Modeling long memory in the EU stock market: Evidence from the STOXX 50 returns

Sonia R. Bentes[*1], Nuno Ferreira[#2]

[*] *ISCAL, Av. Miguel Bombarda 20, 1069-035 Lisbon, Portugal and BRU-IUL, Av. das Forcas Armadas, 1649-026 Lisbon, Portugal*
[1]smbentes@iscal.ipl.pt

[#]*Department of Quantitative Methods, IBS-ISCTE Business School, ISCTE Avenida das Forças Armadas, Lisboa, Portugal*
[1]nuno.ferreira@iscte.pt

Abstract **- This paper examines the persistence behaviour of STOXX 50 returns. To this end, we estimated the GARCH, IGARCH and FIGARCH models based on a data set comprising the daily returns from January 5th, 1987 to December 27th, 2013. The results show that the long-memory in the volatility returns constitutes an intrinsic and empirically significant characteristic of the data and are, therefore, in consonance with previous evidence on the subject. Moreover, our findings reveal that the FIGARCH is the best model to capture linear dependence in the conditional variance of the STOXX 50 returns as given by the information criteria.**

Keywords - *Stock long-memory; persistence; volatility; conditional variance; FIGARCH.*

1. Introduction

A major topic of research in Finance concerns the degree of persistence or long memory in stock returns. By long memory we mean a high degree of persistence in the time series. In particular, this concept can be generally expressed either in the time domain or in the frequency domain. In the time domain, long memory is characterized by a slow decay of the autocorrelation function of time series with larger sample data. This means that observations far from each other are still strongly correlated and decay at a slow rate. On the other hand, in the frequency domain the same information comes in the form of a spectrum. Thus, if the spectral density is unbounded at low frequencies, time series is said to exhibit a long memory process. Though both definitions are not equivalent they are related by the Hurst exponent – H (Beran, 1994), who was the first author to document this property in nature. Motivated by the desire to understand the persistence of the steam flow and therefore the design of reservoirs Hurst (1951) analyzed 900 geophysical time series and found significant long-term correlations among the fluctuation of the river Nile outflow. After his seminal work several other studies documented persistence in very distinct domains of Science, such as, Biology, Geophysics, Climatology and Economics, *inter alia*.

In addition, this phenomenon has also gathered much attention in Finance, which bases on an alternative approach built on the ARCH (Autoregressive Conditional Heteroskedasticity) type models to capture persistence in time series. In particular, the GARCH (General ARCH) process introduced by Bollerslev (1986) has become quite popular in modelling conditional dependence in financial volatility. However, though this constitutes an effective tool to capture volatility clustering it has revealed inappropriate to accommodate for persistence since it assumes that shocks decay at a fast exponential rate. Therefore, it is only suited to account for short-run effects. In an attempt to overcome this limitation Engle and Bolerslev (1986) developed the IGARCH (Integrated GARCH) framework, which allows infinite persistence. However, the infinite memory is a very unrealistic assumption, which motivated the need for an alternative approach. In the light of this, Baillie *et al.*

IJLTFES, E-ISSN: 2047-0916

(1996) formulated the FIGARCH (Fractional Integrated) model, which characterizes an intermediate range of persistence. This is accomplished through the introduction of the fractional difference parameter *d*. Another advantage of this model is that it nets both GARCH, for $d = 0$, and IGARCH, for $d = 1$, processes as special cases.

Bearing on these models we investigate the degree of persistence of the EU stock market using the STOXX 50 index as a proxy. We use a data set comprising daily data from January 5th, 1987 to December 27th, 2013. The empirical results are, therefore, intended to broaden the available evidence to date and to characterize conditional dependence in the volatility process of the EU stock market. Moreover, we report estimates of the GARCH, IGARCH and FIGARCH parameters and use the information criteria to discriminate between models. Our aim is to find which model is more appropriate to characterize the dependence in the volatility process.

The rest of the paper proceeds as follows. Section 2 outlines the conditional volatility models, which is followed by a description and initial analysis of the data in Section 3. Section 4 presents the empirical results and, finally, Section 5 concludes the paper.

2. Model framework

2.1 GARCH model

Following Engle (1982) consider the time series y_t and the associated prediction error $\varepsilon_t = y_t - E_{t-1}[y_t]$, where $E_{t-1}[.]$ is the expectation of the conditional mean on the information set at time *t*-1. The standard GARCH(*p*,*q*) model introduced by Bollerslev (1986) is defined as:

$$\sigma_t^2 = \omega + \alpha(L)\varepsilon_t^2 + \beta(L)\sigma_t^2, \qquad (1)$$

where $\omega > 0$, $\alpha \geq 0$, $\beta \geq 0$, L denotes the lag or backshift operator $\alpha(L) \equiv \alpha_1 L + \alpha_2 L^2 + ... + \alpha_q L^q$ and $\beta(L) \equiv \beta_1 L + \beta_2 L^2 + ... + \beta_p L^p$. Hence, according to the GARCH formulation the conditional variance is a function of (i) the past squared residuals and of (ii) the lagged values of the variance itself (Daly, 2008). In general, a GARCH(1,1) specification is sufficient to describe the vast majority of time series and rarely is any higher order model estimated. The great advantages of this model when compared to the seminal ARCH of Engle (1982) are two-fold: (i) it is more parsimonious and (ii) avoids overfitting. Consequently, the model is less prone to breach non-negativity constraints.

A common empirical finding in most applied work is the apparent persistence implied by the estimates for the conditional variance. This is manifested by the presence of an approximate unit root in the autoregressive polynomial, that is, $\alpha + \beta \approx 1$, meaning that shocks are infinitely persistent (Bollerslev *et al.*, 1992). As the GARCH formulation considers that shocks decay at a fast geometric rate this specification is not appropriate to describe long memory, being only suited to accommodate for short-memory phenomena.

2.2 IGARCH model

In order to overcome this limitation Engle and Bollerslev (1986) derived the IGARCH model, which captures *I*(1) type processes for the conditional variance as infinite persistence remains important for forecasts of all horizons. Assuming that $\upsilon_t \equiv \varepsilon_t^2 - \sigma_t^2$ the GARCH model can be re-written in the form of an ARMA(*m*,*p*) process

$$\Phi(L)(1-L)\varepsilon_t^2 = \omega + [1-\beta(L)]\upsilon_t, \qquad (2)$$

where $\Phi(L) = [1-\alpha(L)-\beta(L)](1-L)^{-1}$ and all roots of $\Phi(L)$ and $[1-\beta(L)]$ lie outside the unit root circle.

However, despite its insight when compared to its predecessor this model is not fully satisfactory in describing long memory in the volatility process as shocks in the IGARCH methodology never die out.

2.3 FIGARCH model

In an attempt to describe the long memory process in a more realistic way Baillie *et al.* ([1] introduced a new class of models called FIGARCH. In contrast to an *I*(0) time series in which shocks die out at a fast geometric rate or an *I*(1) time series where there is no mean reversion, shocks to an *I*(*d*) time series with $0 < d < 1$ decay at a very slow hyperbolic rate.

The FIGARCH(*p*,*d*,*q*) model can be obtained by replacing the difference operator in Eq. (2) with a fractional differencing operator $\left(1-L\right)^d$ as in the following expression:

$$\Phi\left(L\right)\left(1-L\right)^d \varepsilon_t^2 = \omega + \left[1-\beta\left(L\right)\right]\upsilon_t. \quad (3)$$

Rearranging the terms in Eq. (3), one can write the FIGARCH model as follows:

$$\left[1-\beta\left(L\right)\right]\varepsilon_t^2 = \omega + \left[1-\beta\left(L\right)-\Phi\left(L\right)\left(1-L\right)^d\right]\varepsilon_t^2. \quad (4)$$

The conditional variance of ε_t^2 is obtained by:

$$\sigma_t^2 = \frac{\omega}{\left[1-\beta\left(L\right)\right]} + \left[1-\frac{\Phi\left(L\right)}{\left[1-\beta\left(L\right)\right]}\left(1-L\right)^d\right]\varepsilon_t^2, \quad (5)$$

which corresponds to

$$\sigma_t^2 = \frac{\omega}{\left[1-\beta\left(L\right)\right]} + \lambda\left(L\right)\varepsilon_t^2, \quad (6)$$

where $\lambda\left(L\right) = \lambda_1 L + \lambda_2 L^2 \ldots$

The FIGARCH metodology provides greater flexibility for modeling the conditional variance, because it accommodates the covariance stationary GARCH model when $d = 0$ and the IGARCH model when $d = 1$, as special cases. For the FIGARCH model the persistence of shocks to the conditional variance or the degree of long memory is measured by the fractional differencing parameter d. Thus, the attraction of this methodology is that for $0 < d < 1$, it is sufficiently flexible to allow for an intermediate range of long memory. It is worthy to note that the parameters of the FIGARCH model can be estimated by an approximate quasi-maximum likelihood estimation technique (Bollerslev and Wooldridge, 1992).

3. Data and some preliminary statistical results

The data employed in this study consist of the daily closing prices for the STOXX 50 during the period from January 5th, 1987 to December 27th, 2013, which totals 7040 observations. STOXX 50 index was designed to provide a Blue-chip representation of supersector leaders in the Eurozone and covers 50 stocks from 12 Eurozone countries: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal and Spain. Because the STOXX 50 is considered to be a proxy of the overall Eurozone stock market, it is frequently used as the underlying index for several derivative financial instruments such as options, futures, index funds and structured products. Figure 1 plots the STOXX 50 daily returns'.

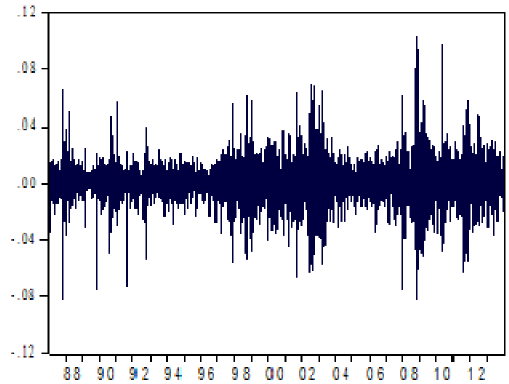

Figure 1. STOXX 50 daily returns'

More specifically, we define the STOXX 50 daily returns' as

$$R_t = \ln P_t - \ln P_{t-1}. \quad (7)$$

where R_t denotes the index returns at time t, and P_t and P_{t-1}, prices at time t and t-1, respectively. The data were collected from Datastream database.

Preliminary analysis for the STOXX 50 returns over the period under consideration is displayed in Table 1. Starting with the descriptive statistics we find that the average daily returns are positive and very small when compared to the standard deviation. The series is also characterized by negative skewness and high levels of positive kustosis, indicative of a heavier tailed distribution than the Normal. Jarque-Bera (J-B) test further confirms departure from normality, which can be graphically observed by looking at the plot of the histogram (Fig. 2). Accordingly, these results encourage the adoption of an alternative distribution, which embodies these features of the data, such as, the GED (Generalized Error Distribution) distribution.

Table 1. Preliminary analysis of STOXX 50 daily returns

Panel A. Summary statistics

Mean	Std. Dev.	Skewness	Kurtosis	J-B	$Q(10)$	BG(10)
0.000176	0.013272	-0.15006	9.031538	10696.21**	56.159**	5.549556**

Panel B. Heteroskedasticity tests

ARCH-LM	$Q^2(10)$
155.5133**	3792.4**

Panel C. Unit root tests

	ADF	KPSS
intercept	-40.0337**	0.178982
trend + intercept	-40.04168**	0.086268

Notes: 1. Denotes significantly at the 1% level. 2. J-B represents the statistics of the Jarque and Bera (1987)'s normal distribution test. 3. $Q(10)$ denotes the Ljiung-Box Q test for the 10th order serial correlation for the null hypothesis of no autocorrelation. 4. BG reveals the Breusch-Godfrey test for the null hypothesis of no serial correlation up to 10 lags. 5. ARCH-LM refers to the ARCH test for the null of no autoregressive conditional heteroskedasticity up to 10 lags. 6. $Q^2(10)$ is the Ljiung-Box Q test for serial correlation in the squared standardized residuals with 10 lags, which is used to check for ARCH effects. 7. ADF defines the Augmented Dickey and Fuller (1979) test for the null of non-stationarity. Critical values: -3.43 (1%) and -2.86 (5%) for constant and -3.96 (1%) and -3.41 (5%) for constant and linear trend. 8. KPSS indicates the Kwiatkowski, Phillips, Schmidt and Shin (1992) test for the null of stationarity. Critical values: 0.739 (1%) and 0.463 (5%) for constant and 0.216 (1%) and 0.146 (5%) for constant and linear trend.

Moreover, the Ljiung-Box Q-statistics ($Q(10)$) and the Breusch-Godfrey statistics both reveal serial correlation. In addition, the results for the ARCH-LM and $Q^2(10)$ tests are highly significant, thus unveiling ARCH effects.

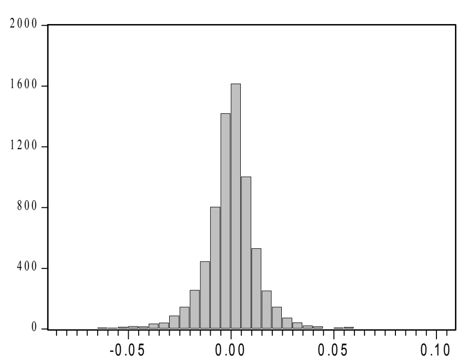

Figure 2. Histogram of the daily returns of the STOXX 50

Turning to the unit root tests (Table 1 – Panel C), an analysis of the ADF and KPSS statistics unfolds that the return series is stationary. In fact, the ADF test rejects the null hypothesis of non-stationarity at the 1% significance level, while for the KPSS the null of stationarity in not rejected. Both tests were estimated considering a constant and a constant and a linear trend in the exogenous regressors.

4. Model estimates

In order to remove any serial correlation present in the data[1] we first fit an AR(p) model (Autoregressive model). In this context we chose an AR(5) specification to account for this feature of the data. Table 2 provides the residual analysis for this specification. Results show that the residuals are non-normally distributed as the Jarque-Bera test is rejected at the 1% level. This is also implied by the negative skewness and excess kurtosis evidenced by the residual series. In addition, the Ljiung-Box and Breusch-Godfrey statistics both indicate no evidence of serial dependence, thus revealing that this specification is adequate in describing the linear dependence in the data. Nevertheless, residuals still exhibit conditional heteroskedasticity given that the ARCH-LM and the Ljiung-Box statistics of the squared residuals are all significant at the 1% level. Therefore, a specification that accounts for this property should be used to model the data.

Table 2. Residual's analysis for the fitted AR(p) model

Panel A. Summary statistics

Mean	Std. Dev.	Skewness	Kurtosis	J-B	$Q(10)$	BG(10)
-2.37E-20	0.013227	-0.247482	8.742039	9735.055**	7.5577	0.984831

Panel B. Heteroskedasticity tests

ARCH-LM	$Q^2(10)$
156.7961**	3825.3**

Notes: 1. Denotes significantly at the 1% level. 2. J-B represents the statistics of the Jarque and Bera (1987)'s normal distribution test. 3. $Q(10)$ denotes the Ljiung-Box Q test for the 10th order serial correlation for the null hypothesis of no autocorrelation. 4. BG reveals the Breusch-Godfrey test for the null hypothesis of no serial correlation up to 10 lags. 5. ARCH-LM refers to the ARCH test for the null of no autoregressive conditional heteroskedasticity up to 10 lags. 6. $Q^2(10)$ is the Ljiung-Box Q test for serial correlation in the squared standardized residuals with 10 lags, which is used to check for ARCH effects. 7. ADF defines the Augmented Dickey and Fuller (1979) test for the null of non-stationarity. Critical values: -3.43 (1%) and -2.86 (5%) for constant and -3.96 (1%) and -3.41 (5%) for constant and linear trend. 8. KPSS indicates the Kwiatkowski, Phillips, Schmidt and Shin (1992) test for the null of stationarity. Critical values: 0.739 (1%) and 0.463 (5%) for constant and 0.216 (1%) and 0.146 (5%) for constant and linear trend.

Given this, we start by fitting the standard GARCH(1,1). Subsequently, in order to account for persistence the IGARCH(1,1) and FIGARCH(1,d,1) specifications are also estimated. All parameters were estimated by quasi maximum likelihood estimation method in terms of the BFGS optimization algorithm using the econometric package of OxMetrics 5.00. Since the returns follow a distribution with thicker tails than the normal, as shown in Section 3, we assumed a GED distribution for estimation purposes. Model estimates and diagnostic tests are provided in Table 3. As one can observe, the parameters ω, α, β of the GARCH equation and the tail coefficient of the GED distribution are all positive and found to be highly significant. Further, there is also evidence of strong persistence in the return series, as $\alpha + \beta \approx 1$, which motivated the estimation of the IGARCH(1,1) and FIGARCH(1,d,1) models. As in the GARCH case, results for these specifications uncover positive and highly significant parameters at the 1% level. A fractional difference parameter of 0.46053 was found for the return series, which shows a moderate level of persistence.

A number of diagnostic tests were then performed in order to assess the adequacy of these models in describing the returns volatility. According to our results, the p-values of the Ljung-Box Q statistic test at lag 10 of the standardized residuals for all models fail to reject the null of no autocorrelation at the 1% significance level. Therefore, all models appear to be adequate in describing the linear dependence in the return series. In addition, for all the three models considered the ARCH-LM(10) test cannot reject the null hypothesis of no ARCH effects. This is corroborated by the Ljiung-Box statistics of the squared residuals, which is significant at the 1% level, thus unfolding that these specifications are sufficient to capture conditional heteroskedasticity in the conditional variance equation. Finally, there is still evidence of non-normality in the residual series as the Jarque-Bera test rejects the null of Gaussianity at the 1% level.

Having estimated these three specifications one question remains to be answered: which one is the best model to describe conditional dependence in the volatility process? In order to discriminate between models we employ the LL (log-likelihood), AIC (Akaike Information Criterion) and SIC (Schwarz Information Criterion) information criteria. According to Sin and White [13] the most appropriate model to describe the data is the one that maximizes the LL function and minimizes the SIC and AIC criteria. In our particular case, the model that fulfills these conditions is the FIGARCH(1,d,1) model. This is not surprising as the results obtained in the GARCH formulation also suggested persistence since the sum of α and β is very close to 1.

Table 3. GARCH(1,1), IGARCH(1,1) and FIGARCH (1,d,1) estimates

Parameter	GARCH (1,1)	IGARCH (1,1)	FIGARCH (1,*d*,1)
ω	0.019349**	0.014945**	1.579819**
	(0.0000)	(0.0001)	(0.0053)
d	---	---	0.46053**
	---	---	(0.0000)
α	0.096188**	0.104375**	0.121130**
	(0.0000)	(0.0000)	(0.0020)
β	0.894336	0.895625**	0.518799**
	(0.0000)	(0.0000)	(0.0000)
GED	1.31078	1.295030**	1.321006**
	(0.0000)	(0.0000)	(0.0000)
LL	22026.727	22023.625	22041.272
AIC	-6.255641	-6.255043	-6.259489
SIC	-6.245896	-6.246273	-6.24877
J-B	14696**	18727**	15000**
	(0.00000)	(0.00000)	(0.00000)
Q(10)	7.49615	7.64307	8.61503
	(0.1862771)	(0.1770350)	(0.1254403)
ARCH-LM	0.13182	0.13897	0.11751
	(0.9994)	(0.9992)	(0.9996)
$Q^2(10)$	1.32262	1.38635	1.18414
	(0.9952764)	(0.9944382)	(0.9967962)

Notes: 1. The p-values are included in parenthesis. 2. GED refers to the tail coefficient of the GED distribution. 3. LL refers to the log-likelihood value 4. SIC designates the Schwarz Information Criterion. 5. AIC denotes the Akaike Information Criterion. 6. ** Indicates the rejection of the null hypothesis at the 1% significance level. 7. J-B represents the statistics of the Jarque and Bera (1987)'s normal distribution test. 8. *Q*(10) denotes the Ljiung-Box *Q* test for the 10th order serial correlation for the null hypothesis of no autocorrelation. 9. ARCH-LM refers to the ARCH test for the null of no autoregressive conditional heteroskedasticity up to 10 lags. 10. $Q^2(10)$ is the Ljiung-Box *Q* test for serial correlation in the squared standardized residuals with 10 lags, which is used to check for ARCH effects.

5. Conclusions

In order to examine the degree of persistence of the STOXX 50 daily returns' we estimated three different models of conditional volatility – GARCH, IGARCH and FIGARCH models. To this end, we collected data from January 5th, 1987 to December 27th, 2013.

Basically, we conducted our study in three steps: first, we provided a preliminary analysis on the data based on the descriptive statistics of the variable under consideration. Our results showed that the STOXX 50 returns did not follow a normal distribution. In addition, we demonstrated that though prices were non-stationary returns were stationary, thus enabling further analysis. Moreover, we also found serial correlation and conditional heteroskedasticity in the return series. Second, and in order to capture the autocorrelation we fitted an AR(5) model, which proved to be sufficient to remove any serial dependence in the series. Nonetheless, heteroskedasticity was still present in the returns. Finally, we estimated the GARCH, IGARCH and FIGARCH parameters and used the information criteria to discriminate between models. Our results showed a moderate level of persistence (*d* = 0.46053). Furthermore, the FIGARCH model was proven to be the best model to describe the data. Finally, residual tests also revealed absence of serial correlation and ARCH effects.

References

[1] Baillie, R.T., T. Bollerslev and H.O. Mikkelsen (1996), Fractionally integrated generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 74, 3-30.

[2] J. Beran (1994), *Statistics for Long Memory Processes*, New York: Chapman and Hall.

[3] Bollerslev, T. (1986), Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307-327.

[4] Bollerslev, T. and J.M. Wooldridge (1992), Quasi-maximum likelihood estimation of dynamic models with time varying covariances, *Econometric Reviews* 11, 143-172.

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

[6] Daly, K. (2008), Financial volatility: Issues and measuring techniques *Physica A* 387, 2377-2393.

[7] Dickey, D.A. and W.A. Fuller (1979), Distribution of the estimator for autoregressive time series with a unit root, *Econometrica* 49, 1057-1072.

[8] Engle, R.F. (1982), Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* 50, 987-1007.

[9] Engle, R.F. and T. Bollerslev (1986), Modelling the persistence of conditional variances, *Econometric Reviews* 5, 1-50.

[10] Hurst, H.E. (1951), Long term storage capacities of reservoirs, *Transactions of the American Society of Civil Engineers* 116, 770-799.

[11] Jarque, C.M. and A.K. Bera (1987), A test of normality of observations and regression residuals, *International Statistics Review* 55, 163-172.

[12] Kwiatkowski, D., P.C.B, Phillips, P. Schmidt, and Y. Shin (1992), Testing the null hypothesis of stationarity against the alternative of a unit root, *Journal of Econometrics* 54, 159-178.

[13] Sin, C.Y. and H. White (1996), Information Criteria for selecting possibly misspecified parametric models, *Journal of Econometrics* 71, 207-325.