



TI 2003-071/4

Tinbergen Institute Discussion Paper

Periodic Heteroskedastic RegARFIMA Models for Daily Electricity Spot Prices

M. Angeles Carnero¹

Siem Jan Koopman^{2,3}

Marius Ooms^{2,3}

1 Dpt. Fundamentos del Análisis Económico, University of Alicante, Alicante, 2 Faculty of Economics and Business Administration, Vrije Universiteit Amsterdam, 3 Tinbergen Institute.

Tinbergen Institute

The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam

Roetersstraat 31

1018 WB Amsterdam

The Netherlands

Tel.: +31(0)20 551 3500

Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam

Burg. Oudlaan 50

3062 PA Amsterdam

The Netherlands

Tel.: +31(0)10 408 8900

Fax: +31(0)10 408 9031

Please send questions and/or remarks of non-scientific nature to wdriessen@few.eur.nl.

Most TI discussion papers can be downloaded at <http://www.tinbergen.nl>.

Periodic Heteroskedastic RegARFIMA models for daily electricity spot prices

September 4, 2003

M. Angeles Carnero

University of Alicante,

Dpt. Fundamentos del Análisis Económico,

03071 Alicante

`acarnero@merlin.fae.ua.es`

Siem Jan Koopman¹ and Marius Ooms

*Tinbergen Institute and Vrije Universiteit Amsterdam, Department of Econometrics and
Operations Research,*

De Boelelaan 1105, NL-1081 HV Amsterdam

`s.j.koopman@feweb.vu.nl` and `mooms@feweb.vu.nl`

¹Corresponding Author

Periodic Heteroskedastic RegARFIMA models for daily electricity spot prices

M. Angeles Carnero, Siem Jan Koopman and Marius Ooms

Abstract

Although the main interest in the modelling of electricity prices is often on volatility aspects, we argue that stochastic heteroskedastic behaviour in prices can only be modelled correctly when the conditional mean of the time series is properly modelled. In this paper we consider different periodic extensions of regression models with autoregressive fractionally integrated moving average disturbances for the analysis of daily spot prices of electricity. We show that day-of-the-week periodicity and long memory are important determinants for the dynamic modelling of the conditional mean of electricity spot prices. Once an effective description of the conditional mean of spot prices is empirically identified, focus can be directed towards volatility features of the time series.

For the older electricity market of Nord Pool in Norway, it is found that a long memory model with periodic coefficients is required to model daily spot prices effectively. Further, strong evidence of conditional heteroskedasticity is found in the mean corrected Nord Pool series. For daily prices at three emerging electricity markets that we consider (APX in The Netherlands, EEX in Germany and Powernext in France) periodicity in the autoregressive coefficients is also established, but evidence of long memory is not found and existence of dynamic behaviour in the variance of the spot prices is less pronounced. The novel findings in this paper can have important consequences for the modelling and forecasting of mean and variance functions of spot prices for electricity and associated contingent assets.

Keywords: Autoregressive fractionally integrated moving average model; Generalised autoregressive conditional heteroskedasticity model; Long memory process; Periodic autoregressive model; Volatility.

1 Introduction

Electricity supply has been the responsibility of public-private companies in many OECD countries until recently. It is anticipated that the private trading of electricity will intensify further in future and eventually move towards fully privatised electricity markets. In such markets large volumes of electricity power will be traded for the short and long term together with future contracts and options. Although similarities with financial markets may exist with respect to its operations, electricity markets generally are subject to the limitations of the underlying asset. Since the characteristics of electricity prices have important implications for derivative pricing and real option analysis, the modelling of such time series has become of interest to researchers and practitioners.

Following the standard practice in modelling volatility of financial returns, we are interested in the conditional mean and in the volatility of the innovations of the pricing process. For electricity prices it is well established that the mean process is not a simple random walk in logarithms but has specific characteristics. For example, Escribano, Peña, and Villaplana (2002) give a good overview of the salient features of electricity prices: (i) Seasonality: the strong dependence of electricity demand on weather conditions and economic and business activities lead to periodic behaviour in the prices, (ii) Mean-reversion: electricity prices are mean-reverting since weather is a dominant factor and influences equilibrium prices through changes in demand; (iii) Jumps and volatility: it is impossible or difficult to store large quantities of electricity so that supply and demand shocks cannot easily be smoothed out.

The literature on analysing electricity prices is scarce but it grows quickly. Some initial contributions are the papers by Knittel and Roberts (2001) and Lucia and Schwartz (2002) who propose several models for electricity prices using Californian hourly data and Norwegian daily data for the Nord Pool market, respectively. Escribano et al. (2002) propose and estimate a more general model for daily spot prices in the electricity markets of Nord Pool, Argentina, Australia, New Zealand and Spain. Wilkinson and Winsen (2002) point out that the pattern of prices varies across day-types. By analysing the sample means of electricity prices in New South Wales (Australia), they find that four type of days should be considered: Monday, Tuesday to Friday, Saturday and Sunday.

We go further and look at periodicities, both in the deterministic function which explains the yearly seasonality of these data, and in the dynamic parameters. We argue that not only the mean and variance of daily electricity prices depend on the day of the week but also skewness, kurtosis and autocorrelation structure. Therefore, the model parameters associated with the dynamics should be different for each day of the week.

Once we have dealt with these periodic features of the time series, it is not taken for granted that volatility clustering exist for the resulting innovations. We show that for some important European electricity markets, the dynamic behaviour in the variance of such innovations depends on the specification for the conditional mean. This means that the empirical work on volatility in electricity prices may have produced spurious results.

However, our empirical results indicate that this does not apply to all markets. First, the model for the mean process can be more complicated than standard autoregressive dynamic structures, especially if longer series are considered. Long memory features seem to be present in market prices for which we have a long data set. Second, after the dynamics in the conditional mean have been identified, the variance process may still have dynamic structure and volatility clustering. For this paper we capture these variance features using a standard generalised autoregressive conditional heteroskedastic (GARCH) model assuming a Student-t distribution for the innovations and including a yearly cycle in the equation of the conditional variance with the objective of capturing possible seasonal patterns in the conditional variance. Simultaneous treatment of dynamics in the conditional mean and variance of electricity price requires further investigation in the future.

The paper is organised as follows. Section 2 describes the markets and data sets and provides the main motivation for the periodic analysis. Section 3 discusses the main Seasonal Periodic RegARIMA models for the conditional mean. Section 4 presents the Seasonal Periodic RegARFIMA model that allows for long memory characteristics. Section 5 reveals the main empirical results obtained using the periodic models. Section 6 presents evidence on conditional heteroskedasticity. Finally, section 7 concludes the paper.

2 European electricity markets and daily spot prices

2.1 Some facts about electricity markets

We analyse time series of daily spot electricity prices from the following European markets: Nord Pool in Norway (www.nordpool.no), Amsterdam Power Exchange (APX) in The Netherlands (www.apx.nl), European Energy Exchange (EEX) in Germany (www.eex.de), and Powernext in France (www.powernext.fr). These markets have started in different years and therefore the four daily time series are of different length. The oldest market is Nord Pool that started in 1991 for the trading of all hydro electricity power generated by Norway. In 1996 Sweden, in 1998 Finland and in 1999 Denmark also joined the Nord Pool market but in this paper we only consider prices for electricity produced by Norway. Most of this electricity (99%) is generated in hydro electric power stations and therefore supply depends heavily on weather conditions. The average production capability of Norway's hydro power plants is about 113 Terawatt hours (TWh= 10^9 KWh) per year. However, this production depends on precipitation levels. For example, in 1998, electricity production was 116.7 TWh while in 1996 it was only 104.7 TWh. The daily spot market APX has been operational since May 1999 and in 2001 a total of 8.24 TWh were traded on this market. The volume traded on the Amsterdam Power Exchange increased by more than 60% in the last years and now it represents more than 9% of the electricity consumption in the Netherlands. EEX is the largest national electricity market in Europe and the volume traded on this market has also increased in the last years. In 2002, about 33 TWh were traded on the EEX spot market, 50% more than in 2001. Finally, the

spot market Powernext in France started in November 2001 and in 2002 volumes traded on this market represented already 2.5% of the France’s electricity consumption. All four markets operate as “day-ahead” markets that concentrate on daily trade for electricity delivered on the next day. Daily series are constructed as the average of 24 price series for the different hours of the day. The resulting prices are referred to as spot prices.

2.2 Time series descriptives of electricity spot prices

The time series of spot prices are observed from January, 1, 2001 to June, 8, 2003, from October, 1, 2001 to June, 8, 2003, from December, 3, 2001 to June, 8, 2003 and from January, 4, 1993 to November, 14, 1999 for the markets of the APX, EEX, Powernext and Nord Pool respectively. Prices are in Euros/MWh in the APX, EEX and Powernext and in NOK/MWh in Nord Pool (1 Euro \approx 8 NOK).

Figure 1 plots the four series of daily spot prices, P_t , computed as the average of the 24 hourly prices. As we can see in the graphs, the dynamic behaviour of the Nord Pool time series is different from the behaviour of the other three series. One of the most important reasons that could explain these differences is the type of electricity traded on the markets. Most of the electricity traded on the Nord Pool market is produced by hydro power generation and therefore, it depends on weather conditions. In the APX, most of the electricity traded is thermal (via the burning of coal) and gas while the EEX and Powernext markets trade electricity produced mainly by nuclear power plants.

As it is usual in the electricity literature, we are interested in modelling the logarithm of prices. Figure 2 plots the log-daily prices, $p_t = \log P_t$, together with a kernel density estimate and the correlogram. As we can see, the log-transformation stabilised the variation in the series. The pattern of the correlations is similar for the three younger markets. The simultaneous time series plots also reveal similar local trends in these markets. Strong weekly patterns are clearly apparent in the correlograms of the series. The correlogram of Nord Pool is dominated by long memory dynamics. This may partly be explained by its relation to river flow levels and, as it is known, mean river flows often display long memory characteristics; see, for example Ooms and Franses (2001) and references therein. Figure 1 at least reveals that the nature of the Nord Pool series is very different from the other markets.

Table 1 contains, for the APX, EEX and Powernext markets, descriptive statistics for the whole (pooled) sample and for the samples of different days of the week, together with the periodic autocorrelation coefficients as defined by McLeod (1994). For example, the fourth column of this table shows $r_{\text{Tue}}(1) = \text{corr}(p_t, p_{t-1})$ with values 0.72, 0.60 and 0.67, for APX, EEX and Powernext, respectively. These values correspond to the sample estimate of

$$\frac{E[(p_{\text{Tue}} - E(p_{\text{Tue}}))(p_{\text{Mon}} - E(p_{\text{Mon}}))]}{\sqrt{\text{Var}(p_{\text{Tue}})}\sqrt{\text{Var}(p_{\text{Mon}})}}.$$

As Wilkinson and Winsen (2002) pointed out, the pattern of prices varies across day-types. Table 1 further shows that the mean and variance of daily electricity prices depend on the day

of the week and so does the skewness, kurtosis and autocorrelation structure.

For all markets, the mean is bigger on Tuesdays and smaller on Sundays and the correlation coefficients depend on the day of the week. For example, the correlation between Wednesdays and Tuesdays (0.83, 0.74 and 0.83 for APX, EEX and Powernext, respectively) is higher than the correlation between Mondays and preceding Sundays (0.40, 0.63 and 0.63, respectively) as one would expect.

The large day-to-day differences in the autocorrelation function strongly motivate a periodic time series modelling approach. This pertains both to the regression part and the dynamic part as we will illustrate in the following sections. In order to save space we mention without showing all the results that the autocorrelation functions remain clearly periodic if apparent nonstationarities due to day-of-the-week effects and yearly weather cycles have been removed by regression or by (seasonally) differencing the data. Table 2 contains the corresponding descriptive statistics for the Nord Pool data, both for log prices and daily changes in log prices. The second panel of this table shows that the correlation function for the daily price changes on the Nord Pool market is extremely periodic, even for longer lags, motivating a seasonal periodic long memory time series model.

Even before specifying a modelling strategy it is relevant to discuss possible consequences of neglecting periodicity when interpreting estimates of nonperiodic autocorrelation functions of levels and squares of periodic processes.

2.3 Autocorrelation, periodic autocorrelation, in levels and squares

First, it is interesting to discuss the slightly puzzling fact that periodic autocorrelations for all periods can be higher than an individual autocorrelation coefficient. For example, this is the case for $r(1)$ of EEX spot prices and for $r(1)$ and $r(2)$ of Powernext spot prices as reported in Table 1. An understanding of this feature can be important for effective identification of periodic time series models. The matter can be illustrated with the following example. Consider the time series generating process

$$y_t = \begin{cases} \varepsilon_j & t = 4j - 3 \\ -\varepsilon_j & t = 4j - 2 \\ \eta_j & t = 4j - 1 \\ \eta_j & t = 4j \end{cases}, \quad \varepsilon_j \stackrel{i.i.d.}{\sim} N(0, 1), \quad \eta_j \stackrel{i.i.d.}{\sim} N(0, 1), \quad (1)$$

for $t = 1, \dots, T$ and $j = 1, \dots, T/4$ with T chosen to be some multiple of 4. Note that ε_j and η_j are assumed independent of each other at all times. It follows that the theoretical

autocorrelation $E(y_t y_{t-1})$ is periodic

$$E(y_t y_{t-1}) = \begin{cases} E(\varepsilon_j \eta_{j-1}) = 0 & t = 4j - 3 \\ E(-\varepsilon_j^2) = -1 & t = 4j - 2 \\ E(-\varepsilon_j \eta_j) = 0 & t = 4j - 1 \\ E(\eta_j^2) = 1 & t = 4j \end{cases}. \quad (2)$$

Computing the sample counterparts of the autocorrelation and the periodic autocorrelations, it is expected that two periodic autocorrelations will be significantly away from zero while the overall first order correlation coefficient will be close to zero. Also higher order sample autocorrelations will be close to zero. Therefore, in this case the standard correlogram will not give evidence of serial correlation in y_t .

A further consequence of the illustration is that the theoretical autocovariance in the squares, $E\{(y_t^2 - 1)(y_{t-1}^2 - 1)\}$, is periodic as well (0, 2, 0, 2 in the order of (2)). However, in this case the sample first order autocorrelation of the squared y_t will be significantly away from zero (approximately 0.5). Given the fact that the correlogram for the levels suggests that y_t is white noise, this example shows that evidence provided by the correlogram of the squared y_t can be misleading, when periodic features are present in the mean. Therefore the possibility of detecting spurious volatility clustering is existent. Further it is important to account for periodic features of time series before dynamic models are fitted to conditional variances.

3 Periodic RegARIMA models

3.1 RegARIMA model

Regression models with autoregressive integrated moving average (ARIMA) disturbances have been used in many areas of time series analysis and were introduced in the seminal book of Box and Jenkins (1970). Most notably this time series model is often used as the basis for model-based seasonal adjustment; see, for example, Bell and Hillmer (1984). Also the current version of the U.S. Census program X-12-Arima allows the identification, estimation and diagnostic checking of RegARIMA models as a precursor to the X-11 seasonal adjustment method; see Findley, Monsell, Bell, Otto, and Chen (1998).

A RegARIMA model is defined as a multiple regression model with ARIMA disturbances. For an observed time series y_t and for k known and fixed time series x_{it} with $i = 1, \dots, k$ and $t = 1, \dots, T$, the RegARIMA model is given by

$$y_t = \sum_{i=1}^k \beta_i x_{it} + u_t, \quad \phi(L)\Delta^d u_t = \theta(L)\eta_t, \quad \eta_t \sim N(0, \sigma^2), \quad (3)$$

where $\beta_1, \dots, \beta_k, \sigma^2$ are fixed unknown coefficients, $\phi(L)$ and $\theta(L)$ are polynomials in the lag operator L ($L^p y_t = y_{t-p}$) and Δ^d is the differencing operator $(1 - L)^d$ for $d = 0, 1, \dots$. In

section 4 of this paper we extend the model to allow for (seasonal) fractionally integrated disturbances, that is “RegARFIMA” models, which allow for a non-integer order of differencing. The coefficients of the lag polynomials

$$\phi(L) = 1 - \phi_1 L - \dots - \phi_p L^p, \quad \theta(L) = 1 + \theta_1 L + \dots + \theta_q L^q, \quad (4)$$

are fixed and unknown for any set of non-negative integers p and q . In short notation, we refer to the model as RegARIMA(p, d, q).

The time series models that have been used for the modelling of spot electricity prices are special cases of the RegARIMA model (3). The basic model for the logarithm of the daily spot price, p_t , can be represented as

$$p_t = f_t + u_t, \quad f_t = \sum_{i=1}^k \beta_i x_{it}, \quad u_t = \phi u_{t-1} + \eta_t, \quad (5)$$

where f_t is a deterministic function to allow for specific seasonal patterns in the prices. The basic model (5) is in effect the RegARIMA(1,0,0) model. The standard explanatory variables are polynomial, trigonometric and other periodic functions of time. Observations and forecasts for (stochastic) weather variables are not included in these types of market models.

For example, we can take $k = 7$ with

$$x_{1t} = 1, \quad x_{2t} = t, \quad x_{3t} = \sin \lambda t, \quad x_{4t} = \cos \lambda t, \quad x_{5t} = \sin 2\lambda t, \quad x_{6t} = \cos 2\lambda t, \quad (6)$$

where $\lambda = 2\pi/365$ and with

$$x_{7t} = \begin{cases} 0 & t \text{ is workday} \\ 1 & t \text{ is not a workday} \end{cases}. \quad (7)$$

Models with such specifications for the mean function for log electricity prices have been used in, for example, Lucia and Schwartz (2002), Escribano, Peña, and Villaplana (2002) and De Jong and Huisman (2002). The first two authors propose the RegARIMA(1,0,0) model (5), with four regressors, x_{1t}, x_{3t}, x_{4t} and x_{7t} in f_t , for modelling spot prices for “next day” electricity. Forecasts from this model are subsequently used for prices for electricity delivered at later days (“future prices”). Other contributions in the electricity market pricing literature have adopted this specification for the mean of log prices in order to model the presumed stochastic volatility in the price series.

The considered RegARIMA model assumes that the regression component f_t is appropriate for the seasonal effects in the time series of prices. We argued in section 2.2 that a model with periodic coefficients for f_t and u_t is required to adequately describe the dynamics of log electricity prices. Before presenting extended models we discuss a reparameterisation of the RegARIMA model that simplifies estimation in some interesting cases.

3.2 Parameterisation and estimation of RegARIMA models

Maximum likelihood estimation of RegARIMA models is straightforward since an explicit expression for the (Gaussian) loglikelihood function is available; see standard textbooks such as Brockwell and Davis (1994) for a time series point of view, or Johnston and Dinardo (1997) for an econometric point of view. The consequence is that the likelihood needs to be maximised numerically with respect to the unknown regression and lag polynomial coefficients. Numerical optimisation can be avoided in models with pure AR processes for u_t , by conditioning on the first observations, leading to conditional least squares, see e.g. Shumway and Stoffer (2000).

In this case with only a stationary autoregressive polynomial u_t , the RegARIMA($p, 0, 0$) model (5) can be written in an equivalent form of an Autoregressive model with exogenous input variables (ARX) as

$$\phi(L)p_t = \sum_{i=1}^k \beta_i z_{it} + \eta_t,$$

where $z_{it} = \phi(L)x_{it}$ for $i = 1, \dots, k$. Furthermore, in cases where x_{it} is a polynomial or periodic function of the time index t such as in (6) and (7), applying a stationary autoregressive lag polynomial to x_{it} to obtain z_{it} , has no effect on the interpretation of the explanatory variables and therefore its effect is void. Since we consider model (5) in which all explanatory variables are polynomial or periodic functions of time t , we can take x_{it} rather than z_{it} for estimation purposes. Of course, this replacement does lead to different regression coefficients. It is mainly for this reason that the RegARIMA($p, 0, 0$) model can be written as a standard regression model with lagged dependent variables, that is

$$p_t = \phi_1 p_{t-1} + \dots + \phi_p p_{t-p} + \sum_{i=1}^k \beta_i^* x_{it} + \eta_t. \quad (8)$$

The coefficients of this model can be estimated by Ordinary Least Squares (OLS) to get consistent and efficient estimates for all coefficients $\phi_1, \dots, \phi_p, \beta_1^*, \dots, \beta_k^*$. From these it is straightforward to derive β_1, \dots, β_k in the AR(p) extension of equation (5), see e.g. Ooms (1994, §A2.2.2).

3.3 Periodic RegARIMA models

Spot electricity prices change on a daily basis and are subject to strong seasonal effects. The nature of the seasonality described in section 2.2 suggests that each day of the week should be described by a different model. We therefore consider the periodic RegARIMA model which in effect allows for different coefficients for different days in the week. If all coefficients are different, including the variance σ^2 , we can isolate the days from each other and estimate separate models for the seven weekly series of prices for Monday until Sunday effectively creating one time invariant multiple RegARIMA model, c.f. Tiao and Grupe (1980). Our focus will also be on more subtle periodic formulations of the RegARIMA model in which only certain mean

coefficients for f_t are different for each day or in which only certain lag polynomial coefficients for u_t are different.

Periodic autoregressive (periodic AR) models were first applied by Jones and Brelsford (1967) and have found widespread use in geophysical time series modelling. Tiao and Grupe (1980) discuss the consequences for traditional ARMA modelling if the underlying process really follows a periodic ARMA model. Vecchia (1985) developed and applied conditional ML estimation for periodic ARMA models, but these models are still not widely used. McLeod (1994) discussed the empirical identification of periodic AR models. Econometric work on periodic models with focus on testing for periodic integration, i.e. the nonstationarity of the periodic AR part of the model, is discussed in various articles and in books by Franses (1996) and Ghysels and Osborn (2001).

We follow the time series approach to periodic modelling by using the sample periodic autocorrelation function as the basis for identifying the periodic RegARIMA model in a similar way as is done for standard ARMA modelling. A definition of the sample periodic autocorrelation function is found in McLeod (1994) and as we have seen in section 2.2, Tables 1 and 2 contain the sample periodic autocorrelation coefficients for the data analysed in this paper. To allow for the regression effect f_t , the sample periodic autocorrelation function can be computed for the OLS residuals of periodic regressions.

In the context of model (3) we formulate the periodic RegARIMA model by

$$y_t = \sum_{i=1}^k \beta_{i,j} x_{it} + u_t, \quad \phi_j(L) \Delta^d u_t = \theta_j(L) \eta_t, \quad \eta_t \sim N(0, \sigma_j^2), \quad (9)$$

for $j = j(t) = 1, \dots, s$, with fixed periodic (or seasonal) length s and where $\phi_j(L)$ and $\theta_j(L)$ are polynomials with coefficients indexed by j . The indices $j(t)$ indicate that coefficients can be different for different “seasons”. They are periodic functions of the time index t . For example, in the case of daily time series with a general day-of-the-week periodic behaviour, we have $j(t) = 1 + (t \bmod 7)$. In the particular case of the electricity price model (5), both the components f_t and u_t are taken as periodic counterparts of their original definitions.

3.4 Estimation of Periodic RegARIMA models

The number of coefficients to be estimated increases when periodic coefficients are introduced in the model. Therefore the method of maximum likelihood that requires direct numerical optimisation of the likelihood is to be avoided whenever possible.

As discussed in section 3.2, a linear regression approach can be taken for the RegARIMA($p,0,0$) model with polynomial and periodic regressors. This also holds for the periodic version of this model for log electricity prices given by

$$p_t = \phi_{1,j} p_{t-1} + \dots + \phi_{p,j} p_{t-p} + \sum_{i=1}^k \beta_{i,j}^* x_{it} + \eta_t, \quad \eta_t \sim N(0, \sigma_j^2),$$

where j is a (modulus) function of t and $j = 1, \dots, s$. This representation effectively is a set of s different equations and estimation can be carried out by the method of least squares, equation by equation. However when particular coefficients are restricted to be the same across days, estimation requires the consideration of all equations jointly, hence the need to estimate all coefficients simultaneously and nonlinear optimisation often becomes necessary. In the particular case $j = 1 + t \bmod s$, we can also specify the model as a single equation by

$$p_t = \sum_{j=1}^s D_{jt} \left(\phi_{1,j} p_{t-1} + \dots + \phi_{p,j} p_{t-p} + \sum_{i=1}^k \beta_{i,j}^* x_{it} + \eta_t \right), \quad \eta_t \sim N(0, \sigma_j^2),$$

where D_{jt} is a dummy variable that is one when $t \bmod j = 0$ and is zero otherwise. Straightforward (iteratively weighted) least squares can be used for the estimation of the coefficients, whether periodic coefficients (or a selection thereof) are restricted to be the same or not, see Oberhofer and Kmenta (1974). In effect one maximises the approximate conditional loglikelihood function

$$c - \frac{T}{2s} \sum_{j=1}^s \ln \hat{\sigma}_j^2 \tag{10}$$

where c is a constant, T is the sample size and $\hat{\sigma}_j^2$ is the average of the squared residuals for “season” j . Periodicity can be tested by comparing loglikelihoods of the models under various restrictions. Note that (10) simplifies to $c - \frac{T}{2} \ln \hat{\sigma}^2$ under homoskedasticity by taking $s = 1$.

Implementation is relatively straightforward. The software programs have been written for this purpose using `0x 3.2`, see Doornik (2002). The `0x` programs are available from the authors upon request.

4 Seasonal periodic heteroskedastic RegARFIMA models

In this section we present two extensions that build on the RegARIMA models discussed in the previous section. First the seasonal RegARIMA will be briefly introduced. Next, the fractional integration extension will be discussed in more detail.

4.1 Seasonal extension

The seasonal Box-Jenkins analysis is based on multiplicative polynomials for dynamics in the level and for dynamics in the seasonal. Seasonal ARIMA models are discussed by Box and Jenkins (1970) and in several textbooks such as Shumway and Stoffer (2000). The seasonal model is given by

$$\phi(L)\Phi(L^s)\Delta^d\Delta_s^D u_t = \theta(L)\Theta(L^s)\eta_t, \quad t = 1, \dots, T, \tag{11}$$

where $\phi(L)$ and $\theta(L)$ are polynomials in the lag operator L of orders p and q respectively (as given by (4)), $\Phi(L^s)$ and $\Theta(L^s)$ are polynomials in L^s of orders P and Q respectively (similar to (4)), $\Delta^d = (1 - L)^d$, $\Delta_s^D = (1 - L^s)^D$, $p < s$ and $q < s$. Short hand notation for this model is seasonal ARIMA(p, d, q) \times (P, D, Q) $_s$. Since the multiplication of two polynomials leads to another single polynomial, the analysis based on seasonal ARIMA models is similar to ARIMA models although the analysis can become more intricate since it involves more parameters and higher dimensions. Tiao and Grupe (1980) showed how seasonal lags in the non-periodic model can arise if one attempts to whiten the residuals of a periodic time series process, i.e. if one attempts to approximate a periodic time series process by a nonperiodic model. It is therefore interesting to include a seasonal ARIMA model in the analysis. The seasonal RegARIMA model is given by

$$y_t = \sum_{i=1}^k \beta_i x_{it} + u_t, \quad t = 1, \dots, T,$$

where u_t is modelled as in (11). In the remainder of this paper we will concentrate on to the seasonal RegARIMA($p, 0, 0$) \times ($1, D, 0$) $_s$ model.

4.2 Seasonal fractional integration extension

The seasonal periodic ARFIMA model was first considered by Ooms and Franses (2001) to model the long memory characteristics of the monthly flows in the Canadian Fraser river. They based their analysis on the seasonal ARIMA($p, 0, 0$) \times (P, D, Q) $_{12}$ model allowing the nonseasonal AR parameters and the seasonal fractional integration parameter D to vary with the month of the year.

The “fractional differencing” model introduced by Adenstedt (1974) has become a standard model for long memory behaviour, see Beran (1994). In econometrics this model is known as an $I(d)$ process. The discrete integration model $(1 - L)^d u_t = \varepsilon_t$, where ε_t is white noise and d is a non-negative integer, becomes a fractional integration model when d is replaced by a real number.

If $|d| < 0.5$ the “fractionally integrated” series u_t is stationary and invertible and can be written as an infinitely long AR process using the binomial expansion of $(1 - L)^d$. For $d > 0$ the autocorrelation function of this process dies out hyperbolically, indicating the process is long memory.

The generalisation towards a mixed autoregressive fractionally integrated moving average (ARFIMA) process, introduced by Granger and Joyeux (1980) and Hosking (1981), is obtained by considering (3), without regressors, with a real number $|d| < 0.5$. Statistical properties and inference for ARFIMA processes and other long memory models are extensively discussed in the monograph by Beran (1994) and more recently in Robinson (2003). The further generalisation towards seasonal RegARFIMA models is evident by adding regressors for the mean and by letting both d and D be continuous variables in (11). Carlin, Dempster, and Jonas (1985) provided an early analysis of ARFIMA models with seasonal fractional integration parameter

D. Porter-Hudak (1990) is probably the best known application to economic data.

As a final step one can let the fractional d and D be periodic. In this paper we do not consider a periodic d , see Franses and Ooms (1997), but we do allow for periodic AR parameters and a periodic seasonal fractional D_j . If one views this parametric model in its multiple equation form, it can be given a fractional cointegration interpretation as in Sowell (1987), see e.g. Dueker and Startz (1998) for an application. In our parametric context fractional cointegration implies that different linear combinations of the separate (fractionally integrated) time series for the different days are allowed to be integrated of different orders. Robinson and Yajima (2002) consider estimation and testing for fractional cointegration in a semiparametric multiple equation context.

As we show below, this turns out to be relevant for the Nord Pool data. For example, we find that the daily difference of p_t between Fridays and preceding Thursdays is approximately white noise and therefore $I(0)$, whereas the difference between Saturdays and preceding Fridays is integrated of order 0.4. See the second panel of Table 2. In this case Thursdays and Fridays share the same long memory property, which can be removed by considering their difference.

For the Nord Pool data, we specify a seasonal periodic ARFIMA(2, 0, 0) \times (0, D , 0) $_7$ model with $|D_j| < 0.5$. We require $D_j > -0.5$ so that the model can be estimated using the nonlinear least squares method following Beran (1995). We require $D_j < 0.5$ so that the regression part of the model can be interpreted in a standard way. Note that this specification with 14 autoregressive parameters, 7 integration parameters and 7 variances is much more parsimonious than a full fractional vector autoregressive model of order 1 for all the days of the week, which would contain 49 autoregressive parameters, 7 integration parameters and 21 parameters for the variance covariance parameters.

For example, the stationary periodic model for log spot prices with $|D_j| < 0.5$ and $p = 2$ is given by

$$p_t = \sum_{i=1}^k \beta_{i,j} x_{it} + u_t, \quad (1 - \phi_{1,j}L - \phi_{2,j}L^2)(1 - L^s)^{D_j} u_t = \eta_t, \quad (12)$$

for $t = 1, \dots, T$ and for $j = 1, \dots, s$ with $j = 1 + (t \bmod s)$, and where η_t is white noise, with a stable periodic AR part of the model.

The same arguments that led to formulation (8) of a RegARIMA($p, 0, 0$) model also hold for the RegARIMA part of a stationary RegARFIMA($p, 0, 0$) \times ($P, D, 0$) model with $|D| < 0.5$. Similarly, the arguments hold for the short memory seasonal and periodic extensions of the RegARIMA part of the model. We therefore can represent model (12) by

$$(1 - \phi_{1,j}L - \phi_{2,j}L^2)p_t = \sum_{i=1}^k \beta_{i,j}^* x_{it} + (1 - L^s)^{-D_j} \eta_t, \quad \eta_t \sim N(0, \sigma_j^2), \quad (13)$$

for $j = 1, \dots, s$, and $t = 1, \dots, T$. The transformed regression coefficients $\beta_{i,j}^*$ are here only defined with respect to the autoregressive polynomial $\phi_j(L)$. Note that the regression coefficients would not be identified if one (also) would try to transform them with respect to the fractional

difference operator. This is easily seen if one considers only the simple nonperiodic 'ARFIX' transformation analogous to (8) and if x_{it} contains only a constant, $x_{1t} = 1$:

$$\phi(L)(1-L)^d p_t = \phi(L)(1-L)^d \beta_1^{**} + \eta_t,$$

where $(1-L)^d \beta_1^{**} = (1-1)^d \beta_1^{**} = 0$ and for all $d > 0$ for all β_1^{**} , since the infinite sum of the coefficients in the $\text{AR}(\infty)$ expansion of $(1-L)^d$ equals one. It is therefore impossible to derive the mean β_1 in the corresponding RegARFIMA model (12) using this transformation.

4.3 Estimating seasonal periodic RegARFIMA models

Although the general model can have potentially a large number of parameters that need to be estimated, the suggested extensions of the RegARFIMA model do not lead to further complexities in the estimation methodology.

Frequency domain methods for estimating long memory regression models are surveyed by Robinson (2003). Important time domain methods were developed by Sowell (1992) for stationary ARFIMA processes and by Beran (1995) for invertible ARFIMA processes. Beran (1995) based his method on nonlinear least squares methods applying the approximate (Gaussian) loglikelihood function similar to (10). Fast computational methods for estimation and analysis of ARFIMA models are implemented in the ARFIMA package of Doornik and Ooms (1999) and available for the `0x` programming environment. Doornik and Ooms (2003) provide a detailed discussion.

Ooms and Franses (2001) showed that under the assumption of covariance stationarity the model can be easily estimated using a periodic version of the nonlinear least squares method of Beran (1995), considering separate equations for each day of the week.

The least squares methods do not impose stationarity for the (periodic) AR polynomials or for the periodic seasonal integration parameters, so one should test for nonstationarity after estimation. Test statistics for nonstationarity of the periodic AR polynomial follow so-called Dickey-Fuller limit distributions under the null of nonstationarity, see Boswijk and Franses (1996). Test statistics for specific values of the fractional integration parameters follow standard t - and chi-squared limit distributions in most relevant cases, see Robinson (2003). Under our stationarity assumption, test statistics on the regression parameters also follow standard limit distributions.

5 Empirical results

5.1 Results for new European markets

In this section we analyse daily spot prices of electricity observed in three European markets: APX Netherlands, EEX Germany and Powernext France. Nord Pool results follow in section 5.2. We provide detailed results for three models which can be summarised as follows.

Model 1 is given by

$$p_t = f_t + u_t, \quad f_t = \beta_1 + \beta_3 x_{3t} + \beta_4 x_{4t} + \beta_7 x_{7t} + u_t, \quad u_t = \phi_1 u_{t-1} + \eta_t, \quad (14)$$

where η_t is white noise with variance σ^2 . The explanatory variables are specified above, see (6) and (7).

Lucia and Schwartz (2002) proposed this RegARIMA(1,0,0) model (5) with four regressors for a yearly cycle and a weekend effect. Other explanatory variables for the mean function f_t , these are x_{2t} , x_{5t} and x_{6t} (trend and half-year cycle), have been considered in the analysis but were not found to be significant. This also applies to the models that are discussed below.

Model 2 is an extended version of Model 1 with separate constants for each day of the week.

The weekend variable x_{7t} in (14) is superseded and deleted. The regression part is therefore given by

$$f_t = \beta_{1,j} + \beta_3 x_{3t} + \beta_4 x_{4t}, \quad j = 1, \dots, 7.$$

This model further involves a more elaborate autoregressive specification for the disturbance u_t , in particular a seasonal ARIMA(2,0,0) \times (1,0,0)₇ model, see (11),

$$(1 - \phi_1 L - \phi_2 L^2)(1 - \Phi_1 L^s)u_t = \eta_t, \quad \eta_t \sim N(0, \sigma_j^2) \quad (15)$$

where the disturbances η_t have periodic variances σ_j^2 for $j = 1, \dots, s$ with $s = 7$.

The dynamic part of Model 2 is obtained using Box-Jenkins methodology, but it includes different constants for different days following the suggestion of Wilkinson and Winsen (2002) and also allows for periodic heteroskedasticity, which is clearly required by the data.

Model 3 is the completely periodic RegARIMA(2,0,0) model which can be written in ARX form, that is

$$p_t = \phi_{1,j} p_{t-1} + \phi_{2,j} p_{t-2} + \beta_{1,j}^* + \beta_{3,j}^* x_{3t} + \beta_{4,j}^* x_{4t} + \eta_t, \quad \eta_t \sim N(0, \sigma_j^2) \quad (16)$$

for $j = 1, \dots, s$, with $s = 7$;

We estimate Models 1, 2 and 3 by (weighted) least squares to maximise the approximate loglikelihood function (10). Tables 3, 4, and 5 report the parameter estimates for the three models for the APX, EEX and Powernext markets, respectively. Models 1 and 3 are estimated by ordinary least squares, where Model 3 is treated as a set of seven equations for each day of the week. Model 2 implies a nonlinear restriction on the AR polynomial and is estimated by nonlinear iterative weighted least squares. We compute standard errors for the parameter estimates by the delta method using (numerical) second order derivatives of the approximate loglikelihood (10). Diagnostic checks are performed on the (appropriately ordered) scaled residuals.

The effective number of observations for the three series, i.e. the number of nonzero residuals we can compute, varies from model to model. The percentage of the number of observations lost by conditioning at the beginning of the sample is maximum $9/553 = 1.6\%$ for Model 2 for the Powernext data, while it is 1.0% and 1.5% for APX and EEX, respectively. The exact likelihood for the first observations can be taken into account if (Gaussian) Exact ML (EML) is applied. However, if one uses EML, the results are very similar to the least squares outcomes for Models 1 and 2, as the number of observations is relatively large.

It is remarkable that the magnitudes and signs of the estimated dynamic parameters of Models 1 and 2 are similar for the three series. As it might not be easy to grasp the main patterns, similarities and differences for daily periodic results of Model 3 in Tables 3, 4 and 5, we also provide a graphical representation. Figures 3, 4 and 5 plot interval estimates for the parameters of periodic Model 3 for the three markets. Regression coefficients, dynamic parameters and residual variances vary from day to day in all three markets, and the periodic patterns in the dynamic parameters vary significantly from market to market. Monday's AR(2) coefficients for the Powernext market deviate from the APX and EEX market. Sunday's AR(2) coefficients for the APX differ from the EEX and Powernext. Note that $\hat{\beta}_{3,j}$ and $\hat{\beta}_{4,j}$ are not directly comparable across markets as the starting dates for the data sets are different. Statistical significance, however, can be compared.

Figure 6 plots the scaled residuals, $\hat{\eta}_t$, of the three fitted models for the APX data together with their correlograms, in levels ($\hat{\eta}_t$) and squares ($\hat{\eta}_t^2$). It is clear that Model 1 is not able to capture the structure of the conditional mean of the series. In particular, its residual correlations at seasonal lags 7, 14, 28, \dots , are clearly significant confirming a strong weekly seasonality in the time series. Model 2 performs better in capturing the dynamics in the mean, since, as we can see, the scaled residuals seem to be uncorrelated. But looking at the correlogram of the squared residuals, we can see some structure in them suggesting the presence of conditional heteroskedasticity. However, it is interesting to note that the structure in the squares is not as strong in the residuals of the periodic Model 3. In this case, it seems that there are no dynamics left in the residuals, not in the levels nor in the squares.

Figure 7 helps to explain what is happening. When we look at the correlation coefficients of the residuals of Model 2, they are not significant at any lag, however, looking at the periodic correlation coefficients, this is not the case. Periodic correlations are significant for some lags and days. For example, $r(1)$ is highly significant and positive for Wednesdays, which indicates a strong correlation between the residuals corresponding to Wednesdays and Tuesdays. In the case of Sundays, $r(1)$ is also highly significant but negative. This is related to our example in section 2.3, where we illustrate how spurious volatility clustering could be detected when looking at the correlogram of the squares if periodic features are present in the data and they are not taken into account. The last seven plots of Figure 7 illustrate this periodic diagnostic check for a periodic time series model. We also advise to make the corresponding periodic scatter plots to check whether the periodic correlations are due to a few influential observations only, see Ooms and Franses (1997) for an application of this periodic diagnostic. For the data sets

in this paper we do not find that the periodic correlations are caused by only a few influential observations.

Figures 8 and 9 show the standard diagnostics for the estimated EEX and Powernext models. Model 3 outperforms Models 1 and 2 in both markets. Still, Model 3 shows remarkable outliers. We experimented by adding dummy variables to Model 3 to take account of the outliers and this improved normality diagnostics. However, this did not change the main results on the other aspects of the models.

Table 6 provides loglikelihood values and Akaike information criteria for Models 1, 2 and 3 for the three markets. It also reports results for a number of intermediate models, which confirm that both periodicity in the regression coefficients, periodicity in the dynamic parameters and periodicity in the variance of the white noise disturbances contribute to goodness-of-fit in a remarkable way. As can be seen in the table, the highest values of the loglikelihood, -135.06 , 39.97 and 71.60 for the APX, EEX and Powernext respectively, correspond to Model 3. Also, Model 3 has for the three markets the lowest values of the AIC.

Finally we note that the scaled residuals are significantly correlated across the three markets, with correlations between 0.3 and 0.6 for the period December 2001-June 2003. This suggests the feasibility of successful joint modelling, but this is outside the scope of this paper.

5.2 Nord Pool results

In this section we present empirical results for the daily Nord Pool data, for which data characteristics are summarised in section 2.2. It is indicated that a seasonal periodic long memory model might be needed to adequately capture the dynamics in the conditional mean of the series. Moreover, the Nord Pool series are substantially longer than the series for the new markets, which makes a parametric long memory analysis feasible.

In addition to Models 1, 2 and 3 we analyse Model 4 as well.

Model 4 is the seasonal periodic RegARFI(MA) model as in (13), in particular

$$(1 - \phi_{1,j}L - \phi_{2,j}L^2)p_t = \beta_{1,j}^* + \beta_{3,j}^*x_{3,t} + \beta_{4,j}^*x_{4,t} + (1 - L^7)^{-D_j}\eta_t, \eta_t \sim N(0, \sigma_j^2) \quad (17)$$

for $t = 1, \dots, T$ and $j = 1, \dots, s$.

Again, the process η_t is assumed to be white noise with zero mean and day-of-the-week variances σ_j^2 , for $j = 1, \dots, s$.

We estimate Model 4 by nonlinear least squares with separate equations for each day of the week. Again we use the approximate loglikelihood function (10) as the basis for inference. Treating

$$p_t = \phi_{1,j}p_{t-1} + \phi_{2,j}p_{t-2} + \beta_{1,j}^* + \beta_{3,j}^*x_{3,t} + \beta_{4,j}^*x_{4,t} + (1 - L^7)^{-D_j}\eta_t$$

as s regression models (for weekly data) with (weekly) ARFIMA(0, d ,0) errors and estimating them using Exact ML delivered similar results. Diagnostic checks are performed on the (appropriately ordered) scaled residuals of these seven equations.

Table 7 presents the estimation results for Model 1, 2, 3 and 4. Figure 10 shows confidence intervals for the main parameters of Model 4. The (periodic) autoregressive parts and the fractional components of all these models are stationary. For example, the largest characteristic root of the periodic AR part of Model 4 equals 0.845 which is well inside the unit circle, whereas the largest D_j is estimated at 0.387 with a standard error of 0.06. The first plot of Figure 10 shows the seasonal long memory parameter of Model 4 that appears to be strongly periodic with large values for Monday and Saturday only. For Wednesdays, Thursdays and Fridays $\phi_{1,j}$ is close to one and $\phi_{2,j}$ close to zero. For Saturdays and Sundays the parameter sums $\phi_{1,j} + \phi_{2,j}$ are close to one. This means that we can interpret the model in terms of daily differences for Wednesdays to Sundays. Most of the mean reversion seems to take place on Mondays as $\phi_{1,1} + \phi_{2,1}$ is clearly below one. Figure 11 shows diagnostics for the scaled residuals $\hat{\eta}_t$ of Models 1, 2, 3 and 4. Model 4 effectively removes serial correlation at weekly lags. Figure 11 provides evidence of nonnormality of $\hat{\eta}_t$. Strong serial correlation in $\hat{\eta}_t^2$ is also detected in all models. We address the latter issue in the next section.

The last column of Table 6 compares the goodness-of-fit of Models 1, 2, 3 and 4 and some intermediate ones. The seasonal periodic long memory Model 4 outperforms Models 2 and 3 as expected from the analysis of the periodic autocorrelation functions in Table 2 above. The estimated equations for Friday and Saturday in Model 4 clearly show the periodic behaviour in the seasonal long memory part. A good approximation for Friday's equation is given in obvious notation by

$$p_{\text{Fri},t^*} = p_{\text{Thu},t^*} + \eta_{\text{Fri},t^*}$$

where t^* is a time index for weeks. This means that $p_{\text{Fri},t^*} - p_{\text{Thu},t^*}$ is approximately white noise. This is confirmed by its autocorrelation function reported in the lower panel of Table 2, which also shows that it is (periodically) uncorrelated with $p_{\text{Thu},t^*} - p_{\text{Wed},t^*}$ and $p_{\text{Wed},t^*} - p_{\text{Tue},t^*}$. On the other hand the equation for Saturday's prices clearly contains seasonal long memory, even after conditioning on Friday's and Thursday's prices. This type of long memory was also indicated by the large values for the autocorrelation function for $p_{\text{Sat},t^*} - p_{\text{Fri},t^*}$ at weekly lags in Table 2. In econometric terms one can say that the weekly series p_{Fri,t^*} and p_{Thu,t^*} are fractionally cointegrated in the context of a (restricted) fractionally integrated Vector Autoregressive Model for weekly data, see e.g. Dueker and Startz (1998) for a recent application of this model introduced by Sowell (1987). Moreover, the time series $p_{\text{Fri},t^*} - p_{\text{Thu},t^*}$ does not seem to contain a yearly cycle as $\beta_{3,\text{Fri}}$ and $\beta_{4,\text{Fri}}$ do not differ significantly from zero. Again, the presence of these common features is also indicated in Table 2 where p_{Fri,t^*} and p_{Thu,t^*} display the same slowly decaying autocorrelation function, whereas all autocorrelations of $p_{\text{Fri},t^*} - p_{\text{Thu},t^*}$ are close to zero. As most of the periodic AR polynomials $\phi_j(L)$ seem to contain a "unit root", we also estimated the nonstationary model with daily differences so that $\phi_j(L) = (1 - L)(1 - \phi_{1,j}^*L)$, $j = 1, \dots, s$. This model is clearly rejected by the data, compare the last rows of Table 6. This rejection is due to the bad fit for the equation for Monday.

Summarising the results of Model 4 for the Nord Pool data we find that the model adequately captures the conditional mean of the process. The parameters are comparatively easy

to interpret. As the estimated periodic AR polynomials $\phi_j(L)$ approximately contain a daily difference factor $1 - L$ we can also interpret separate equations in terms of daily differences. However, the combined periodic model is stationary.

6 Volatility in electricity prices

It is well known that spot electricity prices are very volatile. The price can change very much from one day to the other, with increasing variance on some particular days. In general, the variance of the weekend days is smaller than the variance of the days at the beginning of the working week; see, for example, Table 1. Volatility clustering also seems to be present, implying that days with a high conditional variance tend to be followed by days with high variance and more quiet days tend to be followed by quiet days. However, as we have discussed above, the volatility clustering found in these series could be spurious. For the three younger markets, APX, EEX and Powernext, we have found that once the conditional mean is modelled and weekly periodicity in the variance is taken into account, no clear structure seems to remain in the variance of the innovations. This is not the case for the Nord Pool data, where the correlation in the squared residuals is present even after taking into account the periodic features of the data, see Figure 11.

Whereas the linear time series modelling of the conditional mean has received little attention in the literature, a wide range of nonlinear time series models has been applied for volatility in daily electricity prices. One has considered various modifications for the statistical assumptions for the innovations η_t . For example, Knittel and Roberts (2001) consider a jump diffusion model and an exponential GARCH model. Escribano et al. (2002) take η_t as a GARCH process with jumps and consider periodicity in the volatility with respect to the four seasons of the year.

De Jong and Huisman (2002) propose a stationary two-regime Markov switching model for the APX data consisting of an AR(1) regime and a spike regime with high variance to capture outliers. Byström (2001) considers hourly data for the Nord Pool market and fits fat-tailed distributions in an AR-GARCH model.

We do not distinguish a separate regime for the outliers and confine ourselves to a standard GARCH model to capture the volatility clustering in the Nord Pool data. Detailed analysis of the extreme observations is beyond the scope of our analysis.

GARCH models, proposed by Bollerslev (1986) generalising the ARCH models of Engle (1982), are the most popular to capture volatility clustering in financial markets. An extensive review and discussion of these models can be found, for example, in Bollerslev et al. (1994). If $\eta_t^* = \eta_t/\sigma_j$ follows a GARCH(1,1) model, then $\eta_t^* = \varepsilon_t h_t^{1/2}$ where ε_t is a serially independent and identically distributed process with zero mean, unit variance and finite fourth order moment, independent of the volatility process $h_t^{1/2}$ represented as

$$h_t = \alpha_0 + \alpha_1 \eta_{t-1}^{*2} + \gamma_1 h_{t-1}$$

where α_0 , α_1 and γ_1 are parameters such that $\alpha_0 > 0$ and $\alpha_1, \gamma_1 \geq 0$. We have estimated

GARCH models with Gaussian and Student t -distributions for the scaled residuals $\widehat{\eta}_t^*$ of Model 3 for the APX, EEX and Powernext and for Model 4 for the Nord Pool series. We used Maximum Likelihood as implemented in PcGive, see Doornik and Hendry (2001). Following Byström (2001) we include a yearly cycle in the conditional variance equation, represented by $x_{3,t}$ and $x_{4,t}$. Let γ_3 and γ_4 be the corresponding parameters in the conditional variance equation:

$$h_t = \alpha_0 + \alpha_1 \widehat{\eta}_{t-1}^{*2} + \gamma_1 h_{t-1} + \gamma_3 x_{3,t} + \gamma_4 x_{4,t} \quad (18)$$

The t -distribution with unknown and fixed degrees of freedom ν is clearly preferable, increasing the approximate loglikelihood by more than 100 compared with Gaussian errors. The resulting estimated parameters and the corresponding standard errors are given in Table 8.

As can be seen in the table, GARCH behaviour is not clear for the APX and EEX series of residuals since the estimated γ_1 parameter is not statistically significant. However, the Powernext results provide clear evidence of GARCH behaviour. For the Nord Pool data we find significant estimates for α_1 and γ_1 , close to the IGARCH case for which $\alpha_1 + \gamma_1 = 1$. This indicates high persistence in the conditional variance, a typical finding in many financial applications. Moreover, the resulting scaled residuals $\widehat{\varepsilon}_t$ do not show significant serial correlation in the squares.

The two-step estimation method is not optimal. Simultaneous estimation of the seasonal periodic RegARFIMA Model 4 with GARCH- t -errors, combining equations (17) and (18), is not easy using existing software. However it is possible to estimate (18) with a periodic $\alpha_{0,j}$. This does not significantly improve the fit, an indication that the estimates for σ_j presented in Table 7, are satisfactory although they do not take the GARCH- t process into account.

7 Summary and Conclusions

In this paper we have analysed the dynamic behaviour of daily spot prices in deregulated European electricity markets. We find that the day-of-the-week periodic structure present in these data should not be neglected, otherwise spurious volatility clustering could be detected. In particular, in comparison with existing models, the correlation in the squared residuals is reduced by allowing for periodic heteroskedasticity and by considering a periodic model for the conditional mean. We present an empirical analysis for four European markets, the APX in The Netherlands, the EEX in Germany, the Powernext in France and the Nord Pool in Norway. After estimating a range of models for these markets, we propose to use periodic heteroskedastic seasonal RegARIMA models to explain the dynamics in the conditional mean of log prices for APX, EEX and Powernext.

For the series of Nord Pool, additional seasonal periodic long memory features appear to be present. We fit a periodic long memory model that is able to capture the dynamics in the conditional mean. The daily returns on Wednesdays, Thursdays and Fridays seem to be white noise, whereas the returns on Mondays and Saturdays display long memory behaviour. This model is stationary from year to year. The periodic long memory model is not able to capture

the conditional variance dynamics since the squared residuals are still correlated. We fit a GARCH- t model to capture the volatility clustering assuming a Student t -distribution for the error term. This removes serial correlation in the scaled residuals.

The resulting models allow for dynamic point forecasting and stochastic simulation. Future research should entail simultaneous estimation of the periodic parameters for the models for the conditional mean and variance as well as joint modelling of these related markets. Moreover, outliers could be captured by examining other distributions for the shocks affecting daily electricity prices. This would also allow the construction of realistic confidence intervals for forecasts. These interval forecasts can become more realistic as more observations for these young markets become available to estimate and select models.

Acknowledgements

The initial work on this paper was carried out while the first author was visiting the Free University Amsterdam in the spring of 2003. We would like to thank Peter Boswijk, Esther Ruiz, Neil Shephard and Pablo Villaplana for valuable discussions that have led to many improvements and further insights. We are also grateful to Lorenzo Pascual (Endesa Trading) for providing us with the data and for very helpful comments and suggestions. Finally we are indebted to the participants of the Tinbergen Institute seminar where we presented initial results.

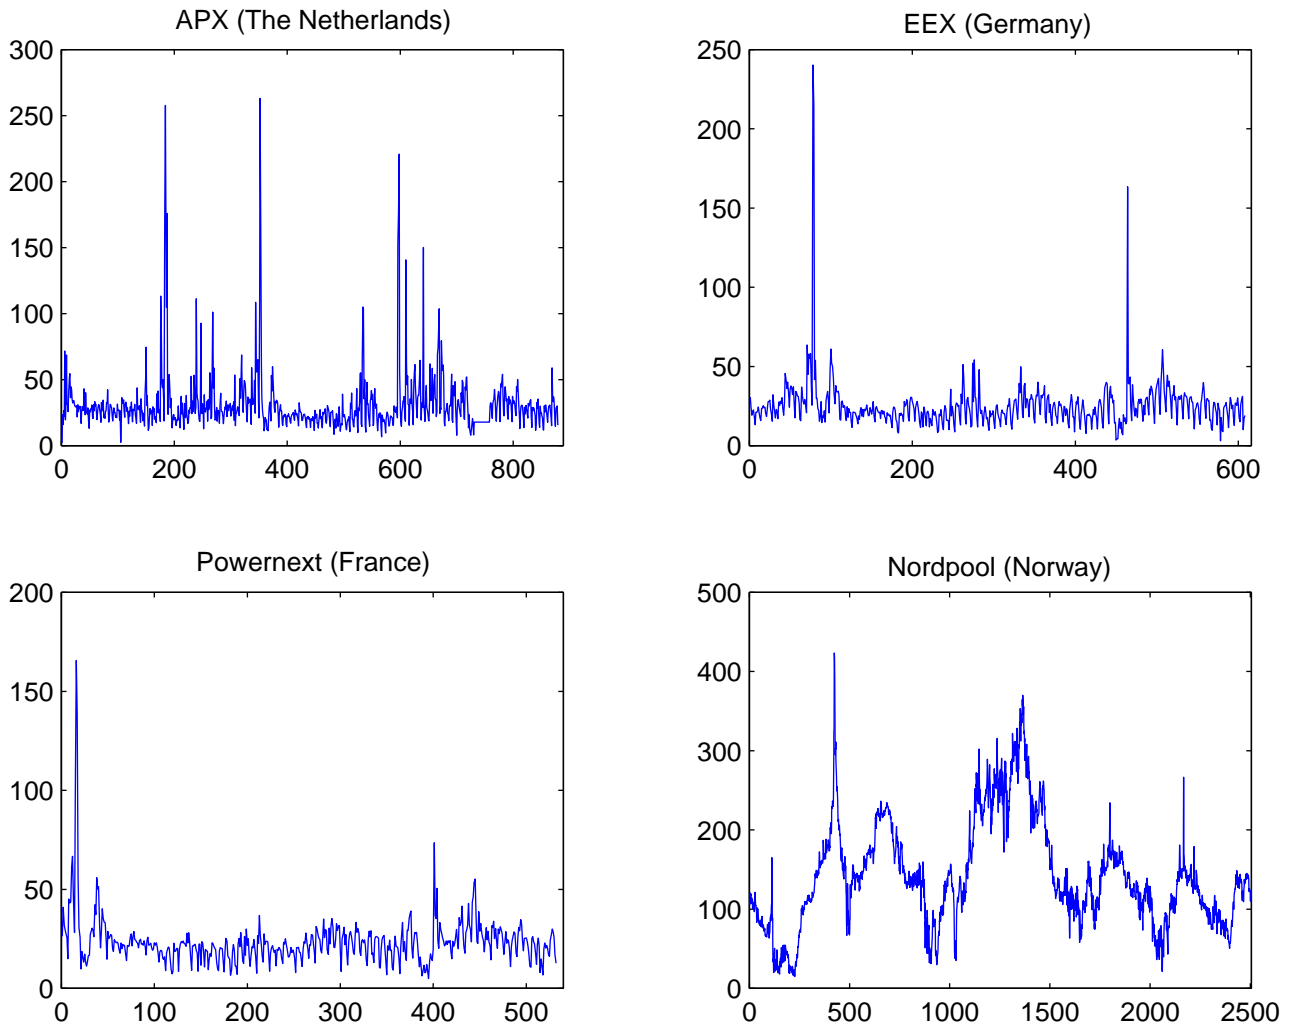
References

- Adenstedt, R. K. (1974). On large-sample estimation for the mean of a stationary random sequence. *The Annals of Statistics*, **2**, 1095–1107.
- Bell, W. R. and S. C. Hillmer (1984). Issues involved with the seasonal adjustment of economic time series. *Journal of Business and Economic Statistics*, **2**, 291–320.
- Beran, J. (1994). *Statistics for Long-Memory Processes*. London: Chapman and Hall.
- Beran, J. (1995). Maximum likelihood estimation of the differencing parameter for invertible short and long memory autoregressive integrated moving average models. *Journal of the Royal Statistical Society*, **57**, 659–672.
- Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, **31**, 307–327.
- Bollerslev, T., R. F. Engle, and D. B. Nelson (1994). ARCH models. *The Handbook of Econometrics*, **4**, 2959–3038.
- Boswijk, H. P. and P. H. Franses (1996). Unit roots in periodic autoregressions. *Journal of Time Series Analysis*, **17**, 221–245.
- Box, George E P and G M Jenkins (1970). *Time Series Analysis: Forecasting and Control*. San Francisco, CA: Holden-Day.

- Brockwell, P. J. and R. A. Davis (1994). *ITSM for Windows, A User's Guide to Time Series Modelling and Forecasting*. Springer Verlag.
- Byström, H. (2001). Extreme value theory and extremely large electricity price changes. Technical Report 19, Department of Economics, Lund University, Forthcoming in *International Review of Economics and Finance*.
- Carlin, J. B., P. Dempster, and A. B. Jonas (1985). On methods and models for Bayesian time series analysis. *Journal of Econometrics*, **30**, 67–90.
- De Jong, C. and R. Huisman (2002). Option Formulas for Mean-Reverting Power Prices with Spikes. Manuscript.
- Doornik, J.A. (2002). *Object-Oriented Matrix Programming using Ox 3.0* (4th ed.). London: Timberlake Consultants Ltd. See <http://www.nuff.ox.ac.uk/Users/Doornik>.
- Doornik, J. A. and D. F. Hendry (2001). *Econometric Modelling Using PcGive, Volume III*. London, U.K.: Timberlake Consultants Press.
- Doornik, J. A. and M. Ooms (1999). A package for estimating, forecasting and simulating Arfima models: Arfima package 1.0 for Ox. Technical report, Nuffield College, Oxford, UK.
- Doornik, J. A. and M. Ooms (2003). Computational aspects of maximum likelihood estimation of autoregressive fractionally integrated moving average models. *Computational Statistics & Data Analysis*, **42**, 333–348.
- Dueker, M. and R. Startz (1998). Maximum-likelihood estimation of fractional cointegration with an application to U.S. and Canadian bond rates. *Review of Economics and Statistics*, **80**, 420–426.
- Engle, R. F. (1982). Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, **50**, 987–1007.
- Escribano, A., J.I. Peña, and P. Villaplana (2002). Modeling Electricity Prices: International Evidence. W.P. 02-27, Universidad Carlos III de Madrid.
- Findley, D.F., B.C. Monsell, W.R. Bell, M.C. Otto, and B-C Chen (1998). New capabilities of the X-12-ARIMA seasonal adjustment program (with discussion). *Journal of Business and Economic Statistics*, **16**, 127–77.
- Franses, P. H. (1996). *Periodicity and Stochastic Trends in Economic Time Series*. Oxford, UK.: Oxford University Press.
- Franses, P. H. and M. Ooms (1997). A periodic long memory model for quarterly UK inflation. *International Journal of Forecasting*, **13**, 119–128.
- Ghysels, E. and D. R. Osborn (2001). *The Econometric Analysis of Seasonal Time Series*. Cambridge, UK: Cambridge University Press.
- Granger, C. W. J. and R. Joyeux (1980). An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis*, **1**, 15–29.
- Hosking, J. R. M. (1981). Fractional differencing. *Biometrika*, **68**, 165–176.

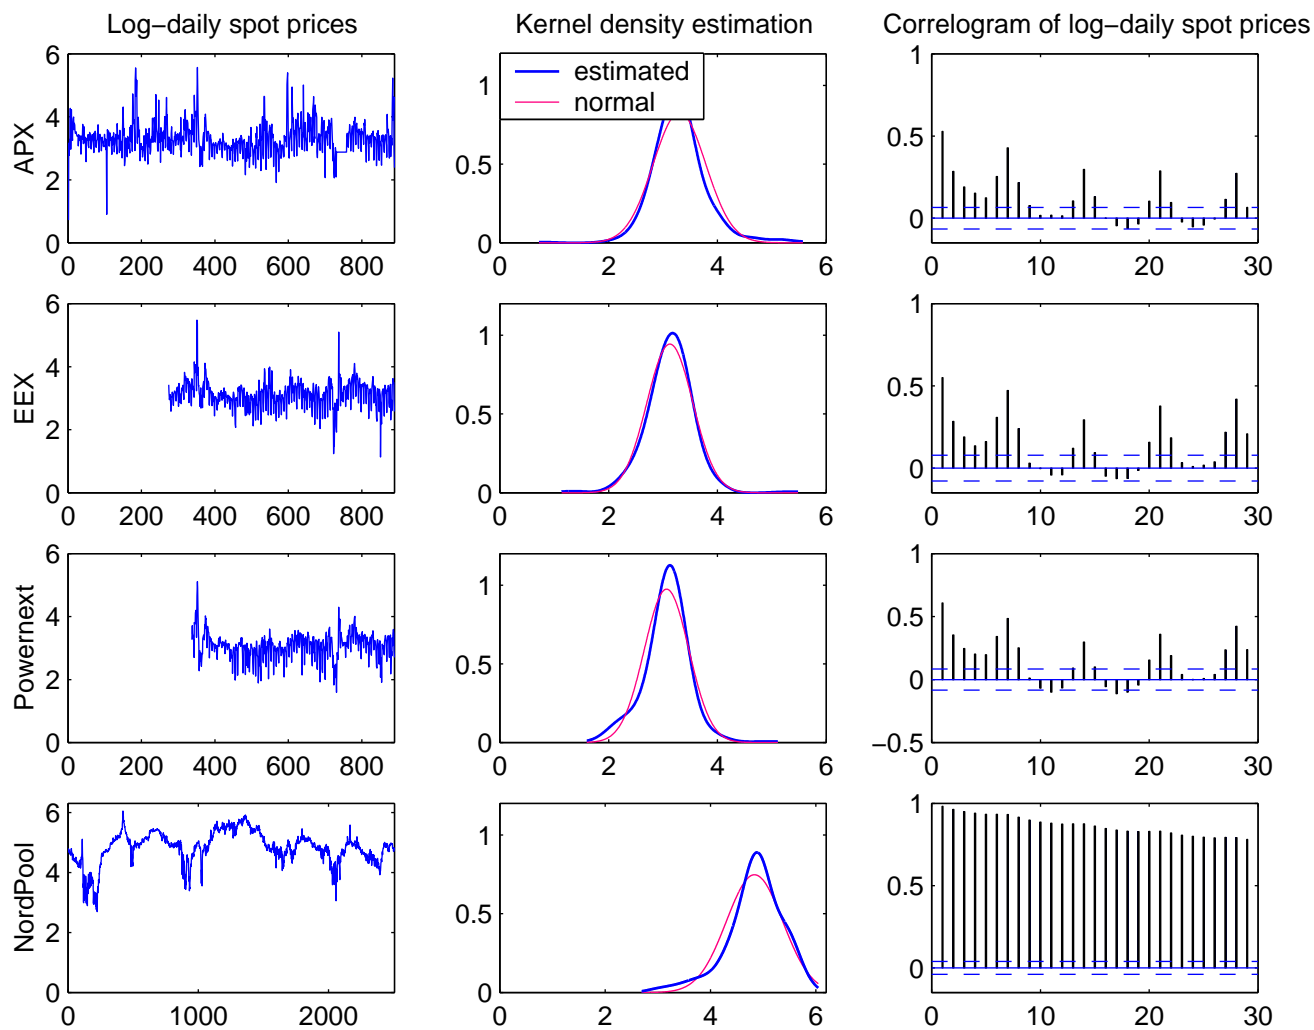
- Johnston, J. and J. Dinardo (1997). *Econometric Methods, fourth edition*. McGraw-Hill, New York etc.
- Jones, R. H. and W. M. Brelsford (1967). Time series with periodic structure. *Biometrika*, **54**, 403–408.
- Knittel, C.R. and M.R. Roberts (2001). An Empirical Examination of Deregulated Electricity Prices. PWP 087, POWER. University of California Energy Institute.
- Lucia, J.J. and E.S. Schwartz (2002). Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research*, **5**, 5–50.
- McLeod, A.I. (1994). Diagnostic Checking of Periodic Autoregression Models with Application. *Journal of Time Series Analysis*, **15**, 221–233.
- Oberhofer, W. and J. Kmenta (1974). A general procedure for obtaining maximum likelihood estimates in generalized regression models. *Econometrica*, **42**, 579–590.
- Ooms, M. (1994). *Empirical Vector Autoregressive Modeling*, Volume 407 of *Lecture Notes in Economics and Mathematical Systems*. Berlin, Germany: Springer-Verlag.
- Ooms, M. and P.H. Franses (2001). A Seasonal Periodic Long Memory Model for Monthly River Flows. *Environmental Modelling & Software*, **16**, 559–569.
- Ooms, M. and P. H. Franses (1997). On periodic correlations between estimated seasonal and nonseasonal components in German and U.S. unemployment. *Journal of Business and Economic Statistics*, **15**, 470–481.
- Porter-Hudak, S. (1990). An application of the seasonal fractionally differenced model to the monetary aggregates. *Journal of the American Statistical Association*, **85**, 338–344.
- Robinson, P. M. (Ed.) (2003). *Time Series with Long Memory*, Oxford, UK. Oxford University Press.
- Robinson, P. M. and Y. Yajima (2002). Determination of cointegrating rank in fractional systems. *Journal of Econometrics*, **106**, 217–241.
- Shumway, R. H. and D. S. Stoffer (2000). *Time Series Analysis and Its Applications*. New York: Springer-Verlag.
- Sowell, F. (1987). Maximum likelihood estimation of fractionally integrated time series models. Technical Report 87-07, Department of Economics, Duke University.
- Sowell, F. (1992). Maximum likelihood estimation of stationary univariate fractionally integrated time series models. *Journal of Econometrics*, **53**, 165–188.
- Tiao, G.C. and M.R. Grupe (1980). Hidden Periodic Autoregressive Moving Average Models in Time Series Data. *Biometrika*, **67**, 365–373.
- Vecchia, A. V. (1985). Maximum likelihood estimation for periodic autoregressive moving average models. *Technometrics*, **27**, 375–384.
- Wilkinson, L. and J. Winsen (2002). What we can Learn from a Statistical Analysis of Electricity Prices in New South Wales. *The Electricity Journal*, **April**, 60–69.

Figure 1: Daily spot prices for four European electricity markets



NOTES: APX: January, 1, 2001 - June, 8, 2003, EEX: October, 1, 2001 - June, 8, 2003,
Powernext: December, 3, 2001 - June, 8, 2003, price in Euros/MWh,
Nord Pool: January, 4, 1993 - November, 14, 1999, price in NOK/MWh .

Figure 2: Log daily spot prices for four European electricity markets



NOTES: APX: January, 1, 2001 - June, 8, 2003, EEX: October, 1, 2001 - June, 8, 2003,
 Powernext: December, 3, 2001 - June, 8, 2003, log price in Euros/MWh,
 Nord Pool: January, 4, 1993 - November, 14, 1999, log price in NOK/MWh .

Table 1: Descriptive statistics of daily log-prices for new markets

APX (The Netherlands)								
Day	All	Mon	Tue	Wed	Thu	Fri	Sat	Sun
T	889	127	127	127	127	127	127	127
Mean	3.30	3.44	3.47	3.45	3.44	3.37	3.12	2.82
S.D.	0.48	0.55	0.53	0.45	0.41	0.35	0.29	0.33
Skewness	0.73	-0.06	1.51	1.51	0.59	1.43	1.16	-1.41
Kurtosis	7.31	8.21	7.17	7.73	5.44	8.17	5.47	11.37
$r(1)$	0.53	0.40	0.72	0.83	0.71	0.69	0.30	0.46
$r(2)$	0.28	0.34	0.25	0.74	0.63	0.64	0.33	0.22
$r(7)$	0.43	0.23	0.25	0.26	0.34	0.30	0.42	0.21
$r(14)$	0.30	0.09	0.03	0.05	0.14	0.06	0.34	0.13
EEEX (Germany)								
Day	All	Mon	Tue	Wed	Thu	Fri	Sat	Sun
T	616	88	88	88	88	88	88	88
Mean	3.12	3.24	3.31	3.27	3.25	3.20	2.95	2.65
S.D.	0.42	0.40	0.47	0.38	0.42	0.29	0.24	0.26
Skewness	0.03	1.60	0.48	-1.43	-2.17	0.30	0.29	0.04
Kurtosis	7.72	13.80	12.24	9.36	11.94	2.86	3.21	2.67
$r(1)$	0.55	0.63	0.60	0.74	0.67	0.71	0.70	0.83
$r(2)$	0.28	0.52	0.34	0.54	0.58	0.73	0.58	0.64
$r(7)$	0.47	0.32	0.10	0.35	0.15	0.38	0.40	0.62
$r(14)$	0.29	0.14	-0.30	-0.01	0.03	0.09	0.33	0.47
Powernext (France)								
Day	All	Mon	Tue	Wed	Thu	Fri	Sat	Sun
T	553	79	79	79	79	79	79	79
Mean	3.06	3.14	3.23	3.23	3.20	3.16	2.90	2.60
S.D.	0.41	0.35	0.36	0.40	0.37	0.30	0.28	0.36
Skewness	-0.03	0.37	1.72	-0.32	-0.54	0.46	-0.14	-0.14
Kurtosis	5.75	7.66	12.39	10.76	5.04	4.58	4.59	2.18
$r(1)$	0.61	0.63	0.67	0.83	0.69	0.82	0.79	0.62
$r(2)$	0.35	0.72	0.47	0.70	0.63	0.64	0.66	0.62
$r(7)$	0.48	0.30	0.14	0.26	0.26	0.32	0.42	0.40
$r(14)$	0.30	0.13	-0.09	-0.05	-0.01	-0.09	0.14	0.21

NOTES: T: Sample size.

$r(\tau)$: Periodic autocorrelation of y_t for a lag of τ days

Table 2: Descriptive statistics of daily log-prices Nord Pool

Day	All	Mon	Tue	Wed	Thu	Fri	Sat	Sun
T	2506	358	358	358	358	358	358	358
Mean	4.84	4.86	4.88	4.87	4.87	4.85	4.78	4.76
S.D.	0.53	0.52	0.50	0.50	0.51	0.52	0.57	0.58
Skewness	-0.90	-0.84	-0.78	-0.80	-0.82	-0.84	-0.99	-0.98
Kurtosis	4.41	4.35	4.33	4.21	4.31	4.24	4.47	4.23
$r(1)$	0.98	0.98	0.98	0.99	0.99	0.99	0.98	0.99
$r(2)$	0.96	0.98	0.97	0.97	0.98	0.98	0.97	0.98
$r(7)$	0.93	0.93	0.94	0.94	0.93	0.93	0.93	0.92
$r(14)$	0.87	0.86	0.88	0.88	0.88	0.88	0.87	0.86
First difference of log-prices								
Day	All	Mon	Tue	Wed	Thu	Fri	Sat	Sun
T	2506	358	358	358	358	358	358	358
Mean	0.00	0.10	0.02	-0.00	-0.01	-0.02	-0.06	-0.02
S.D.	0.10	0.12	0.10	0.07	0.07	0.06	0.11	0.06
Skewness	0.43	1.75	2.36	-1.97	-0.25	-1.12	-2.64	-0.97
Kurtosis	14.44	7.14	18.90	26.66	14.17	8.60	12.69	8.01
$r(1)$	0.06	-0.27	0.07	0.09	0.04	-0.01	0.26	0.16
$r(2)$	-0.19	-0.58	0.09	0.03	-0.01	0.05	0.11	0.13
$r(7)$	0.41	0.39	0.02	-0.13	-0.01	0.04	0.57	0.19
$r(14)$	0.38	0.38	-0.01	0.05	-0.01	0.06	0.38	0.05

NOTES: T: Sample size.

$r(\tau)$: Periodic autocorrelation of y_t for a lag of τ days

Table 3: Estimation results for daily log-prices APX (The Netherlands)

Model	Day	$\hat{\beta}_1$	$\hat{\beta}_3$	$\hat{\beta}_4$	$\hat{\phi}_1$	$\hat{\phi}_2$	$10\hat{\sigma}^2$
1		3.46 ^{**} _(0.03)	-0.15 ^{**} _(0.04)	0.02 _(0.04)	0.54 ^{**} _(0.03)		11.76
	Mon	3.47 ^{**} _(0.05)					17.98
	Tue	3.48 ^{**} _(0.05)					13.05
	Wed	3.46 ^{**} _(0.04)					6.03
2	Thu	3.45 ^{**} _(0.04)	-0.11 [*] _(0.04)	0.06 _(0.04)	0.53 ^{**} _(0.03)	0.10 ^{**} _(0.03)	7.68
	Fri	3.39 ^{**} _(0.04)					6.43
	Sat	3.12 ^{**} _(0.04)					7.50
	Sun	2.83 ^{**} _(0.04)					8.42
Model	Day	$\hat{\beta}_1^*$	$\hat{\beta}_3^*$	$\hat{\beta}_4^*$	$\hat{\phi}_1$	$\hat{\phi}_2$	$10\hat{\sigma}^2$
	Mon	0.57 _(0.46)	-0.10 _(0.05)	-0.18 ^{**} _(0.06)	0.60 ^{**} _(0.13)	0.39 ^{**} _(0.14)	16.80
	Tue	0.91 ^{**} _(0.29)	-0.05 _(0.04)	-0.06 _(0.05)	0.87 ^{**} _(0.07)	-0.15 _(0.11)	11.10
	Wed	0.68 ^{**} _(0.15)	0.04 _(0.03)	0.05 _(0.03)	0.54 ^{**} _(0.06)	0.26 ^{**} _(0.05)	5.60
3	Thu	1.39 ^{**} _(0.19)	0.13 ^{**} _(0.04)	0.05 _(0.03)	0.53 ^{**} _(0.09)	0.07 _(0.08)	7.50
	Fri	1.08 ^{**} _(0.20)	0.02 _(0.03)	-0.03 _(0.03)	0.45 ^{**} _(0.08)	0.22 ^{**} _(0.07)	6.10
	Sat	2.29 ^{**} _(0.24)	-0.08 [*] _(0.04)	0.11 [*] _(0.03)	0.18 _(0.09)	0.07 _(0.08)	6.90
	Sun	1.32 ^{**} _(0.33)	-0.06 _(0.04)	0.12 ^{**} _(0.04)	0.37 ^{**} _(0.10)	0.10 _(0.07)	8.10

NOTES: In Model 1, see (14): $\hat{\beta}_7 = -0.45^{**}_(0.03)$ and in Model 2, see (15): $\hat{\Phi}_1 = 0.09[*]_(0.03)$. Model 3 is defined in (16). * and ** indicate significance at 5% and 1% respectively.

Table 4: Estimation results for daily log-prices EEX (Germany)

Model	Day	$\hat{\beta}_1$	$\hat{\beta}_3$	$\hat{\beta}_4$	$\hat{\phi}_1$	$\hat{\phi}_2$	$10\hat{\sigma}^2$
1		3.24 ^{**} _(0.03)	0.10 [*] _(0.04)	0.02 _(0.04)	0.58 ^{**} _(0.03)		8.71
	Mon	3.22 ^{**} _(0.04)					10.52
	Tue	3.30 ^{**} _(0.06)					15.21
	Wed	3.26 ^{**} _(0.04)					6.32
2	Thu	3.23 ^{**} _(0.05)	0.15 ^{**} _(0.04)	-0.01 _(0.04)	0.48 ^{**} _(0.04)	0.16 ^{**} _(0.03)	9.52
	Fri	3.19 ^{**} _(0.04)					3.56
	Sat	2.93 ^{**} _(0.03)					2.64
	Sun	2.63 ^{**} _(0.03)					4.56
Model	Day	$\hat{\beta}_1^*$	$\hat{\beta}_3^*$	$\hat{\beta}_4^*$	$\hat{\phi}_1$	$\hat{\phi}_2$	$10\hat{\sigma}^2$
	Mon	0.56 _(0.46)	-0.06 _(0.06)	0.05 ^{**} _(0.05)	0.95 ^{**} _(0.23)	0.06 _(0.25)	10.17
	Tue	1.10 [*] _(0.45)	-0.04 _(0.06)	0.01 _(0.06)	0.74 ^{**} _(0.13)	-0.07 _(0.21)	14.63
	Wed	1.04 ^{**} _(0.23)	-0.02 _(0.04)	-0.01 _(0.04)	0.51 ^{**} _(0.07)	0.17 _(0.08)	6.51
3	Thu	0.74 ^{**} _(0.27)	0.11 [*] _(0.05)	0.04 _(0.04)	0.61 ^{**} _(0.12)	0.15 _(0.10)	9.16
	Fri	1.21 ^{**} _(0.17)	0.07 [*] _(0.03)	0.01 _(0.03)	0.23 [*] _(0.06)	0.38 ^{**} _(0.07)	3.07
	Sat	1.24 ^{**} _(0.19)	0.12 ^{**} _(0.02)	-0.01 _(0.02)	0.44 ^{**} _(0.08)	0.09 _(0.05)	2.46
	Sun	-0.03 _(0.21)	0.02 _(0.03)	0.01 _(0.02)	0.80 ^{**} _(0.10)	0.10 _(0.08)	2.17


NOTES: In Model 1, see (14): $\hat{\beta}_7 = -0.43^{**}_(0.03)$ and in Model 2, see (15): $\hat{\Phi}_1 = 0.10[*]_(0.04)$. Model 3 is defined in (16). * and ** indicate significance at 5% and 1% respectively. Standard errors in parenthesis.

Table 5: Estimation results for daily log-prices Powernext (France)

Model	Day	$\hat{\beta}_1$	$\hat{\beta}_3$	$\hat{\beta}_4$	$\hat{\phi}_1$	$\hat{\phi}_2$	$10\hat{\sigma}^2$
1		3.17** _(0.03)	0.05 _(0.05)	0.13** _(0.04)	0.64** _(0.03)		7.08
	Mon	3.11** _(0.05)					6.64
	Tue	3.20** _(0.05)					6.99
	Wed	3.20** _(0.05)					5.18
2	Thu	3.18** _(0.05)	0.09 _(0.06)	0.07 _(0.05)	0.59** _(0.04)	0.15** _(0.04)	7.06
	Fri	3.13** _(0.04)					2.83
	Sat	2.87** _(0.04)					6.43
	Sun	2.58** _(0.05)					4.70
Model	Day	$\hat{\beta}_1^*$	$\hat{\beta}_3^*$	$\hat{\beta}_4^*$	$\hat{\phi}_1$	$\hat{\phi}_2$	$10\hat{\sigma}^2$
	Mon	0.52 _(0.28)	-0.10* _(0.04)	0.08* _(0.04)	0.23* _(0.10)	0.71** _(0.13)	5.48
	Tue	1.04** _(0.29)	0.01 _(0.05)	0.01 _(0.04)	0.63** _(0.11)	0.08 _(0.11)	7.27
	Wed	-0.17 _(0.24)	-0.04 _(0.03)	-0.06 _(0.03)	0.75** _(0.08)	0.32** _(0.09)	4.30
3	Thu	1.14** _(0.27)	0.01 _(0.04)	0.10* _(0.04)	0.53** _(0.13)	0.11 _(0.15)	6.89
	Fri	0.98** _(0.18)	0.04 _(0.03)	0.05 _(0.03)	0.55** _(0.07)	0.13 _(0.06)	2.95
	Sat	0.59** _(0.20)	0.07** _(0.03)	-0.01 _(0.03)	0.67** _(0.11)	0.06 _(0.08)	2.73
	Sun	-0.02 _(0.32)	0.05 _(0.04)	0.06 _(0.04)	0.76** _(0.17)	0.12 _(0.16)	6.39

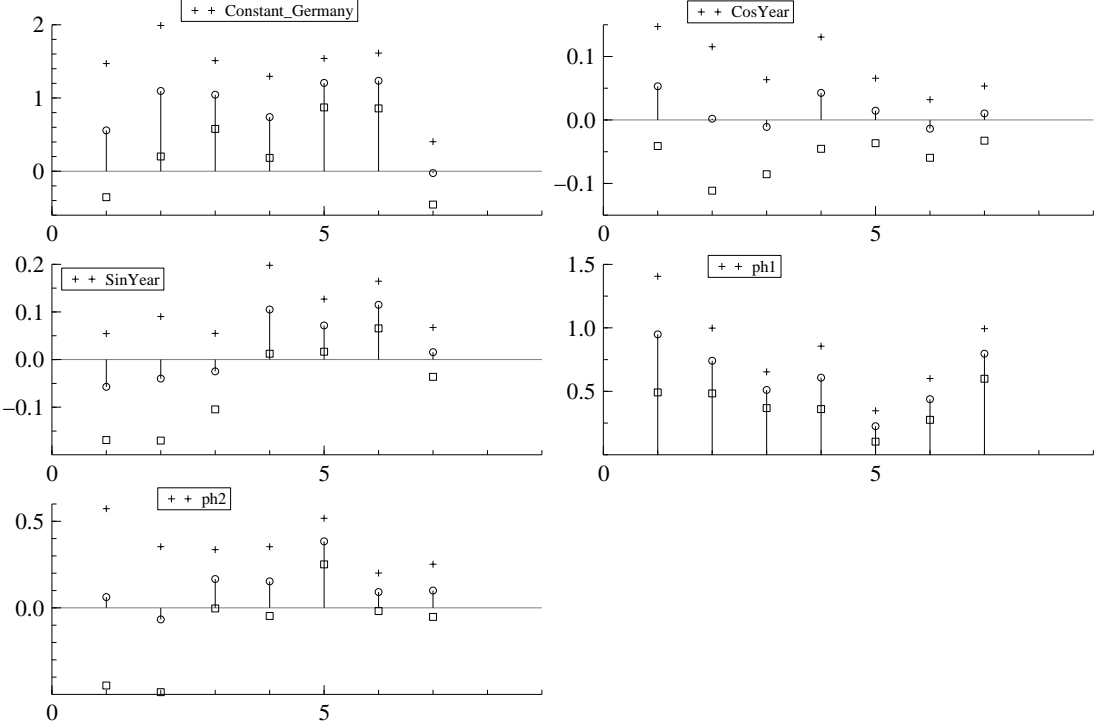
NOTES: In Model 1, see (14): $\hat{\beta}_7 = -0.40$ **_(0.02) and in Model 2, see (15): $\hat{\Phi}_1 = 0.05$ _(0.04). Model 3 is defined in (16). * and ** indicate significance at 5% and 1% respectively. Standard errors in parenthesis.

Figure 3: Estimated parameters of the periodic Model 3 for the APX




NOTES: See equation (16). 95% confidence intervals for $\beta_{1,j}$, $\beta_{3,j}$, $\beta_{4,j}$, $\phi_{1,j}$, $\phi_{2,j}$ respectively. $j = 1$: Monday, \dots , $j = 7$: Sunday.

Figure 4: Estimated parameters of the periodic Model 3 for the EEX



NOTES: See equation (16). 95% confidence intervals for $\beta_{1,j}$, $\beta_{3,j}$, $\beta_{4,j}$, $\phi_{1,j}$, $\phi_{2,j}$ respectively. $j = 1$: Monday, \dots , $j = 7$: Sunday.

Figure 5: Estimated parameters of the periodic Model 3 for Powernext



NOTES: See equation (16). 95% confidence intervals for $\beta_{1,j}$, $\beta_{3,j}$, $\beta_{4,j}$, $\phi_{1,j}$, $\phi_{2,j}$ respectively. $j = 1$: Monday, \dots , $j = 7$: Sunday.

Figure 6: Residuals of the three models for the APX data


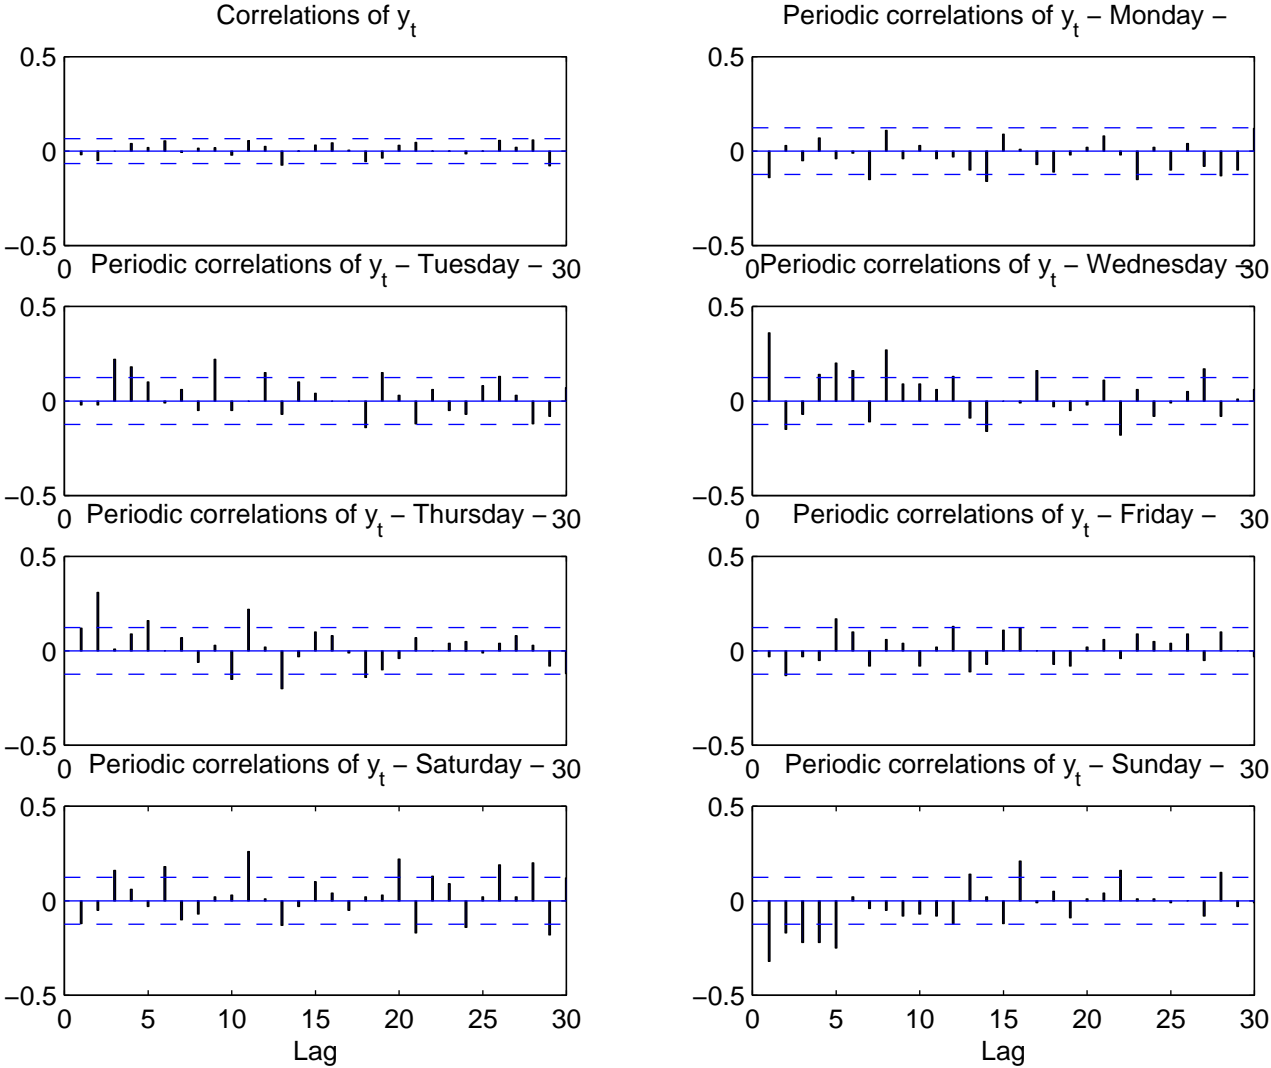


Figure 7: Correlations and Periodic correlations of the residuals of Model 2 for the APX



NOTES: Model 2 is defined in (15).

Figure 8: Residuals of the three models for the EEX data




Figure 9: Residuals of the three models for the Powernext data




Table 6: Goodness of fit of different models

Model	k	T	APX 889		EEX 616		Pownext 553		Nord Pool 2506	
			Llik	AIC	Llik	AIC	Llik	AIC	Llik	AIC
1	6		-308.06	0.7066	-120.45	0.4105	-50.76	0.2053	2533.4	-2.0171
1s	12		-274.97	0.6456	-86.35	0.3193	-19.24	0.1130	2690.5	-2.1376
2ns	13		-212.80	0.5080	-53.15	0.2148	26.29	-0.0481	2658.7	-2.1115
2	19		-181.47	0.4510	11.58	0.0241	43.46	-0.0885	2779.6	-2.2032
2b	25		-186.92	0.4767	-45.02	0.2273	38.92	-0.0503	2714.6	-2.1465
2bs	31		-157.32	0.4237	18.61	0.0402	56.42	-0.0919	2824.6	-2.2296
3	42		-135.06	0.3983	39.97	0.0066	71.60	-0.1071	2896.6	-2.2779
4d	42								2892.01	-2.27455
4	49								2944.5	-2.3108

NOTES: T is the number of observations. The effective number of observations slightly varies from model to model depending on the (periodic) AR order of the model. Llik: the loglikelihood used for estimation: see (10). AIC: Akaike Information criterion defined as $= (1/T) \cdot (-2 \cdot \text{Llik} + 2 \cdot k)$, with k the number of parameters. Models 1, 2 and 3 (reported in bold) are defined in section 5.1, see (14), (15), see (16). Model 4 (also reported in bold) is defined in section 5.2, see (17). The other models are defined as follows:

Model 1s is Model 1 allowing for periodic variances: $\sigma_j^2, j = 1, \dots, s$,

Model 2ns is Model 2 restricting the variance to be constant: $\sigma_j^2 = \sigma^2$,

Model 2b is Model 2ns allowing for periodic coefficients for the yearly cycle: $\beta_{3,j}$ and $\beta_{4,j}, j = 1, \dots, s$,

Model 2bs is Model 2b allowing for periodic variances: $\sigma_j^2, j = 1, \dots, s$,

Model 4d is Model 4 with “unit roots“ imposed in the PAR part: $\phi_j(L) = (1 - L)(1 - \phi_{1,j}^* L), j = 1, \dots, s$.

Table 7: Estimation results for daily log-prices Nord Pool (Norway)


Model	Day	$\hat{\beta}_1$	$\hat{\beta}_3$	$\hat{\beta}_4$	$\hat{\phi}_1$	$\hat{\phi}_2$	\hat{D}	$10\hat{\sigma}^2$
1		4.87** _(0.11)	0.01 _(0.01)	0.32** _(0.10)	0.98** _(0.01)			0.08
	Mon	4.88** _(0.13)						0.12
	Tue	4.90** _(0.13)						0.09
	Wed	4.89** _(0.13)						0.06
2	Thu	4.89** _(0.13)	0.17 _(0.13)	0.52** _(0.12)	1.05** _(0.02)	-0.06** _(0.02)		0.05
	Fri	4.87** _(0.13)						0.04
	Sat	4.80** _(0.13)						0.10
	Sun	4.78** _(0.13)						0.03
Model	Day	$\hat{\beta}_1^*$	$\hat{\beta}_3^*$	$\hat{\beta}_4^*$	$\hat{\phi}_1$	$\hat{\phi}_2$	\hat{D}	$10\hat{\sigma}^2$
	Mon	0.54** _(0.05)	-0.01 _(0.01)	-0.03** _(0.01)	0.56** _(0.08)	0.34** _(0.08)		0.09
	Tue	0.29** _(0.05)	0.01 _(0.01)	0.01 _(0.01)	0.92** _(0.05)	0.03 _(0.04)		0.09
	Wed	0.06 _(0.04)	-0.01* _(0.01)	0.01 _(0.01)	1.05** _(0.04)	-0.06 _(0.04)		0.05
3	Thu	0.01 _(0.04)	-0.01 _(0.01)	0.01 _(0.01)	1.02** _(0.05)	-0.03 _(0.05)		0.04
	Fri	-0.06 _(0.04)	-0.01 _(0.01)	0.01 _(0.01)	0.98** _(0.05)	0.03 _(0.05)		0.04
	Sat	-0.41** _(0.05)	-0.01 _(0.01)	0.02* _(0.01)	1.36** _(0.08)	-0.28** _(0.08)		0.09
	Sun	0.08* _(0.03)	-0.01* _(0.01)	0.02** _(0.01)	1.09** _(0.03)	-0.11** _(0.03)		0.03
	Mon	0.71** _(0.09)	0.01 _(0.01)	-0.02 _(0.01)	0.65** _(0.08)	0.22* _(0.09)	0.22** _(0.05)	0.09
	Tue	0.29** _(0.07)	0.01 _(0.01)	0.01 _(0.01)	0.92** _(0.05)	0.03 _(0.05)	0.00 _(0.06)	0.09
	Wed	0.03 _(0.03)	-0.01** _(0.01)	0.01 _(0.01)	1.04** _(0.04)	-0.05 _(0.04)	-0.10* _(0.04)	0.05
4	Thu	0.01 _(0.04)	-0.01 _(0.01)	0.01 _(0.01)	1.02** _(0.05)	-0.02 _(0.05)	-0.02 _(0.04)	0.04
	Fri	-0.06 _(0.04)	-0.01 _(0.01)	0.01 _(0.01)	0.98** _(0.05)	0.03 _(0.05)	-0.01 _(0.04)	0.04
	Sat	-0.25* _(0.11)	-0.01 _(0.01)	0.03* _(0.02)	1.25** _(0.07)	-0.21** _(0.07)	0.39** _(0.06)	0.08
	Sun	0.08 _(0.04)	-0.01 _(0.01)	0.02** _(0.01)	1.09** _(0.03)	-0.11** _(0.04)	0.07 _(0.05)	0.03

NOTES: In Model 1: $\hat{\beta}_7 = -0.08$ **_(0.01) and in Model 2: $\hat{\Phi}_1 = 0.17$ **_(0.02).

* and ** indicate significance at 5% and 1% nominal level respectively.

Standard errors in parenthesis.

Figure 10: Estimated parameters of the seasonal periodic long memory Model 4 for the Nord Pool data



NOTES: See equation (17). 95% confidence intervals for D_j , $\beta_{1,j}$, $\beta_{3,j}$, $\beta_{4,j}$, $\phi_{1,j}$, $\phi_{2,j}$ respectively. $j = 1 : \text{Monday}, \dots, j = 7 : \text{Sunday}$.

Table 8: GARCH- t models applied to $\hat{\eta}_t/\sigma_j$

Market	Model	$\hat{\alpha}_0$	$\hat{\alpha}_1$	$\hat{\gamma}_1$	$\hat{\gamma}_3$	$\hat{\gamma}_4$	$\hat{\nu}$
APX	3	0.44** _(0.16)	0.25** _(0.07)	0.32 _(0.19)	-0.13* _(0.05)	-0.25* _(0.10)	5.86** _(1.09)
EEX	3	0.43** _(0.10)	0.36** _(0.12)	0.23 _(0.12)	0.02 _(0.06)	0.08 _(0.06)	4.25** _(0.73)
Powernext	3	0.06* _(0.03)	0.13** _(0.04)	0.82** _(0.05)	-0.01 _(0.01)	-0.01 _(0.01)	4.80** _(0.97)
Nord Pool	4	0.04** _(0.01)	0.24** _(0.03)	0.74** _(0.03)	-0.01* _(0.01)	0.02** _(0.01)	5.06** _(0.51)

NOTES: Estimates for equation (18) applied to scaled residuals of Model 3 and 4; see (16) and (17) respectively. Standard errors in parenthesis.

Figure 11: Scaled Residuals of the fitted models for the Nord Pool data

