Department of Economics and Management

DEM Working Paper Series

Long memory and Periodicity in Intraday Volatility

Eduardo Rossi
(Università di Pavia)

Dean Fantazzini
(Moscow School of Economics)

 # 15 (11-12)

November 2012

Long memory and Periodicity in Intraday Volatility [*]

Eduardo Rossi [†]

University of Pavia

Dean Fantazzini[‡]

Moscow School of Economics

October 26, 2012

Abstract

Intraday return volatilities are characterized by the contemporaneous presence of periodicity and long memory. This paper proposes two new parameterizations of the intraday volatility: the Fractionally Integrated Periodic EGARCH and the Seasonal Fractional Integrated Periodic EGARCH, which provide the required flexibility to account for both features. The periodic kurtosis and periodic autocorrelations of power transformations of the absolute returns are computed for both models. The empirical application shows that volatility of the hourly Emini S&P 500 futures returns are characterized by a periodic leverage effect coupled with a statistically significant long-range dependence. An out-of-sample forecasting comparison with alternative models shows that a constrained version of the FI-PEGARCH provides superior forecasts. A simulation experiment is carried out to investigate the effects that sample frequency has on the fractional differencing parameter estimate.

JEL classification: C22, C58, G13.

Keywords: Intraday volatility, Long memory, FI-PEGARCH, SFI-PEGARCH, Periodic models.

1 Introduction

The studies on intraday return volatilities show that the intraday dynamics is characterized by the contemporaneous presence of periodicity and long memory. Andersen and Bollerslev (1997b) observe that conclusions on return volatility and market microstructure variables at the intraday frequencies are likely to be subject to severe distortions due to the strong periodicity in volatility. Therefore, a model for intraday volatility to be used in prediction requires that intraday periodicities and long memory are jointly modeled. Intraday volatility forecasts are an important input for algorithmic trading, and, in general, it is widely recognized that high-frequency volatility forecasting is of substantial practical importance for the finance industry (see Engle and Sokalska, 2012). However, specifying and estimating a GARCH-type model for the conditional variance of intraday return can be a challenging task, see for example Andersen and Bollerslev (1997b).

The main contribution of our article are two new parameterizations of the intraday volatility that combine both long memory and periodicity. The proposed models are flexible enough to account for the presence of different persistent periodic components. Our main purpose is to model and forecast intraday volatility. Moreover, using an asymmetric long memory periodic model allows us to estimate the intraday patterns of the leverage effects (asymmetric effects) and to analyze if the observed periodicities are the outcome of trades originated in different geographical areas.

The extant approaches in the literature can be classified into two groups. In the first one, pioneered by Andersen and Bollerslev (1997a,b, 1998), the volatility process is modeled as the outcome of the simultaneous interaction of a long memory factor and an intraday periodic component. A number of studies have adopted this framework, e.g. Beltratti and Morana (1999), Baillie, Han, Myers, and Song (2007) and Martens, Chang, and Taylor (2002).

The second approach, instead, is based on seasonal fractional differencing. Examples are the long memory stochastic volatility model of Arteche (2004), the Periodic Long-Memory GARCH model of Bordignon, Caporin, and Lisi (2009) and the Gegenbauer-GARCH model of Bordignon, Caporin, and Lisi (2007).

The models presented here are closer to the second approach. The *Fractionally Integrated Periodic EGARCH* (FI-PEGARCH), an extension of the Fractionally Integrated EGARCH of

Baillie, Bollerslev, and Mikkelsen (1996), Bollerslev and Mikkelsen (1996), is periodic with long memory in each period, while the *Seasonally Fractionally Integrated Periodic EGARCH* (SFI-PEGARCH), is periodic with seasonal long memory. In both models the parameters that determine the short-run dependence are periodic, which in turn implies that the return kurtosis and the autocorrelations of power transformations of absolute returns are periodic. Their flexibility allows us to accommodate different intraday volatility patterns as long as different long range dependencies. Moreover, the FI-PEGARCH and SFI-PEGARCH nest alternative models. For example, restricting the parameters in the FI-PEGARCH to be constant across the periods we obtain the Periodic Long Memory EGARCH (PLM-EGARCH) model of Bordignon, Caporin, and Lisi (2009). The flexibility in modeling periodicity and long memory, provided by the FI-PEGARCH and SFI-PEGARCH models, is particularly important for the estimate of the fractional differencing parameter. The empirical results, corroborated by a Monte Carlo experiment, show that, with high-frequency data the modeling of periodicity has an impact on the estimate of long memory, which is null when using daily data. In fact, in simulations, the fractional integration parameter estimate is close to the true one when daily data are employed to estimate a long memory volatility model. Instead, when high-frequency data are employed, underspecifying the periodicity produces a biased estimate in finite samples.

We assume that the periodicity is known but allowed to vary, like in the Periodic GARCH models of Bollerslev and Ghysels (1996). Periodic models have found application in the stochastic volatility framework too. Tsiakas (2006) studies the day-of-the-week, the non-trading day, and the month-of-the-year seasonal effects in daily S&P500 returns and in volatility adopting a periodic stochastic volatility model.

Not differently from what other studies have shown, we find that squared and absolute returns of the E-mini SP500 futures contracts quoted at the Chicago Mercantile Exchange are characterized by long memory and periodicity. The empirical results illustrate that during the Asian and European trading time the volatility is much lower than during the American trading time when we observe a sharp increase. Furthermore, our estimates show that the leverage effect varies across the day. Meaningfully, it turns out that long memory estimates obtained in a non-periodic long memory model are greater than those obtained with FI-PEGARCH and SFI-PEGARCH models which include both periodicity and long memory. This confirms that the long memory estimation with high-frequency data is strictly connected to the periodicity

specification. Restricting the latter increases the former, and viceversa. The FI-PEGARCH with a constrained specification of the parameter periodicity turns out to provide the best out-of-sample forecasting performance among a wide range of competing models. This model can well represent the best compromise between model complexity and parsimony.

The remainder of the paper is organized as follows. Section 2 illustrates the FI-PEGARCH and the SFI-PEGARCH models and their features. Section 3 presents the data and the preliminary analysis. Section 4 reports the empirical results of estimation, while Section 5 discusses the in-sample and out-of-sample forecasting performances. In Section 6, we present some simulation results on the estimation of the long memory in parametric volatility models with high-frequency generated data. Section 7 concludes.

2   Long memory in periodic models

In this section we illustrate two models for the logarithm of the conditional variance that are characterized by long memory and periodicity: the FI-PEGARCH and the SFI-PEGARCH, both can be thought of as a generalization of the PEGARCH, which is illustrated in the next subsection. In the FI-PEGARCH the log conditional variance has nonseasonal and nonperiodic long memory, while the SFI-PEGARCH is characterized by seasonal long memory.

2.1   The Periodic EGARCH model

Periodic volatility models are a straightforward representation of the seasonal patterns observed in intraday volatility. The Periodic GARCH introduced by Bollerslev and Ghysels (1996) and used for the analysis of periodicity in volatilities by Franses and Paap (2000), Taylor (2004), among others, can be considered as the first example of a periodic volatility model. An analogous extension of the stochastic volatility model can be found in Tsiakas (2006). Let P_t be the asset price at time t, and by y_t the percent continuously compounded return $y_t = 100 \times [\ln(P_t) - \ln(P_{t-1})]$, where y_t is observed S times intradaily, for a total number of observations which is T. The Periodic EGARCH(p, q) process (Bollerslev and Ghysels, 1996) is a time varying coefficient model for the log-conditional variance of the returns:

$$y_t = \eta_t \sqrt{h_t} \qquad t = 1, \dots, T \tag{1}$$

$$h_t = Var[y_t | \Phi_{t-1}^s] \tag{2}$$

where $\eta_t \sim i.i.d.(0,1)$. Φ_{t-1}^s is a modified Borel σ-field filtration in which the Borel σ-field filtration based on the realization of the $\{y_t\}$ process up to time $t-1$ is augmented by a process defining the stage of the periodic cycle at each point in time. The log-conditional variance process is modeled as:

$$\ln h_t = \omega_s + \sum_{i=1}^p \delta_{is} \ln h_{t-i} + g_s(\eta_{t-1}) + \sum_{j=1}^q \alpha_{js} g_s(\eta_{t-1-j}) \ \ t=1,\dots,T \ \ s=1,\dots,S \quad (3)$$

or using the lag operator (L):

$$(1 - \delta_s(L)) \ln h_t = \omega_s + \alpha_s(L) g_s(\eta_{t-1}), \quad (4)$$

where

$$
\begin{aligned}
g_s(\eta_t) &= \psi_s \left[|\eta_t| - E\left(|\eta_t|\right) \right] + \gamma_s \eta_t \\
\alpha_s(L) &= 1 + \alpha_{1s}L + \dots + \alpha_{qs}L^q \\
\delta_s(L) &= \delta_{1s}L + \dots + \delta_{ps}L^p.
\end{aligned}
$$

2.2 The FI-PEGARCH model

When we assume that each period is long-range dependent, that is $(1 - \delta_s(L)) = (1 - \beta_s(L))(1 - L)^d$ we obtain the FI-PEGARCH(p,d,q) model:

$$(1 - \beta_s(L))(1 - L)^d (\ln h_t - \omega_s) = \alpha_s(L) g_s(\eta_{t-1}) \quad (5)$$

with the roots of $(1 - \beta_s(z)) = 0$ strictly outside the unit circle, which excludes the possibility of periodic integration. All long memory properties of the model are captured in $(1 - L)^d$. The process is, by analogy to the ARFIMA models, covariance stationary and invertible for $-1/2 \le d \le 1/2$, and strictly stationary and ergodic for $d < 1/2$ (from Theorem 2.1 in Nelson, 1991), see also Bollerslev and Mikkelsen (1996). In the FI-PEGARCH model the long memory component is applied at the non-periodic frequency, it is a periodic process which features nonseasonal and nonperiodic long memory. The parameter d determines the long-term behavior, whereas the parameters in ω_s, $\beta_s(L)$, $g_s(\eta_t)$ and $\alpha_s(L)$ allow for periodic modeling of short-range properties. With $d = 0$ the FI-PEGARCH becomes a PEGARCH model where there is no long-

range dependence and, as we have seen above, the short-range dependence is determined by the periodically varying parameters. When $d > 0$ the process is long memory. When the parameters are constant, i.e. $\omega_s = \omega, \psi_s = \psi, \gamma_s = \gamma$, and $\beta_s = \beta, \forall s$, the FI-PEGARCH becomes a FIEGARCH model. It is particularly useful, especially in the estimation process, to express the FI-PEGARCH(p, d, q) process in infinite ARCH form,

$$\ln h_t = \omega_s + (1 - \beta_s(L))^{-1}(1 - L)^{-d}\alpha_s(L)g_s(\eta_{t-1}). \tag{6}$$

For instance, the infinite ARCH representation of the FI-PEGARCH(1,d,1) with $-1/2 < d < 1/2$ is

$$
\begin{aligned}
\ln h_t &= \omega_s + (1 - \beta_{1s}L)^{-1}(1 - L)^{-d}(1 + \alpha_{1s}L)g_s(\eta_{t-1}) \\
&= \omega_s + \lambda_s(L)(1 + \alpha_{1s}L)g_s(\eta_{t-1})
\end{aligned} \tag{7}
$$

where

$$(1 - \beta_{1s}L)^{-1}(1 - L)^{-d} = 1 + \sum_{i=1}^{\infty}\lambda_{i,s}L^i = \lambda_s(L) \tag{8}$$

with λ's given by

$$
\begin{aligned}
\lambda_{0,s} &= 1 \\
\lambda_{i,s} &= \sum_{j=0}^{i-1}\frac{\Gamma(j+d)}{\Gamma(j+1)\Gamma(d)}\beta_s^{i-j-1}
\end{aligned} \tag{9}
$$

where $\Gamma(\cdot)$ is the Gamma function. Finally, the fractional filter can be modeled as periodic, i.e., $(1 - L)^{d_j}, j = 1, \ldots, S$, as in the Seasonal periodic ARFIMA of Koopman, Ooms, and Carnero (2007).

2.3 The SFI-PEGARCH model

If we consider the possibility that each regime is seasonally persistent, that is $(1 - \delta_s(L)) = (1 - \beta_s(L))(1 - L^S)^d$, (see Ooms and Franses, 2001, for a seasonal periodic long memory conditional mean model), we have the SFI-PEGARCH(p, d, q):

$$(1 - \beta_s(L))(1 - L^S)^d(\ln h_t - \omega_s) = \alpha_s(L)g_s(\eta_{t-1}) \qquad t = 1, 2, \ldots, T \tag{10}$$

As for FI-PEGARCH model the process is covariance stationary and invertible for $-1/2 \leq d \leq 1/2$, and strictly stationary and ergodic for $d < 1/2$, with $d > 0$ the process is characterized by long memory. The SFI-PEGARCH model is a periodic process with a seasonal long-range dependence. The long memory component is applied at the seasonal frequencies, e.g., $(1 - L^S)^d$, together with periodic modeling of short-range features. As in the case of FI-PEGARCH, the PEGARCH model is nested in the SFI-PEGARCH. When the parameters are restricted to be constant across the periods, the SFI-PEGARCH reduces to the PLM-EGARCH of Bordignon, Caporin, and Lisi (2009). The infinite ARCH representation of SFI-PEGARCH(p, d, q) is given by:

$$\ln h_t = \omega_s + (1 - \beta_s(L))^{-1}(1 - L^S)^{-d}\alpha_s(L)g_s(\eta_{t-1}) \qquad t = 1, 2, \ldots, T \qquad (11)$$

The inverse of the fractional differential operator is $(1 - L^S)^{-d} = 1 + \sum_{k=1}^{\infty} \pi_k L^{kS}$ with $\pi_k \equiv \prod_{j=1}^{k} \frac{j-1+d}{j}$. For instance the infinite ARCH representation of the SFI-PEGARCH$(1, d, 1)$ is:

$$\ln h_t = \omega_s + (1 - \beta_{1s}L)^{-1}(1 - L^S)^{-d}(1 + \alpha_{1s}L)g_s(\eta_{t-1})$$

A compact expression for $(1 - \beta_{1s}L)^{-1}(1 - L^S)^{-d}$ is given by

$$(1 - \beta_{1s}L)^{-1}(1 - L^S)^{-d} = 1 + \beta_{1s}L + \ldots + (\beta_{1s}^S + \pi_1)L^S + (\beta_{1s}\pi_1 + \beta_{1s}^{S+1})L^{S+1} + \ldots$$
$$= 1 + \sum_{k=1}^{\infty} \lambda_{k,s}L^k = \lambda_s(L).$$

where the coefficients of the polynomial $\lambda_s(z)$, are defined as follows

$$\lambda_{i \cdot S + j, s} = \begin{cases} \beta_{1s}^{iS+j} + \beta_{1s}^{(i-1)S+j}\pi_1 + \ldots + \beta_{1s}^{S+j}\pi_{i-1} & i = 1, 2, \ldots \quad j = \ldots, -2, -1, 1, 2, \ldots \\ \beta_{1s}^{iS} + \beta_{1s}^{(i-1)S}\pi_1 + \ldots + \beta_{1s}^{S}\pi_{i-1} + \pi_i & i = 1, 2, \ldots \quad j = 0 \end{cases}$$
$$(12)$$

with $k = i \cdot S + j$. From (12) is clear that whenever $s = S$ the λ_k coefficient is increased by π_S, with $\pi_S > 0$, when $d > 0$.

The coefficients in the polynomial $\lambda(z)$ for the FI-PEGARCH$(1, d, 0)$ show a smooth decaying, while, as expected, those for the SFI-PEGARCH$(1,d,0)$ are characterized by a jagging pattern, see Figure 1. In the next subsection we show that FI and SFI-PEGARCH models induce periodic kurtosis and periodic autocorrelations of absolute and squared returns. These

properties are particularly important because, as it is shown in Section 3, the hourly Emini returns feature periodic sample kurtosis and autocorrelations.

We want to remark that the number of periods S does not necessarily coincide with the sampling frequency, that is the numbers of returns observed during the day: for example, we can have $S = 24$ intraday periods with returns sampled every 5 minutes or with higher frequency. This flexibility avoids that the number of parameters necessarily increases with the sampling frequency. Furthermore, when S is large, the number of periodic parameters can be limited assuming that their seasonal variation is described by periodic functions like combinations of cosine and sinus, as we will see in more details in Section 4.2.
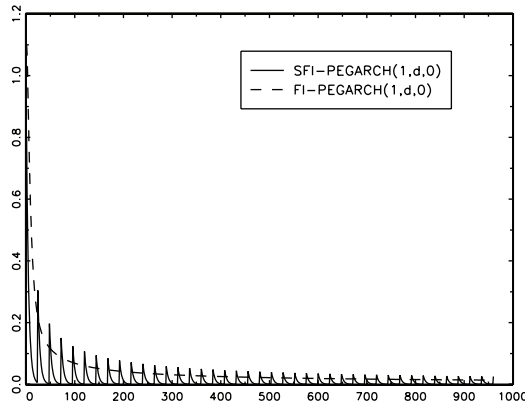

Figure 1: Series expansion of $(1 - L)^{-d}(1 - \beta_1 L)^{-1}$ and $(1 - L^S)^{-d}(1 - \beta_1 L)^{-1}$ with $d = 0.3$, $\beta_1 = 0.8$

2.4 The periodic kurtosis

He, Teräsvirta, and Malmsten (2002) derived the unconditional kurtosis and the autocorrelations of absolute and squared observations for the short memory first order EGARCH process, Ruiz and Vega (2008) extended the analysis to the FIEGARCH model. From Theorem 2 in He, Teräsvirta, and Malmsten (2002, p.870), we can compute the unconditional periodic kurtosis for season s.

Proposition 2.1 *Consider the FI-PEGARCH(1,d,0) and the SFI-PEGARCH(1,d,0) process and assume that the fourth moment $\nu_4 \equiv E(\eta_t^4) < \infty$, that $E[e^{g_s(\eta_t)}]$ and $E[e^{2g_s(\eta_t)}]$ exist and that $|\beta_{1s}| < 1, \forall s$ holds. Then the kurtosis of y_t for period s exists and is given by:*

$$k_{y,s} = \nu_4 \frac{\prod_{i=0}^{\infty} E\left[\exp\{2\lambda_{i,s}g_s(\eta_t)\}\right]}{\left\{\prod_{i=0}^{\infty} E\left[\exp\{\lambda_{i,s}g_s(\eta_t)\}\right]\right\}^2} \quad s = 1,\ldots,S \quad (13)$$

7

where $\lambda_{i,s}$ is defined in (9) and (12) for the FI-PEGARCH(1,d,0) and the SFI-PEGARCH(1,d,0) respectively.

Proof. See Appendix A.

It is evident from the expression (13) that for both models the kurtosis results to be periodic. In Figure 2, the kurtosis of the FIPEGARCH$(1, d, 0)$ and SFI-PEGARCH$(1, d, 0)$ are displayed as a function of each parameter that enters in expression (13). The kurtosis of the FI-PEGARCH is slightly larger than that of the SFI-PEGARCH for each level of the parameters. The tail thickness of the unconditional distribution of the process increases with the short-memory parameter β_{1s} as well as with the parameters ψ_s which capture the effect that large shocks have on the log-conditional variance (i.e., *magnitude effect*, see Nelson, 1991). On the contrary, the effect of the parameters γ_s (i.e., *sign or leverage effect*) is negligible. Moreover, larger values for the long memory parameter d (in the covariance stationarity interval $0 < d < 1/2$) do not affect significantly the kurtosis of both processes.
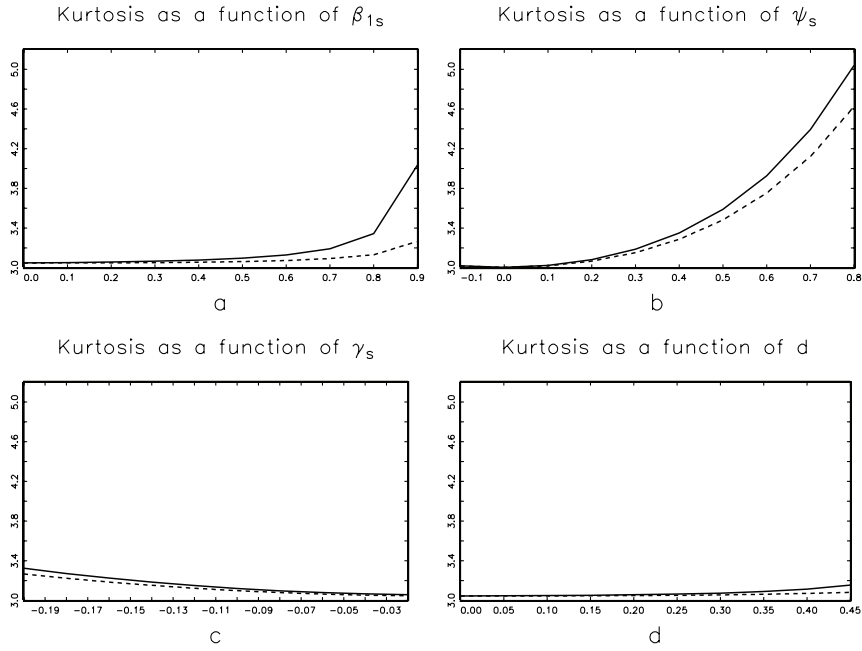

Figure 2: Kurtosis as a function of model parameters of FI-PEGARCH$(1,d,0)$ (solid line) and SFI-PEGARCH$(1,d,0)$ (dashed line). Each graph is obtained considering the following parameters values $\beta_{1s} = 0.2710$, $\psi_s = 0.1748$, $\gamma_s = -0.0337$, $d = 0.2441$

2.5 The periodic autocorrelation function of $|y_t|^{2m}$

It is possible to derive the periodic autocorrelation function of $|y_t|^{2m}$ for season s by applying Proposition 2.1 and a recursion similar to that of Theorem 3 (p.872) in He, Teräsvirta, and Malmsten (2002).

Proposition 2.2 *Consider the FI-PEGARCH(1,d,0) or the SFI-PEGARCH(1,d,0) and assume that $\nu_{4m} \equiv E(\eta_t^{4m}) < \infty$, $E[|\eta_t|^{2m} e^{mg_s(\eta_t)}] < \infty$ and $E[e^{2mg_s(\eta_t)}] < \infty$ for $0 < m < \infty$, and that $|\beta_{1s}| < 1, \forall s$ holds. Then the autocorrelation function of $|y_t|^{2m}$ for season s and lag n equals*

$$\rho_{m,s}(n) = \frac{E[|y_t|^{2m}|y_{t-n}|^{2m}] - [E(|y_t|^{2m})][E(|y_{t-n}|^{2m})]}{(E[|y_t|^{4m}] - [E(|y_t|^{2m})]^2)^{1/2} (E[|y_{t-n}|^{4m}] - [E(|y_{t-n}|^{2m})]^2)^{1/2}} \quad for \quad n \geq 1, \quad (14)$$

where the numerator is

$$E[|y_t|^{2m}|y_{t-n}|^{2m}] - [E(|y_t|^{2m})][E(|y_{t-n}|^{2m})] = \nu_{2m} \exp\{m[(\omega_r + \omega_s) + \beta_{1s}^n(\omega_r - \omega_s)]\} \times$$

$$E\left[\eta_t^{2m} \exp\{m\chi_{n-1,s}g_s(\eta_t)\}\right] \prod_{i=0}^{n-2} E\left[\exp\{m\chi_{i,s}g_s(\eta_t)\}\right] \times$$

$$\prod_{i=n}^{\infty} E\left[\exp\{m\chi_{i,s}g_s(\eta_t)\}\right] \prod_{i=1}^{\infty} E\left[\exp\{m(1+\beta_{1s}^n)\lambda_{i,r}g_r(\eta_t)\}\right] -$$

$$\nu_{2m}^2 \exp\{m(\omega_s + \omega_r)\} \left(\prod_{i=0}^{\infty} E\left[\exp\{m\lambda_{i,s}g_s(\eta_t)\}\right]\right) \times$$

$$\left(\prod_{i=0}^{\infty} E\left[\exp\{m\lambda_{i,r}g_r(\eta_t)\}\right]\right) \quad (15)$$

s denotes the season in time t, whereas r denotes the season which occurs at $t - n$. If $n = 1$ then $\prod_{i=0}^{n-2} E\left[\exp\{m\chi_{i,s}g_s(\eta_t)\}\right] = 1$. The denominator is

$$\left(E[y_t^{4m}] - [E(y_t^{2m})]^2\right)^{1/2} \left(E[y_{t-n}^{4m}] - [E(y_{t-n}^{2m})]^2\right)^{1/2} =$$

$$\exp\{m(\omega_r + \omega_s)\} \left[\nu_{4m} \prod_{i=1}^{\infty} E\left[\exp\{2m\lambda_{i,s}g_s(\eta_t)\}\right] - \nu_{2m}^2 \left(\prod_{i=0}^{\infty} E\left[\exp\{m\lambda_{i,s}g_s(\eta_t)\}\right]\right)^2\right]^{1/2} \times$$

$$\left[\nu_{4m} \prod_{i=1}^{\infty} E\left[\exp\{2m\lambda_{i,r}g_r(\eta_t)\}\right] - \nu_{2m}^2 \left(\prod_{i=0}^{\infty} E\left[\exp\{m\lambda_{i,r}g_r(\eta_t)\}\right]\right)^2\right]^{1/2} \quad (16)$$

with $\lambda_s(L) = (1-L)^{-d}(1-\beta_{1s}L)^{-1}$ and $\chi_s(L) = (1-L)^{-d} \sum_{i=1}^{n} \beta_{1s}^{i-1}L^{i-1}$ for the FI-PEGARCH(1,d,0)

9

and $\lambda_s(L) = (1 - L^S)^{-d}(1 - \beta_{1s}L)^{-1}$, and $\chi_s(L) = (1 - L^S)^{-d} \sum_{i=1}^{n} \beta_{1s}^{i-1} L^{i-1}$ for the SFI-PEGARCH(1,d,0).

Proof. See Appendix A.

Combining (15) and (16) we obtain the expression of the periodic autocorrelation function of $|y_t|^{2m}$. It is clear that the periodic autocorrelation varies with s. This is consistent with the behavior of sample periodic autocorrelograms discussed in Section 3. Using theorem A1.1 in Nelson (1991) we can further simplify the previous expressions for the kurtosis and the periodic autocorrelation function in case of normal standardized innovation, see Appendix A for more details.

3   The data

The data set consists of transaction prices of the E-mini SP500 stock index futures, recorded every 60 minutes on the electronic market Chicago Mercantile Exchange, from March 14, 2004 through January 25, 2008, comprising 23,952 observations (data from Computer Information Systems 101 Holly Ridge Monroe, LA 71203 U.S.A.). This future contract, besides being the most important stock index futures, has the characteristic of being quoted almost 24 hours a day: to be more precise, there is normal "pit" trading from 9.30 to 16.15 - New York time - (named *Regular Trading Hour, RTH*), and an electronic session from 16.45 until 9.15 the day after (*Globex*). The Monday morning session in Asian markets starts at 17.00 NY time, that corresponds to Sunday afternoon in the US. The index contains the most liquid stocks from the corresponding market and hence the problem of spurious autocorrelation induced by non-synchronous trading should not arise. The prices of the futures contracts used have been built with the futures contracts closest to expiration, while missing data were substituted with the last available price, see e.g. Wasserfallen and Zimmermann (1985), Dacorogna, Gencay, Muller, Olsen, and Pictet (2001) and Shephard and Sheppard (2010). The analysis uses continuously compounded hourly percent returns, y_t. The sampling frequency is chosen to avoid any problem linked to the market liquidity. With periodic models the choice of the sampling frequency is relevant for the model specification. Thus an increase in the sampling frequency is reflected in a dramatic increase in the number of parameters. However, as stressed before, periodic models can be restricted such that the parameter space dimension does not explode. Moreover, employing

hourly returns hides the seasonalities which are present in returns sampled at higher frequency, e.g., 5-minutes, but this does not invalidate the analysis of periodic patterns which are present in hourly sampled returns. The first 23,451 observations are used for the estimation while the last 500 are employed for the out-of-sample forecasting exercise. The main descriptive statistics are plotted in Figure 3. The sample kurtosis per hour varies across the day, and the night time
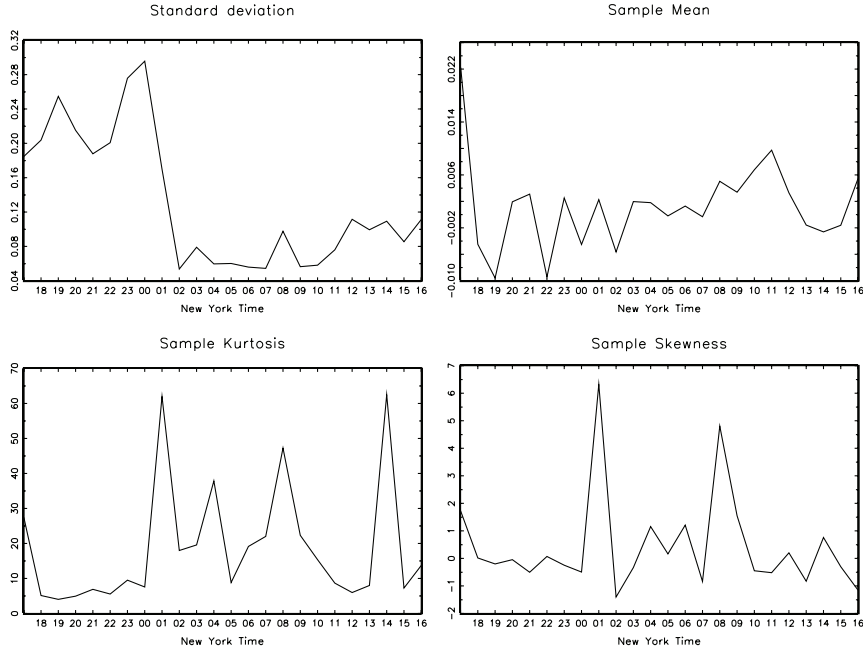

Figure 3: Descriptive statistics for E-mini SP500 continuously compounded hourly percent returns (Eastern Standard Time, E.S.T.).

returns are characterized by extremely leptokurtic distributions. There is also some evidence of asymmetry. This confirms that the E-mini stock index futures SP500 hourly returns are not normally distributed. Moreover, the series of Emini-SP500 returns exhibits an increase in the fraction of zero returns in the night time.

Periodicity in returns. The sample autocorrelation function (ACF) of y_t (not reported) does not display any periodic or seasonal pattern, such as those evidenced in Figure 4 for the squared and absolute returns. To specify the conditional mean of y_t, we adopt a simple-to-general strategy, as suggested by Franses and Paap (2004). First, we run an auxiliary regression, with the fitted residuals $(\widehat{\epsilon}_t)$ of an AR(1) model for y_t, to test the presence of periodicity:

$$\widehat{\epsilon}_t = \sum_{s=1}^{24} \delta_s D_{s,t} + \gamma_1 y_{t-1} + \sum_{s=1}^{24} \zeta_s D_{s,t} \widehat{\epsilon}_{t-1} + u_t \quad s = 1, \dots, S \quad t = 2, \dots, T$$

11

where $D_{s,t}$ are seasonal dummy variables. The asymptotic F-test does not reject the null $\zeta_1 = \ldots = \zeta_{24} = 0$ (F is 1.382, with a p-value of 0.1), namely no periodic autocorrelation of order 1. However, given that the AR(1) component in the return regression is insignificant, we also test for periodic intercepts:

$$y_t = c_s + \epsilon_t \quad s = 1, \ldots, S \quad t = 1, \ldots, T$$

and accept the null, e.g., $c_s = c$ for $s = 1, \ldots, S$ (LR test statistic is 27.880 with p-value 0.22). Further, to analyze the presence of day-of-the-week effect in the returns, we simply run a regression of y_t on a constant and six day-of-the-week dummies. The F-test for the joint nullity of the dummies is equal to 1.239 with a p-value of 0.287, which suggests the absence of any lower frequency seasonality in the returns. Thus, in the following analysis we do not make any attempt to correct for the lower frequency interday patterns.

Periodicity in squared returns. In Section 2.5 we have shown that the ACF of the powers of absolute returns, generated by both FI and SFI-PEGARCH processes, are periodically varying. When periodic patterns are present in the data, like in the case of squared and absolute hourly returns, a simple though useful tool of analysis is represented by the periodic autocorrelogram, i.e. an ACF for each hour of the day (see McLeod, 1994 for a definition of the sample periodic ACF). The sample periodic ACF of squared hourly returns, displayed in Figure 5, are substantially different from hour to hour. For instance, the intervals 23-24, 24-1 and 5-6, during the Asia-Europe trading time, display a low autocorrelation when compared to the others. This stylized fact can hardly be accounted for by non-periodic volatility models. In Section 4, we provide some empirical evidence on the performances of FI and SFI-PEGARCH models, whose population periodic ACFs are known in closed form, in replicating the pattern observed in sample periodic autocorrelations.

Long-range dependence in squared returns. Along with periodicity, the long memory (or long-range dependence) is a distinctive feature of squared and absolute hourly returns. The autocorrelograms of squared and absolute hourly returns in Figure 4 clearly show the presence of a strong persistent periodic behavior. In particular, they seem to be characterized by a decay at a slow mean-reverting hyperbolic rate, which is a characteristic of long-range dependence. The same pattern is evident in the sample periodic ACF in Figure 5. This motivates the adoption of
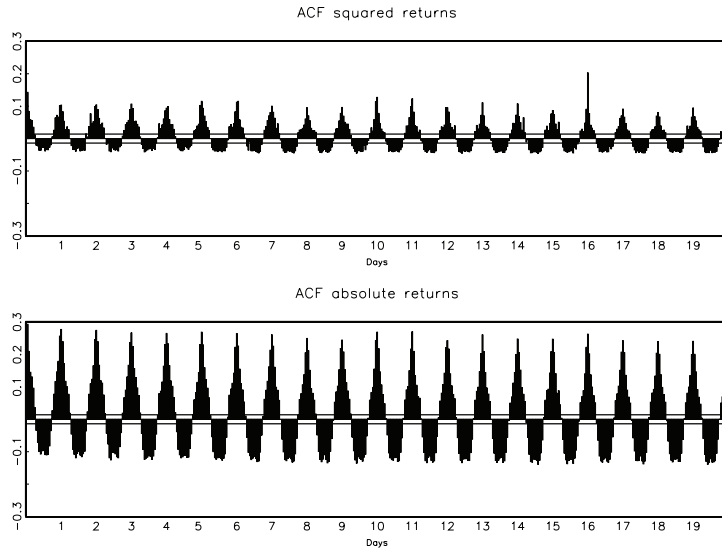
12

Figure 4: Autocorrelograms of squared and absolute Emini-SP500 continuously compounded hourly percent returns (from March 14, 2004 through January 25, 2008)

long memory periodic models in the study of intraday volatilities. The presence of long memory in the squared hourly returns at frequency zero is confirmed by the Gaussian semiparametric estimate, with a bandwidth m of 150, of the long memory parameter d, that is equal to 0.346 with a standard error of 0.041.[1] However, the spectrum estimates (obtained as the smoothed sample periodogram based on a Bartlett's window equal to 200) of absolute and squared hourly returns, displayed in Figure 6, have a peak at the daily frequency of $\pi/12$; this finding has previously been documented by Andersen and Bollerslev (1997a). This evidence suggests that long memory can be present at frequency $\pi/12$, too. We estimate the degree of long memory at $\pi/12$ by means of the Gaussian semiparametric estimator proposed by Arteche and Robinson (2000), which requires that the spectral density satisfies mild technical conditions (see Arteche and Robinson, 2000 and Arteche, 2004 for more details). The estimate of d at frequency $\pi/12$, with $m = 150$, is 0.269 with a standard error of 0.041, confirming the presence of long memory beyond frequency zero.

4   Estimation results

From the preliminary analysis the hourly return on Emini SP500 is modeled as

$$y_t = c + \eta_t \sqrt{h_t}, \quad \eta_t \sim i.i.d.N(0,1) \quad t = 1, \dots, T$$

[1]For a description of the Gaussian semiparametric estimator see Section 4.4.3 in Palma (2007).
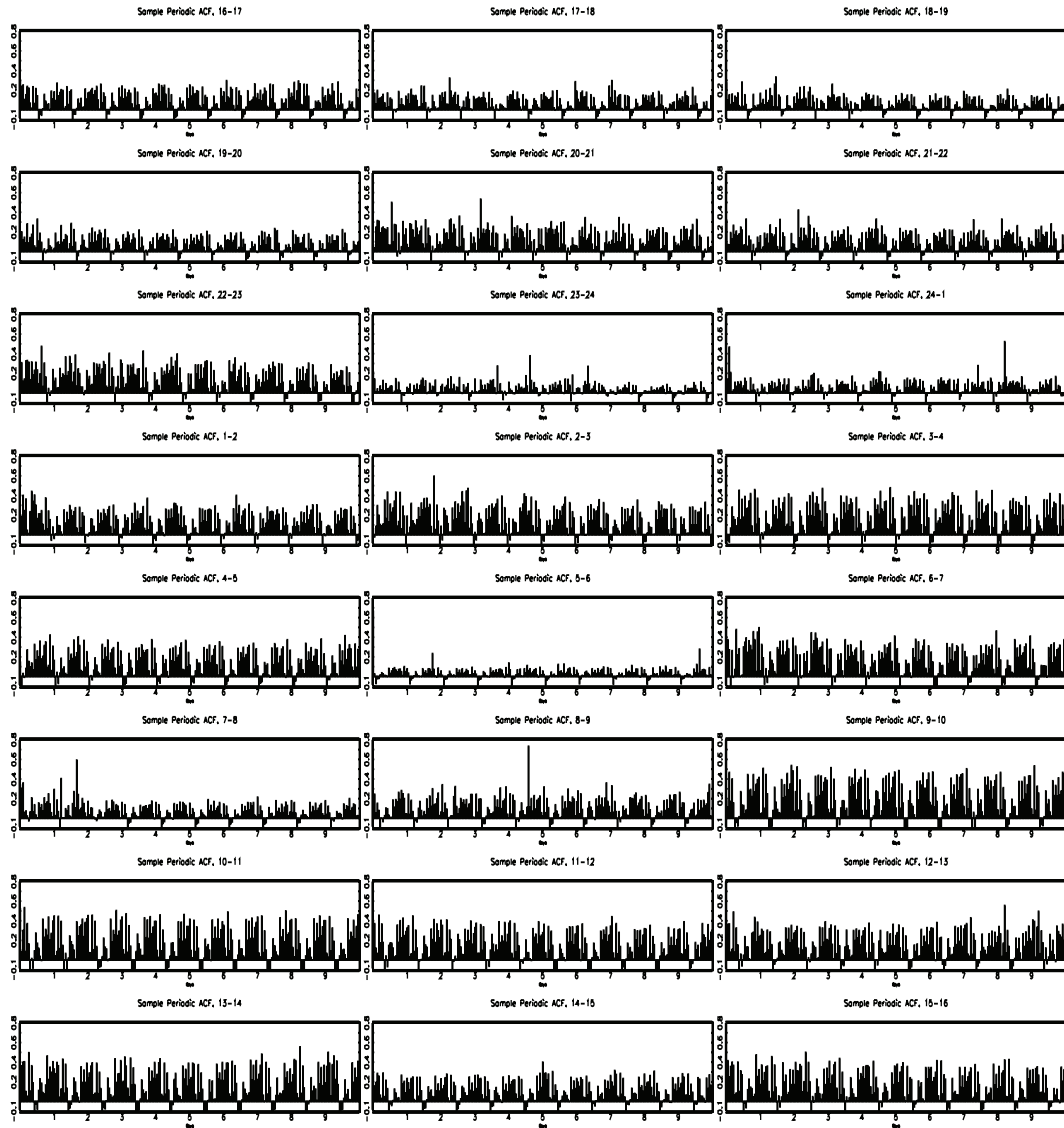
Figure 5: Periodic autocorrelograms of squared Emini-SP500 continuously compounded hourly percent returns (from March 14, 2004 through January 25, 2008).
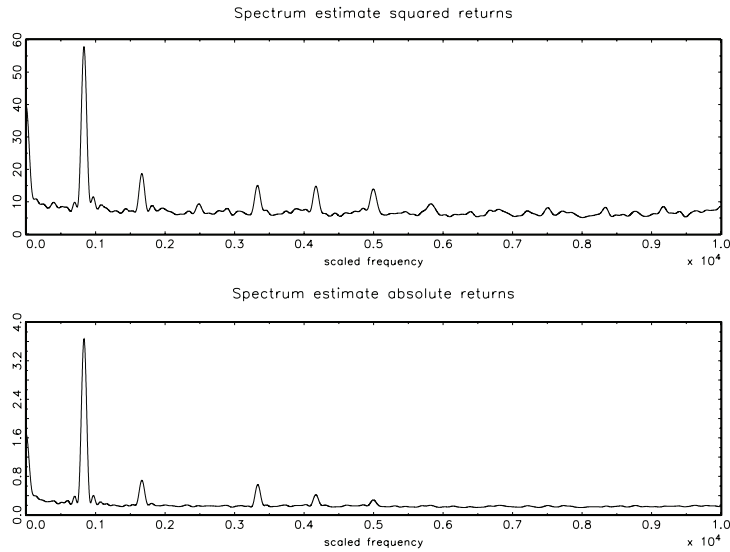
Figure 6: Bartlett estimate of the spectrum of squared and absolute Emini-SP500 continuously compounded hourly percent returns. The estimate is obtained using a bandwidth of 200.

with $y_t = 1000\Delta \ln P_t$. The model is estimated considering alternative specifications for h_t, as it will be shown in the following sections. The model parameters θ are estimated by maximizing the conditional Gaussian log-likelihood function. The computation of the conditional log-likelihood necessitates the truncation of the infinite lag polynomials in (6) and (11). One well known problem in the estimation of long-memory ARCH models is the treatment of initial conditions required to start up the recursions for the conditional variance function. Baillie, Bollerslev, and Mikkelsen (1996) observe that truncating at too low a lag may destroy important long-run dependencies. They adopt, in a simulation study with $1,500$ and $3,000$ observations, a truncation lag of $1,000$. Since there are no theoretical criteria to choose the truncation lag, the models were estimated considering different truncations. These experiments suggested that to retain the long-run dependencies in the conditional variance process, a truncation lag at $1,008$, which is equal to 42 days, can be considered a reasonable choice given the sample size. We use as starting values in the recursion for the calculation of $\log h_t$ the first $1,008$ standardized residuals, $\widehat{\eta}_t$, where the conditional variances are replaced with the sample variance. Thus in the estimation $22,442$ observations were used.

In general, when the assumption of conditional normality is violated, asymptotically valid inference can be carried out on the basis of the results of the theory of quasi-maximum likelihood estimation (QMLE) (Bollerslev and Wooldridge, 1992). Unfortunately, there is a lack of theoretical results on the properties of the QMLE method when used with FIGARCH (and FIE-

GARCH) model, and this is also true for FI and SFI-PEGARCH models. The interested reader is referred to Baillie, Bollerslev, and Mikkelsen (1996) for the discussion of this point. Compared to the sample sizes adopted in the simulation experiments of Baillie, Bollerslev, and Mikkelsen (1996) our data set is fifteen time longer. The second concern is related to the relatively short span of calendar time, e.g., three years and a half, of the hourly Emini returns data set. However, as Andersen and Bollerslev (1997a, p.987) showed "the mechanism responsible for the fractional integration in volatility is generic to the returns process, and thus ever present, so that with high frequency data it may be feasible to identify the manifestation of the phenomenon even over relatively short spans of calendar time". Motivated by these two observations, we think that the long memory parameter can be reliably estimated with this data set.

In order to evaluate the proposed models we report the estimation results of alternative models: the multiplicative component GARCH of Engle and Sokalska (2012), the Periodic Long Memory EGARCH$(1, d, 0)$ (PLM-EGARCH) and the k-factor PLM-EGARCH of Bordignon, Caporin, and Lisi (2009), the Log-Gegenbauer GARCH (Log-G-GARCH) of Bordignon, Caporin, and Lisi (2007), the PEGARCH model. First, we present models based on 24 seasons, then we illustrate the results of restricted models.[2]

4.1   Models with twenty-four periods

The main results can be summarized as follows:

Multiplicative component GARCH:

Engle and Sokalska (2012), building on the work of Andersen and Bollerslev (1997b, 1998), propose a multiplicative model for the high-frequency intraday financial returns volatility. The conditional variance is expressed as a product of daily, diurnal, and stochastic intraday volatility. The model includes two intradaily components: a deterministic diurnal pattern and a stochastic intradaily GARCH. Intraday returns are supposed to be given by

$$y_{g,i} = (h_g \cdot s_i \cdot q_{g,i})^{1/2} z_{g,i}, \quad z_{g,i} \sim N(0,1) \quad g = 1, \dots, N \quad i = 1, 2, \dots, S \qquad (17)$$

where h_g is the daily variance component, s_i is the diurnal variance pattern, $q_{i,g}$ is the intraday variance component. The intraday conditional variance varies with i but it is assumed to be

---

a short memory process, excluding the possibility of modeling the interactions between long memory and seasonalities in the conditional variance.

$$Var_{g,i-1}[y_{g,i}] = h_g \, s_i \, q_{g,i}.$$

Assuming that the intradaily returns are serially uncorrelated, the daily conditional variance is the sum of the conditional variances in each of the hours,

$$E\Big[\frac{\sum_{i=1}^{N} y_{g,i}^2}{h_g}\Big] = \lambda. \tag{18}$$

If the daily variance estimator is unbiased and if overnight returns are included in the analysis $\lambda = 1$. The s_i is calculated as the variance of returns in each hour after deflating by the daily variance

$$E\Big[\frac{y_{g,i}^2}{h_g}\Big] = s_i E[q_{g,i}] = s_i$$

with $\sum_i s_i = \lambda$. The diurnal component for each hour is estimated as

$$\hat{s}_i = \frac{1}{N} \sum_{g=1}^{N} \frac{y_{g,i}^2}{h_g}$$

The residual volatility is modeled as a GARCH(1,1):

$$\frac{y_{g,i}}{\sqrt{h_g s_i}} = z_{g,i}\sqrt{q_{g,i}}$$

$$q_{g,i} = \omega + \alpha\Big(\frac{y_{g,i-1}}{\sqrt{h_g s_{i-1}}}\Big)^2 + \beta q_{g,i-1}$$

The forecast of intraday conditional variance is readily obtained as

$$E_{g,i}[y_{g,i+1}^2] = h_g \, s_{i+1} \, q_{g,i+1}. \tag{19}$$

The model is estimated with a two-step procedure: first, daily conditional variance and diurnal volatility, then, using normalized returns, the intraday variance component ($q_{g,i}$). The model does not consider long memory neither in the daily conditional variance nor in the intraday variance component. The estimation results confirm that this model is not able to catch the long-run dependence, as shown by the Ljung-Box test on standardized residuals ($\hat{\eta}_t$) and standardized

squared residuals ($\hat{\eta}_t^2$). More surprisingly, the seasonality seems not to be fully accounted for by the model, as testified by the seasonality test, based on a simple regression of the standardized squared returns on a set of seasonal dummies.

PLM-EGARCH(1,d,0):

$$(1 - \beta L)(1 - L^S)^d(\ln h_t - \omega) \quad = \quad \psi(|\eta_{t-1}| - E\,|\eta_{t-1}|) + \gamma\eta_{t-1}, \quad s = 1, \ldots, 24 \quad (20)$$

This model, based on fractional seasonal filter of the log-conditional variance with $S = 24$, is a simple alternative to the FI-PEGARCH and SFI-PEGARCH. The parameters (ψ, γ) of the news impact function are strongly significant. In particular, the asymmetry parameter γ takes a negative value, corresponding to a leverage effect. The estimated fractional differencing parameter $\hat{d} = 0.589$ is strongly significant (with a robust standard error of 0.009). It turns out to be larger than the estimates obtained with models with periodic parameters. The diagnostic tests in Table 1 show that the PLM-EGARCH (Bordignon, Caporin, and Lisi, 2009) provides an adequate fit of the long memory component of the log-volatility while is clearly unable to catch the periodic patterns in volatility, as it is well shown by the rejection of the null hypothesis of the seasonality test.

k-Factor PLM-EGARCH$(1, \{d_1, \ldots, d_k\}, 0, \{S_1, \ldots, S_k\})$:

$$\left[\prod_{i=1}^{k}(1 - L^{S_i})^{d_i}\right](\ln h_t - \omega) = \psi(|\eta_{t-1}| - E\,|\eta_{t-1}|) + \gamma\eta_{t-1}, \quad s = 1, \ldots, 24 \quad (21)$$

The k-Factor PLM-EGARCH is a generalization of the PLM-EGARCH presented above. The model allows for long memory at different frequencies. The parameter k is a non-negative integer indicating the number of cycles; S_i $(i = 1, \ldots, k)$ are positive integers defining the length of cycles; d_i specify the degree of memory of the cyclical components S_i. The spectrum of squared returns in Figure 6 has two main peaks, one in zero ($S_1 = 0$) and a second one in $\pi/12$ ($S_2 = 24$). Thus, we choose a 2-Factor PLM-EGARCH model. The estimate of d_1 and d_2 are -0.001 with a standard error of 0.013 and 0.603 with a standard error of 0.011, respectively. Only the second factor is significant and close to the one estimated with the PLM-EGARCH. The diagnostic tests confirm the results of PLM-EGARCH, that is the model is not able to fit the strong periodicities present in the data.

Log-Gegenbauer-GARCH(p, d_h, q):

$$(1 - \beta(L))(\ln h_t - \omega) = \left[1 - \beta(L) - P(L)\phi(L)\right]\left[\ln(\eta_t^2 h_t) - \tau\right], \tag{22}$$

$$P(L) = (1-L)^{d_0}(1+L)^{d_h I(E)} \prod_{j=1}^{h-1}\left(1 - 2\cos(\lambda_j)L + L^2\right)^{d_j}.$$

where $\phi(L) = \phi_1 L + \ldots + \phi_6 L^6$ and $\beta(L) = \beta_1 L^1 + \ldots + \beta_{10} L^{10}$ and $\tau = E[\ln(\eta_t^2)]$, (in the Gaussian case $\tau = -1.27$); $P(L)$ is the Gegenbabuer polynomial that models the long memory periodic behavior at frequencies $\lambda_j = \left(\frac{2\pi j}{S}\right)$, $(j = 0, \ldots, h)$ through the long memory parameters d_j, $(j = 0, \ldots, h)$, where h is an integer and $h + 1 = [S/2] + 1$, where $[\cdot]$ stands for the integer part, see Bordignon, Caporin, and Lisi (2007) for more details. Therefore, each periodic frequency in the log-G-GARCH model is modeled by means of a specific long memory parameter d_j. Bordignon, Caporin, and Lisi (2007) show that when $d_0 = d_1 = \ldots = d_h$ and the relevant frequencies are associated to a single periodic component, the Gegenbabuer-GARCH model becomes the PLM-GARCH.[3] The long memory parameters are all significant but d_8 and d_{13}. In terms of SIC, the Log-Gegenbauer-GARCH is better than 2-Factor PLM-EGARCH but worse than PLM-EGARCH. The diagnostic tests in Table 1 also show that the Log-G-GARCH is clearly unable to catch the periodic patterns in volatility.

PEGARCH(1,0):

$$(1 - \beta_s L)\ln h_t = \omega_s + \psi_s(|\eta_{t-1}| - E|\eta_{t-1}|) + \gamma_s \eta_{t-1}, \quad s = 1, \ldots, 24 \tag{23}$$

Figure 7 displays the estimated parameters of PEGARCH along with those of FIPEGARCH and SFI-PEGARCH. The estimates of the periodic coefficients are markedly different, across the 24 hours, with the exception of the ψ's, which are fairly stable across the day. The estimated β's satisfy the condition of stationarity for a periodic EGARCH model, in fact $\left|\prod_s \widehat{\beta}_s\right| < 1$. For what concerns the γ's, only nine parameters are statistically significant at the 1% level. One corresponds to the European market hour, during which American macroeconomic data are released, the others are in the American trading periods. The seasonality test in Table 1 suggests that the estimated PEGARCH is able to fit the periodic features of the volatility process but, as expected, is unable to fit the long memory components, as the rejection of the

---

[3]We are grateful to Massimiliano Caporin for providing us with the GAUSS code to compute the log-likelihood function.
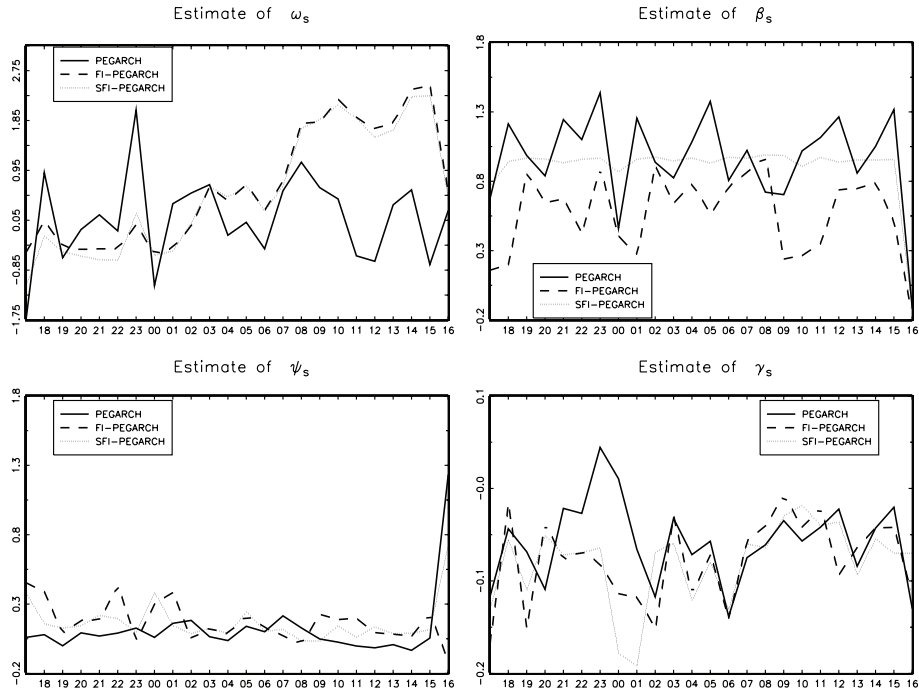
Figure 7: Estimates of $\omega_s, \beta_s, \psi_s, \gamma_s$ in PEGARCH, FI-PEGARCH, and SFI-PEGARCH models.

null of the Ljung-Box test statistic of long autocorrelation in the squared standardized residuals testifies. Moreover, the estimated fractional integration parameter of the squared standardized residuals is significantly different from zero.

We also estimate the model imposing the restriction $\gamma_s = \gamma$. In this case, the leverage effect, given by γ, is strongly statistically significant and the sign is negative confirming the presence of an asymmetric effect of returns shocks on volatility. The diagnostic tests, reported in Table 1, are analogous to those obtained with the unrestricted specification.

FI-PEGARCH(1,d,0):

$$(1 - \beta_s L)(1 - L)^d (\ln h_t - \omega_s) = \psi_s(|\eta_{t-1}| - E|\eta_{t-1}|) + \gamma_s \eta_{t-1}, \quad s = 1, \ldots, 24 \quad (24)$$

The estimates of ω's and β's are markedly different across the 24 hours, while the ψ's do not display a strong periodic pattern (see Figure 7). Differently from the PEGARCH model, now the periodic leverage effect γ_s is statistically significant at 5% level across different hours. In particular, significant estimated γ_s are found in the evening hours after the closing of the American market and in the European market hours just before the opening of American markets. This is not by chance, since in the American "after-hours" quoted companies disclose their financial

20

results (or other sensitive data), whereas the most important American macroeconomic data are released just before the markets openings. The estimate of the fractional integration parameter is equal to 0.427. The Ljung-Box test statistic reported in Table 1 for squared standardized residuals, e.g., $\widehat{\eta}_t^2$, confirm that the estimated model is able to give full account of the dynamics in log-conditional variance. The estimate of fractional integration parameter d of the squared standardized residuals (see Table 1), using the Gaussian semiparametric estimator, indicates that there is no persistence that survives after the filtering. Furthermore, the seasonality test accepts the null, namely there is no evidence of periodic effects in the squared standardized residuals, $\widehat{\eta}_t^2$. The model-implied periodic ACF's provide interesting clues to the ability of the estimated model to approximate the observed behavior. To compute the periodic ACF implied by the estimate of FI-PEGARCH, we replace unknown parameters by QML estimates in (15) and (16). The infinite products in (15) and (16) are truncated at lag $1,008$ while the expected values are computed using the expressions in (35)-(38). It is noticeable in Figure 8 that the model-implied periodic ACF's display an overall slow hyperbolic rate of decay and are different across the hours, which is the kind of pattern observed in the sample periodic ACF in Figure 5. When we adopt a restricted specification, i.e., with $\gamma_s = \gamma, \forall s$, γ is strongly statistically significant, as in the case of PEGARCH. The diagnostics are similar to those of the unrestricted model. The null hypothesis of constant β across periods is strongly rejected.

SFI-PEGARCH(1,d,0):

$$(1 - \beta_s L)(1 - L^S)^d(\ln h_t - \omega_s) \quad = \quad \psi_s(|\eta_{t-1}| - E|\eta_{t-1}|) + \gamma_s \eta_{t-1}, \quad s = 1, \ldots, 24 \quad (25)$$

The estimates of ω's and ψ_s are fairly close to those of FI-PEGARCH. Differences arise in the estimated β_s (see Figure 7) which are larger than those of FIPEGARCH. Again the γ_s's are strongly statistically significant and vary across seasons, like those obtained with the FI-PEGARCH. The estimate of d is 0.25, within the stationary bounds, and close to the semi-parametric estimate (0.269) at frequency $\pi/12$. According to the seasonality test in Table 1, the $\widehat{\eta}_t^2$ have no seasonal components. Moreover, like in the FI-PEGARCH case, the estimate of the long-range dependence of the squared standardized residuals is not statistically different from zero. The model-implied periodic ACF's, computed exactly as in the FI-PEGARCH case, show the same slow hyperbolic rate of decay[4], generated by the long-memory features of

---

[4]Not shown here to save space, but available upon request from the authors.

Figure 8: Estimated FI-PEGARCH periodic autocorrelogram computed using expressions (14)-(16), where the unknown parameters are replaced by the estimates.

SFI-PEGARCH, and periodically varying correlations, implied by the periodic parameters in the log-conditional variance process. The results suggest that during the Asian and European trading time the volatility is much lower than during the American trading time when we observe a sharp increase. Moreover, this evidence seems to confirm that hourly Emini returns are characterized by different seasonal patterns with a statistically significant persistence. The restricted model with $\gamma_s = \gamma, \forall s$, presents similar features. The Wald test statistics for the null of $d = 0$ and constant β's are, like the FI-PEGARCH, strongly rejected (see Table 1).

Finally, it is worth noticing that the long memory estimates are quite heterogeneous across different models. In this case, the estimate of d turns out to be smaller than the one obtained in the PLM-EGARCH case, where the same fractional integration filter is used. This exemplifies how the fractional integration parameter estimates are connected to the specification of the periodicity.

Skewness and kurtosis of the standardized residuals: In general, most of the models are able to capture the differences in the hourly skewness and kurtosis during European and American hours, while some (weak) excess skewness and (strong) kurtosis remain during American "after-hours" and when American macroeconomic news are released before American market openings. This suggests that different distributional assumptions may be needed for these particular market periods. However, this goes beyond the scope of the paper and we leave it as avenue for further research.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
$Q_{\hat\eta_t}(5)$	3.276	3.587	3.382	4.194	13.470	9.957	11.340	10.526	11.865	12.317	10.902	16.624	7.380	7.591
p-value	0.657	0.610	0.641	0.522	0.019	0.076	0.045	0.062	0.037	0.031	0.053	0.005	0.194	0.180
$Q_{\hat\eta_t}(25)$	18.745	49.768	51.390	57.503	48.977	48.475	50.177	48.079	48.331	48.202	52.223	61.353	46.089	56.969
p-value	0.809	0.002	0.001	0.000	0.003	0.003	0.002	0.004	0.003	0.004	0.001	0.000	0.006	0.000
$Q_{\hat\eta_t}(500)$	787.177	532.471	532.990	548.757	529.570	527.989	548.541	541.113	538.654	536.822	520.038	559.377	501.375	533.364
p-value	0.000	0.152	0.149	0.065	0.174	0.187	0.066	0.099	0.113	0.123	0.259	0.034	0.474	0.146
$Q_{\hat\eta_t^2}(5)$	0.981	0.996	3.634	44.889	51.395	70.516	12.323	11.303	22.918	11.741	12.418	0.412	16.500	0.700
p-value	0.731	0.360	0.603	0.000	0.000	0.000	0.031	0.046	0.000	0.039	0.029	0.995	0.006	0.984
$Q_{\hat\eta_t^2}(25)$	23.799	7.736	34.418	78.975	62.615	82.476	21.952	19.925	33.654	21.322	15.481	32.621	25.115	15.117
p-value	0.531	1.000	0.099	0.000	0.000	0.000	0.638	0.751	0.116	0.675	0.929	0.141	0.456	0.939
$Q_{\hat\eta_t^2}(500)$	1973.690	204.076	616.955	407.910	673.512	635.034	451.626	431.432	494.540	471.645	310.369	731.934	378.228	323.071
p-value	0.000	1.000	0.000	0.999	0.000	0.000	0.941	0.988	0.560	0.814	1.000	0.000	1.000	1.000
Seasonality test	0.000	0.000	0.000	0.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	0.000	0.000	0.000
$\hat d\,(\hat\eta_t^2)$	0.223	-0.009	-0.021	0.044	0.101	0.188	-0.005	0.054	0.079	0.072	-0.078	0.063	-0.031	0.072
p-value	0.041	0.041	0.041	0.041	0.041	0.041	0.041	0.041	0.041	0.041	0.041	0.041	0.041	0.041
Log-lik function	-32696.57	-34597.16		-32888.859	-30883.11	-30953.45	-30825.96	-30839.939	-30893.41	-30976.71	-32367.345	-33755.433	-32812.065	-33363.94
SIC	2.791	2.954		2.832	2.795	2.791	2.671	2.662	2.677	2.674	2.767	2.885	2.804	2.852
$H_0: d = 0$	4132.415						1074.387	6791.595	148.523	808.264				
p-value	0.000						0.000	0.000	0.000	0.000				
$H_0: \beta_s = \beta$				2573.436		1733.656		196.2448						
p-value		0.000	0.000	0.000	0.000		0.000		0.000		0.000	0.000	0.000	0.000

Table 1: Diagnostic tests of estimated models. 1=Multiplicative component GARCH, 2 = PLM-EGARCH, 3= k-Factor PLM-EGARCH, 4=Log-G-GARCH, 5 = PEGARCH(1,0), 6 = PEGARCH(1,0) with $\gamma_s = \gamma$, $\forall s$, 7 = FI-PEGARCH(1,d,0) with $\gamma_s = \gamma$, $\forall s$, 8 = FI-PEGARCH(1,d,0), 9 = SFI-PEGARCH(1,d,0) with $\gamma_s = \gamma$, $\forall s$, 10 = SFI-PEGARCH(1,d,0) with constrained periodicity, 12 = SFI-PEGARCH(1,d,0) with constrained periodicity, 13 = FI-PEGARCH(1,d,0) with three regimes, 14 = SFI-PEGARCH(1,d,0) with three regimes. In the three regimes case the parameters vary over three periods (the Asia-Europe trading period (17.00−09.00), the morning trading in the US (10.00−13.00) and the US afternoon (14.00−16.00)). Constrained periodic parameters are defined in (26). Ljung-Box test statistics for standardized residuals ($Q_{\hat\eta_t}$) and square standardized residuals ($Q_{\hat\eta_t^2}$) are computed for the number of lags reported in parentheses. The Schwarz Information Criterion (SIC) is computed as $SIC = -L(\hat{\boldsymbol{\theta}})/T + p\frac{\ln T}{T}$, where p is the number of parameters and T is the number of observations. The table reports the p-value of the seasonality test, namely a F test of the regression of the standardized squared residual ($\hat\eta_t^2$) on a set of 23 dummies, one for each hour. $\hat d$ is the estimated fractional integration parameter of the standardized squared residuals by the Gaussian semiparametric estimator with $m = 150$, asymptotic standard errors equal to $1/\sqrt{4m}$ are in parentheses. Wald test statistics for the null hypotheses in the Table are reported.

4.2 Models with constrained patterns of periodicity

Models with twenty-four periods can be heavily overparameterized, and this can cause some loss of efficiency. When S is large, like the case we are examining, the number of periodic parameters can be limited assuming that their seasonal variation is described by periodic functions like combinations of cosine and sinus. Following Jones and Brelsford (1967), we assume that the periodic parameters in the FI-PEGARCH(1,d,0) and SFI-PEGARCH(1,d,0) are specified as:

$$\omega_s = \widetilde{\omega}_0 + \widetilde{\omega}_1 \cos\left(\frac{2\pi s}{S} - \widetilde{\omega}_2\pi\right) \quad \beta_s = \widetilde{\beta}_0 + \widetilde{\beta}_1 \cos\left(\frac{2\pi s}{S} - \widetilde{\beta}_2\pi\right)$$

$$\psi_s = \widetilde{\psi}_0 + \widetilde{\psi}_1 \cos\left(\frac{2\pi s}{S} - \widetilde{\psi}_2\pi\right) \quad \gamma_s = \widetilde{\gamma}_0 + \widetilde{\gamma}_1 \cos\left(\frac{2\pi s}{S} - \widetilde{\gamma}_2\pi\right) \quad s = 1,\dots,S \qquad (26)$$

For parameter identification we restrict $\widetilde{\omega}_2, \widetilde{\beta}_2, \widetilde{\psi}_2, \widetilde{\gamma}_2 \in [0,1)$ as $\cos(x + k\pi) = (-1)^k \cos(x)$ for $k \in \mathbb{Z}$ and $x \in \mathbb{R}$. This specification is close to the Flexible Fourier Form GARCH adopted by Taylor (2004). Figure 9 plots the estimates of γ's and ψ's based on the specification in (26). The γ's and ψ's in the FI-PEGARCH model have a more pronounced variability than that obtained with the SFI-PEGARCH model. This is also clear from Figure 10 where the estimated periodic news impact functions $g_s(\eta)$ for both models are displayed. The negative shocks during the last hours of the Asia-Europe trading period have a larger effect on volatility than those which can occur in the rest of the day. It should be also noted that the US morning period is characterized by news impact function values larger than those estimated for the US afternoon trading hours. The early morning trading hours can be the heaviest trading hours of the day. During the US morning trading session, the quantity and the quality of information that investors receive is such that to provoke large portfolio reallocations. In the FI-PEGARCH case the asymmetric effect of shocks on log-volatility is more pronounced than that obtained with the SFI-PEGARCH. This comes along with different long memory estimates. The $\widehat{d}$ in the FI-PEGARCH model is 0.707 while in the SFI-PEGARCH is 0.438, both strongly significant and higher than those estimated with unconstrained models. This finding confirms that restricting the periodicity pattern leads to an increase in the estimate of the fractional differencing parameter. The Ljung-Box tests of autocorrelation of different orders in the squared standardized residuals, reported in Table 1, indicate that the SFI-PEGARCH is worse than the FI-PEGARCH.

4.3 Models with three periods

An alternative strategy to the constrained periodic structure of model's parameters, explored in Section 4.2, is to exploit some a priori knowledge on the trading activity. The division of the 24-hour daily period in two sessions suggests two possible distinct trading cycles for a US investor. In this regard, Taylor (2007) finds that the overnight information flow, as revealed in the E-mini SP500 futures market during the Globex session, which broadly coincides with the trading times of the major European and Asian securities markets, has a significant impact upon the information flow in the regular SP500 futures market. In addition to a geographic motivation, the specification of models with three periods (morning, afternoon and overnight) also has a daytime vs. overnight motivation, due to the different institutional arrangements for trading in the daytime and overnight periods. In this regard, Tsiakas (2008) finds that daytime returns and volatility originate from a different data generating process than overnight returns and volatility, and that financial information that accumulates during nontrading hours has a superior predictive ability than that which accumulates during trading hours.

We can identify three distinct regimes: the Asia-Europe trading period $(17.00 - 09.00)$, the morning trading in the US $(10.00 - 13.00)$ and the US afternoon $(14.00 - 16.00)$. The estimated γ's, plotted in Figure 9, are markedly different between FI-PEGARCH and SFI-PEGARCH while the estimated ψ are much closer. The estimated $g_s(\eta)$, in Figure 10 are similar to those obtained with the constrained periodicity parameter specification in (26), that is the impact of negative shocks is larger during the Asia-Europe trading session. In this case, the asymmetric effects in the first period $(17.00 - 09.00)$ are larger than those in the other two. The estimates of d are significantly different from zero but smaller than those in the constrained periodicity models, 0.458 and 0.37 for the FI-PEGARCH and SFI-PEGARCH, respectively. The Ljung-Box tests in Table 1 suggest that the two models are close in terms of the autocorrelation in the squared standardized disturbances. However both models do not catch the periodicity, as the seasonality test clearly highlights.
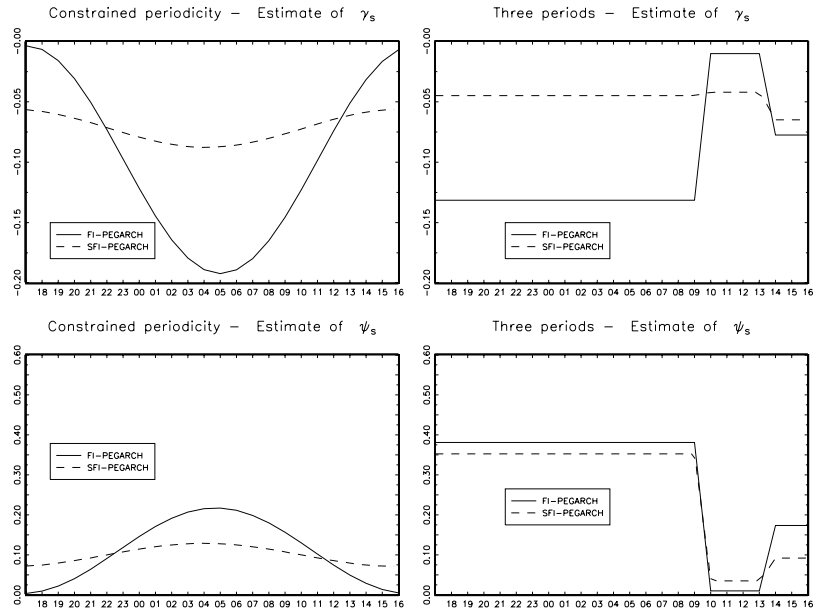
Figure 9: Estimated γ's and ψ's with constrained periodicity (left plot) and three regimes (right plot) FI-PEGARCH and SFI-PEGARCH models.
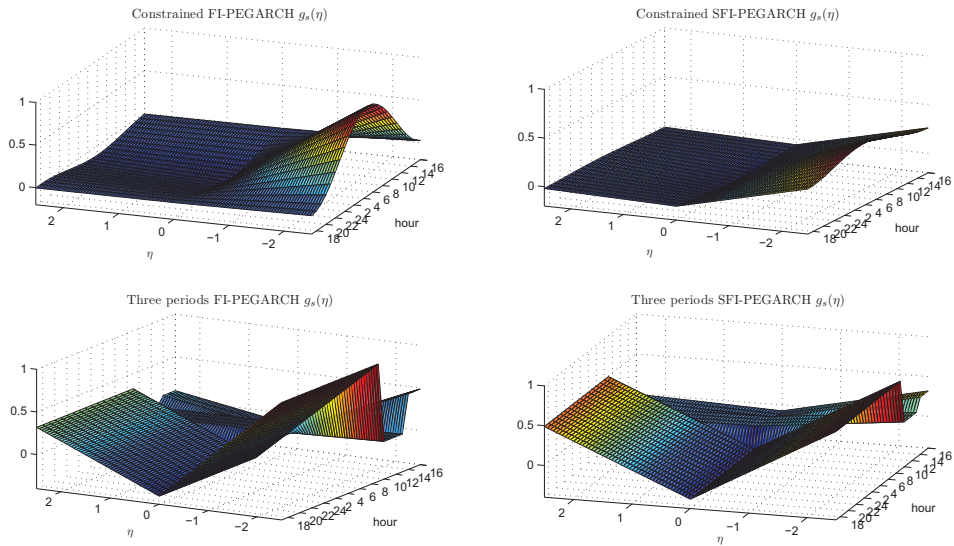

Figure 10: Estimated $g_s(\eta)$ (the z-coordinate of the surface points) of restricted FI-PEGARCH (left plots) and SFI-PEGARCH (right plots) models.

27

5 Forecasts

5.1 In-sample forecasts

In this section we present the in-sample forecasting performances of the models presented in the previous section. The evaluation and comparison of volatility forecasts is made difficult by the fact that the object of interest is unobservable, even ex post. Thus the evaluation and comparison of volatility forecasts must rely on direct or indirect methods to overcome this difficulty. Direct methods use a volatility proxy, i.e. some observable variable that is related to the latent variable of interest. To this end we compute the *hourly realized volatility* (RV_t), an ex-post volatility measure, namely the aggregated squared 5-minutes returns over the hour. Even though this estimate can be biased, in this context it represents a first attempt to describe the intraday pattern of the volatility process.

The models used in forecasting one-step ahead volatility are shown in Table 2. The forecasts are obtained by first estimating the parameters of the model on the full sample and performing a series of static one-step ahead forecasts. It is important to remark that, for in-sample forecasting, the constant mean parameter is estimated using the full data sample, while for out-of-sample forecasting this parameter is re-estimated every time a volatility forecast is generated, by using only the information contained in the rolling samples of 23,450 observations (see the next subsection). This is done in order to avoid any problem of look-ahead bias.

The forecast accuracy is evaluated employing the bias, the Mean Square Error (MSE) and the $QLIKE = \ln h_t + RV_t/h_t$. The choice of these loss functions is motivated by Patton (2011) who shows that MSE and QLIKE provide consistent ranking when there is noise in the volatility proxy used in place of the unobserved volatility.

The results in Table 2 shows that the Multiplicative component GARCH is the worst in terms of bias, MSE and QLIKE, while the best is the FI-PEGARCH with constant γ. FI-PEGARCH and SFI-PEGARCH models outperform the PLM-EGARCH, 2-Factor PLM-EGARCH, and Log-G-GARCH models. The SFI-PEGARCH models behave, in terms of MSE and QLIKE, always worse than the corresponding FI-PEGARCH models while for the bias the results are mixed.

	In-sample			Ouf-of-sample		
	bias	MSE	QLIKE	bias	MSE	QLIKE
Multiplicative GARCH	-1.865	1.812	183.332	-100.491	667.208	4.114
PLM-EGARCH	0.329	1.355	64.740	7.635	14.285	4.146
2-Factor PLM-EGARCH	0.018	1.462	80.774	6.672	14.965	4.185
Log-G-GARCH	0.384	1.379	66.598	9.188	13.987	3.969
PEGARCH	0.308	1.319	62.530	1.358	32.479	8.528
PEGARCH constant γ	0.331	1.321	62.826	5.896	13.232	3.699
FI-PEGARCH	0.391	1.330	63.818	8.698	14.169	3.861
FI-PEGARCH constant γ	0.350	1.317	61.620	7.079	14.522	3.733
FI-PEGARCH constrained	0.419	1.405	66.022	6.113	13.061	3.478
FI-PEGARCH 3 periods	0.294	1.376	64.058	-2.867	131.333	4.385
SFI-PEGARCH	0.403	1.334	63.692	7.888	13.782	3.985
SFI-PEGARCH constant γ	0.372	1.331	62.261	6.466	13.923	3.885
SFI-PEGARCH constrained	0.380	1.381	66.806	5.035	14.777	3.850
SFI-PEGARCH 3 periods	0.003	1.385	118.647	7.758	14.048	4.470

Table 2: Measures of forecast accuracy. The bias is the sample mean of the ex-post hourly realized volatility minus the forecast.

5.2 Out-of-sample forecasts

While in-sample forecasting provides a general indication of the predictive ability, it is only the evaluation of the out-of-sample forecasts that can assess the models capability of providing predictions that are relevant for financial decision making. Hence, in order to closely mimic a real-world forecast situation we also report on the out-of-sample forecast results. The one-hour-ahead forecasts are obtained from parameter estimates based on rolling samples with fixed sizes of $23,450$ observations. For every date $t > 23,450$ in the sample, we estimate the parameters of each specification over the $23,450$ data points up to including date t. We calculate 500 out-of-sample forecasts, corresponding to the period from December 26, 2007 to January 25, 2008.

It turns out that the best model in terms of bias is the PEGARCH, while the worst is the Multiplicative component GARCH. While, looking at the MSE, the best model is the FI-PEGARCH with constrained periodicity, on the contrary the SFI-PEGARCH models with unconstrained periodicity are better than the constrained ones. FI-PEGARCH and SFI-PEGARCH models perform better than the PLM-EGARCH and 2-Factor PLM-EGARCH. A similar picture emerges when using the QLIKE loss function. The best is still the FI-PEGARCH with constrained periodicity while the PEGARCH with constant γ turns out to be the second best model. There is also a quite strong evidence against the FI-PEGARCH with three periods. The FI-PEGARCH

with constrained periodicity can represent the best compromise between model complexity and over-parametrization.

To test for significant differences in forecast accuracy between competing models we apply the Model Confidence Set (MCS) approach developed by Hansen, Lunde, and Nason (2011). The MCS is a sequential test of equal predictive ability. Given an initial set of forecasts it allows to test the null that no forecast is distinguishable from any other and discard any inferior forecasts if they exist. The MCS procedure yields a model confidence set that can be interpreted as a confidence interval for the forecasts in that it is the set containing the best forecasts at some confidence level. The starting hypothesis is that all models considered have equal forecasting performances.

The equivalence test statistics employed in the procedure are the chi-square (T_Q), the maximum deviation of the sample losses (T_{max}), and the semi-quadratic T_{SQ}. The loss functions used here are the MSE and the QLIKE. The asymptotic distributions of the test statistics of T_{max} and T_{SQ} are non-standard and the p-values are obtained by means of stationary block bootstrap with a block length of 48 hours. The MCS p-value is the threshold level that determines whether the model belongs to the MCS or not.

Results in Table 3, for all models considered, indicate that based on the MSE loss function almost all models are statistically indistinguishable given the relatively high p-values.[5] On the contrary, given the QLIKE loss function, the results of the three test statistics are quite different. Choosing a significance level of 0.1, the FI-PEGARCH with constrained periodicity is the sole model in the MCS. Patton and Sheppard (2009) show that QLIKE, relative to MSE, exhibits significantly more power in differentiating between forecasts. Thus overall the FI-PEGARCH with constrained periodicity provides forecast performance gains relative to a number of alternative models.

To summarize, the empirical evidence suggests that fully unrestricted fractional models describe the data very well in-sample, whereas the constrained (for the FI-PEGARCH) and restricted (in the SFI-PEGARCH case) versions are better or equivalent in terms of out-of-sample forecasting accuracy. It may be the case that the reduction in forecast uncertainty and accuracy gains from less-parameterized models outweigh the possible misspecification problems associated with simpler model structures. Moreover, the issue of overfitting noisy data such as our

---

[5]The MCS p-values have been computed with the *MULCOM 2.00 Econometric Toolkit for Multiple Comparisons* by Hansen and Lunde (2010).

intraday data set may well explain why more complex formulations tend not to predict well out-of-sample.

	MSE			QLIKE		
	T_Q	T_{max}	T_{SQ}	T_Q	T_{max}	T_{SQ}
Multiplicative GARCH	0.000	0.051	0.044	0.003	0.000	0.019
PLM-EGARCH	**0.289**	**0.532**	**0.462**	0.003	0.000	0.019
2-Factor PLM-EGARCH	0.026	**0.305**	**0.171**	0.003	0.000	0.019
Log-G-GARCH	**0.745**	**0.778**	**0.776**	0.026	0.035	0.056
PEGARCH	0.000	0.091	0.045	0.003	0.000	0.016
PEGARCH constant γ	**0.745**	**0.532**	**0.531**	0.026	0.035	0.056
FI-PEGARCH	**0.289**	**0.532**	**0.380**	0.003	0.000	0.024
FI-PEGARCH constant γ	**0.289**	**0.400**	**0.380**	0.026	0.035	0.056
FI-PEGARCH constrained	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**
FI-PEGARCH 3 periods	0.000	0.091	0.044	0.003	0.000	0.012
SFI-PEGARCH	**0.289**	**0.400**	**0.380**	0.003	0.000	0.019
SFI-PEGARCH constant γ	**0.190**	**0.323**	**0.239**	0.010	0.035	0.044
SFI-PEGARCH constrained	0.000	**0.247**	**0.254**	0.003	0.000	0.035
SFI-PEGARCH 3 periods	**0.289**	**0.424**	**0.382**	0.003	0.000	0.012

Table 3: MCS of out-of-sample forecasts. The last column reports the MCS p-values according to the test statistics T_{max}, i.e. maximum deviation, computed with 10,000 bootstrap resamples, with block length set to 48. In bold p-values larger than 0.1 are reported in bold.

6 High frequency or low frequency data to estimate long memory?

So far we have seen how the estimate of long memory is influenced by the specification of the intraday patterns observed in volatility. This should not be the case with daily data. In fact, at daily level the intraday periodic patterns are obviously absent and the long-range dependence can be estimated with a long memory volatility model (e.g. FIEGARCH). In this situation, one would expect that the fractional differencing parameter estimate based on daily data should be close to the true one. On the contrary, the presence of periodicity at higher frequencies should lead, when using an underspecified long memory volatility model, to a biased estimate of d. In this section, we investigate this issue by simulation. The data-generating process employed in simulations is a FI-PEGARCH(1,d,0) with 12 periods; the number of simulations is 1,000. The parameters values used in simulations are displayed in Table 4.

Daily returns are obtained as $x_g = \sum_{s=1}^{12} y_{s,g}$, $g = 1, \ldots, N$, where $y_{s,g}$ is the log-return for period s and day g. The criteria adopted for measuring finite sample performance are bias and

31

	1	2	3	4	5	6	7	8	9	10	11	12
ω_s	-1.500	-1.500	-1.500	-1.500	-1.500	-1.500	-1.500	-1.500	-1.500	-1.500	-1.500	-1.500
ψ_2	0.010	0.011	0.012	0.020	0.022	0.024	0.030	0.031	0.032	0.010	0.011	0.012
β_s	0.830	0.830	0.830	0.600	0.600	0.600	0.800	0.800	0.800	0.590	0.590	0.590
γ_s	-0.290	-0.290	-0.290	-0.410	-0.410	-0.410	-0.210	-0.210	-0.210	-0.290	-0.290	-0.290

Table 4: Parameters used in simulation of FI-PEGARCH(1,d,0), with $S = 12$.

root mean squared error (RMSE). The results in Table 5(a) show how the presence of periodic patterns at intraday level are essentially irrelevant for the estimation of d with the FIEGARCH with daily data. Aggregating the intraday returns annihilates the periodicity and thus the estimation of long memory turns out to be only slightly biased. Further, the bias decreases when the value of the fractional differencing parameter, i.e. $d = 0.35$, increases. With larger sample sizes we obtain an improvement in the RMSE.

When the FIEGARCH(1,d,0) is estimated with intradaily data the results are completely different. The estimate of d is severely biased, as shown from Table 5(b), no matter what the sample size is. There is also some evidence that, differently from the daily data case, the larger the fractional integration parameter, the bigger the bias in the estimate of d.

This simple experiment shows how the long memory estimate can be biased when periodic patterns are not properly modeled.

(a) Daily data

N	$\bar{\hat{d}}$	0.25 Bias	RMSE	$\bar{\hat{d}}$	0.35 Bias	RMSE
2,500	0.2625	0.0125	0.0777	0.3578	0.0078	0.0560
4,000	0.2716	0.0216	0.0544	0.3611	0.0111	0.0382
6,000	0.2671	0.0171	0.0402	0.3524	0.0024	0.0307
8,000	0.2664	0.0164	0.0382	0.3507	0.0007	0.0269

(b) Hourly data

T	$\bar{\hat{d}}$	0.25 Bias	RMSE	$\bar{\hat{d}}$	0.35 Bias	RMSE
12,000	0.3482	0.0982	0.1011	0.6235	0.2735	0.2742
14,400	0.3491	0.0991	0.1010	0.6266	0.2766	0.2771
16,800	0.3479	0.0979	0.1002	0.6285	0.2785	0.2790

Table 5: Monte Carlo results on the estimate of d in a FIEGARCH(1,d,0) using data generated with FIPEGARCH(1,d,0) with 12 periods. $d = 0.25$ and $d = 0.35$ are used in simulation. N is the number of days and T is the number of intraday observations.

7 Conclusions

This paper presents a new approach to the modeling of intraday financial return volatility that takes persistence, periodicity and asymmetry into account. The proposed models are used to investigate the features of the returns volatility of the E-mini SP500 futures contracts. The empirical application shows that the proposed models provide an adequate description of high frequency return volatility. We find that the FI-PEGARCH specification with constrained patterns of periodicity provides the best out-of-sample forecasting performance, representing a good compromise between flexibility and parsimony. Furthermore, the results confirm that the volatility is characterized by long memory, periodicity and asymmetric responses to return shocks. Moreover, we found out that asymmetric effects vary during the day, with negative shocks having a larger effect on volatility during the US trading hours than those which occur in the rest of the day. The empirical findings also suggest that restricting the periodicity patterns leads to a larger estimate of the long memory parameter. A simulation experiment corroborates this conjecture: that is a model that misspecifies the periodic behavior of the intraday volatility leads to a biased estimate of d.

Finally, we want to remark that we did not consider the presence of price jumps to avoid additional model complexity: in this regard, modeling the volatility jump and continuous components at the intraday level for algorithmic trading is certainly an interesting avenue for further research.

References

ANDERSEN, T. G., AND T. BOLLERSLEV (1997a): "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns," *Journal of Finance*, 52, 975–1005.

——— (1997b): "Intraday Periodicity and Volatility Persistence in Financial Markets," *Journal of Empirical Finance*, 4, 115–158.

——— (1998): "Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies," *Journal of Finance*, 53, 219–265.

ARTECHE, J. (2004): "Gaussian semiparametric estimation in long memory in stochastic volatility and signal plus noise models," *Journal of Econometrics*, 119, 131–154.

ARTECHE, J., AND P. ROBINSON (2000): "Semiparametric inference in seasonal and cyclical long memory processes," *Journal of Time Series Analysis*, 21, 1–25.

BAILLIE, R., T. BOLLERSLEV, AND H. O. MIKKELSEN (1996): "Fractionally integrated generalized autoregressive conditional heteroskedasticity," *Journal of Econometrics*, 74, 3–30.

BAILLIE, R., Y.-W. HAN, R. J. MYERS, AND J. SONG (2007): "Long memory models for daily and high frequency commodity futures returns," *Journal of Futures Markets*, 27, 643–668.

BELTRATTI, A., AND C. MORANA (1999): "Computing value at risk with high frequency data,," *Journal of Empirical Finance*, 6(5), 431–455.

BOLLERSLEV, T., AND E. GHYSELS (1996): "Periodic autoregressive conditional heteroscedasticity," *Journal of Business and Economic Statistics*, 14, 139–151.

BOLLERSLEV, T., AND H. MIKKELSEN (1996): "Modeling and pricing long memory in stock market volatility," *Journal of Econometrics*, 74, 3–30.

BOLLERSLEV, T., AND J. WOOLDRIDGE (1992): "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-varying Covariances," *Econometric Reviews*, 11, 143–172.

BORDIGNON, S., M. CAPORIN, AND F. LISI (2007): "Generalised long-memory GARCH models for intra-daily volatility," *Computational Statistics and Data Analysis*, 51, 5900–5912.

——— (2009): "Periodic Long-Memory GARCH models," *Econometric Reviews*, 28, 60–82.

DACOROGNA, M., R. GENCAY, U. MULLER, R. B. OLSEN, AND O. V. PICTET (2001): *An Introduction to High-Frequency Finance*. Academic Press.

ENGLE, R. F., AND M. E. SOKALSKA (2012): "Forecasting intraday volatility in the US equity market. Multiplicative component GARCH," *Journal of Financial Econometrics*, 10(1), 54–83.

FRANSES, P., AND R. PAAP (2000): "Modelling day-of-the-week seasonality in the SP 500 index," *Applied Financial Economics*, 10, 483–488.

——— (2004): *Periodic Time Series Models*. Oxford University Press.

GRADSHTEYN, I., AND I. RYZHIK (1980): *Tables of Integrals, Series, and Products*. Academic Press.

HANSEN, P. R., A. LUNDE, AND J. M. NASON (2011): "The Model Confidence Set," *Econometrica*, 79(2), 453–497.

HE, C., T. TERÄSVIRTA, AND H. MALMSTEN (2002): "Moment Structure of a Family of First-Order Exponential GARCH Models," *Econometric Theory*, 18, 868–885.

JONES, R. H., AND M. BRELSFORD (1967): "Time Series with Periodic Structure," *Biometrika*, 54, 403–407.

KOOPMAN, S., M. OOMS, AND M. CARNERO (2007): "Periodic Seasonal Reg-ARFIMA-GARCH Models for Daily Electricity Spot Prices," *Journal of American Statistical Association*, 102, 16–27.

MARTENS, M., Y. CHANG, AND S. J. TAYLOR (2002): "A comparison of seasonal adjustment methods when forecasting intraday volatility,," *Journal of Financial Research*, 15(2), 283–299.

MCLEOD, A. I. (1994): "Diagnostic checking of periodic autoregression models with application," *Journal of Time Series Analysis*, 15, 221–233.

NELSON, D. (1991): "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59(2), 347–370.

OOMS, M., AND P. FRANSES (2001): "A Seasonal Periodic Long Memory Model for Monthly River Flows," *Environmental Modelling & Software*, 16, 559–569.

PALMA, W. (2007): *Long-memory time series: theory and methods.* J. Wiley and Sons.

PATTON, A., AND K. SHEPPARD (2009): "Evaluating volatility and correlation forecasts," in *Handbook of Financial Time Series*, ed. by T. Andersen, R. Davis, J. Kreiss, and T. Mikosch. Springer-Verlag.

PATTON, A. J. (2011): "Volatility forecast comparison using imperfect volatility proxies," *Journal of Econometrics*, 160(1), 246 – 256.

RUIZ, E., AND H. VEGA (2008): "Modelling long-memory volatilities with leverage effect: A-LMSV versus FIEGARCH," *Computational Statistics and Data Analysis*, 52, 2846 – 2862.

SHEPHARD, N., AND K. SHEPPARD (2010): "Realising the future: forecasting with high-frequency-based volatility (HEAVY) models," *Journal of Applied Econometrics*, 25, 197 – 231.

TAYLOR (2004): "Modeling discontinuous periodic conditional volatility: evidence from the commodity futures market," *Journal of Futures Markets*, 24, 805–834.

TAYLOR, N. (2007): "A Note on the Importance of Overnight Information in Risk Management Models," *Journal of Banking and Finance*, 31, 161–180.

TSIAKAS, I. (2006): "Periodic Stochastic Volatility and Fat Tails," *Journal of Financial Econometrics*, 4, 90–135.

——— (2008): "Overnight Information and Stochastic Volatility: A Study of European and US Stock Exchanges," *Journal of Banking and Finance*, 32, 251–268.

WASSERFALLEN, W., AND H. ZIMMERMANN (1985): "The behaviour of intraday exchange rates," *Journal of Banking and Finance*, 9, 55–72.

A PROOFS

Proof of Proposition 2.1.
The FI-PEGARCH(1,d,0) can be written as:

$$(1 - \beta_{1s}L) \log h_t = (1 - \beta_{1s})\omega_s + (1 - L)^{-d}g_s(\eta_{t-1})$$

Take the exponential and raise both sides of the equation to the power m

$$h_t^m = \exp\{m(1 - \beta_{1s})\omega_s\} \exp\{m(1 - L)^{-d}g_s(\eta_{t-1})\}h_{t-1}^{m\beta_{1s}} \tag{27}$$

If we solve (27) recursively we get

$$h_t^m = \exp\left\{m(1 - \beta_{1s})\omega_s \sum_{i=1}^{n} \beta_{1s}^{i-1}\right\} \exp\{m(1 - L)^{-d}\sum_{i=1}^{n} \beta_{1s}^{i-1}g_s(\eta_{t-i})\}h_{t-n}^{m\beta_{1s}^n} \tag{28}$$

Assuming that $|\beta_{1s}| < 1, \ \forall s$ and letting $n \to \infty$ then

$$h_t^m = \exp\{m\,\omega_s\} \exp\left\{m(1 - L)^{-d}\sum_{i=1}^{\infty} \beta_{1s}^{i-1}g_s(\eta_{t-i})\right\}$$

using (8)

$$h_t^m = \exp\{m\,\omega_s\} \exp\left\{m\lambda_s(L)g_s(\eta_{t-1})\right\} \tag{29}$$

Now, given that h_t and η_t are uncorrelated, the even moments of y_t are equal to

$$E\left[|y_t|^{2m}\right] = E\left[|\eta_t\sqrt{h_t}|^{2m}\right] = E[|\eta_t|^{2m}]\,E[h_t^m]$$

$$= \nu_{2m}\exp\{m\,\omega_s\}\prod_{i=0}^{\infty} E\left[\exp\{m\lambda_{i,s}g_s(\eta_{t-1-i})\}\right].$$

where $\nu_{2m} = E[|\eta_t|^{2m}]$. Since the kurtosis of y_t is defined as $k_{y,s} = \frac{E[(y_t)^4]}{E[(y_t)^2]^2}$, setting $m = 1$ for the denominator and $m = 2$ for the numerator, respectively, we obtain

$$k_{y,s} = \nu_4 \frac{\prod_{i=0}^{\infty} E\left[\exp\{2\lambda_{i,s}g_s(\eta_{t-1-i})\}\right]}{\{\prod_{i=0}^{\infty} E\left[\exp\{\lambda_{i,s}g_s(\eta_{t-1-i})\}\right]\}^2}. \tag{30}$$

The same procedure is followed for the SFI-PEGARCH(1,d,0), where $\lambda_s(L) = (1 - \beta_{1s}L)^{-1}(1 - L^S)^{-d}$.

Proof of Proposition 2.2
The autocorrelation function is defined as

$$\rho_{m,s}(n) = \frac{E[|y_t|^{2m}|y_{t-n}|^{2m}] - [E(|y_t|^{2m})][E(|y_{t-n}|^{2m})]}{(E[|y_t|^{4m}] - [E(|y_t|^{2m})]^2)^{1/2}\,(E[|y_{t-n}|^{4m}] - [E(|y_{t-n}|^{2m})]^2)^{1/2}}. \tag{31}$$

Multiply both sides of (28) by h_{t-n}^m and $\eta_t^{2m}\eta_{t-n}^{2m}$, knowing that $\sum_{i=0}^{n-1} \beta_{1s}^i = (1 - \beta_{1s}^n)(1 - \beta_{1s})^{-1}$,

$$|y_t|^{2m}|y_{t-n}|^{2m} = |\eta_t|^{2m}|\eta_{t-n}|^{2m}\exp\{m(1 - \beta_{1s}^n)\omega_s\}\times$$

$$\exp\left\{m(1 - L)^{-d}\sum_{i=1}^{n} \beta_{1s}^{i-1}g_s(\eta_{t-i})\right\}h_{t-n}^{m(1+\beta_{1s}^n)}. \tag{32}$$

and taking expectations yields

$$E[|y_t|^{2m}|y_{t-n}|^{2m}] = \nu_{2m} \exp\left\{m(1-\beta_{1s}^n)\omega_s\right\} E\left[|\eta_{t-n}|^{2m} \exp\left\{m(1-L)^{-d}\sum_{i=1}^{n}\beta_{1s}^{i-1}g_s(\eta_{t-i})\right\}\right] \times$$

$$E\left[h_{t-n}^{m(1+\beta_{1s}^n)}\right]. \tag{33}$$

By (29), we know that

$$E\left[h_{t-n}^{m(1+\beta_{1s}^n)}\right] = \exp\{m\omega_r(1+\beta_{1s}^n)\}\prod_{i=1}^{\infty} E\left[\exp\{m(1+\beta_{1s}^n)\lambda_{i-1,r}g_r(\eta_{t-n-i})\}\right] \tag{34}$$

where r denotes the period corresponding to time $t-n$. Then, replacing $E\left[h_{t-n}^{m(1+\beta_{1s}^n)}\right]$ in (33) with the expression in (34) we obtain

$$E[|y_t|^{2m}|y_{t-n}|^{2m}] = \nu_{2m}\exp\left\{m[(\omega_r+\omega_s)+\beta_{1s}^n(\omega_r-\omega_s)]\right\}$$

$$E\left[\eta_{t-n}^{2m} \times \exp\left\{m(1-L)^{-d}\sum_{i=1}^{n}\beta_{1s}^{i-1}g_s(\eta_{t-i})\right\}\right] \times$$

$$\prod_{i=1}^{\infty} E\left[\exp\{m(1+\beta_{1s}^n)\lambda_{i,r}g_r(\eta_{t-n-i})\}\right]$$

denoting $\chi_s(L) = (1-L)^{-d}\sum_{i=1}^{n}\beta_{1s}^{i-1}L^{i-1} = 1 + \chi_{1,s}L + \ldots$, then

$$E[|y_t|^{2m}|y_{t-n}|^{2m}] = \nu_{2m}\exp\left\{m[(\omega_r+\omega_s)+\beta_{1s}^n(\omega_r-\omega_s)]\right\} E\left[|\eta_{t-n}|^{2m}\prod_{i=0}^{\infty}\exp\left\{m\chi_{i,s}g_s(\eta_{t-1-i})\right\}\right] \times$$

$$\prod_{i=1}^{\infty} E\left[\exp\{m(1+\beta_{1s}^n)\lambda_{i,r}g_r(\eta_{t-n-i})\}\right]$$

$$= \nu_{2m}\exp\left\{m[(\omega_r+\omega_s)+\beta_{1s}^n(\omega_r-\omega_s)]\right\} E\left[|\eta_{t-n}|^{2m}\exp\left\{m\chi_{n-1,s}g_s(\eta_t)\right\}\right]$$

$$\prod_{i=0}^{n-2} E\left[\exp\{m\chi_{i,s}g_s(\eta_t)\}\right] \prod_{i=n}^{\infty} E\left[\exp\{m\chi_{i,s}g_s(\eta_t)\}\right] \prod_{i=0}^{\infty} E\left[\exp\{m(1+\beta_{1s}^n)\lambda_{i,r}g_r(\eta_t)\}\right]$$

Furthermore, we know that

$$E[y_t^{2m}] = \nu_{2m}\exp\{m\,\omega_s\}\prod_{i=0}^{\infty} E\left[\exp\{m\lambda_{i,s}g_s(\eta_t)\}\right]$$

The numerator in (14) is equal to

$$E[|y_t|^{2m}|y_{t-n}|^{2m}]-[E(|y_t|^{2m})][E(|y_{t-n}|^{2m})] = \nu_{2m}\exp\left\{m[(\omega_r+\omega_s)+\beta_{1s}^n(\omega_r-\omega_s)]\right\}\times$$

$$E\left[|\eta_t|^{2m}\exp\left\{m\chi_{n-1,s}g_s(\eta_t)\right\}\right]\prod_{i=0}^{n-2}E\left[\exp\left\{m\chi_{i,s}g_s(\eta_t)\right\}\right]\times$$

$$\prod_{i=n}^{\infty}E\left[\exp\left\{m\chi_{i,s}g_s(\eta_t)\right\}\right]\prod_{i=0}^{\infty}E\left[\exp\{m(1+\beta_{1s}^n)\lambda_{i,r}g_r(\eta_t)\}\right]-$$

$$\nu_{2m}^2\exp\{m(\omega_s+\omega_r)\}\left(\prod_{i=0}^{\infty}E\left[\exp\{m\lambda_{i,s}g_s(\eta_t)\}\right]\right)\times$$

$$\left(\prod_{i=0}^{\infty}E\left[\exp\{m\lambda_{i,r}g_r(\eta_t)\}\right]\right).$$

The variance of $|y_t|^{2m}$ is given by

$$E[|y_t|^{4m}]-(E[|y_t|^{2m}])^2 = \nu_{4m}\exp\{2m\omega_s\}\prod_{i=0}^{\infty}E\left[\exp\{2m\lambda_{i,s}g_s(\eta_t)\}\right]$$

$$-\nu_{2m}^2\exp\{2m\,\omega_s\}\left(\prod_{i=0}^{\infty}E\left[\exp\{m\lambda_{i,s}g_s(\eta_t)\}\right]\right)^2$$

then the denominator is

$$\exp\{m\,(\omega_r+\omega_s)\}\left[\nu_{4m}\prod_{i=0}^{\infty}E\left[\exp\{2m\lambda_{i,s}g_s(\eta_t)\}\right]-\nu_{2m}^2\left(\prod_{i=0}^{\infty}E\left[\exp\{m\lambda_{i,s}g_s(\eta_t)\}\right]\right)^2\right]^{1/2}\times$$

$$\left[\nu_{4m}\prod_{i=0}^{\infty}E\left[\exp\{2m\lambda_{i,r}g_r(\eta_t)\}\right]-\nu_{2m}^2\left(\prod_{i=0}^{\infty}E\left[\exp\{m\lambda_{i,r}g_r(\eta_t)\}\right]\right)^2\right]^{1/2}$$

The same procedure can be followed for the SFI-PEGARCH(1,d,0) simply setting $\lambda_s(L)\equiv(1-L^S)^{-d}(1-\beta_{1s}L)^{-1}$, and $\chi_s(L)=(1-L^S)^{-d}\sum_{i=1}^{n}\beta_{1s}^{i-1}L^{i-1}$.

Using theorem A1.1 in Nelson (1991) we can further simplify the expressions for the kurtosis and the periodic autocorrelation function in case of normality for the error term. Under the hypothesis of $\eta_t\sim i.i.d.N(0,1)$, we have that:

$$E\left[\exp\left\{m\lambda_{i,s}g_s(\eta_t)\right\}\right]=\left\{\Phi[m(\psi_s+\gamma_s)\lambda_{i,s}]\exp\left[m^2\lambda_{i,s}^2(\gamma_s+\psi_s)^2/2\right]+\right.$$

$$\Phi[m(\psi_s-\gamma_s)\lambda_{i,s}]\exp\left[m^2\lambda_{i,s}^2(\psi_s-\gamma_s)^2/2\right]\right\}\times$$

$$\left.\exp\left[-m\lambda_{i,s}\psi_s\sqrt{(2/\pi)}\right]<\infty. \tag{35}$$

for finite real scalars m and $\lambda_{i,s}$. For $m>0$ and any finite real $\chi_{n-1,s}$ (see Lemma 1 in He,

Teräsvirta, and Malmsten, 2002) we have that

$$
\begin{aligned}
E\left[|\eta_t|^{2m}\exp\{m\chi_{n-1,s}g_s(\eta_t)\}\right] &= \frac{1}{\sqrt{2\pi}}\Gamma(2m+1)\exp\left\{-\sqrt{\frac{2}{\pi}}m\psi_s\chi_{n-1,s}\right\}\times \\
&\quad \exp\left\{\frac{1}{4}m^2(\gamma_s+\psi_s)^2\chi_{n-1,s}^2\right\}\times \\
&\quad \left\{D_{-(2m+1)}[-m\chi_{n-1,s}(\gamma_s+\psi_s)]+\exp\left\{-m^2\chi_{n-1,s}^2(\gamma_s\psi_s)\right\}\times\right. \\
&\quad \left. D_{-(2m+1)}[-m\chi_{n-1,s}(\psi_s-\gamma_s)]\right\}
\end{aligned}
\tag{36}
$$

where $D_{-(2m+1)}(\cdot)$ is the parabolic cylinder function (see Gradshteyn and Ryzhik, 1980), defined as

$$
D_{(-p)}[q] = \frac{\exp\{-q^2/4\}}{\Gamma(p)}\int_0^\infty x^{p-1}\exp\{-qx-x^2/2\}dx \quad p > 0
\tag{37}
$$

and $\Gamma(\cdot)$ is the gamma function. When $p = 2$, He, Teräsvirta, and Malmsten (2002) show that

$$
\int_0^\infty x^2\exp\left(-qx-\frac{x^2}{2}\right)dx = -q+\sqrt{2\pi}(1+q^2)\Phi(-q)\exp\left(q^2/2\right).
\tag{38}
$$