Equation-by-Equation Estimation of Multivariate Periodic Electricity Price Volatility *

Álvaro Escribano† and Genaro Sucarrat‡

This version: 29th January 2018
[First version: 22 July 2016]

Abstract

Electricity prices are characterised by strong autoregressive persistence, periodicity (e.g. intraday, day-of-the week and month-of-the-year effects), large spikes or jumps, GARCH and – as evidenced by recent findings – periodic volatility. We propose a multivariate model of volatility that decomposes volatility multiplicatively into a non-stationary (e.g. periodic) part and a stationary part with log-GARCH dynamics. Since the model belongs to the log-GARCH class, the model is robust to spikes or jumps, allows for a rich variety of volatility dynamics without restrictive positivity constraints, can be estimated equation-by-equation by means of standard methods even in the presence of feedback, and allows for Dynamic Conditional Correlations (DCCs) that can – optionally – be estimated subsequent to the volatilities. We use the model to study the hourly day-ahead system prices at Nord Pool, and find extensive evidence of periodic volatility and volatility feedback. We also find that volatility is characterised by (positive) leverage in one third of the hours, and that a DCC model provides a better fit of the conditional correlations than a Constant Conditional Correlation (CCC) model.

JEL Classification: C22, C32, C51, C58
Keywords: Electricity prices, financial return, volatility, ARCH, exponential GARCH, log-GARCH, Multivariate GARCH, Dynamic Conditional Correlations, inverse leverage, Nord Pool

---

*We are grateful to the Editor, two anonymous reviewers, Juan Ignacio Peña, and participants at the CFE 2017 conference (London), the 2017 CATE workshop (Oslo), the International Conference in Honour of Luc Bauwens (Brussels), the 2017 SDNDE conference (Paris), the GREQAM seminar (Marseille) and ECOMFIN2016 (Paris) for useful comments, suggestions and questions.

†Department of Economics, Universidad Carlos III de Madrid (Spain).

‡Corresponding author. Department of Economics, BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway. Email genaro.sucarrat@bi.no, phone +47+46410779, fax +47+23264788. Webpage: http://www.sucarrat.net/

1 Introduction

Modelling the uncertainty or volatility of electricity prices is of great importance for energy market participants. On the supply side, producers of electricity need estimates of the time-varying price volatility in order to determine the risks of future production levels. On the demand side, consumers of electricity need the same type of information in order to ascertain the risks associated with decisions about when and where to produce goods, and in order to hedge against adverse price changes.

It is well known that electricity prices are characterised by autoregressive persistence, periodicity effects (e.g. hour-of-the-day, day-of-the-week and month-of-the-year effects) in the conditional mean, see e.g. Bunn (2000), Knittel and Roberts (2005), Janczura et al. (2013), and Weron (2014). It is also well known that the volatility of electricity prices is characterised by Autoregressive Conditional Heteroscedasticity (ARCH) and large spikes or jumps, see e.g. Escribano et al. (2002, 2011), Koopman et al. (2007), and Hellström et al. (2012). Since the periodicity effects in the conditional mean usually account for a considerable proportion of the conditional mean dynamics, it is reasonable to conjecture that the same may also be the case for volatility. Recently, this line of research has received increasing attention. Bauwens et al. (2013, Section 4.2), for example, in a three-dimensional multivariate model of monthly, quarterly and yearly Phelix baseload futures at the European Energy Exchange, find that volatility depends on the number of days-to-delivery, i.e. that the volatility increases as the future in question approaches maturity. Sucarrat et al. (2016, Section 4), in a two-dimensional multivariate model of peak and off-peak day-ahead prices in the Oslo region (Nord Pool), find that day-of-the-week effects matter for volatility, and that peak volatility dynamics is less persistent than off-peak. Dupuis (2017), in a fifteen-dimensional multivariate model of electricity prices in the New York area, includes dummies in the volatility equations to accommodate hour-of-the-day and day-of-the-week effects.

There are two main challenges in the multivariate modelling of electricity price volatility. The first is the socalled "curse of dimensionality": As the multivariate dimension grows, joint estimation of the full model becomes infeasible in practice due to the number of parameters that has to be estimated. This problem is not specific to electricity prices, but it is more severe. The reason is that volatility is likely to depend on additional covariates, e.g. weather and market specific stochastic conditioning variables, in addition to periodicity effects similar to those that often characterise the conditional mean dynamics. Moreover, if standard or non-exponential GARCH models are used, then the curse of dimensionality problem is compounded, since the covariates and/or their parameters need to be restricted in estimation in order to ensure the positivity of fitted volatility. An example in which such a parameter restriction is needed in electricity price markets is the socalled "inverse leverage effect", as coined by Knittel and Roberts (2005), whereby negative shocks in one period leads to a reduction in volatility in the next period.[1] Knittel and Roberts (2005) avoid the need for a restriction by using Nelson's (1991) Exponential GARCH (EGARCH). However, as is well-known, the EGARCH is not robust to spikes.[2] This leads to the second main challenge in the modelling of electricity prices: The occurrence of price spikes. It is well-known that the ordinary GARCH model is not robust to such spikes. This is because the spikes affect estimation and inference inadvertently (Carnero et al. (2007), Gregory and Reeves (2010)), and because it makes the model propensive to volatility forecast failure subsequent to the spikes, see e.g. Harvey and Sucarrat (2014, Introduction). One multivariate model specification that has been put forward as being able to accommodate fat-tailed standardised errors, is the exponential version of the Generalised Autoregressive Score (GAS) model, see e.g. Creal et al. (2011). However, even univariate versions of this model can be very difficult to estimate due to its nature (see the section on "Computational challenges" in Sucarrat (2013, p. 142)), and the problem is compounded even further in the multivariate case.

We propose a multivariate model of electricity price volatility that is robust to spikes, that sidesteps the curse of dimensionality through equation-by-equation estimation, and which can include both deterministic and stochastic covariates to accommodate periodicity effects, leverage, the effect of weather-related variables, and so on. The model we propose is a multivariate multiplicative component log-GARCH-X model that decomposes volatility multiplicatively into a non-stationary deterministic part of arbitrary form, and a stationary stochastic part. In order to enable equation-by-equation estimation, we make use of recent ideas developed formally in Francq and Zakoïan (2016), and in Francq and Sucarrat (2017). In particular, our model allows for feedback volatility effects among the equations, and Dynamic Conditional Correlations (DCCs) that – optionally – can be estimated subsequent to the volatility equations. As long as the DCC specification is appropriately chosen, this will ensure positive definiteness of the conditional covariance matrix. The model we propose can be viewed as a generalisation of Sucarrat et al. (2016, Section 4) in two ways. First, the deterministic component is much more general, since it can be of arbitrary form (i.e. it needs not be a linear combination of non-stochastic covariates). Second, we set up the estimation problem in such a way that the deterministic and stationary parts can be estimated separately, each by common methods that are

[1] In stock markets, by contrast, a negative shock is usually followed by an increase. Arguably, the inverse leverage effect should instead be referred to as negative asymmetry, since the effect is not due to leverage in many markets (e.g. electricity and currency markets), and because a negative parameter value is not obtained as the mathematical inverse of a positive parameter.

[2] This is the reason why Nelson proposed his model in combination with the Generalised Error Distribution (GED) rather than with the standardised Student's t, since the unconditional variance will generally not exist if the standardised error is distributed as the latter, see Nelson (1991, p. 365).

widely available. In particular, in many cases the deterministic part will be estimable by an Ordinary Least Squares (OLS) regression, and the stochastic part will be estimable via an ARMA-regression. The equation-by-equation estimation procedure that we propose is thus readily implemented in software that is widely available. We use the model to study the multivariate volatility of hourly day-ahead system prices at Nord Pool. We find extensive evidence of periodicity in the volatility in that it depends on the day-of-the-week, and in that volatility dynamics varies intradaily. We also find extensive evidence of volatility feedback from adjacent hours. Leverage (of positive type), however, is only present in about one third of the instances, and mostly between 1am and 5am. In only a single instance – at 7pm – does a plain log-GARCH(1,1) without periodicity provide a better fit of the volatility. Finally, we also find that the corrected DCC (cDCC) of Aielli (2013) provides a better fit of the conditional correlations than a Constant Conditional Correlation (CCC) specification. Interestingly, the correlations are found to be at their strongest among adjacent hours, and that the strength is inversely related to the degree of adjacency: The further away, the weaker the correlation. This has implications for risk-management, since it implies that portfolio risk is reduced if the degree of adjacency among the portfolio components is reduced.

The rest of the paper is organised as follows. The next section, Section 2, outlines the model and the equation-by-equation estimation procedure. Section 3 contains our study of hourly day-ahead price volatility at Nord Pool. Section 4 contains the conclusions, whereas tables and figures are located at the end after the references.

2 Model and estimator

2.1 The model

Let $\boldsymbol{r}_t = (r_{1t}, \ldots, r_{Mt})'$ denote an M-dimensional vector of price returns at t. A generic model of $\boldsymbol{r}_t$ can be written as (see e.g. Engle (2002))

$$\begin{aligned} \boldsymbol{r}_t &= \boldsymbol{\mu}_t + \boldsymbol{\epsilon}_t, \quad t \in \mathbb{Z}, & (1)\\ \boldsymbol{\epsilon}_t &= (\epsilon_{1t}, \ldots, \epsilon_{Mt})', \quad \boldsymbol{H}_t = E_{t-1}(\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t'), \quad \boldsymbol{D}_t^2 = \text{diag}(\boldsymbol{H}_t), & (2)\\ \boldsymbol{\eta}_t &= \boldsymbol{D}_t^{-1} \boldsymbol{\epsilon}_t, \quad \boldsymbol{R}_t = E_{t-1}(\boldsymbol{\eta}_t \boldsymbol{\eta}_t'), & (3) \end{aligned}$$

where $\boldsymbol{\mu}_t$ is the conditional mean (say, a VAR-X as in the empirical section, see Section 3.2), $\boldsymbol{\epsilon}_t = (\epsilon_{1t}, \ldots, \epsilon_{Mt})'$ is the error term, $\boldsymbol{H}_t$ is an $M \times M$ covariance matrix conditional on the past information set $\mathcal{F}_{t-1}$, $E_{t-1}(\cdot)$ is shorthand notation for $E(\cdot|\mathcal{F}_{t-1})$, $\boldsymbol{D}_t^2$ is a diagonal $M \times M$ matrix with the conditional variance or volatility $\boldsymbol{\sigma}_t^2 = (\sigma_{1t}^2, \ldots, \sigma_{Mt}^2)'$ on the diagonal, $\boldsymbol{\eta}_t = (\eta_{1t}, \ldots, \eta_{Mt})'$ is the standardised error, i.e. $E(\boldsymbol{\eta}_t) = \boldsymbol{0}$ and $E(\boldsymbol{\eta}_t^2) = \boldsymbol{1}$ where $\boldsymbol{0}$ and $\boldsymbol{1}$ are $M \times 1$ vectors, $\boldsymbol{D}_t^{-1}$ is a diagonal $M \times M$ matrix with $(1/\sigma_{1t}, \ldots, 1/\sigma_{Mt})'$ on the diagonal and $\boldsymbol{R}_t$ is the correlation matrix conditional on the past. The relationships between $\boldsymbol{H}_t$ and $\boldsymbol{R}_t$ are given by $\boldsymbol{H}_t = \boldsymbol{D}_t \boldsymbol{R}_t \boldsymbol{D}_t$ and $\boldsymbol{R}_t = \boldsymbol{D}_t^{-1} \boldsymbol{H}_t \boldsymbol{D}_t^{-1}$. The return vector $\boldsymbol{r}_t$ can be replaced with a price vector $\boldsymbol{S}_t = (S_{1t}, \ldots, S_{Mt})'$, albeit – as is well-known – any vector of prices can be obtained via a straightforward transformation of $\boldsymbol{r}_t$. For example, if $\boldsymbol{r}_t$ is log-return, then $\boldsymbol{S}_t = \exp(\boldsymbol{\mu}_t + \boldsymbol{\epsilon}_t) \odot \boldsymbol{S}_{t-1}$, where $\odot$ is the elementwise

(Hadamard) matrix product.[3] Alternatively, if $\boldsymbol{r}_t$ is relative return (this approach is preferable in markets where negative prices are possible), then $\boldsymbol{S}_t = (\boldsymbol{r}_t + \mathbf{1}) \odot \boldsymbol{S}_{t-1}$. To accommodate that our model belongs to the log-GARCH class of models, and in order to enable equation-by-equation estimation, we add the two assumptions

$$m = 1, \ldots, M : P_{t-1}(\eta_{mt} = 0) = 0, \tag{4}$$

$$\eta_{mt} \text{ is independent of } \mathcal{F}_{t-1}, \tag{5}$$

where $P_{t-1}(\cdot)$ denotes a probability conditional on the past. The first assumption is standard in log-GARCH models, but can be relaxed via the modifications suggested in Sucarrat and Escribano (2017), and in Grønneberg and Sucarrat (2017). The second assumption enables equation-by-equation estimation of $\boldsymbol{\sigma}_t^2$ in the case where the conditional correlations (i.e. the off-diagonals of $\boldsymbol{R}_t$) are dynamic and dependent on the past, i.e. they are DCCs, see Francq and Zakoïan (2016), and Francq and Sucarrat (2017). It should be noted, however, that the estimation of $\boldsymbol{R}_t$ is optional. In other words, estimation of $\boldsymbol{\sigma}_t^2$ does not require the specification nor the estimation of $\boldsymbol{R}_t$. Accordingly, estimation of $\boldsymbol{R}_t$ – if needed – proceeds subesquent to the estimation of $\boldsymbol{\sigma}_t^2$ by means of the standardised residuals (i.e. estimates of $\boldsymbol{\eta}_t$). In Section 2.3 we provide a brief survey of potential specifications of $\boldsymbol{R}_t$, and in the empirical section we estimate the DCC of Aielli (2013) and a CCC.

Periodic volatility means volatility is not covariance-stationary, since then the unconditional variance $E(\boldsymbol{\epsilon}_t^2)$ depends on t. The most common approach to non-stationary volatility is to decompose $\boldsymbol{\sigma}_t^2$ multiplicatively, see (amongst other) Van Bellegem and Von Sachs (2004), Engle and Rangel (2008), Mazur and Pipien (2012), and Amado and Terasvirta (2014a, 2014b). This means

$$\boldsymbol{\sigma}_t^2 \;\; = \;\; \boldsymbol{g}_t \odot \boldsymbol{h}_t = (g_{1t}h_{1t}, \ldots, g_{Mt}h_{Mt})', \tag{6}$$

where $\boldsymbol{g}_t$ is the non-stationary component and $\boldsymbol{h}_t$ is the stationary component (typically a GARCH-like process). In our model, the non-stationary component is specified as

$$\ln \boldsymbol{g}_t \;\; = \;\; \left(\ln g_1(\boldsymbol{\lambda}_1, \boldsymbol{x}_{1t}), \ldots, \ln g_M(\boldsymbol{\lambda}_M, \boldsymbol{x}_{Mt})\right)', \tag{7}$$

where $\ln g_1, \ldots, \ln g_M$ are known functions (linear or nonlinear), $\boldsymbol{x}_{1t}, \ldots, \boldsymbol{x}_{Mt}$ are known, non-stochastic regressors, and $\boldsymbol{\lambda}_1, \ldots, \boldsymbol{\lambda}_M$ are unknown parameters to be estimated. We do not restrict the $\boldsymbol{x}_{mt}$'s nor the functions $\ln g_m$ to be equal across equations, and the $\ln g_m$'s can assume a variety of shapes. In the simplest case the $\ln g_m$'s are linear functions made up of time dummies (calendar effects, etc.), but it can also assume the shape of an exponential spline as in Engle and Rangel (2008), the Fourier Flexible Form (FFF) as in Mazur and Pipien (2012), or smooth threshold models as in Amado and Terasvirta (2014a, 2014b). Under appropriate assumptions, the functions may also be estimated nonparametrically, as in Van Bellegem and Von Sachs (2004).

[3] For example, if $\boldsymbol{a}$ and $\boldsymbol{b}$ are two equally sized $M \times 1$ vectors, say, $\boldsymbol{a} = (a_1, \ldots, a_M)'$ and $\boldsymbol{b} = (b_1, \ldots, b_M)'$, then $\boldsymbol{a} \odot \boldsymbol{b} = (a_1b_1, \ldots, a_Mb_M)'$.

The stationary component is specified as

$$\ln \boldsymbol{h}_t = \boldsymbol{\omega} + \sum_{i=1}^{p} \boldsymbol{\alpha}_i \ln \widetilde{\boldsymbol{\epsilon}}_{t-i}^2 + \sum_{j=1}^{q} \boldsymbol{\beta}_j \ln \boldsymbol{h}_{t-j} + \boldsymbol{\delta} \boldsymbol{x}_{t-1}^s, \tag{8}$$

where $\ln \boldsymbol{h}_t = (\ln h_{1,t}, \ldots, \ln h_{M,t})'$, $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_M)'$, $\ln \widetilde{\boldsymbol{\epsilon}}_t^2 = (\ln \boldsymbol{\epsilon}_t^2 - \ln \boldsymbol{g}_t) = (\ln h_{1t}\eta_{1t}^2, \ldots, \ln h_{Mt}\eta_{Mt}^2)'$, $\boldsymbol{x}_t^s = (x_{1t}^s, x_{2t}^s, \ldots)'$ is a vector of strictly stationary stochastic – hence the superscript s – variables that satisfy $E(\boldsymbol{x}_t^s) = \mathbf{0}$,

$$\boldsymbol{\alpha}_i = \begin{pmatrix} \alpha_{11.i} & \cdots & \alpha_{1M.i} \\ \vdots & \ddots & \vdots \\ \alpha_{M1.i} & \cdots & \alpha_{MM.i} \end{pmatrix}, \quad \boldsymbol{\beta}_j = \text{diag}(\beta_{11.j}, \ldots, \beta_{MM.j})$$

and $\boldsymbol{\delta}$ is a parameter matrix of appropriate size. The non-diagonality of $\boldsymbol{\alpha}_i$ enables feedback among equations, whereas the diagonality of $\boldsymbol{\beta}_j$ enables equation-by-equation estimation. The motivation for $\boldsymbol{x}_t^s$ is that electricity price volatility may also depend on a range of stochastic factors, e.g. leverage, weather-related quantities and other market specific variables. Contrary to non-exponential GARCH models, we do not need to impose any non-negativity constraints on $\boldsymbol{\delta}$, nor on the variables in $\boldsymbol{x}_t^s$. The model is stable (in $\ln \boldsymbol{h}_t$) if all eigenvalues of $\sum_{i=1}^{p}(\boldsymbol{\alpha}_i + \boldsymbol{\beta}_i)$ are strictly less than 1 in modulus, and $\ln \boldsymbol{h}_t$ is invertible if all the eigenvalues of $\sum_{j=1}^{q} \boldsymbol{\beta}_j$ are strictly less than 1 in modulus.

2.2 Equation-by-equation estimation of $\boldsymbol{\sigma}_t^2$

In our model, given by (1)-(8), the mth. log-volatility equation can be written as

$$\begin{aligned} \ln \sigma_{mt}^2 &= \ln g_{mt} + \ln h_{mt}, && (9) \\ \ln g_{mt} &= \ln g_m(\boldsymbol{\lambda}_m, \boldsymbol{x}_{mt}), && (10) \\ \ln h_{mt} &= \omega_m + \sum_{i=1}^{p} \boldsymbol{\alpha}_{m.i} \ln \widetilde{\boldsymbol{\epsilon}}_{t-i}^2 + \sum_{j=1}^{q} \beta_{mm.j} \ln h_{m,t-j}^2 + \boldsymbol{\delta}_m \boldsymbol{x}_{t-1}^s, && (11) \end{aligned}$$

where $\boldsymbol{\alpha}_{m.i}$ is the mth. row of $\boldsymbol{\alpha}_i$, i.e. $\boldsymbol{\alpha}_{m.i} = (\alpha_{m1.i}, \ldots, \alpha_{mM.i})$, and $\boldsymbol{\delta}_m$ is the mth. row of $\boldsymbol{\delta}$. Estimation of equation m proceeds in three steps:

1. Estimate $\boldsymbol{\lambda}_m$ by means of the auxiliary regression

   $$\ln \epsilon_{mt}^2 = \lambda_{m0} + \ln g_m(\boldsymbol{\lambda}_m, \boldsymbol{x}_{mt}) + y_{mt}, \tag{12}$$

   where λ_{m0} is the intercept and y_{mt} the error-term. Below, we show that $\lambda_{m0} = E(\ln \widetilde{\epsilon}_{mt}^2)$ and that y_{mt} is a zero-mean stationary error under standard assumptions. If $\boldsymbol{\lambda}_m$ enters linearly in $\ln g_m$, then it can be estimated by OLS.

2. Fit an ARMA model to the residuals $\widehat{y}_{mt}$ from the first step. Below, we show that the error-term y_{mt} from Step 1 is in fact governed by a mean-corrected ARMA representation of $\ln h_{mt}$. Due to the relationships between the parameters of the log-GARCH model and the parameters of the mean-corrected ARMA-representation, this provides consistent estimates of all the log-GARCH parameters apart from the

intercept ω_m. As we will show, however, an estimate of ω_m is not needed in order to estimate σ^2_{mt}. Nevertheless, ω_m can – if needed – be estimated subsequently in a fourth step (see below).

3. Estimate the log-moment $E(\ln \eta^2_{mt})$ needed to complete the estimate of σ^2_{mt}. As we show below, estimation of $E(\ln \eta^2_{mt})$ is straightforward by means of a simple formula made up of the residuals from Step 2.

We now provide the details of this three step estimator.

Step 1 consists of estimating an auxiliary regression whose error-term follows the mean-corrected ARMA-representation of $\ln h_{mt}$. If $E|\ln \eta^2_{mt}| < \infty$, then the ARMA-representation of $\ln h_{mt}$ is

$$\ln \widetilde{\epsilon}^2_{mt} \quad = \quad \phi_{m0} + \sum_{i=1}^{p} \boldsymbol{\phi}_{m.i} \ln \widetilde{\boldsymbol{\epsilon}}^2_{t-i} + \sum_{j=1}^{q} \theta_{mm.j} u_{m,t-j} + \boldsymbol{\delta}_m \boldsymbol{x}^s_{t-1} + u_{m,t}, \tag{13}$$

where

$$\phi_{m0} = \omega_m + (1 - \sum_{j=1}^{q} \beta_{mm.j}) E(\ln \eta^2_{mt}), \tag{14}$$

$$\boldsymbol{\phi}_{m.i} = (\alpha_{m1.i}, \ldots, \alpha_{mm.i} + \beta_{mm.i}, \ldots, \alpha_{mM.i}), \qquad \theta_{mm.j} = -\beta_{mm.j}, \tag{15}$$

$$u_{m,t} = \ln \eta^2_{m,t} - E(\ln \eta^2_{m,t}), \qquad u_{m,t} \sim WN(0, \sigma^2_{u_m}). \tag{16}$$

In other words, $\boldsymbol{\phi}_{m.i}$ is the mth. row of $\boldsymbol{\phi}_i$. If $E|\ln \widetilde{\boldsymbol{\epsilon}}^2_t| < \boldsymbol{\infty}$, then the mean-corrected ARMA representation is

$$y_{m,t} = \sum_{i=1}^{p} \boldsymbol{\phi}_{m.i} \boldsymbol{y}_{t-i} + \sum_{j=1}^{q} \theta_{mm.j} u_{m,t-j} + \boldsymbol{\delta}_m \boldsymbol{x}^s_{t-1} + u_{m,t}, \qquad u_{m,t} \sim WN(0, \sigma^2_{u_m}), \tag{17}$$

where $y_{m,t} = \ln \widetilde{\epsilon}^2_{m,t} - E(\ln \widetilde{\epsilon}^2_{m,t})$ and $\boldsymbol{y}_t = (y_{1,t}, \ldots, y_{M,t})'$. To obtain the auxiliary regression in (12), we simply add $\ln \eta^2_{mt}$ to $\ln \sigma^2_{mt}$ in (9), which gives

$$\begin{aligned} \ln \epsilon^2_{mt} &= \ln g_{mt} + \ln h_{mt} + \ln \eta^2_{mt} \\ &= \ln g_{mt} + \ln \widetilde{\epsilon}^2_{mt} \\ &= \lambda_{m0} + \ln g_{mt} + y_{mt}, \end{aligned}$$

where $\lambda_{m0} = E(\ln \widetilde{\epsilon}^2_{mt})$ and $y_{mt} = \ln \widetilde{\epsilon}^2_{mt} - E(\ln \widetilde{\epsilon}^2_{mt})$. In other words, (12) is a standard regression model in which the error-term follows a zero-mean stationary error. In particular, if $\boldsymbol{\lambda}_m$ enters $\ln g_{mt}$ linearly, then $\lambda_{m0} + \ln g_{mt}$ can be estimated by OLS.

Step 2 consists of estimating (17) using the residuals $\widehat{y}_{mt}$ from Step 1. This is an ARMA-X estimation problem that provides estimates of all the ARCH and GARCH parameters – except ω_m – due to the relationships in (15). An estimate of ω_m, however, is not needed if the aim is to estimate σ^2_{mt}. The reason for this is that the fitted values from the first two steps provide estimates of $E(\ln \widetilde{\epsilon}^2_{mt}) + \ln g_{mt}$ and $E_{t-1}(y_{mt})$, respectively.

Adding these gives

$$\begin{aligned} E(\ln \tilde{\epsilon}^2_{mt}) + \ln g_{mt} + E_{t-1}(y_{mt}) &= \ln g_{mt} + E_{t-1}(\ln \tilde{\epsilon}^2_{mt}) \\ &= \ln g_{mt} + \ln h_t + E(\ln \eta^2_{mt}), \end{aligned} \tag{18}$$

since $\ln \tilde{\epsilon}^2_{mt} = \ln h_{mt} + \ln \eta^2_{mt}$. So only an estimate of $E(\ln \eta^2_{mt})$ is needed to complete the estimate of σ^2_{mt}.

Step 3 thus consists of estimating $E(\ln \eta^2_{mt})$. Sucarrat et al. (2016) noted that, if $E|\ln \eta^2_{mt}| < \infty$ and $E(\eta^2_{mt}) = 1$, then it follows straightforwardly that $-\ln E(\exp(u_{mt})) = E(\ln \eta^2_{mt})$. This suggests

$$-\ln \left[T^{-1} \sum_{t=1}^{T} \exp(\widehat{u}_{mt}) \right] \tag{19}$$

is a consistent estimator of $E(\ln \eta^2_{mt})$, where the $\widehat{u}_{mt}$'s are the residuals from Step 2.[4] Sucarrat et al. (2016) provide conditions under which this indeed holds, whereas Francq and Sucarrat (2017) prove that this holds when the ARMA-X representation of equation m in a first order multivariate log-GARCH-X model – where the X-part refers to stochastic conditioning variables – is estimated by Least Squares.

Summarised, then, the estimate of σ^2_{mt} is given by

$$\widehat{\sigma}^2_{mt} = \exp\left(\widehat{E}(\ln \tilde{\epsilon}^2_{mt}) + \ln \widehat{g}_{mt} + \widehat{E}_{t-1}(y_{mt}) - \widehat{E}(\ln \eta^2_{mt}) \right), \tag{20}$$

where $\widehat{E}(\ln \tilde{\epsilon}^2_{mt}) + \ln \widehat{g}_{mt}$ is the fitted value of the auxiliary regression in Step 1, $\widehat{E}_{t-1}(y_{mt})$ is the fitted value of the mean-corrected ARMA representation in Step 2 and $\widehat{E}(\ln \eta^2_{mt})$ is the estimate of $E(\ln \eta^2_{mt})$ in Step 3.

An estimate of ω_m requires estimation of the other equations, in addition to equation m. This is because the expression for $E(\ln \tilde{\epsilon}^2_{mt})$, which can be written as $E(\ln \tilde{\epsilon}^2_{mt}) = \phi_{m0} + \sum_{i=1}^{p} \boldsymbol{\phi}_{m.i} E(\ln \tilde{\boldsymbol{\epsilon}}^2_t)$, depends on the unconditional expectations of the other equations. Solving for ω_m in the expression for $E(\ln \tilde{\epsilon}^2_{mt})$ gives

$$\omega_m = (1 - \sum_{j=1}^{q} \beta_{mm.j}) E(\ln \tilde{\epsilon}^2_{mt}) - \sum_{i=1}^{p} \boldsymbol{\alpha}_{m.i} E(\ln \tilde{\boldsymbol{\epsilon}}^2_t) - (1 - \sum_{j=1}^{q} \beta_{mm.j}) E(\ln \eta^2_{mt}), \tag{21}$$

where we have used the expression for ϕ_{m0} in (14), and that $\sum_{i=1}^{p} \boldsymbol{\phi}_{m.i} E(\ln \tilde{\boldsymbol{\epsilon}}^2_t) = \sum_{i=1}^{p} \boldsymbol{\alpha}_{m.i} E(\ln \tilde{\boldsymbol{\epsilon}}^2_t) + \sum_{j=1}^{q} \beta_{mm.j} E(\ln \tilde{\epsilon}^2_{mt})$. It should be noted that only the elements in $E(\ln \tilde{\boldsymbol{\epsilon}}^2_t)$, apart from the mth. entry, comes from the other equations. In other words, if there is no feedback effects (i.e. all entries in the $\boldsymbol{\alpha}_{m.i}$'s apart from the mth. entry are zero), then there is no need to estimate the other equations in order to estimate ω_m.

For inference an estimate of the covariance matrix of the coefficients is needed. To this end an estimate of the covariance matrix of the mean-corrected ARMA specification can be used, since the log-GARCH parameters – apart from the log-volatility intercept ω_m – is a linear combination of the ARMA-parameters, see (15). For inference on ω_m the delta method can be used. Finally, for inference involving $\boldsymbol{\lambda}_m$ a "sandwich" approach can be used, where the autocovariance structure of y_{mt} provides the "meat" of the sandwich.

[4] The expression in square brackets in (19) is the smearing estimate of Duan (1983).

2.3 Dynamic Conditional Correlations (DCCs)

Assumption (5) enables DCCs. The three-step estimation procedure described above, however, does not provide estimates of the DCCs. Nevertheless, they can – if needed – be estimated in a subsequent step. The estimates $\widehat{\sigma}_{1t}^2, \ldots, \widehat{\sigma}_{Mt}^2$ from the three-step procedure above lead to the standardised residuals $\widehat{\boldsymbol{\eta}}_t = (\widehat{\eta}_{1t}, \ldots, \widehat{\eta}_{Mt})'$, where $\widehat{\eta}_{mt} = \epsilon_{mt}/\widehat{\sigma}_{mt}$. These residuals can be used to estimate either a CCC or DCC specification of $\boldsymbol{R}_t = E(\boldsymbol{\eta}_t\boldsymbol{\eta}_t'|\mathcal{F}_{t-1})$. An example is the DCC of Engle (2002), or alternatively the corrected version of Aielli (2013), see e.g. the empirical section of Francq and Sucarrat (2017). Another option is the robust (to spikes) DCC model proposed for electricity prices by Dupuis (2017). In the empirical section (see Section 3.3) we estimate the DCC of Aielli (2013) and a CCC.

3 The volatility of hourly day-ahead system prices at Nord Pool

3.1 Data

Nord Pool Spot AS is one of the largest energy exchanges in the world measured in traded volume of terrawatts per hour (TWh). Currently 380 members operate on the exchange, and these include public and private energy producers, energy intensive industries, large consumers, distributors, funds, investment companies, banks, brokers, utility companies and financial institutions. Arguably, the most important price at the exchange is the "system price". This is because it constitutes some sort of aggregate or equilibrium price (it is determined by the intersection of the aggregate supply and demand curves of all bids and offers), and because it is used as reference in financial contracts – used for hedging and risk management – traded at Nasdaq Commodities.[5]

Our rawdata consist of the hourly day-ahead system prices in Euros per kw/h from 1 January 2010 to 20 May 2014. This amounts to $T = 1601$ daily observations for each m before differencing and lagging. The price at day t in hour m we denote S_{mt}, where $m = 1, \ldots, 24$. Note that S_{1t} should be interpreted as the price from midnight to 1am in day t, S_{2t} is the price from 1am to 2am in day t, and so on. The daily log-return for hour m, denoted r_{mt}, is defined as $\ln S_{m,t} - \ln S_{m,t-1}$, i.e. the daily log-return for hour m. Graphs of S_{mt} and r_{mt} are contained in Figures 1 and 2, respectively, whereas the top graph in Figure 3 contains the average hourly price. The prices and returns exhibit the usual characteristics, i.e. that return variability is substantially larger than those of stocks, stock indices and exchange rates, and that big spikes or jumps occur relatively frequently. On average, the price is highest at 9am and lowest at 4am. There are no negative prices in our data, but five spurious zeros due to daylight saving time.[6] These zeros we replace by the average of the two adjacent values.

[5] See https://www.nordpoolspot.com/About-us/, http://www.nordpoolspot.com/How-does-it-work/Financial-market/ and http://www.nordpoolspot.com/TAS/Day-ahead-market-Elspot/Price-calculation/. All accessed 14 July 2016.

[6] The five zeros all occurred for $m = 3$, one in each year: 28 March 2010, 27 March 2011, 25 March 2012, 31 March 2013 and 30 March 2014.

3.2 Models of the mean and volatility

We start by fitting a conditional mean specification to the vector $\boldsymbol{r}_t$ of daily returns. A restricted Vector Autoregressive (VAR) model is formulated and estimated, where each equation contains its own AR-lags from 1 to 7 and six day-of-the-week dummies (Tuesday to Sunday). The total number of estimated parameters in each conditional mean equation is thus fourteen: One intercept + seven AR-parameters + six dummy-parameters. We do not explicitly model the spikes in the mean, since our volatility model is robust to them due to the log-transform. However, if a user wishes do to so, our model is – of course – compatible with specifications that model them explicitly. The second column of Table 1 and the second graph in Figure 3 contain the R-squareds of the twenty-four conditional mean equations. As is clear, predictability varies substantially across the day, since the R-squareds range from only 5% when $m = 1$ to a peak of 50% when $m = 9$. As a whole, the graph clearly indicates that the explanatory power is higher in peak hours, i.e. from about $m = 7$ to about $m = 19$.

We fit five different multivariate volatility models to the vector of errors $\boldsymbol{\epsilon}_t$. The models we label (a) – (e), and in each of the five models equation m is given by

$$
\begin{aligned}
\text{(a)}\quad \ln\sigma_{mt}^2 &= \omega_m + \alpha_{mm.1}\ln\widetilde{\epsilon}_{m,t-1}^2 + \beta_{mm.1}\ln h_{m,t-1},\\
\text{(b)}\quad \ln\sigma_{mt}^2 &= \sum_{i=1}^{6}\lambda_{mi}d_{it},\\
\text{(c)}\quad \ln\sigma_{mt}^2 &= \sum_{i=1}^{6}\lambda_{mi}d_{it} + \omega_m + \alpha_{mm.1}\ln\widetilde{\epsilon}_{m,t-1}^2 + \beta_{mm.1}\ln h_{m,t-1},\\
\text{(d)}\quad \ln\sigma_{mt}^2 &= \sum_{i=1}^{6}\lambda_{mi}d_{it} + \omega_m + \alpha_{mm.1}\ln\widetilde{\epsilon}_{m,t-1}^2 + \alpha_{mm(1).1}\ln\widetilde{\epsilon}_{m(1),t-1}^2\\
&\quad + \alpha_{mm(2).1}\ln\widetilde{\epsilon}_{m(2),t-1}^2 + \beta_{mm.1}\ln h_{m,t-1},\\
\text{(e)}\quad \ln\sigma_{mt}^2 &= \sum_{i=1}^{6}\lambda_{mi}d_{it} + \omega_m + \alpha_{mm.1}\ln\widetilde{\epsilon}_{m,t-1}^2 + \alpha_{mm(1).1}\ln\widetilde{\epsilon}_{m(1),t-1}^2\\
&\quad + \alpha_{mm(2).1}\ln\widetilde{\epsilon}_{m(2),t-1}^2 + \beta_{mm.1}\ln h_{m,t-1} + \delta_m x_{m,t-1}.
\end{aligned}
$$

Model (a) is a plain log-GARCH(1,1) and serves as benchmark. The variables $d_{1t},\ldots,d_{6t}$ are dummies for Tuesday to Sunday, respectively. So (b)–(e) all contain periodicity. In (e) and (d), $\ln\widetilde{\epsilon}_{m(1),t-1}^2$ and $\ln\widetilde{\epsilon}_{m(2),t-1}^2$ are the two most adjacent (within the same day) log-ARCH lags in equation m. For example, if $m = 1$, then $m(1) = 2$ and $m(2) = 3$. Similarly, if $m = 2$, then $m(1) = m - 1 = 1$ and $m(2) = m + 1 = 3$. And so on. The idea is to include those log-ARCH terms that are most likely to have a feedback effect on the volatility of equation m.[7] Finally, the variable $x_{m,t-1}$ is a lagged mean-corrected asymmetry or "leverage" term, where the lagged asymmetry term is given by $I_{\{\epsilon_{m,t-1}<0\}}$ before mean-correction.

[7]For simplicity we only include lags from the same day. However, as pointed out by one of the reviewers, for $m = 1$, it is not obvious that $m(2) = 2$ for $t - 1$ is more adjacent that $m(2) = 24$ in $t - 2$. Indeed, additional estimation results (not reported) suggests the $m = 1$ equation can be improved further by including the ARCH-lag for $m = 24$ in $t - 2$. A detailed study of the optimal adjacency structure we leave for future research.

The number of estimated parameters (denoted k_m) in equation $\ln\sigma_m^2$ is 3, 7, 9, 11 and 12, respectively, for the five specifications (a) – (e). The total number of estimated parameters in each of the five multivariate volatility specifications of $\boldsymbol{\sigma}_t^2$ is thus $3\times24 = 72$ for (a), $7\times24 = 168$ for (b), $9\times24 = 216$ for (c), $11\times24 = 264$ for (d) and $12\times24 = 288$ for (e). Estimation of the five multivariate models all together takes about thirty seconds on an average laptop, and we experience no numerical issues.[8] Both the deterministic and stationary parts are estimated by Least Squares (LS), and a summary of the estimation results are contained in Table 1. For comparison we use the Schwarz (1978) information criterion (BIC), which favours parsimony. The best model in hour m according to the BIC is identified with an asterisk (*) to the right of its BIC-value. In all but two cases the best model is either (d) or (e). In other words, in all but two cases the best model contains periodicity and feedback terms. The exceptions occur for $m = 16$ and $m = 19$, in which models (c) and (a), respectively, are the best according to the BIC. If we only compare (a), (b) and (c) against each other to obtain a more detailed idea of the importance of periodicity, then we see that either (b) or (c) performs better in 17 out of the 24 hours. Moreover, the periods in which (a) performs better are clustered in the evening, since they are $m = 17, 18, 19, 20, 21, 22$ and 24. Finally, with respect to leverage, equation (e) performs better than (d) in 8 instances, whereas the opposite is the case in 16 instances. In other words, there is evidence of leverage in about one third of the hours. Interestingly, whenever present the leverage is always positive except once ($m = 12$) – i.e. we find little evidence of the socalled inverse (i.e. negative) leverage effect. Summarised, then, our results provide extensive evidence of periodicity and feedback effects in the volatility, and in about one third of the hours there is leverage. In only a single instance ($m = 16$) did the plain log-GARCH(1,1) perform better than the other specifications.

The third, fourth and fifth graphs in Figure 3 provide a more detailed picture of the best specifications in each m. The third graph contains the ARCH(1) estimates, i.e. $\widehat{\alpha}_{mm.1}$, for $m = 1, \ldots, 24$. All estimates are positive and lie between 0.06 and 0.12. In other words, once periodicity and feedback effects are controlled for, then the (own) ARCH effect becomes substantially smaller than commonly found in electricity prices, and much closer to those usually found in stock and currency markets. The fourth graph in the figure contains the GARCH(1) parameter estimates, i.e. $\widehat{\beta}_{mm.1}$. The values are in the 0.8 to 0.95 range, and interestingly the values about 0.9 are mostly clustered around the morning hours, i.e. $m = 1, \ldots, 8$. Finally, the bottom graph contains the leverage estimates, i.e. $\widehat{\delta}_m$. The estimated value is zero if the best model in hour m does not contain leverage, and it is notable that half of the non-zero values occur at night, from $m = 2$ to $m = 5$.

3.3 Dynamic Conditional Correlations (DCCs)

To obtain estimates of the off-diagonals of $\boldsymbol{H}_t$, an estimate of $\boldsymbol{R}_t = E_{t-1}(\boldsymbol{\eta}_t'\boldsymbol{\eta}_t)$ is needed. To this end we fit the corrected DCC (cDCC) model of Aielli (2013), which is a modified version of Engle's (2002) DCC. It should be noted that the cDCC is a covariance-stationary model of $\boldsymbol{\eta}_t'\boldsymbol{\eta}_t$ even though $\boldsymbol{\epsilon}_t'\boldsymbol{\epsilon}_t$ is not covariance stationary.

[8]The computations were undertaken with R code (R Core Team (2014)) on a Lenovo X250 with an Intel Core i7-5600U-2.60 Ghz processor running Win7-64bit.

The cDCC model is given by

$$\boldsymbol{R}_t = \boldsymbol{Q}_t^{*-1/2}\boldsymbol{Q}_t\boldsymbol{Q}_t^{*-1/2}, \quad \boldsymbol{Q}_t = (1-\gamma_1-\gamma_2)\boldsymbol{R} + \gamma_1\boldsymbol{Q}_{t-1}^{*1/2}\boldsymbol{\eta}_{t-1}\boldsymbol{\eta}_{t-1}'\boldsymbol{Q}_{t-1}^{*1/2} + \gamma_2\boldsymbol{Q}_{t-1}, \quad (22)$$

where $\gamma_1, \gamma_2 \geq 0$ are scalars such that $\gamma_1 + \gamma_2 < 1$, $\boldsymbol{R}$ is a correlation matrix (which in general is *not* equal to $E(\boldsymbol{R}_t)$), $\boldsymbol{Q}_t^*$ is a diagonal matrix with elements from the diagonal of $\boldsymbol{Q}_t$ and $\boldsymbol{\eta}_t = \boldsymbol{D}_t^{-1}\boldsymbol{\epsilon}_t$. So estimation of the cDCC entails estimation of only two parameters, γ_1 and γ_2. The standardised error $\boldsymbol{\eta}_t$ is made up of the errors from the best model in each m according to the BIC (see above). In other words, if $m = 1$, then $\widehat{\eta}_{1t}$ is that of model (e). If $m = 8$, then $\widehat{\eta}_{8t}$ is that of model (d). And so on.

Recalling that $\boldsymbol{R}_t$ is both the conditional correlation and covariance matrix of $\boldsymbol{\eta}_t$, estimation of γ_1 and γ_2 by Gaussian Quasi Maximum Likelihood (QML) leads to the estimator

$$(\widehat{\gamma}_1, \widehat{\gamma}_2) = \arg\max_{(\widehat{\gamma}_1,\widehat{\gamma}_2)} \sum_{t=1}^{T}\left(-M\ln 2\pi - \ln|\widehat{\boldsymbol{R}}_t| - \widehat{\boldsymbol{\eta}}_t'\widehat{\boldsymbol{R}}_t^{-1}\widehat{\boldsymbol{\eta}}_t\right)/2, \quad (23)$$

where $|\boldsymbol{R}_t|$ is the determinant of $\boldsymbol{R}_t$, $\widehat{\boldsymbol{\eta}}_t$ are the standardised residuals of the best models,

$$\widehat{\boldsymbol{R}}_t = \widehat{\boldsymbol{Q}}_t^{*-1/2}\widehat{\boldsymbol{Q}}_t\widehat{\boldsymbol{Q}}_t^{*-1/2}, \quad \widehat{\boldsymbol{Q}}_t = (1-\widehat{\gamma}_1-\widehat{\gamma}_2)\widehat{\boldsymbol{R}} + \widehat{\gamma}_1\widehat{\boldsymbol{Q}}_{t-1}^{*1/2}\widehat{\boldsymbol{\eta}}_{t-1}\widehat{\boldsymbol{\eta}}_{t-1}'\widehat{\boldsymbol{Q}}_{t-1}^{*1/2} + \widehat{\gamma}_2\widehat{\boldsymbol{Q}}_{t-1}$$

$$\widehat{\boldsymbol{R}} = \frac{1}{T}\sum_{t=1}^{T}\widehat{\boldsymbol{Q}}_t^{*1/2}\widehat{\boldsymbol{\eta}}_t\widehat{\boldsymbol{\eta}}_t'\widehat{\boldsymbol{Q}}_t^{*1/2}, \qquad \widehat{\boldsymbol{Q}}_t^* = \text{diag}(\widehat{q}_{11t}, \ldots, \widehat{q}_{MMt})$$

$$\widehat{q}_{mmt} = (1-\widehat{\gamma}_1-\widehat{\gamma}_2) + \widehat{\gamma}_1\widehat{\eta}_{m,t-1}^2 + \widehat{\gamma}_2\widehat{q}_{mm,t-1} \quad \text{for } m = 1, \ldots, M.$$

The estimates of γ_1 and γ_2 are 0.005 and 0.868, respectively, which suggests the dynamic correlations are very persistent, indeed almost constant. An estimate of the unconditional correlation $E(\boldsymbol{R}_t)$ is obtained as $\widehat{\boldsymbol{R}}^{*-1/2}\widehat{\boldsymbol{R}}\widehat{\boldsymbol{R}}^{*-1/2}$, where $\widehat{\boldsymbol{R}}^*$ is a diagonal matrix containing the diagonal elements of $\widehat{\boldsymbol{R}}$. Figure 4 depicts the evolution of the unconditional correlations over the day depending on m. The general tendency is clear: The strongest unconditional correlations of $\widehat{E}(\eta_{it}\eta_{mt})$ are always those closest to hour i. The top left graph, for example, depicts $\widehat{E}(\eta_{1t}\eta_{mt})$ for $m = 1, \ldots 24$. Naturally, when $m = 1$, then $\widehat{E}(\eta_{1t}\eta_{1t}) = 1$. Next, the unconditional correlations fall gradually until they reach their lowest point at $m = 23$, in which the estimate is about 0.2. More generally, whenever $i < m$, then almost without exception $\widehat{E}(\eta_{it}\eta_{m,t})$ is stronger than $\widehat{E}(\eta_{it}\eta_{m+1,t})$, which is stronger than $\widehat{E}(\eta_{it}\eta_{m+2,t})$, and so on. Similarly, whenever $i > 1$, $\widehat{E}(\eta_{it}\eta_{m-1,t})$ is stronger than $\widehat{E}(\eta_{it}\eta_{m-2,t})$, which is stronger than $\widehat{E}(\eta_{it}\eta_{m-3,t})$, and so on. So just as in the case of volatility feedback (see above), there is a clear adjacency effect among correlations.

Estimation of the cDCC entails fitting a total of $24 \cdot (24-1)/2 = 276$ distinct (off-diagonal) conditional correlation paths. Figure 5 contains graphs of the first 23 together with its own correlation, i.e. of $\widehat{E}_{t-1}(\eta_{1t}\eta_{mt})$ for $m = 1, \ldots, 24$. Graphically, the subset of paths exhibit relatively little variation around their unconditional values, and the near-zero estimate of γ_1 also suggests this is the case. So one may ask whether the cDCC provides a better fit than a Constant Conditional Correlation (CCC) specification. The Gaussian log-likelihood of the cDCC is given by the formula in (23) at the estimated values $\widehat{\gamma}_1$, $\widehat{\gamma}_2$, whereas the log-likelihood of the CCC specification is obtained by replacing $\widehat{\boldsymbol{R}}_t$

with the sample correlation matrix of $\{\widehat{\boldsymbol{\eta}}_t\}$ in the same formula (we use the function *cor* in R to compute the sample correlation matrix). The two values are -24863.82 (cDCC) and -25064.15 (CCC), so the former produces a better fit in terms of the (quasi) log-likelihood. In terms of the BIC, computed in terms of the average (quasi) log-likelihood with $T = 1592$, and defining the cDCC to be characterised by 2 parameters and the CCC by 0, the values are 31.25 (cDCC) and 31.49 (CCC), respectively. In other words, the DCC specification is also warranted according to the BIC.

4 Conclusions

We propose a multivariate model of electricity price volatility that decomposes volatility multiplicatively into a non-stationary part (e.g. periodic) of arbitrary form, and a stationary part with log-GARCH dynamics. The model is robust to spikes or jumps, a common characteristic of electricity prices, the model allows for a rich variety of volatility dynamics without restrictive positivity constraints on the parameters, it can be estimated equation-by-equation by means of standard methods in widely available software, and Dynamic Conditional Correlations (DCCs) can – optionally – be estimated subsequent to the volatilities. In a study of the hourly day-ahead system prices at Nord Pool, we find extensive evidence of periodic volatility and volatility feedback, and that about one third of the hours exhibit (positive) leverage. The strength of the ARCH, GARCH and leverage effects depend on the hour of the day. In only a single instance (at 7pm) does the plain log-GARCH(1,1) perform better than the other specifications. We also find that time-varying conditional correlations provide a better fit than constant correlations, and that the correlations are at their strongest in adjacent hours. This may have implications for risk-management, since it implies that portfolio risk is reduced if the degree of adjacency among the components of a portfolio is reduced.

References

Aielli, G. P. (2013). Dynamic conditional correlations: On properties and estimation. *Journal of Business and Economic Statistics 31*, 282–299. http://dx.doi.org/10.1080/07350015.2013.771027.

Amado, C. and T. Terasvirta (2014a). Modelling changes in the unconditional variance of long stock return series. *Journal of Empirical Finance 25*, 15–35.

Amado, C. and T. Terasvirta (2014b). Modelling volatility by variance decomposition. *Journal of Econometrics 175*, 142–153.

Bauwens, L., C. Hafner, and D. Pierret (2013). Multivariate volatility modelling of electricity futures. *Journal of Applied Econometrics 28*, 743–761.

Bunn, D. W. (2000). Forecasting loads and prices in competitive power markets. *Proceedings of the IEEE 88*, 163–169.

Carnero, M. A., D. Pena, and E. Ruiz (2007). Effects of outliers on the identification and estimation of GARCH models. *Journal of Time Series Analysis 28*, 471–497.

Creal, D., S. J. Koopmans, and A. Lucas (2011). A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations. *Journal of Business and Economic Statistics 29*, 552–563.

Duan, N. (1983). Smearing estimate: A nonparametric retransformation method. *Journal of the Americal Statistical Association 78*, 605–610.

Dupuis, D. J. (2017). Electricity price dependence in new york state zones: A robust detrended correlation approach. *Annals of Applied Statistics 11*, 248–273.

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics 20*, 339–350.

Engle, R. F. and J. G. Rangel (2008). The spline GARCH model for low frequency volatility and its global macroeconomic causes. *Review of Financial Studies 21*, 1187–1222.

Escribano, Á., J. I. Peña, and P. Villaplana (2002). Modelling electricity prices: International evidence. UC3M Working Paper 02-27 in the Economic Series. Available as http://docubib.uc3m.es/WORKINGPAPERS/WE/we022708.pdf.

Escribano, A., J. I. Peña, and P. Villaplana (2011). Modelling electricity prices: International evidence. *Oxford Bulletin of Economics and Statistics 73*, 622–650. working paper version (2002): UC3M Working Paper 02-27 in the Economic Series.

Francq, C. and G. Sucarrat (2017). An equation-by-equation estimator of a multivariate Log-GARCH-X model of financial returns. *Journal of Multivariate Analysis 153*, 16–32.

Francq, C. and J.-M. Zakoïan (2016). Estimating multivariate volatility models equation by equation. *The Journal of the Royal Statistical Society. Series B 78*, 613–635. Working paper version: MPRA Paper No. 54250: http://mpra.ub.uni-muenchen.de/54250/.

Gregory, A. and J. Reeves (2010). Estimation and inference in ARCH models in the presence of outliers. *Journal of Financial Econometrics 8*, 547–569.

Grønneberg, S. and G. Sucarrat (2017). Risk estimation when the zero probability of financial return is time-varying. MPRA Paper No. 81882. Online at https://mpra.ub.uni-muenchen.de/81882/.

Harvey, A. C. and G. Sucarrat (2014). EGARCH models with fat tails, skewness and leverage. *Computational Statistics and Data Analysis 76*, 320–338.

Hellström, J., J. Lundgren, and H. Yu (2012). Why do electricity prices jump? Empirical evidence from the nordic electricity market. *Energy Economics 34*, 1774–1781.

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

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

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

Mazur, B. and M. Pipien (2012). On the empirical importance of periodicity in the volatility of financial returns – time varying GARCH as a second order APC(2) process. *Central European Journal of Economic Modelling and Econometrics 4*, 95–116.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica 59*, 347–370.

R Core Team (2014). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics 6*, 461–464.

Sucarrat, G. (2013). betategarch: Simulation, estimation and forecasting of Beta-Skew-t-EGARCH models. *The R Journal 5*, 137–147. http://journal.r-project.org/archive/2013-2/sucarrat.pdf.

Sucarrat, G. and Á. Escribano (2017). Estimation of log-GARCH models in the presence of zero returns. *European Journal of Finance*. http://dx.doi.org/10.1080/1351847X.2017.1336452.

Sucarrat, G., S. Grønneberg, and Á. Escribano (2016). Estimation and inference in univariate and multivariate log-GARCH-X models when the conditional density is unknown. *Computational Statistics and Data Analysis 100*, 582–594.

Van Bellegem, S. and R. Von Sachs (2004). Forecasting economic time-series with unconditional time-varying variance. *International Journal of Forecasting 20*, 611–627.

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.

Table 1: Estimation results of multivariate models (a)–(e) for each equation $m = 1, \ldots, 24$, see Section 3. Only selected parameter estimates reported

m	R^2_m	*Spec.*	$\widehat{\alpha}_{mm.1}$	$\widehat{\beta}_{mm.1}$	$\widehat{\delta}_m$	$LogL_m$	k_m	BIC_m	T
1	0.05	(a)	0.145	0.785		1696.179	3	-2.1157	1593
		(b)				1223.525	7	-1.5037	1593
		(c)	0.150	0.779		1761.886	9	-2.1704	1593
		(d)	0.073	0.905		1965.803	11	-2.4187*	1592
		(e)	0.095	0.870	0.177	1965.393	12	-2.4135	1592
2	0.08	(a)	0.142	0.800		1366.772	3	-1.7021	1593
		(b)				703.540	7	-0.8509	1593
		(c)	0.151	0.787		1443.756	9	-1.7710	1593
		(d)	0.045	0.945		1539.855	11	-1.8835	1592
		(e)	0.076	0.908	0.251	1599.195	12	-1.9535*	1592
3	0.11	(a)	0.148	0.803		972.495	3	-1.2071	1593
		(b)				377.325	7	-0.4413	1593
		(c)	0.166	0.780		1169.506	9	-1.4266	1593
		(d)	0.062	0.924		1282.371	11	-1.5601	1592
		(e)	0.080	0.903	0.155	1312.374	12	-1.5931*	1592
4	0.13	(a)	0.145	0.808		803.453	3	-0.9948	1593
		(b)				184.655	7	-0.1994	1593
		(c)	0.168	0.778		1016.636	9	-1.2347	1593
		(d)	0.067	0.920		1099.937	11	-1.3309	1592
		(e)	0.078	0.912	0.186	1140.255	12	-1.3769*	1592
5	0.18	(a)	0.192	0.731		591.917	3	-0.7293	1593
		(b)				214.012	7	-0.2363	1593
		(c)	0.225	0.690		1068.156	9	-1.2994	1593
		(d)	0.087	0.895		1112.289	11	-1.3464	1592
		(e)	0.090	0.893	0.064	1124.681	12	-1.3573*	1592
6	0.34	(a)	0.408	0.023		448.025	3	-0.5486	1593
		(b)				221.996	7	-0.2463	1593
		(c)	0.210	0.695		953.960	9	-1.1560	1593
		(d)	0.074	0.909		1143.808	11	-1.3860*	1592
		(e)	0.071	0.912	-0.036	1138.104	12	-1.3742	1592
7	0.46	(a)	0.342	0.026		536.814	3	-0.6601	1593
		(b)				569.559	7	-0.6827	1593
		(c)	0.206	0.654		1078.526	9	-1.3124	1593
		(d)	0.089	0.879		1180.251	11	-1.4318*	1592
		(e)	0.085	0.882	-0.065	1165.374	12	-1.4085	1592
8	0.49	(a)	0.230	0.247		473.888	3	-0.5811	1593
		(b)				582.832	7	-0.6993	1593
		(c)	0.146	0.705		800.839	9	-0.9638	1593
		(d)	0.072	0.887		890.887	11	-1.0683*	1592
		(e)	0.074	0.880	-0.100	889.508	12	-1.0619	1592
9	0.50	(a)	0.112	0.795		702.853	3	-0.8685	1593

Table continues on next page. Explanatory note at the end of table.

m	R^2_m	$Spec.$	$\widehat{\alpha}_{mm.1}$	$\widehat{\beta}_{mm.1}$	$\widehat{\delta}_m$	$LogL_m$	k_m	BIC_m	T
		(b)				636.460	7	-0.7667	1593
		(c)	0.120	0.784		841.824	9	-1.0152	1593
		(d)	0.105	0.822		859.360	11	-1.0287*	1592
		(e)	0.106	0.820	-0.058	852.266	12	-1.0151	1592
10	0.48	(a)	0.106	0.820		1169.681	3	-1.4546	1593
		(b)				958.053	7	-1.1704	1593
		(c)	0.111	0.813		1218.310	9	-1.4879	1593
		(d)	0.105	0.825		1264.594	11	-1.5377*	1592
		(e)	0.105	0.825	-0.032	1262.402	12	-1.5304	1592
11	0.44	(a)	0.087	0.848		1436.812	3	-1.7900	1593
		(b)				1230.942	7	-1.5130	1593
		(c)	0.092	0.840		1478.363	9	-1.8144	1593
		(d)	0.086	0.857		1517.217	11	-1.8551*	1592
		(e)	0.083	0.862	-0.066	1518.019	12	-1.8515	1592
12	0.40	(a)	0.096	0.837		1652.884	3	-2.0613	1593
		(b)				1348.482	7	-1.6606	1593
		(c)	0.103	0.827		1675.852	9	-2.0624	1593
		(d)	0.080	0.880		1737.819	11	-2.1322	1592
		(e)	0.074	0.890	-0.139	1745.956	12	-2.1378*	1592
13	0.39	(a)	0.108	0.817		1750.552	3	-2.1839	1593
		(b)				1463.161	7	-1.8046	1593
		(c)	0.115	0.805		1786.314	9	-2.2010	1593
		(d)	0.086	0.870		1836.079	11	-2.2557*	1592
		(e)	0.084	0.875	-0.060	1836.131	12	-2.2511	1592
14	0.42	(a)	0.116	0.801		1736.312	3	-2.1660	1593
		(b)				1497.177	7	-1.8473	1593
		(c)	0.127	0.786		1794.275	9	-2.2110	1593
		(d)	0.108	0.829		1851.949	11	-2.2756	1592
		(e)	0.109	0.827	0.067	1856.392	12	-2.2766*	1592
15	0.41	(a)	0.104	0.800		1647.765	3	-2.0549	1593
		(b)				1448.164	7	-1.7858	1593
		(c)	0.114	0.787		1698.284	9	-2.0905	1593
		(d)	0.105	0.811		1731.209	11	-2.1239*	1592
		(e)	0.106	0.808	0.049	1734.441	12	-2.1234	1592
16	0.40	(a)	0.103	0.819		1655.581	3	-2.0647	1593
		(b)				1428.192	7	-1.7607	1593
		(c)	0.113	0.804		1734.833	9	-2.1364*	1593
		(d)	0.106	0.823		1735.811	11	-2.1297	1592
		(e)	0.106	0.822	-0.018	1733.501	12	-2.1222	1592
17	0.35	(a)	0.119	0.792		1574.735	3	-1.9632	1593
		(b)				1326.363	7	-1.6328	1593
		(c)	0.127	0.782		1583.191	9	-1.9460	1593
		(d)	0.098	0.851		1690.942	11	-2.0734	1592
		(e)	0.098	0.852	0.047	1699.919	12	-2.0800*	1592
18	0.30	(a)	0.118	0.833		1649.697	3	-2.0573	1593

Table continues on next page. Explanatory note at the end of table.

m	R^2_m	*Spec.*	$\widehat{\alpha}_{mm.1}$	$\widehat{\beta}_{mm.1}$	$\widehat{\delta}_m$	$LogL_m$	k_m	BIC_m	T
		(b)				1198.336	7	-1.4721	1593
		(c)	0.123	0.827		1663.243	9	-2.0465	1593
		(d)	0.110	0.847		1718.433	11	-2.1079*	1592
		(e)	0.110	0.847	-0.023	1716.875	12	-2.1013	1592
19	0.23	(a)	0.114	0.830		1784.389	3	-2.2264*	1593
		(b)				1320.064	7	-1.6249	1593
		(c)	0.118	0.825		1783.915	9	-2.1980	1593
		(d)	0.118	0.823		1807.227	11	-2.2194	1592
		(e)	0.118	0.823	-0.009	1807.172	12	-2.2147	1592
20	0.20	(a)	0.130	0.778		2051.643	3	-2.5619	1593
		(b)				1558.624	7	-1.9244	1593
		(c)	0.135	0.770		2057.749	9	-2.5418	1593
		(d)	0.119	0.811		2094.393	11	-2.5802*	1592
		(e)	0.117	0.815	-0.058	2093.004	12	-2.5738	1592
21	0.16	(a)	0.109	0.807		2260.769	3	-2.8245	1593
		(b)				1815.880	7	-2.2474	1593
		(c)	0.113	0.801		2269.980	9	-2.8083	1593
		(d)	0.075	0.886		2331.225	11	-2.8777*	1592
		(e)	0.075	0.887	-0.006	2331.454	12	-2.8734	1592
22	0.13	(a)	0.132	0.781		2441.839	3	-3.0518	1593
		(b)				2022.447	7	-2.5068	1593
		(c)	0.136	0.775		2446.624	9	-3.0301	1593
		(d)	0.104	0.846		2478.759	11	-3.0631*	1592
		(e)	0.103	0.848	-0.034	2476.671	12	-3.0558	1592
23	0.10	(a)	0.137	0.761		2448.579	3	-3.0603	1593
		(b)				2014.287	7	-2.4965	1593
		(c)	0.141	0.755		2478.169	9	-3.0697	1593
		(d)	0.103	0.840		2539.517	11	-3.1394*	1592
		(e)	0.096	0.854	-0.096	2539.433	12	-3.1347	1592
24	0.10	(a)	0.131	0.818		2048.593	3	-2.5581	1593
		(b)				1356.737	7	-1.6710	1593
		(c)	0.133	0.815		2001.696	9	-2.4715	1593
		(d)	0.099	0.869		2085.080	11	-2.5685	1592
		(e)	0.103	0.864	0.071	2099.393	12	-2.5819*	1592

R^2_m, the R-squared of conditional mean equation m. *Spec.*, the log-GARCH specification in question, see Section 3. $\widehat{\alpha}_{mm.1}$, $\widehat{\beta}_{mm.1}$ and $\widehat{\delta}_m$, the ARCH, GARCH and asymmetry/leverage estimates, respectively. $LogL_m$, the Gaussian log-likelihood of equation m in question. k_m, the number of parameters for equation m in the log-GARCH specification in question. BIC_m, the value of the Schwarz (1978) information criterion for the equation m in question in terms of the average log-likelihood. T, the number of observations. All computations in R (R Core Team (2014)).

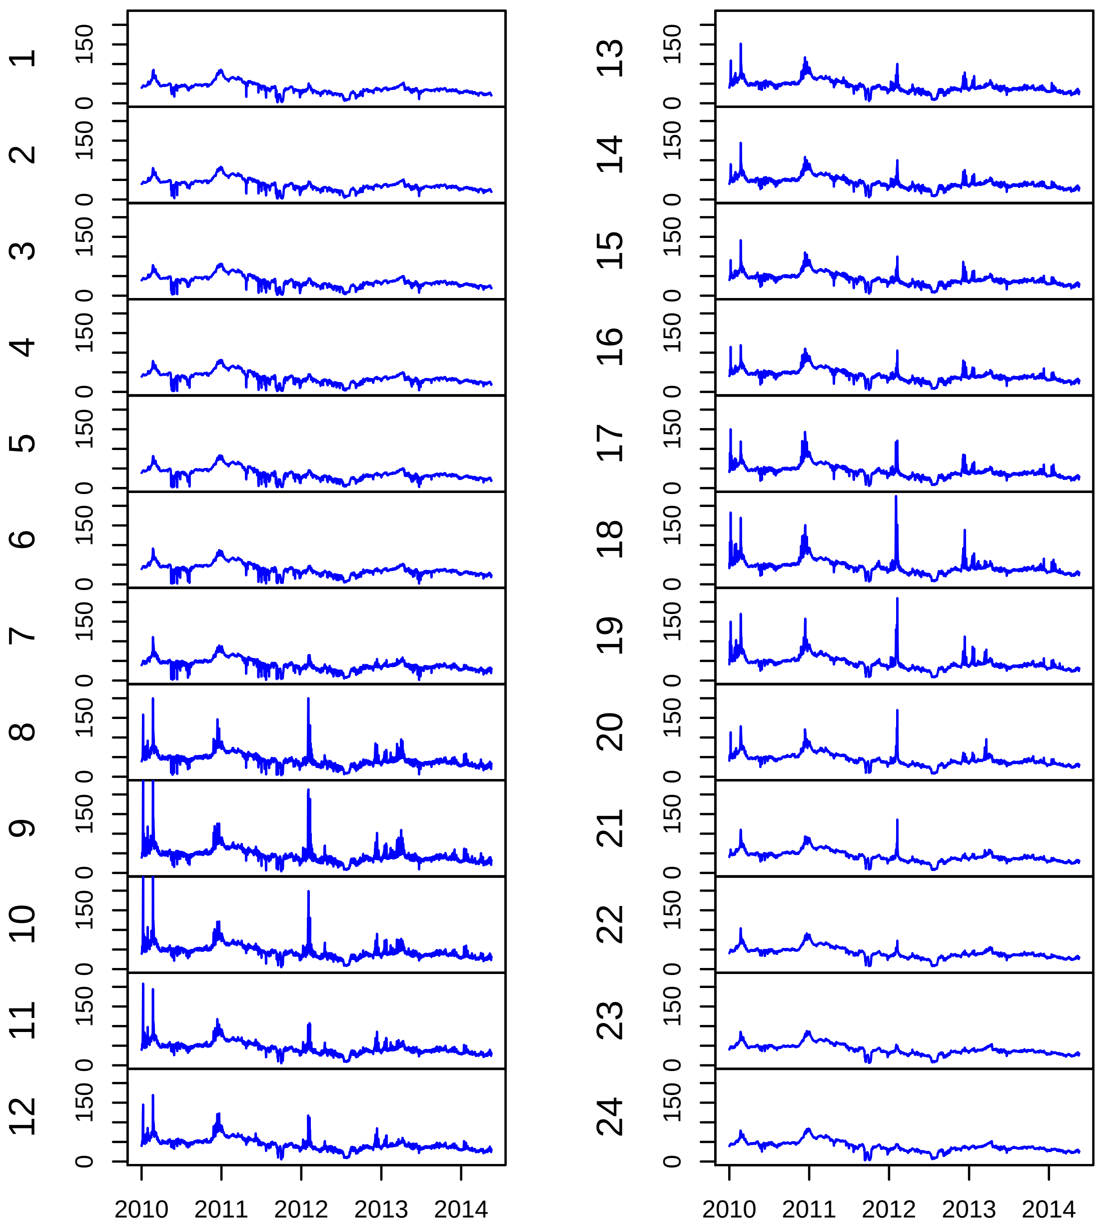

Figure 1: Hourly ($m = 1, \ldots, 24$) day-ahead system prices in Euros per kw/h at Nord Pool from 1 January 2010 to 20 May 2014 ($T = 1601$ observations before differencing and lagging), see Section 3.1.

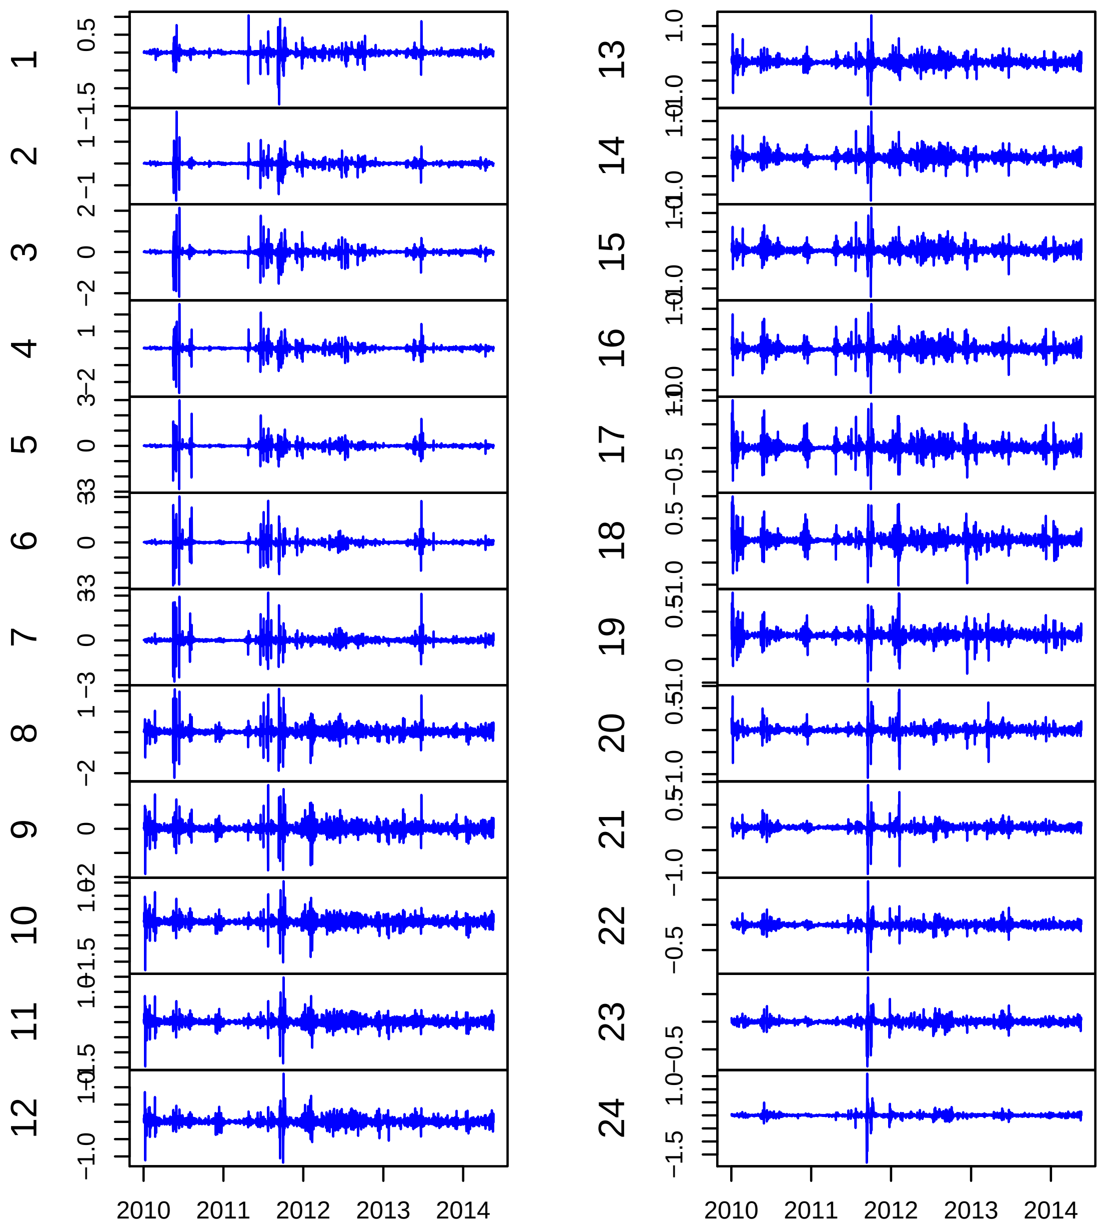

Figure 2: Hourly ($m = 1, \ldots, 24$) day-ahead log-returns of system prices at Nord Pool from 1 January 2010 to 20 May 2014, see Section 3.1.

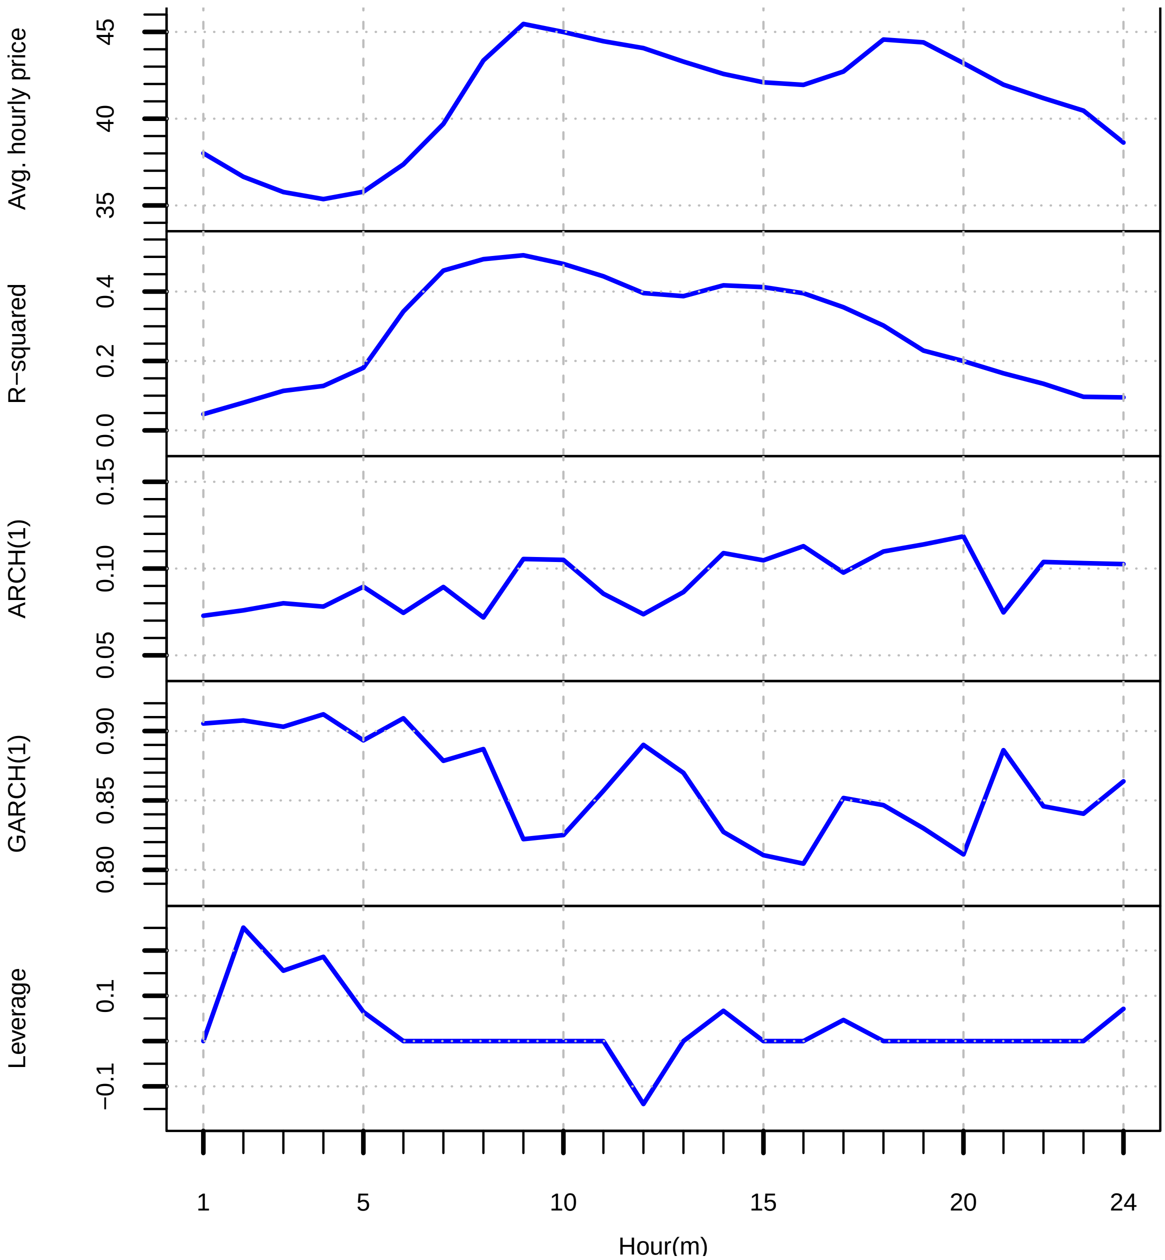

Figure 3: The average hourly price in Euros (top graph), the R-squared of the conditional mean equations (second graph), and the estimated ARCH(1), GARCH(1) and leverage parameters of the best specification (according to BIC) in each m (third, fourth and bottom graphs). The best specification in hour m is identified by an asterisk * to the right of its BIC value in Table 1.

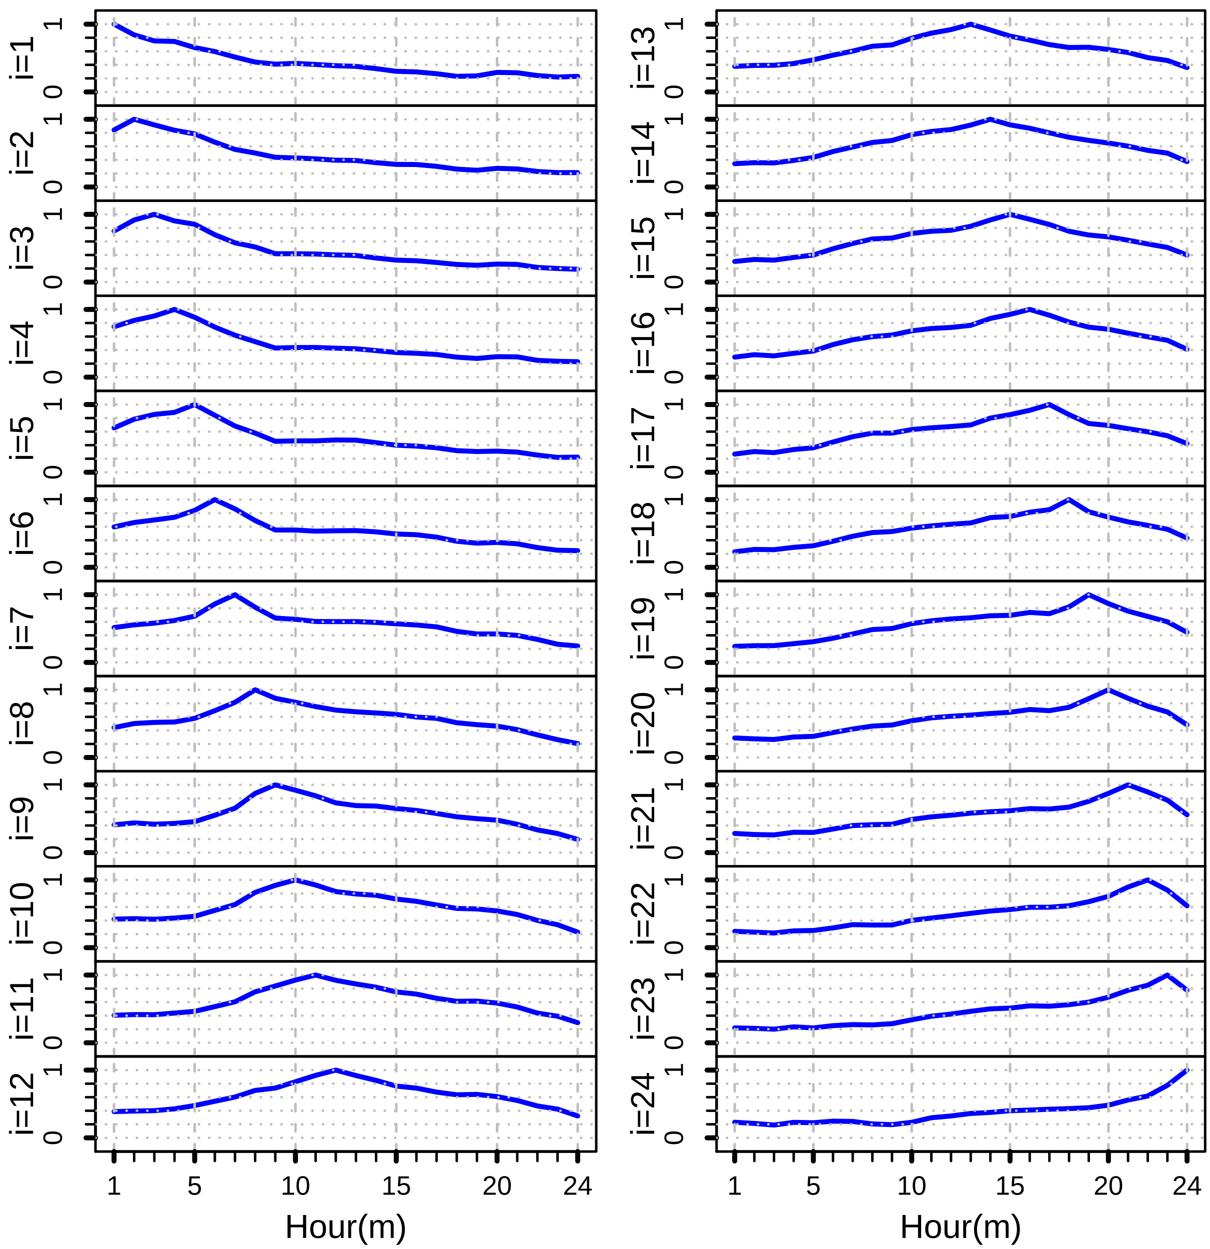

Figure 4: Estimates of the unconditional correlations $E(\eta_{it}\eta_{mt})$, $i = 1, \ldots, 24$ and $m = 1, \ldots, 24$, of the cDCC, see Section 3.3.

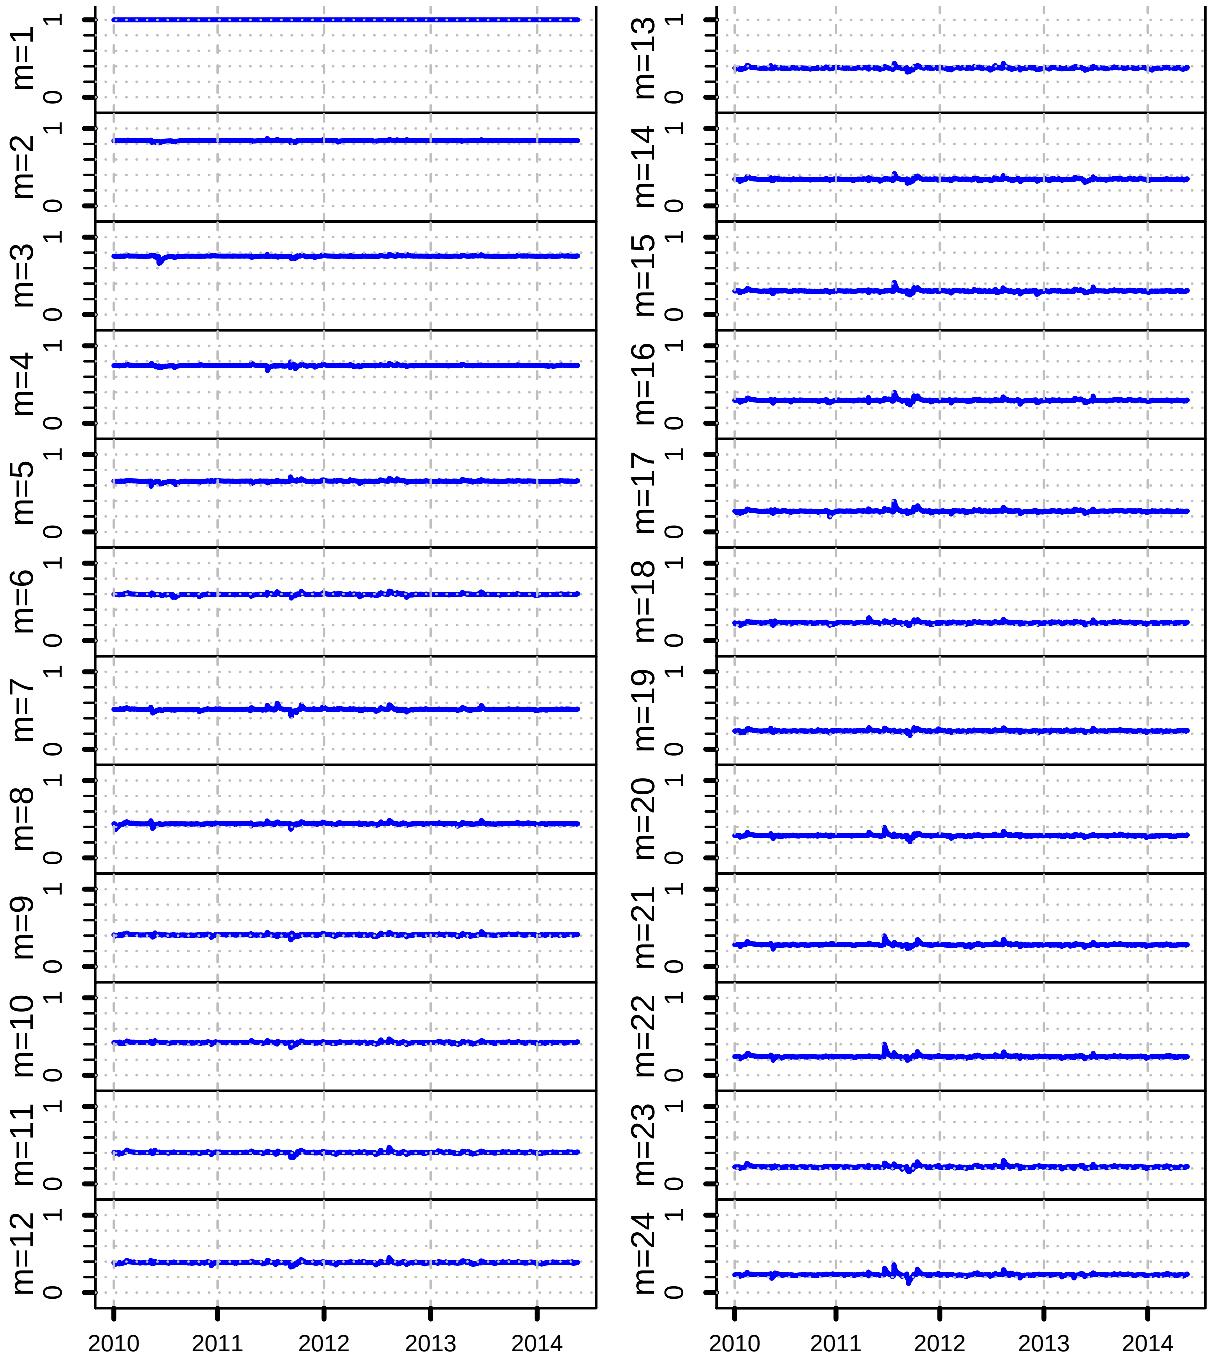

Figure 5: Estimates of a subset (the first 24) of the cDCC conditional correlations paths $\{\widehat{E}_{t-1}(\eta_{1t}\eta_{mt})\}$ for $m = 1, \ldots, 24$, see Section 3.3.