple 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 the presence of both stochastic volatility (ARX-SV versus ARX) and especially regime switching (MS-ARX versus ARX) is very strong. We conclude that the overall forecasting performance of the MS-ARX, ARX-SV, and MS-ARX-SV models is much better than that of the ARX model.

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.

[3] 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 3.2. Predictive performance measures in the empirical study.

		Log predictive Bayes factor against…			
Model	Predictive log density	MS-ARX-SV	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				

Subsample results

In order to better understand the differences in performance between the four models under consideration, as well as to highlight the ease of formally obtaining predictive densities within our framework, we repeat the analysis that we just performed restricted to four specific days in 2015. These days were selected to be the ones where each of the four models performed best relative to its competitors: Monday 12 January (ARX performs best on this day), Sunday 18 January (MS-ARX-SV), Sunday 10 May (MS-ARX), and Tuesday 25 August (ARX-SV).

Tables 3.3–3.6 below are analogous to the full-sample Table 3.2. We observe that, except in the special case where ARX performs best (Table 3.3), this simple linear benchmark performs far worse than all nonlinear competitors (Tables 3.4–3.6). When our full MS-ARX-SV model performs best (Table 3.4), it does so by a very wide margin; note that the smallest Bayes factor is $\exp(12.3544) \approx 2 \times 10^5$ already.

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 3.5 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 3.6), where ARX-SV is the preferred model, the other nonlinear models MS-ARX-SV and MS-ARX are virtually indistinguishable, with a predictive Bayes factor of $\exp(0.8272) \approx 2$. These results confirm our intuition based on the full-sample results: leaving out Markov switching when we need it has a larger negative impact on forecast accuracy than leaving out stochastic volatility when we need it.

Table 3.3. Predictive performance measures in the empirical study, Monday 12 January 2015.

		Log predictive Bayes factor against...			
Model	Predictive log density	MS-ARX-SV	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				

Table 3.4. Predictive performance measures in the empirical study, Sunday 18 January 2015.

		Log predictive Bayes factor against...			
Model	Predictive log density	MS-ARX-SV	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 3.5. Predictive performance measures in the empirical study, Sunday 10 May 2015.

		Log predictive Bayes factor against...			
Model	Predictive log density	MS-ARX-SV	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				

Table 3.6. Predictive performance measures in the empirical study, Tuesday 25 August 2015.

		Log predictive Bayes factor against...			
Model	Predictive log density	MS-ARX-SV	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				

To further investigate what drives these results, Figures 3.3–3.6 show the predictive densities obtained for two selected hours on these four days, one in the afternoon ($h = 15$) and one in the evening ($h = 21$). The ex-post realized price is also included in each of these figures. 12 January (Figure 3.3) was a relatively uneventful day, for which the linear ARX model was "good enough" and its nonlinear extensions turned out to be needless complications.

The MS-ARX-SV model performed best on 18 January (Figure 3.4). 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 10 May (Figure 3.5), which explains the good performance of Markov switching models on this date. Finally, 25 August (Figure 3.6) saw large price fluctuations but no jumps, so that stochastic volatility was an important model component on that day.

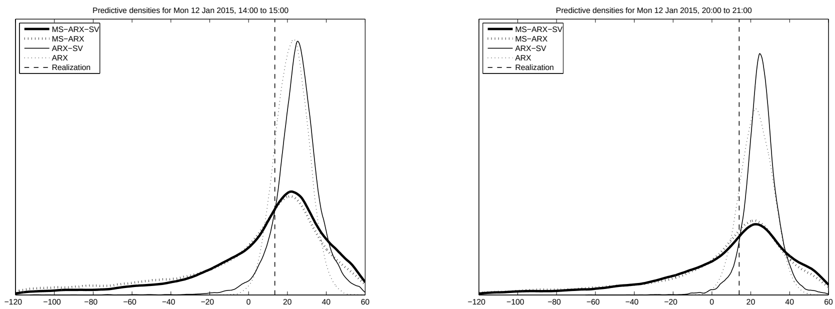

(a) Predictive density for $h = 15$ (b) Predictive density for $h = 21$

Figure 3.3. Predictive densities for two selected hours on Monday 12 January 2015.

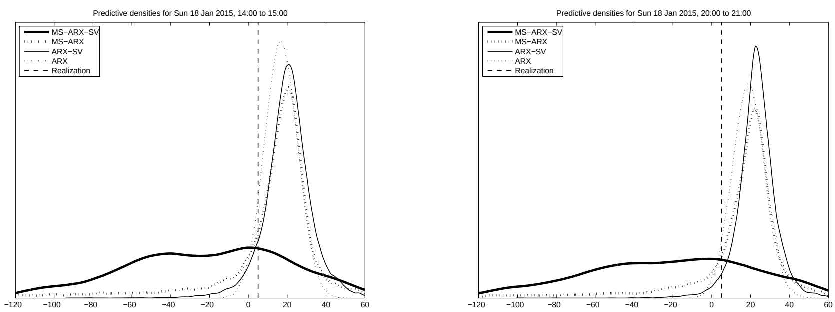

(a) Predictive density for $h = 15$ (b) Predictive density for $h = 21$

Figure 3.4. Predictive densities for two selected hours on Sunday 18 January 2015.

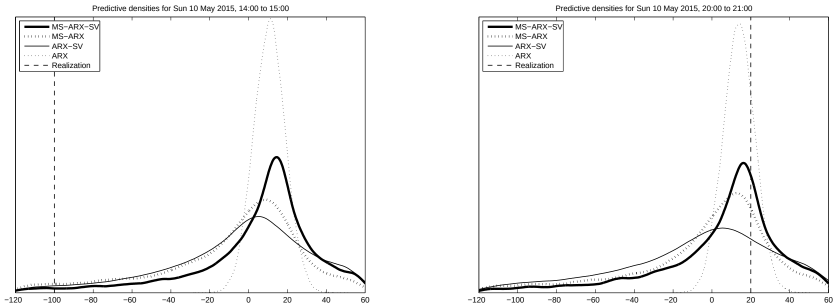

(a) Predictive density for $h = 15$ (b) Predictive density for $h = 21$

Figure 3.5. Predictive densities for two selected hours on Sunday 10 May 2015.

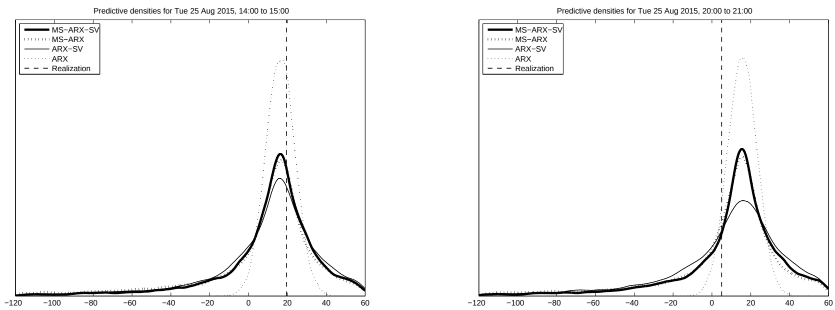

(a) Predictive density for $h = 15$ (b) Predictive density for $h = 21$

Figure 3.6. Predictive densities for two selected hours on Tuesday 25 August 2015.

3.5 Conclusions

The goal of this research was to address fundamental questions within the area of forecasting electricity prices. We proposed a new regime-switching stochastic volatility model with three regimes which takes into account fundamental price drivers. We show how the predictive densities of future electricity prices can be formally constructed via Bayesian inference. Moreover, we introduced a universal method (within the Bayesian framework) for model comparisons, predictive Bayes factors, to the electricity price forecasting literature. Based on this measure, we showed that the Markov switching structure and the stochastic volatility component of our model both contribute to its improved forecasting performance in terms of short-time density forecasting in the Nord Pool intraday market, relative to a model that lacks such features.

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 3.1 confirms this intuition, but a more thorough investigation is required.

3.6 References

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

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

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

Bierbrauer, M., Menn, C., Rachev, S., 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. In: Bubak, M., van Albada, G., Sloot, P., Dongarra, J. (Eds.), Proceedings of the Fourth International Conference on Computational Science. Springer, Berlin/Heidelberg, pp. 859–867.

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

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

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

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

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

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

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

Filardo, A., Gordon, S., 1998. Business cycle durations. Journal of Econometrics 85, 99–123.

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

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

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

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

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

Haldrup, N., Nielsen, M., 2006b. A regime switching long memory model for electricity prices. Journal of Econometrics 135, 349–376.

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

Heydari, S., Siddiqui, A., 2010. 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.

Hong, T., Pinson, P., Fan, S., Zareipour, H., Troccoli, A., Hyndman, R., 2016. Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond. International Journal of Forecasting 32 (3), 896 – 913.

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

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

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

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

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

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

Janczura, J., Weron, R., 2009. 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.

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

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

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

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

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

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

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, 53–84.

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

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

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

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

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

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

Maciejowska, K., Nowotarski, J., Weron, R., 2016. Probabilistic forecasting of electricity spot prices using factor quantile regression averaging. International Journal of Forecasting 32 (3), 957–965.

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

Mount, T., 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.

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

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

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

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

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

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

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

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

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

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

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

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

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

DEPARTMENT OF ECONOMICS AND BUSINESS ECONOMICS

AARHUS UNIVERSITY
SCHOOL OF BUSINESS AND SOCIAL SCIENCES
www.econ.au.dk

PhD dissertations since 1 July 2011

2011-4	Anders Bredahl Kock: Forecasting and Oracle Efficient Econometrics
2011-5	Christian Bach: The Game of Risk
2011-6	Stefan Holst Bache: Quantile Regression: Three Econometric Studies
2011:12	Bisheng Du: Essays on Advance Demand Information, Prioritization and Real Options in Inventory Management
2011:13	Christian Gormsen Schmidt: Exploring the Barriers to Globalization
2011:16	Dewi Fitriasari: Analyses of Social and Environmental Reporting as a Practice of Accountability to Stakeholders
2011:22	Sanne Hiller: Essays on International Trade and Migration: Firm Behavior, Networks and Barriers to Trade
2012-1	Johannes Tang Kristensen: From Determinants of Low Birthweight to Factor-Based Macroeconomic Forecasting
2012-2	Karina Hjortshøj Kjeldsen: Routing and Scheduling in Liner Shipping
2012-3	Soheil Abginehchi: Essays on Inventory Control in Presence of Multiple Sourcing
2012-4	Zhenjiang Qin: Essays on Heterogeneous Beliefs, Public Information, and Asset Pricing
2012-5	Lasse Frisgaard Gunnersen: Income Redistribution Policies
2012-6	Miriam Wüst: Essays on early investments in child health
2012-7	Yukai Yang: Modelling Nonlinear Vector Economic Time Series
2012-8	Lene Kjærsgaard: Empirical Essays of Active Labor Market Policy on Employment
2012-9	Henrik Nørholm: Structured Retail Products and Return Predictability
2012-10	Signe Frederiksen: Empirical Essays on Placements in Outside Home Care
2012-11	Mateusz P. Dziubinski: Essays on Financial Econometrics and Derivatives Pricing

2012-12 Jens Riis Andersen: Option Games under Incomplete Information

2012-13 Margit Malmmose: The Role of Management Accounting in New Public Management Reforms: Implications in a Socio-Political Health Care Context

2012-14 Laurent Callot: Large Panels and High-dimensional VAR

2012-15 Christian Rix-Nielsen: Strategic Investment

2013-1 Kenneth Lykke Sørensen: Essays on Wage Determination

2013-2 Tue Rauff Lind Christensen: Network Design Problems with Piecewise Linear Cost Functions

2013-3 Dominyka Sakalauskaite: A Challenge for Experts: Auditors, Forensic Specialists and the Detection of Fraud

2013-4 Rune Bysted: Essays on Innovative Work Behavior

2013-5 Mikkel Nørlem Hermansen: Longer Human Lifespan and the Retirement Decision

2013-6 Jannie H.G. Kristoffersen: Empirical Essays on Economics of Education

2013-7 Mark Strøm Kristoffersen: Essays on Economic Policies over the Business Cycle

2013-8 Philipp Meinen: Essays on Firms in International Trade

2013-9 Cédric Gorinas: Essays on Marginalization and Integration of Immigrants and Young Criminals – A Labour Economics Perspective

2013-10 Ina Charlotte Jäkel: Product Quality, Trade Policy, and Voter Preferences: Essays on International Trade

2013-11 Anna Gerstrøm: World Disruption - How Bankers Reconstruct the Financial Crisis: Essays on Interpretation

2013-12 Paola Andrea Barrientos Quiroga: Essays on Development Economics

2013-13 Peter Bodnar: Essays on Warehouse Operations

2013-14 Rune Vammen Lesner: Essays on Determinants of Inequality

2013-15 Peter Arendorf Bache: Firms and International Trade

2013-16 Anders Laugesen: On Complementarities, Heterogeneous Firms, and International Trade

2013-17 Anders Bruun Jonassen: Regression Discontinuity Analyses of the Disincentive Effects of Increasing Social Assistance

2014-1 David Sloth Pedersen: A Journey into the Dark Arts of Quantitative Finance

2014-2 Martin Schultz-Nielsen: Optimal Corporate Investments and Capital Structure

2014-3 Lukas Bach: Routing and Scheduling Problems - Optimization using Exact and Heuristic Methods

2014-4 Tanja Groth: Regulatory impacts in relation to a renewable fuel CHP technology: A financial and socioeconomic analysis

2014-5 Niels Strange Hansen: Forecasting Based on Unobserved Variables

2014-6 Ritwik Banerjee: Economics of Misbehavior

2014-7 Christina Annette Gravert: Giving and Taking – Essays in Experimental Economics

2014-8 Astrid Hanghøj: Papers in purchasing and supply management: A capability-based perspective

2014-9 Nima Nonejad: Essays in Applied Bayesian Particle and Markov Chain Monte Carlo Techniques in Time Series Econometrics

2014-10 Tine L. Mundbjerg Eriksen: Essays on Bullying: an Economist’s Perspective

2014-11 Sashka Dimova: Essays on Job Search Assistance

2014-12 Rasmus Tangsgaard Varneskov: Econometric Analysis of Volatility in Financial Additive Noise Models

2015-1 Anne Floor Brix: Estimation of Continuous Time Models Driven by Lévy Processes

2015-2 Kasper Vinther Olesen: Realizing Conditional Distributions and Coherence Across Financial Asset Classes

2015-3 Manuel Sebastian Lukas: Estimation and Model Specification for Econometric Forecasting

2015-4 Sofie Theilade Nyland Brodersen: Essays on Job Search Assistance and Labor Market Outcomes

2015-5 Jesper Nydam Wulff: Empirical Research in Foreign Market Entry Mode

2015-6 Sanni Nørgaard Breining: The Sibling Relationship Dynamics and Spillovers

2015-7 Marie Herly: Empirical Studies of Earnings Quality

2015-8 Stine Ludvig Bech: The Relationship between Caseworkers and Unemployed Workers

2015-9 Kaleb Girma Abreha: Empirical Essays on Heterogeneous Firms and International Trade

2015-10 Jeanne Andersen: Modelling and Optimisation of Renewable Energy Systems

2015-11 Rasmus Landersø: Essays in the Economics of Crime

2015-12 Juan Carlos Parra-Alvarez: Solution Methods and Inference in Continuous-Time Dynamic Equilibrium Economies (with Applications in Asset Pricing and Income Fluctuation Models)

2015-13 Sakshi Girdhar: The Internationalization of Big Accounting Firms and the Implications on their Practices and Structures: An Institutional Analysis

2015-14 Wenjing Wang: Corporate Innovation, R&D Personnel and External Knowledge Utilization

2015-15 Lene Gilje Justesen: Empirical Banking

2015-16 Jonas Maibom: Structural and Empirical Analysis of the Labour Market

2015-17 Sylvanus Kwaku Afesorgbor: Essays on International Economics and Development

2015-18 Orimar Sauri: Lévy Semistationary Models with Applications in Energy Markets

2015-19 Kristine Vasiljeva: Essays on Immigration in a Generous Welfare State

2015-20 Jonas Nygaard Eriksen: Business Cycles and Expected Returns

2015-21 Simon Juul Hviid: Dynamic Models of the Housing Market

2016-1 Silvia Migali: Essays on International Migration: Institutions, Skill Recognition, and the Welfare State

2016-2 Lorenzo Boldrini: Essays on Forecasting with Linear State-Space Systems

2016-3 Palle Sørensen: Financial Frictions, Price Rigidities, and the Business Cycle

2016-4 Camilla Pisani: Volatility and Correlation in Financial Markets: Theoretical Developments and Numerical Analysis

2016-5 Anders Kronborg: Methods and Applications to DSGE Models

2016-6 Morten Visby Krægpøth: Empirical Studies in Economics of Education

2016-7 Anne Odile Peschel: Essays on Implicit Information Processing at the Point of Sale: Evidence from Experiments and Scanner Data Analysis

2016-8 Girum Dagnachew Abate: Essays in Spatial Econometrics

2016-9 Kai Rehwald: Essays in Public Policy Evaluation

2016-10 Reza Pourmoayed: Optimization Methods in a Stochastic Production Environment

2016-11 Sune Lauth Gadegaard: Discrete Location Problems – Theory, Algorithms, and Extensions to Multiple Objectives

2016-12 Lisbeth Palmhøj Nielsen: Empirical Essays on Child Achievement, Maternal Employment, Parental Leave, and Geographic Mobility

2016-13 Louise Voldby Beuchert-Pedersen: School Resources and Student Achievement: Evidence From Social and Natural Experiments

2016-14 Mette Trier Damgaard: Essays in Applied Behavioral Economics

2016-15 Andrea Barletta: Consistent Modeling and Efficient Pricing of Volatility Derivatives

2016-16 Thorvardur Tjörvi Ólafsson: Macrofinancial Linkages and Crises in Small Open Economies

2016-17 Carlos Vladimir Rodríguez Caballero: On Factor Analysis with Long-Range Dependence

2016-18 J. Eduardo Vera-Valdés: Essays in Long Memory

2016-19 Magnus Sander: Returns, Dividends, and Optimal Portfolios

2016-20 Ioana Daniela Neamtu: Wind Power Effects and Price Elasticity of Demand for the Nordic Electricity Markets

2016-21 Anne Brink Nandrup: Determinants of Student Achievement and Education Choice

2016-22 Jakob Guldbæk Mikkelsen: Time-Varying Loadings in Factor Models: Theory and Applications

2016-23 Dan Nguyen: Formidability and Human Behavior: An Interdisciplinary Approach

2016-24 Martin Petri Bagger: Attention and Decision-Making: Separating Top-Down from Bottom-Up Components

2016-25 Samira Mirzaei: Optimization Algorithms for Multi-Commodity Routing and Inventory Routing Problems

2017-1 Viktoryia Buhayenko: Determining Dynamic Discounts for Supply Chain Coordination

2017-2 Mikkel Bennedsen: Rough Continuous-Time Processes: Theory and Applications

2017-3 Max Weiss Dohrn: Business Cycle Costs for Finitely Living Individuals

2017-4 Oskar Knapik: Essays on econometric modelling and forecasting of electricity prices

ISBN: 9788793195585