Volatility of electricity day-ahead prices: Evidence from the French Powernext exchange*

Steve Lawford† and Spyridon Liarmakopoulos

RALM, Electrabel, and Department of Economics and Finance, Brunel University

This version: 30 March 2005.

Abstract

European electricity markets have undergone significant deregulation, and time series are now available on various day-ahead and futures prices. Despite some distributional similarities with asset prices (e.g. high kurtosis of returns, and clear volatility clustering), electricity has dramatically different stochastic properties to standard financial products, and even other commodities. Electricity prices are often mean-reverting, and contain strong seasonal components, which reflect demand and supply factors. Spot prices can make extreme short-term jumps, due to physical considerations such as weather and transmission-line overloads. In this paper, we study the stylized features of the French Powernext daily spot and log return. In order to analyze the volatility properties of the returns series, we employ a multiplicative deseasonalization that was introduced recently in studies of intraday financial data. Returns volatility is split into a deterministic seasonal component, and a stochastic part that can be modelled using conventional univariate GARCH techniques. We find that there are no leverage effects in the daily returns (unlike e.g. asset returns and interest rates), and show that a simple AR(2)-ARCH(1) model, combined with a deterministic intraweek pattern, provides a good representation of the volatility of deseasonalized returns.

*JEL classification: C10, L94. Keywords: Electricity markets, GARCH, Heteroscedasticity, Seasonality, Volatility.

†Corresponding author. Steve Lawford, RALM, Electrabel, Place de l'Université 16 (4ème étage), 1348 Louvain-la-Neuve, Belgium. Email: steve_lawford@yahoo.co.uk. Fax: +32 10 48 51 09.

1 Introduction

Partial or full deregulation of electricity markets has taken place recently in many European countries, with a view to increasing competition and price transparency (see Joskow, 1997, and Mork, 2001, for a clear overview of the issues involved). The creation of new energy exchanges (such as the Dutch APX, German EEX, Scandinavian Nord Pool, and French Powernext) has made accurate modelling of the stochastic behaviour of electricity spot and forward prices, and the construction of reliable forecasts and derivative pricing tools, very important topics for producers, traders, risk managers, and end-users of electrical energy. However, it is not advisable to borrow models developed for standard financial securities (e.g. stocks) and storable commodities (e.g. metals, foods, fuels), and apply them in these new settings, since electricity prices possess stylized characteristics that such models cannot capture, and that are not often seen in other financial markets.

Electricity is a "physical" consumption product, or flow commodity, which results from conversion of various forms of energy (e.g. nuclear, thermal). Electricity is continuously generated and consumed, and meaningful quantities virtually cannot be stored efficiently at reasonable cost – hydroelectric storage schemes are especially expensive, for instance. This leads to a delicate load-generation balancing act, whereupon relatively small changes in demand (e.g. unexpected periods of hot weather, with increased use of cooling systems), which cannot be instantaneously satisfied due to production constraints; or supply (e.g. technical problems, such as power plant failure and transmission-line overload, combined with inelastic demand), lead to sudden large spikes in the spot price, that cannot be smoothed out using inventories. These spikes are typically short-lived and act over several days, contributing to periods of very high volatility (which are periodically recurrent in some markets, e.g. Geman and Roncoroni, 2006), after which the spot price returns rapidly towards its pre-spike level, rather than leading to sustainable higher prices. The degree of volatility in electricity spot markets far exceeds that observed in very volatile stock or commodities (e.g. natural gas, crude oil) markets.

Major load and generation problems often arise simultaneously, creating high risks and exacerbating spikes. These can be extreme, e.g. between 10 and 11 August 2003, spot prices on the Dutch APX exchange rose by over 3,000% (from 19.18 euros/MWh to 660.34 euros/MWh), only to return to original levels 5 days later.[1] Less dramatic spikes tend to be observed in those markets that have efficient transmission networks, and diversified methods of generation. The spike behaviour poses particular identification problems for simple stochastic mean-reverting jump diffusion models, which commonly combine an Ornstein-Uhlenbeck process with an additive upward Poisson jump (usually with lognormally-distributed jump magnitudes), and maximum-likelihood estimation. These models generally require a very high level of mean-reversion following a price spike, which distorts the model performance in non-spike periods (see Huisman and Mathieu, 2003, for

[1] 1 euro/MWh equals the cost of running approximately 16,700 standard 60W light-bulbs for 1 hour.

a critique), while paucity of historical data can result in imprecise parameter estimates. Applications of this class of models to electricity markets include Escribano et al (2002), Lucia and Schwartz (2002), Cartea and Figueroa (2004) and Weron et al (2004). See Borovkova and Permana (2004) for further discussion, and for a generalization of the drift term, in which the mean-reversion rate can be a nonlinear function of the distance from the seasonal average price. In a comprehensive paper, Geman and Roncoroni (2006) address many issues related to electricity spot price modelling (and spike behaviour in particular), using a level-dependent signed-jump diffusion model, with time-varying jump intensity.

Electricity generation is subject to natural phenomena (most often, climate conditions such as temperature or rainfall), which is manifested in a marked seasonal dependence in the data, the pattern of which is market-specific. Demand-driven annual seasonality can arise from fluctuating temperature, or number of daylight hours. Intraweek seasonality can be caused by a lower industrial demand at night and weekends, relative to weekday working hours (Wilkinson and Winsen, 2002, illustrate the dependence of average intraday spot prices on day-type, for the Australian New South Wales market). In some markets, supply-side seasonality is also important, e.g. hydroelectric generation, which has an impact on the Scandinavian Nord Pool exchange, is affected by rainfall and snowmelt. It is common practice to model deterministic periodic patterns such as annual seasonality using a linear trend and one or more sinusoidal functions (with distinct periods to reflect, say, annual and six-month periodicity, due to cold winters and hot summers – more generally, a Fourier series expansion could be used), and intraweek seasonality using a piecewise-constant function (e.g. Pilipović, 1998).

Electricity usually also shows notable intraday seasonality, e.g. weekend hourly prices may be affected by demand surges at midday and evening mealtimes, while weekday and weekend intraday patterns may also differ. Bottazzi et al (2004) examine Nord Pool intraday patterns, and model the distribution of log returns using the general Subbotin distribution, which nests the Laplace, Gaussian and uniform, among others. See also Longstaff and Wang (2004), for a detailed empirical analysis of real-time hourly spot and day-ahead intraday behaviour on the US Pennsylvania–New Jersey–Maryland (PJM) market.

Electricity spot prices are often found to be mean-reverting (Weron and Przybyłowicz, 2000, Atkins and Chen, 2002, Lucia and Schwartz, 2002, Simonsen, 2003, and Resta, 2004, document this empirically for electricity, and more generally mean-reversion of energy commodities is examined by Schwartz, 1997). Simply, this suggests that there is an equilibrium level to which spot prices will return, following a temporary jump, and which reflects economic and fundamental factors such as the marginal cost of production and seasonal weather conditions – the equilibrium may be constant, periodic, or periodic with a trend. Some markets, such as Nord Pool, are characterized by hourly spots with a high degree of long memory, and hyperbolically decaying autocorrelation functions (Carnero et al, 2003, and Haldrup and Nielsen, 2005). Time series from electricity spot markets generally exhibit volatility clustering, that is, periods of high volatility tend to be followed by periods of high volatility, and quiet days tend to be followed by quiet days (e.g. Atkins and Chen, 2002, Resta and Sciutti, 2003, and Knittel and Roberts, 2004).

The unbundling of distribution and generation leads to close links between spot prices and the underlying demand-supply mechanism. Routledge et al (2001) develop a competitive rational expectations framework for non-storable electricity and storable potential fuels. Various empirical features follow directly from their equilibrium analysis, including skewed spot price distributions, price-dependent heteroscedasticity – with higher (lower) volatility when prices are higher (lower), and unstable electricity-fuel correlations. They show that skewness and heteroscedasticity are immediate consequences of inelastic demand, combined with a possibly discontinuous supply curve (the "supply stack") – which is flat at low outputs (corresponding to baseload generation using efficient and cheaper power sources) but becomes very steep at high outputs (corresponding to peak generation using costly and less efficient standby power sources). Capacity and transmission limits dictate that the supply stack will be vertical at maximum output. Hence, demand changes at low output (and low price levels) will have a relatively small impact upon price (low volatility), while demand changes at high output (and high price levels) will result in high price changes (high volatility). We may then expect the proximity of a system to generation capacity, and/or the spot price level, to drive volatility – and this will of course be determined by the physical considerations discussed above.

The outline of the paper is as follows. Section 2 briefly reviews the functioning of the French Powernext day-ahead market. Section 3 describes the stylized features of the Powernext day-ahead spot and log returns series, and reports some preliminary descriptive statistics and tests for unit roots and autocorrelation. This exchange is of some interest, since the market price is derived from a blind auction procedure, and the resultant time series is virtually unstudied in the literature. We see that there is unusual volatility clustering in the returns, and we focus on this effect rather than attempting to explicitly model the electricity spot price (see references in text, and footnote 2 below). In Section 4, we proceed to develop a univariate GARCH model of returns volatility. We first employ a deterministic volatility deseasonalization that was made popular in the context of high-frequency intraday data. The stochastic part of returns volatility can then be modelled successfully with standard symmetric GARCH techniques. We use an AR(2)-ARCH(1) model for the (volatility) deseasonalized returns, with a deterministic (sinusoidal) intraweek pattern, to construct one-step-ahead volatility forecasts. Robust misspecification tests indicate that the conditional mean and variance functions provide satisfactory representations of the data, and volatility is shown to be only mildly persistent. Section 5 concludes.

A number of studies (in a rapidly expanding literature) have now examined models for electricity spot prices, forward curve dynamics, and derivative pricing.[2] However, much

[2]Aside from (Brownian, possibly multi-factor) jump diffusion spot models, other approaches include regime switching models (e.g. Huisman and Mathieu, 2003, and Weron et al, 2004), which permit explicit modelling of "normal" and "spike" regimes and their associated speeds of mean-reversion, Lévy-based models (e.g. Deidersen and Trück, 2002), and models based upon economic fundamentals, with explicit modelling of the aggregate demand and/or supply curves (e.g. Barlow, 2002, who derives a jump-free

less attention has been given to modelling the volatility observed on electricity returns, and this paper is among the first empirical illustrations to use the French Powernext market. We extend the literature by applying a deterministic annual/intraweek volatility deseasonalization to electricity returns, and by showing – surprisingly – that the stochastic part of returns volatility can be modelled using a simple ARCH(1) process. Our method also appears to be robust to removal of the largest spikes from the spot/returns series, in contrast to some existing literature, which models the spike generation mechanism explicitly.

2 The market and data

Powernext SA was created on 26 November 2001, as the first French power exchange, following the 10 February 2000 application into French law of the 1996 European directive 96/92/EC on restructuring of previously monopolistic European electricity markets. It represents a centralized, organized, anonymous exchange for a variety of day-ahead electricity contracts (Powernext Day-Ahead$^{\text{TM}}$, or PDA) and medium-term futures (the Powernext Futures$^{\text{TM}}$ market was incorporated on 18 June 2004). In this paper, we focus on the PDA market. Market participants trade in standardized hourly contracts that commit them to provide to the French transmission network (the restricted geographical area is referred to as a "hub") a specified volume of physical electricity throughout a given hour, or to withdraw it, at a common and recognized market price. The PDA is essentially open to market operators that have been authorized to buy/sell electricity under their national law, and who pay fixed entrance and annual fees, and variable delivery/trading and clearing fees per volume traded. As of July 2004, there were 41 active members, including energy companies and investment banks.

The underlying is physical electricity, traded for next-day (day $t + 1$) delivery over 24 hourly intervals, and this gives rise to the hourly "spot" market clearing price. The daily baseload spot is calculated as the mean of the 24 distinct hourly spots, and is the series studied here. Trading takes place during a "pre-auction" phase, which runs from 17:00 CET on Wednesday the previous week to 11:00 CET on the trading day itself (day t). Trading occurs 24 hours a day, 7 days a week, and including statutory holidays, unlike e.g. NYSE and Nasdaq.[3] Members submit an electronic order form via the internet, which contains a maximum of 64 price/quantity combinations, for each of the 24 one-hour periods of the delivery day, within a minimum and maximum price of 0 euros and 3,000 euros. Bids for purchase (sale) are represented by positive (negative) volumes respectively, and only

diffusion model that can exhibit spikes, and Bessembinder and Lemmon, 2002). Forward price dynamics and joint spot/forward models have been examined by, inter alia, Koekebakker and Ollmar (2001), Manoliu and Tompaidis (2002), Borovkova (2004) and Cartea and Figueroa (2004). Electricity derivatives are discussed by Pilipović (1998), Eydeland and Geman (2000), Barone-Adesi and Gigli (2002) and Vahviläinen (2002).

[3] See Bauwens and Giot (2001, pp 11-17) for discussion of the hybrid (market-maker and order book) New York Stock Exchange (NYSE) and the automated quote-driven National Association of Securities Dealers Automated Quotation (Nasdaq). For extensive details, consider www.nyse.com and www.nasdaq.com.

integer volumes are traded. Volume and price ticks are 1 MWh and 0.01 euros/MWh. Between any two prices entered into the order form, all volumes are assumed to be linear. Order prices are entered to 2 decimal places, while traded quantities are allocated based on the exact equilibrium price. The monthly volume on the PDA exceeded 1.25 TWh (1 terawatt hour=1 million MWh) in July 2004.

Each daily auction (day t) is followed by an approximately 15 minute validation phase, during which the purchase/sale bids are aggregated for each hourly period, and the market clearing price (or PDA "spot") and volume are calculated by linear interpolation. The aggregation is performed by SAPRI, a trading system that is operated by Nord Pool, and acts as a centralized order book, and the concentration of orders via auction is designed to maintain liquidity. Following validation, the settlement transaction amount is the spot price multiplied by the volume traded, taken as a linear interpolation between the two price/volume combinations that contain the market clearing price, on the most recent order form sent by the member for that day. The "quoted" market clearing price is rounded to 3 decimal places. The aggregated buy/sell curves are communicated to members, although the individual order books are withheld, which ensures anonymity. LCH.Clearnet SA acts as a financial counterparty between the member and Powernext, and physical delivery or withdrawal from the French hub is managed by the transmission system operator (RTE).

The above is a slight simplification of actual events, and serves to illustrate the blind auction procedure. While hourly products are a benchmark in electricity markets, it is also possible to trade standardized block bids on the PDA (e.g. hours 1–4, hours 5–8, hours 1–24, hours 9–20), subject to some restrictions. The market clearing price is in fact calculated from both single and block bids, where the block bids are either entirely executed or entirely rejected, according to a decomposition and iterative method. Of crucial importance from the above discussion is the observation that a classical spot market cannot exist for electricity, since the feasibility of the schedule must be verified in advance of delivery. We note that considerable volumes of spot electricity are also traded in bilateral agreements between generation companies and buyers.

The data under study consist of the average PDA daily baseload electricity spot prices in euros/MWh, which was collected and provided by Powernext SA. The in-sample estimation period runs across every day from 27.11.2001 (when the exchange first opened) through to 21.06.2004, including weekends and holidays, which yields a total of $T = 938$ daily observations. For further details on the specifics of the Powernext market and available products, the reader is referred to www.powernext.fr.

3 Empirical features of electricity daily spot and returns

Time series of daily spot prices S_t and one-day log returns $R_t = \ln S_t - \ln S_{t-1}$ are plotted in Figures 1 and 2, which suggest a high level of mean-reversion, and a number of extreme spikes. Summary statistics are reported in Table 1, which indicates highly volatile, positively-skewed and leptokurtic (unconditional) spot and returns, while Jarque-Bera tests strongly reject normality at all conventional significance levels (see Figure 3). The standard deviation for the spot price exceeds 15 euros/MWh, which is almost 60% of the mean value. We find that 28 (8) spot prices exceed 50 (75) euros/MWh. Of these 28, all but one occur in winter (December, January, February) or summer (June, July, August), and all occur on weekdays, when demand is high. The 10% and 90% empirical tail quantiles of S_t are 14.91 and 35.54 euros/MWh respectively. The highly right-skewed spot, and leptokurtic returns, are consistent with many empirical findings (we verified this using APX, EEX and Nord Pool daily spot data over similar time periods), and reflect the spike component in the price process. Daily mean spot prices are 21.21, 29.28 and 30.01 euros/MWh over the 1–year subsamples 01.01.2002–31.12.2002, 01.01.2003–31.12.2003 and 01.06.2003–31.05.2004, suggesting a general rise in average prices over the sample period. Cramér-von Mises and Anderson-Darling empirical distribution tests reject the nulls that the spot and return are well approximated by common distributions such as chi-squared, exponential, gamma, logistic and Pareto.

	mean	median	max.	min.	std. dev.	skew.	kurt.	AC_1
S_t	26.04	24.41	310.37	4.93	15.12	9.12	148.42	0.439
R_t	-1.4×10^{-4}	-0.032	2.78	-1.92	0.37	0.94	8.55	-0.185

Table 1 – Summary descriptive statistics for daily spot S_t and daily log return R_t

The returns and squared returns are clearly not independent, and the first-order autocorrelation coefficients of R_t and R_t^2 equal -0.185 and 0.284 respectively, and are statistically significant (this is confirmed by application of Brock-Dechert-Scheinkman tests).[4] Seasonality is present in both series, and the autocorrelation function attains a local maximum every 7 days (for instance, there is a positive relationship between log returns and returns lagged 7 days, with a simple correlation of 0.43). Moreover, the Ljung-Box portmanteau statistics for up to fourteenth-order serial correlation in R_t and R_t^2 are 559.97 and 117.77, and are highly significant when compared to the limiting chi-squared distribution. The pronounced periodicity and persistence is also evident in the correlograms of R_t and R_t^2 (Figures 4 and 5; see also Weron, 2000, Figure 4). These results strongly suggest both volatility clustering in daily returns, and a seasonal volatility effect.

Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) unit root tests on S_t and R_t are significant at the 1% level, while the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test

[4]First-order autocorrelations AC_1 in Table 1 are similar to those for some other European markets, including the Dutch APX over 02.01.2001–21.06.2004 (spot 0.538, return -0.188) and German EEX over 17.06.2000–21.06.2004 (spot 0.535, return -0.172). Other markets present quite different results, e.g. Scandinavian Nord Pool over 01.01.1999–13.06.2004 (spot 0.977, return -0.130), that are rather similar to those for other financial time series such as stocks, e.g. consider spot AC_1s over an identical sample period to Powernext, computed on the mean of open–close prices, for Microsoft (Nasdaq) 0.995, Yahoo! (Nasdaq) 0.995, Continental Airlines (NYSE) 0.996, and Metro-Goldwyn-Mayer (NYSE) 0.990. Note that weekends and holidays reduce the sample size for these NYSE/Nasqaq stocks to approximately 645.

for stationarity rejects the null at the 5% level for S_t and strongly does not reject the null for R_t.[5] This evidence suggests that the returns series is stationary, while conflicting spot unit root and stationarity test results are perhaps symptomatic of long-range dependence.

4 Volatility modelling

Longstaff and Wang (2004, especially Table 7 and Figure 5) report an "inverted-U" intraday pattern of unconditional volatility (standard deviation) of daily changes in day-ahead PJM prices, while Weron (2000, especially Figure 3) illustrates seasonal volatility for the California Power Exchange spot market. In a different context, the empirical market microstructure literature contains many examples of intraday seasonal volatility patterns in forex and equity markets. For instance, Baillie and Bollerslev (1991) examine returns volatility on four exchange rates (vis-à-vis the US dollar) using six months of continuously recorded hourly data across different world markets. They report a pattern of intraday volatility that corresponds to the open and close of the international financial markets upon which the exchange rates are quoted. Similarly, Andersen and Bollerslev (1997b, especially Figure 2), and Barndorff-Nielsen and Shephard (2002, especially Figure 1(a)), document strong intraday seasonal volatility of the US dollar/German deutschemark using long very-high-frequency time series, which is clearly visible in the autocorrelation function of absolute or squared 5–minute returns. The "U-shape" of equity market intraday volatility was discussed by Wood et al (1985) and Harris (1986). In electricity markets, this periodicity has a natural as well as an institutional basis.

When analyzing volatility, it is important to account appropriately for such strong intraday (or more generally, seasonal) patterns. Some techniques of univariate deseasonalization, based upon periodically-varying volatility weights, were implemented by, inter alia, Andersen and Bollerslev (1997a), and are further discussed by Breymann et al (2003). We choose a simple volatility deseasonalization, based on work by Giot (2000). We then model volatility in the context of a standard GARCH framework, although various extensions are possible. Pronounced volatility clustering is a salient feature of many financial series, and has sparked a huge literature. For excellent surveys of GARCH techniques and applications to finance, see Bollerslev et al (1992), Bera and Higgins (1993), Bollerslev et al (1994), Engle (2001, 2002, 2004), Engle and Patton (2001) and Li et al (2002).

4.1 Seasonal volatility pattern and deseasonalized log returns

Given the autocorrelation found in squared log returns, we investigate a deterministic annual/intraweek volatility pattern. We calculate the mean value $\varphi_t(i;j)$ of the squared log returns observed on day i and month j, where $i = 1$ (Mon), 2 (Tue),…, 7 (Sun), $j = 1$

[5] We use the Schwarz Information Criterion to select lag length for the ADF test, and include a constant in the regression. We use Bartlett kernel spectral estimation when implementing the PP and KPSS tests.

(Jan), 2 (Feb),…, 12 (Dec). For example, there are 12 datapoints that are Mondays in January. Summing the squared log returns observed on these dates and dividing by 12, we obtain $\varphi_t(1;1)$. Repeating this procedure for all i, j, we obtain Table 3. It is apparent that average squared returns on Mondays are significantly greater than on other days, followed by returns on Sundays (in some cases, Saturdays). There is some tendency for average squared returns to be higher during summer months, although this is not conclusive. The intuition is that business activities follow a weekly cycle, with lower levels at weekends, and hence a large negative return on Saturdays/Sundays, and higher levels on working days, giving a high positive return on Mondays. These significant calendar effects must be taken into account, and we define the (volatility) deseasonalized log price return at time t as

$$\widetilde{R}_t(i;j) := R_t(i;j)\left[\varphi_t(i;j)\right]^{-1/2} \tag{1}$$

(see Figure 6). This approach to computing deseasonalized log returns was discussed by Giot (2000) and Bauwens and Giot (2001), and applied to high-frequency financial data. Interestingly, (1) renders the unconditional deseasonalized returns $\widetilde{R}_t$ approximately normal (see Figure 7, and also Andersen et al, 2001, especially Figure 1, for an analogous finding when daily equity returns are standardized by realized volatility).

	mean	median	max.	min.	std. dev.	skew.	kurt.	AC_1
$\widetilde{R}_t$	−0.11	−0.14	3.50	−3.33	0.99	0.14	3.12	−0.104

Table 2 – Summary descriptive statistics for deseasonalized daily log return $\widetilde{R}_t$

4.2 An ARCH model

Although the autocorrelation function in Figure 9 indicates that "obvious" seasonality has been removed from the squared returns (similarly, see Breymann et al, 2003, Figures 3 and 5), the Ljung-Box statistics for up to fourteenth-order serial correlation in $\widetilde{R}_t$ and $\widetilde{R}_t^2$ are 368.13 and 73.92, and are again highly significant. Several authors have indicated the importance of correctly modelling the conditional mean function when applying GARCH techniques to electricity spot prices, e.g. Escribano et al (2002) and Karakatsani and Bunn (2004). Failure to address this issue can lead to erroneous GARCH results, such as near-integrated GARCH (IGARCH), due to the impact of extreme price spikes.

In this paper, we do not model the spike process explicitly, and instead follow two directions. Firstly, we do not treat spikes, but deseasonalize volatility prior to modelling periodicity in (volatility) deseasonalized returns and estimating GARCH. Secondly, we remove large spikes from the spot price process, before following the same procedure as before. We found that it was very difficult to specify a sensible GARCH model if the periodicity in the spot (return) was not explicitly modelled (compare this to simple financial applications, where it is adequate to assume a constant conditional mean, e.g. Engle and Patton, 2001, using 12 years of daily returns data on the Dow Jones Industrial Index).

There is little evidence of a significant annual periodicity in the level of deseasonalized returns. However, spot prices are clearly higher on weekdays than at the weekend, and there is evidence that Friday prices are lower than those on Monday to Thursday. This deterministic weekly periodicity is also present in the returns series (Figure 8). We model this flexibly using two sinusoidal functions (Figure 10), and the correlogram of the volatility deseasonalized returns, filtered to remove weekly periodicity, is given in Figure 11. The remaining autocorrelation in $\widetilde{R}_t$ is captured by including an AR(2) term in $\widetilde{R}_t$. An ARCH-LM(7) test gives the highly significant $F_{\text{ARCH}} = 12.05$, and hence we specify the model

$$\widetilde{R}_t = \theta_0 + \theta_1 \cos\left[\frac{2\pi}{7}(t-\theta_2)\right] + \theta_3 \cos\left[\frac{4\pi}{7}(t-\theta_4)\right] + \psi_1 \widetilde{R}_{t-1} + \psi_2 \widetilde{R}_{t-2} + \varepsilon_t, \qquad (2)$$

where θ_0 is the level of the deseasonalized returns, θ_1 and θ_3 are amplitudes of a 7-day and 3 1/2-day sinusoid, θ_2 and θ_4 are timeshifts, and $\varepsilon_t \sim \text{GARCH}(p, q)$.

We estimated model (2) for $p \in \{0, 1, 2, 3, 4, 5\}$ and $q \in \{0, 1, 2, 3, 4, 5\}$, using both the Berndt–Hall–Hall–Hausman and Marquardt optimization routines. Robust QMLE standard errors are not available. We chose the GARCH model that gave sensible, and significant parameter estimates, and which minimized the Schwarz Criterion (SC = 2.3303). Surprisingly, given the prevalence of GARCH-type behaviour in applied finance studies, we found ARCH(1) to be the order which satisfied these conditions, and which generally converged after the least iterations, with the more stable Marquardt algorithm. Hence, the conditional variance for deseasonalized returns is

$$h_t = \omega + \alpha_1 \varepsilon_{t-1}^2, \qquad (3)$$

where $\varepsilon_t = \epsilon_t h_t^{1/2}$ are residuals from (2), and ϵ_t are Student's $t(r)$ innovations. The implication is that the volatility of the deseasonalized returns process is a function of the error made in forecasting the previous day's return. Note that volatility persistence increases as α_1 approaches unity. Estimated parameters and standard errors are:

$\widehat{\omega}$	$\widehat{\alpha}_1$	$\widehat{r}$
0.4472	0.3273	5.9220
(0.0399)	(0.0799)	(1.3963)

$\widehat{\theta}_0$	$\widehat{\theta}_1$	$\widehat{\theta}_2$	$\widehat{\theta}_3$	$\widehat{\theta}_4$	$\widehat{\psi}_1$	$\widehat{\psi}_2$
−0.1591	−0.7335	52.9226	−0.4249	−2.3518	−0.3189	−0.0767
(0.0233)	(0.0358)	(0.0627)	(0.0350)	(0.0475)	(0.0348)	(0.0301)

and $R^2 = 0.36$. The estimated conditional variance (and other) parameters are all highly significant. The adequacy of (2) and (3) was examined using an ARCH–LM(7) test, which gave an insignificant value of $F_{\text{ARCH}} = 1.29$, indicating that all GARCH effects have successfully been removed, and that (2) and (3) track the short-run interdaily volatility dependencies quite well.

We calculate the unconditional variance as $h_u = \widehat{\omega}(1 - \widehat{\alpha}_1)^{-1} = 0.6648$, which corresponds to an average annualized (stochastic component of) volatility of approximately $((365)(0.6648))^{1/2} \approx 15.6\%$. The volatility half-life is 0.62 days, which is very mildly persistent (a t-test of a unit root in the conditional variance equation gives $|\tau| = 8.42 > 1.96$, and we reject the IGARCH null). This makes intuitive sense, since volatility is driven by short-term imbalances in demand and supply. It is straightforward to combine (1) and (3) to calculate volatility for the raw returns R_t as $(h_t \varphi_t(i;j))^{1/2}$. The in-sample forecasting ability of this model is illustrated in Figure 12, where h_t is plotted against $\widetilde{R}_t^2$ (the latter acts as a proxy for ex post volatility).

4.3 The absence of leverage effects

The symmetric relationship between lagged shocks and conditional variance can be illustrated conveniently using a news impact curve (Engle and Ng, 1993), given directly from (3) as $h_t^{\mathrm{NIC}} = 0.4472 + 0.3273\varepsilon_{t-1}^2$. We check for the presence of asymmetric GARCH effects using the Engle-Ng sign bias (SB) test, negative and positive size bias (NSB and PSB) tests, and a joint LM test. The SB test examines the possibly asymmetric impacts that shocks of different sign have on volatility. The NSB (PSB) test investigates the possibly different impacts that negative (positive) shocks of different magnitude have on volatility. We define a dummy variable

$$S_t^- = \begin{cases} 1 & \text{if} \quad \widehat{\varepsilon}_t < 0 \\ 0 & \text{if} \quad \widehat{\varepsilon}_t \geq 0 \end{cases}$$

and define negative and positive residuals as $\widehat{\varepsilon}_t^- = S_t^- \widehat{\varepsilon}_t$ and $\widehat{\varepsilon}_t^+ = \left(1 - S_t^-\right)\widehat{\varepsilon}_t$ respectively. We then estimate (using standardized and ordinary residuals)

$$\widehat{\epsilon}_t^2 = \kappa_0 + \kappa_1 S_{t-1}^- + \kappa_2 \widehat{\varepsilon}_{t-1}^- + \kappa_3 \widehat{\varepsilon}_{t-1}^+ + \text{ error}, \tag{4}$$

and test SB, NSB and PSB by using the t-ratios on κ_1, κ_2, and κ_3 in (4),

$\widehat{\kappa}_0$	$\widehat{\kappa}_1$	$\widehat{\kappa}_2$	$\widehat{\kappa}_3$	$\widehat{\kappa}_2 + \widehat{\kappa}_3$
1.0740	$-0.0718^{\star\star}$	$0.0511^{\star\star}$	$-0.1446^{\star\star}$	$-0.0935^{\star\star}$
(0.1254)	(0.1852)	(0.1312)	(0.1263)	(0.1821)

where $\star\star$ indicates insignificance at the 10% level, and White's heteroscedasticity-consistent standard errors are given in parentheses. The SB, NSB and PSB tests are highly insignificant, which suggests that there are no asymmetric effects. Moreover, a joint LM test of no asymmetry, computed from (4) as $TR^2 = 0.95 < 7.81 = \chi_{0.95}^2(3)$, does not reject the null.

We corroborate these findings using asymmetric GARCH. A threshold TGARCH(1,0,1) extension of (3) was selected, following the same procedure as above, using conditional mean specification (2), and with Student's $t(r)$ innovations, i.e. $h_t = \omega + \alpha_1\varepsilon_{t-1}^2 + \lambda_1 S_{t-1}^- \varepsilon_{t-1}^2$. Estimated parameter values change slightly, and are not reported here. The "asymmetry" parameter $\widehat{\lambda}_1 = 0.0188^{\star\star}(0.1348)$ is highly insignificant, suggesting that there is no (first-order) leverage effect. Similarly, we selected exponential EGARCH(1,0,1) and power PGARCH(1,1,1) models. Conditional variance specifications are, with asymmetry parameters γ_1 and ϕ_1, $\ln(h_t) = \omega + \alpha_1|\varepsilon_{t-1}h_{t-1}^{-1/2}| + \gamma_1\varepsilon_{t-1}h_{t-1}^{-1/2}$ and $h_t^{\delta/2} = \omega + \alpha_1(|\varepsilon_{t-1}| - \phi_1\varepsilon_{t-1})^\delta + \beta_1 h_{t-1}^{\delta/2}$. We find that $\widehat{\delta} = 3.8315(1.6510)$, while $\widehat{\gamma}_1 = 0.0381^{\star\star}(0.0593)$ and $\widehat{\phi}_1 = -0.0038^{\star\star}(0.0715)$ are again highly insignificant. The conditional variance of the daily PDA deseasonalized returns does not respond asymmetrically to positive and negative shocks, unlike e.g. asset prices (Nelson, 1991) and interest rates (Chan et al, 1992).

4.4 Spike filtering, volatility forecasts, and long-memory

We identified the 8 spot price dates that exceeded 75 euros/MWh, as Mon 17 – Wed 19 Dec 2001, Tue 15 – Wed 16 July 2003, Tue 22 – Wed 23 July 2003, and Mon 11 Aug 2003 (the latter represents the extremely large value of 310.37 euros/MWh). For each of these dates, we replaced the “spikes” with the mean of the spot prices given on the day prior to and following the spike period. We then followed the volatility deseasonalization procedure above to give a large-spike-filtered $\varphi_t(i;j)$, where the bold values in Table 3 are changed. Recalculation of $\widetilde{R}_t$ and estimation of (2) and (3) gave:

$\widehat{\omega}$	$\widehat{\alpha}_1$	$\widehat{r}$
0.4621	0.2624	6.2780
(0.0402)	(0.0707)	(1.5762)

$\widehat{\theta}_0$	$\widehat{\theta}_1$	$\widehat{\theta}_2$	$\widehat{\theta}_3$	$\widehat{\theta}_4$	$\widehat{\psi}_1$	$\widehat{\psi}_2$
−0.1508	−0.7636	−3.0668	−0.4351	−2.3337	−0.3411	−0.0838
(0.0235)	(0.0362)	(0.0620)	(0.0353)	(0.0470)	(0.0345)	(0.0308)

and $R^2 = 0.38$. An ARCH–LM(7) test gave a borderline value of $F_{\text{ARCH}} = 2.11$. The unconditional variance and half-life of the large-spike-filtered series were calculated to be 0.6265 (corresponding to an average annualized - stochastic component of - volatility of approximately 15.1%), and 0.52 days. We see that these values are quantitatively very similar to those above, and that removal of (even very large) spikes from the spot series does not have the same importance as in other studies. We conclude that our deseasonalization of the returns volatility, and correct specification of the conditional mean, serve to alleviate the problems that could otherwise require explicit modelling of the spike process.

Andersen and Bollerslev (1998a, 1998b) show that there is no contradiction between the ability of a GARCH model to provide accurate volatility forecasts, and poor predictive power for daily squared returns. We assess this here, by estimating the Mincer-Zarnowitz equation $\widetilde{R}_t^2 = \pi_1 + \pi_2 h_t + \text{error}$, for both (a) deseasonalized $\widetilde{R}_t^2$ and (b) large-spike-filtered deseasonalized $\widetilde{R}_t^2$. Estimated parameters follow, with White's standard errors in parentheses:

$\widehat{\pi}_1$(a)	$\widehat{\pi}_2$(a)	$\widehat{\pi}_1$(b)	$\widehat{\pi}_2$(b)
0.4345	0.8542	0.3253	1.0686
(0.1243)	(0.2016)	(0.1579)	(0.2692)

Estimated R^2 is 0.052(a) and 0.048(b), suggesting that the AR(2)-ARCH(1) model has low explanatory power, and explains at best 5% of the variability of squared deseasonalized returns (a similar result is obtained by computing the squared sample correlation, $\text{corr}(\bullet, \bullet)^2$, between the one-step-ahead daily volatility forecast and an ex post measure of volatility, such as $|\widetilde{R}_t|$ or $\widetilde{R}_t^2$ – for example, $100 \times \text{corr}(h_t, |\widetilde{R}_t|)^2 \approx$ (a) 3.9% and (b) 3.2% respectively, confirming the previous findings). Essentially, the difficulty is that daily $|\widetilde{R}_t|$ and $\widetilde{R}_t^2$ are noisy estimates of underlying latent volatility. Andersen and Bollerslev (1998a, 1998b) show that use of high-frequency data to construct an improved ex post measure of latent volatility (such as cumulative absolute intraday returns) explains the paradox, and greatly increases explanatory power. We do not examine hourly returns data here, since these exhibit very complicated intraday season- (and time-of-week-) dependent periodicities that would need to be accounted for in a similar spirit to the deseasonalization of daily returns, but leave this as an issue for future research.

We also follow Andersen and Bollerslev (1997b, 1998b), in checking whether it is reasonable to assume hyperbolic decay of the sample autocorrelation function $\text{AC}_m(\bullet)$ of $\widetilde{R}_t^2$, in which case fractionally-integrated or long-memory volatility processes may be appropriate. It is clear from Figures 4, 5 and 8 that the autocorrelation functions of R_t, R_t^2 and $\widetilde{R}_t$ are dominated by intraweek periodicities. This is not the case for the autocorrelation function of deseasonalized squared returns (Figure 9). Hence, we estimate the degree of fractional integration d simply by fitting a hyperbolic decay to $\text{AC}_m(\widetilde{R}_t^2)$ via the regression

$$\ln(\text{AC}_m) = \xi_0 + \xi_1 \ln(m) + \text{error}, \qquad m = 1, \ldots, M,$$

where $\widehat{d} = (\widehat{\xi}_1 + 1)/2$. We use $M = 200$ lags, although a more detailed analysis would require a significantly larger number. Estimated parameters are as follows, with White's standard errors in parentheses, and $\widehat{\text{se}}(\widehat{d}) = \widehat{\text{se}}(\widehat{\xi}_1)/2$:

$\widehat{\xi}_0$	$\widehat{\xi}_1$	$\widehat{d}$
−2.8114	−0.2111	0.3945
(0.2924)	(0.0673)	(0.0337)

An integrated process of order I(d), where $0 < d < 1/2$, will eventually have all positive autocorrelations AC_m, and these will decay at a hyperbolic rate. Since $\widehat{d} = 0.3945$, this suggests that $\text{AC}_m(\widetilde{R}_t^2) \sim m^{2d-1}$, for m large. This seems to be in contrast to the low ARCH(1) volatility persistence found above. Hence, future work could consider the case for fractionally-integrated FIGARCH-type extensions of the results in this paper (see e.g. Baillie et al, 1996).

5 Conclusions

In this paper, we combine a deterministic multiplicative annual/intraweek pattern for electricity volatility with a sinusoidal intraweek level, and an AR(2)-ARCH(1) model for the stochastic part of returns volatility. This procedure builds upon recent work on intraday financial data (which has not previously been applied to the unusual volatility that is present in electricity markets), while taking account of certain stylized features of the market price (notably, weekly seasonality). We find that all estimated parameters are highly significant, and that Powernext PDA electricity volatility is only mildly persistent. Moreover, there are no leverage effects in the stochastic part of volatility, which contrasts sharply with many results on "standard" financial series such as asset prices and interest rates. Our work provides some new insights into the nature, and modelling of, electricity market volatility.

Various extensions of this work are possible, and include (a) modelling the volatility of hourly spot (returns) data, which will involve treatment of the very complicated (season and time-of-week dependent) intraday patterns, (b) further investigation of the volatility deseasonalization in this context, including assessment of whether transformation of the unconditional returns to approximate normality is common across different markets, (c) assessment of the aggregational properties of GARCH-type techniques using this data (especially between daily and hourly data), (d) gauging the impact of exogenous variables on electricity volatility, including temperature and volume, see e.g. Lamoureux and Lastrapes (1990), (e) use of hourly data to provide improved ex post volatility measurements, (f) detailed assessment of possible long-memory dependence in both the spot series, and in volatility. We leave these issues for future research.

Acknowledgements The authors are very grateful to Karim Abadir, Michel Culot, Valérie Limpens, Sébastien de Menten, Michael Monoyios and Nikolaos Roupakas for several discussions and helpful comments during this research. We thank Powernext SA for providing the daily PDA spot price data. Steve Lawford acknowledges support from Electrabel SA, Jacqueline Boucher and André Bihain. Spyridon Liarmakopoulos acknowledges support from the Department of Economics and Finance, Brunel University. The authors take full responsibility for all information expressed herein, and for any remaining errors, and the work does not necessarily reflect the views or interests of Electrabel SA, Powernext SA, or Brunel University. This paper was compiled using MiKTeX, and numerical results were calculated using EViews.

6 References

Andersen T G and Bollerslev T, 1997a, Intraday periodicity and volatility persistence in financial markets, Journal of Empirical Finance, 4, 115-158.

Andersen T G and Bollerslev T, 1997b, Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high-frequency returns, Journal of Finance, 52, 975-1005.

Andersen T G and Bollerslev T, 1998a, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, International Economic Review, 39, 885-905.

Andersen T G and Bollerslev T, 1998b, Deutsche Mark-Dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer-run dependencies, Journal of Finance, 53, 219-265.

Andersen T G, Bollerslev T, Diebold F X and Ebens H, 2001, The distribution of realized stock return volatility, Journal of Financial Economics, 61, 43-76.

Atkins F J and Chen J, 2002, Some statistical properties of deregulated electricity prices in Alberta, Department of Economics Discussion Paper 2002–06, University of Calgary.

Baillie R T and Bollerslev T, 1991, Intra-day and inter-market volatility in foreign exchange rates, Review of Economic Studies, 58, 565-585.

Baillie R T, Bollerslev T and Mikkelsen H O, 1996, Fractionally integrated generalized autoregressive conditional heteroscedasticity, Journal of Econometrics, 74, 3-30.

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

Barndorff-Nielsen O E and Shephard N, 2002, Econometric analysis of realized volatility and its use in estimating stochastic volatility models, Journal of the Royal Statistical Society Series B, 64, 253-280.

Barone-Adesi G and Gigli A, 2002, Electricity derivatives, mimeo: Faculty of Economics (Finance), USI.

Bauwens L and Giot P, 2001, Econometric Modelling of Stock Market Intraday Activity (Dordrecht: Kluwer), Advanced Studies in Theoretical and Applied Econometrics, #38.

Bera A and Higgins M L, 1993, ARCH models: properties, estimation and testing, Journal of Economic Surveys, 7, 305-366.

Bessembinder H and Lemmon M, 2002, Equilibrium pricing and optimal hedging in electricity forward markets, Journal of Finance, 57, 1347-1382.

Bollerslev T, Chou R Y and Kroner K F, 1992, ARCH modeling in finance: A review of the theory and empirical evidence, Journal of Econometrics, 52, 5-59.

Bollerslev T, Engle R F and Nelson D B, 1994, ARCH models, Handbook of Econometrics vol IV, ed R F Engle and D McFadden (Amsterdam: North-Holland) pp 2959-3038.

Borovkova S, 2004, The forward curve dynamic and market transition forecasts, Modelling Prices in Competitive Electricity Markets, ed D W Bunn (Chichester: John Wiley and Sons) pp 267-284.

Borovkova S and Permana F J, 2004, Modelling electricity prices by the potential jump-diffusion, paper presented at Stochastic Finance 2004, ISEG, Lisbon.

Bottazzi, G, Sapio S and Secchi A, 2004, A statistical analysis of electricity price fluctuations, paper presented at Stochastic Finance 2004, ISEG, Lisbon.

Breymann W, Dias A and Embrechts P, 2003, Dependence structures for multivariate high-frequency data in finance, Quantitative Finance, 3, 1-14.

Carnero M, Koopman S and Ooms M, 2003, Periodic heteroskedastic RegARFIMA models for daily electricity spot prices, Discussion Paper TI 2003–071/4, Tinbergen Institute.

Cartea A and Figueroa M G, 2004, Pricing in electricity markets: A mean reverting jump diffusion model with seasonality, mimeo: Birkbeck College, University of London.

Chan K C, Karolyi G A, Longstaff F A and Sanders A B, 1992, An empirical comparison of alternative models of the short-term interest rate, Journal of Finance, 47, 1209-1227.

Deidersen J and Trück S, 2002, Energy price dynamics: Quantitative studies and stochastic processes, Technical Report TR–ISWM–12/2002, Universität Karlsruhe.

Engle R F, 2001, GARCH 101: The use of ARCH/GARCH models in applied econometrics, Journal of Economic Perspectives, 15, 157-168.

Engle R F, 2002, New frontiers for ARCH models, Journal of Applied Econometrics, 17, 425-446.

Engle R F, 2004, Risk and volatility: Econometric models and financial practice, American Economic Review, 94, 405-420.

Engle R F and Ng V K, 1993, Measuring and testing the impact of news on volatility, Journal of Finance, 48, 1749-1778.

Engle R F and Patton A J, 2001, What good is a volatility model?, Quantitative Finance, 1, 237-245.

Escribano A, Peña J I and Villaplana P, 2002, Modelling electricity prices: International evidence, Economics Series Working Paper 08: 02–27, Universidad Carlos III de Madrid.

Eydeland A and Geman H, 2000, Fundamentals of Electricity Derivatives, in: Energy Modeling and the Management of Uncertainty (London: Risk Publications).

Geman H and Roncoroni A, 2006, Understanding the fine structure of electricity prices, Journal of Business, 79, forthcoming.

Giot P, 2000, Time transformations, intraday data, and volatility models, Journal of Computational Finance, 4, 31-62.

Haldrup N and Nielsen M, 2005, A regime switching long memory model for electricity prices, Journal of Econometrics, forthcoming.

Harris L, 1986, A transaction data study of weekly and intradaily patterns in stock returns, Journal of Financial Economics, 16, 99-117.

Huisman R and Mathieu R, 2003, Regime jumps in electricity prices, Energy Economics, 25, 425-434.

Joskow P L, 1997, Restructuring, competition and regulatory reform in the U.S. electricity sector, Journal of Economic Perspectives, 11, 119-138.

Karakatsani N V and Bunn D W, 2004, Modelling stochastic volatility in high-frequency spot electricity prices, mimeo: Department of Decision Sciences, London Business School.

Knittel C R and Roberts M R, 2004, An empirical examination of restructured electricity prices, mimeo: Department of Economics, University of California–Davis.

Koekebakker S and Ollmar F, 2001, Forward curve dynamics in the Nordic electricity market, Discussion Paper 2001–21, Department of Finance and Management Science, Norwegian School of Economics and Business Administration.

Lamoureux C G and Lastrapes W D, 1990, Heteroskedasticity in stock return data: Volume versus GARCH effects, Journal of Finance, 45, 221-229.

Li W K, Ling S and McAleer M, 2002, Recent theoretical results for time series models with GARCH errors, Journal of Economic Surveys, 16, 245-269.

Longstaff F A and Wang A W, 2004, Electricity forward prices: A high-frequency empirical analysis, Journal of Finance, 59, 1877-1900.

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

Manoliu M and Tompaidis S, 2002, Energy futures prices: Term structure models with Kalman filter estimation, Applied Mathematical Finance, 9, 21-43.

Mork E, 2001, Emergence of financial markets for electricity: A European perspective, Energy Policy, 29, 7-15.

Nelson D B, 1991, Conditional heteroscedasticity in asset returns: A new approach, Econometrica, 59, 347-370.

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

Resta M, 2004, Multifractal analysis of power markets: Some empirical evidence, mimeo: School of Economics (Financial Mathematics), University of Genova.

Resta M and Sciutti D, 2003, Spot price dynamics in deregulated power markets, mimeo: School of Economics (Financial Mathematics), University of Genova.

Routledge B R, Seppi D J and Spatt C S, 2001, The "spark spread": An equilbrium model of cross-commodity price relationship in electricity, mimeo: Graduate School of Industrial Administration, Carnegie Mellon University.

Schwartz E S, 1997, The stochastic behaviour of commodity prices: Implications for valuation and hedging, Journal of Finance 52, 923-973.

Simonsen I, 2003, Measuring anti-correlations in the Nordic electricity spot market by wavelets, Physica A, 322, 597-606.

Vahviläinen I, 2002, Basics of electricity derivatives pricing in competitive markets, Applied Mathematical Finance, 9, 45-60.

Weron R, 2000, Energy price risk management, Physica A, 285, 127-134.

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

Weron R and Przybyłowicz B, 2000, Hurst analysis of electricity price dynamics, Physica A, 283, 462-468.

Wilkinson L and Winsen J, 2002, What we can learn from a statistical analysis of electricity prices in New South Wales, Electricity Journal, April, 60-69.

Wood R A, McInish T H and Ord J K, 1985, An investigation of transaction data for NYSE stocks, Journal of Finance, 25, 723-739.

	Jan	Feb	Mar	Apr	May	Jun
Mon	0.172549	0.153134	0.203704	0.362829	0.410702	0.634206
Tue	0.150761	0.021786	0.012664	0.167280	0.076427	0.346457
Wed	0.093718	0.006296	0.014559	0.010824	0.037801	0.065023
Thu	0.084144	0.018015	0.015025	0.007921	0.280272	0.064096
Fri	0.039971	0.033769	0.017683	0.014563	0.094877	0.067420
Sat	0.102417	0.048805	0.049265	0.048285	0.105214	0.190565
Sun	0.062880	0.034773	0.074696	0.255434	0.142691	0.173772

	Jul	Aug	Sep	Oct	Nov	Dec
Mon	0.544370	**1.301646**	0.358908	0.355290	0.466470	**0.451605**
Tue	**0.122955**	**0.516680**	0.075058	0.017599	0.030139	**0.067990**
Wed	**0.043689**	0.014402	0.031409	0.007924	0.032166	**0.052622**
Thu	**0.115656**	0.052249	0.008192	0.019330	0.017346	**0.146491**
Fri	0.046879	0.029561	0.003446	0.035347	0.039659	0.073754
Sat	0.100539	0.078190	0.109275	0.094008	0.220945	0.081298
Sun	0.125789	0.117001	0.116502	0.077310	0.143587	0.190638

Table 3. Mean value $\varphi_t(i; j)$ of squared log daily returns on day i and month j, where values in bold change following removal of the 8 largest spikes from the spot (returns) series – new values are not reported here.

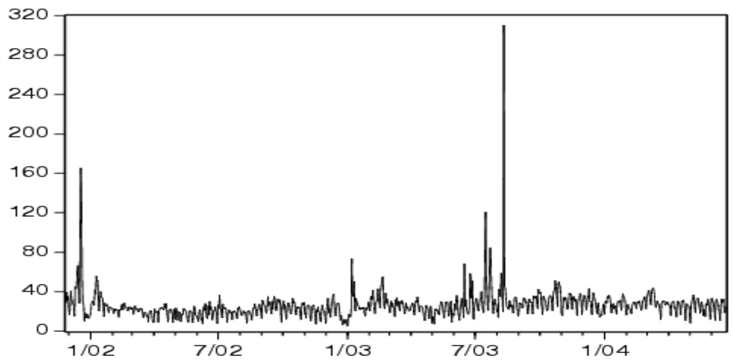

Figure 1: Spot S_t (27.11.01-21.06.04)

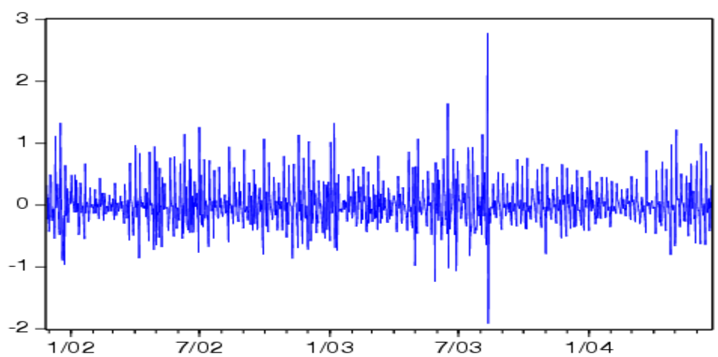

Figure 2: Return R_t (28.11.01-21.06.04)

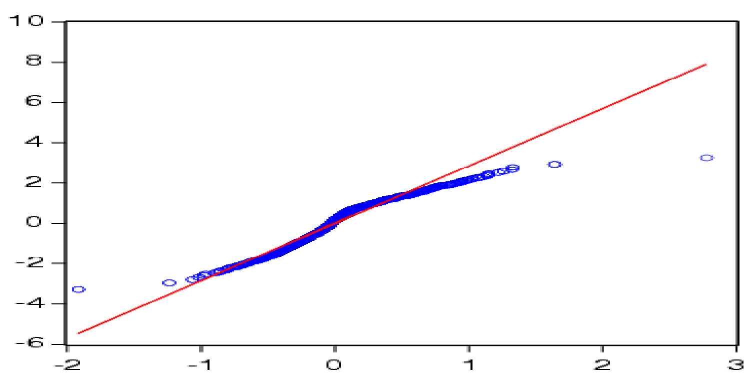

Figure 3: QQ-plot, R_t against normal

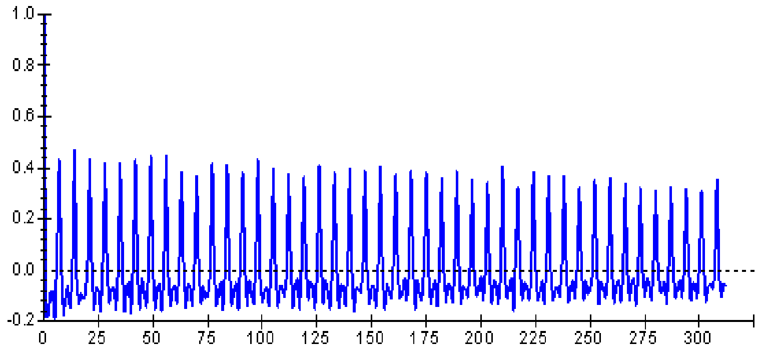

Figure 4: Autocorrelogram of R_t

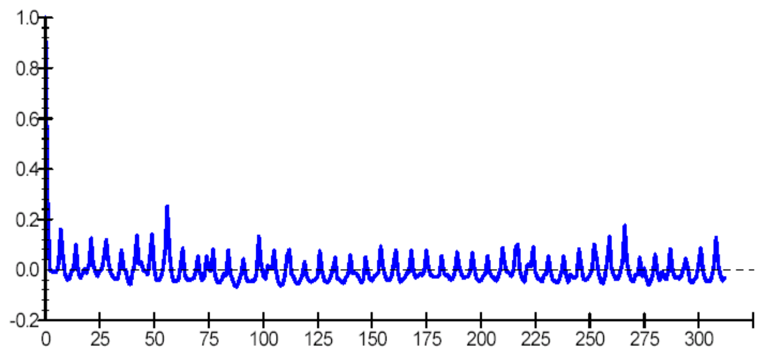

Figure 5: Autocorrelogram of R_t^2

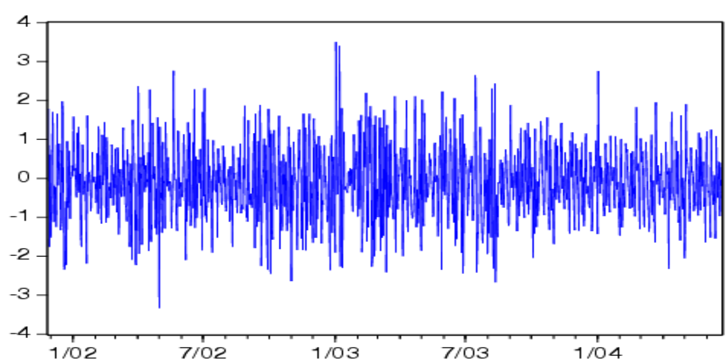

Figure 6: Return $\widetilde{R}_t$ (28.11.01-21.06.04)

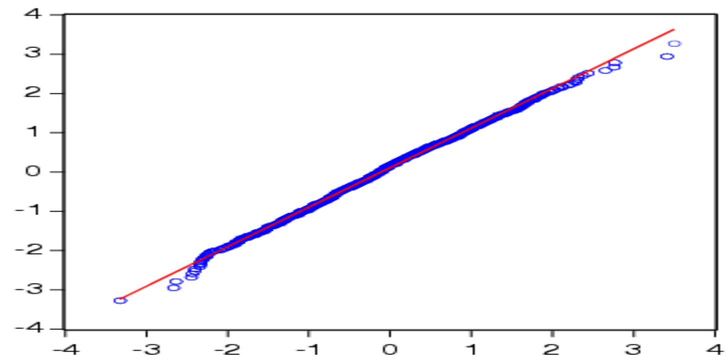

Figure 7: QQ-plot, $\widetilde{R}_t$ against normal

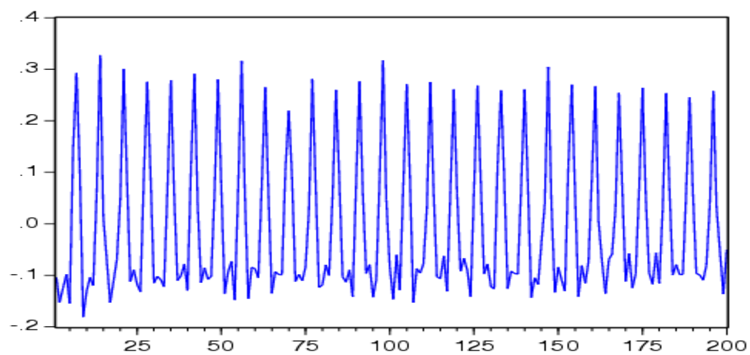

Figure 8: Autocorrelogram of $\widetilde{R}_t$

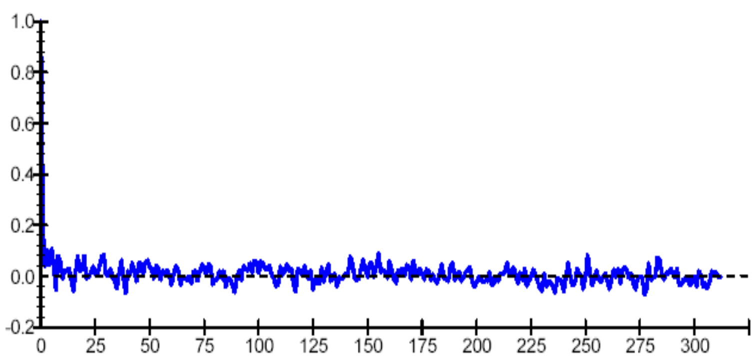

Figure 9: Autocorrelogram of $\widetilde{R}_t^2$

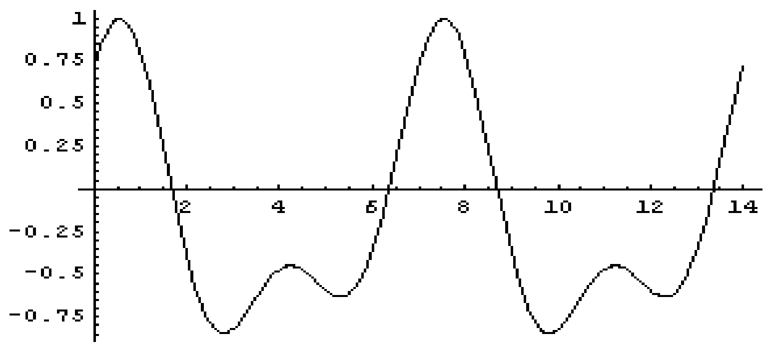

Figure 10: Estimated weekly periodicity in $\widetilde{R}_t$, over 14 days

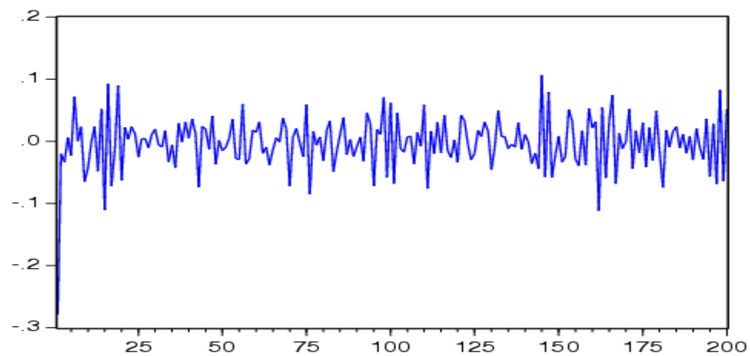

Figure 11: Autocorrelogram of $\widetilde{R}_t$, weekly periodicity removed using sinusoids

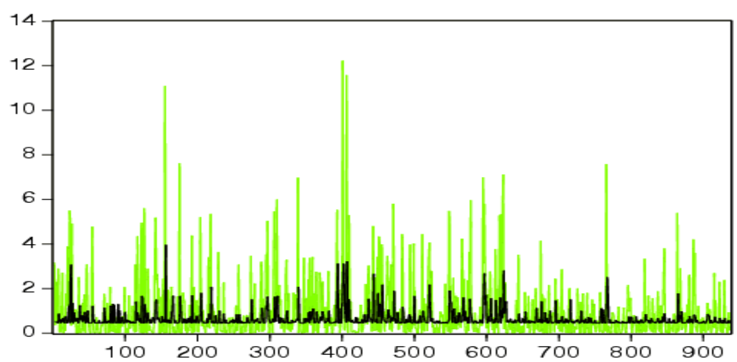

Figure 12: Static AR(2)-ARCH(1) forecast, against $\widetilde{R}_t^2$ (28.11.01-21.06.04)