ELSEVIER

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Journal of Empirical Finance 12 (2005) 187–215

Journal of EMPIRICAL FINANCE

www.elsevier.com/locate/econbase

European exchange rate volatility dynamics: an empirical investigation

Ali Khalil Malik*

Manchester School of Management, University of Manchester Institute of Science and Technology (UMIST), P.O. Box 88, Manchester M60 1QD, UK

Accepted 25 September 2003
Available online 25 March 2004

Abstract

The hourly and daily dummy variables in the conditional variance functions of the two European currencies, the British pound and the euro are estimated. The conditional variance functions are specified as GARCH models to capture the time-dependent conditional heteroskedasticty at both the hourly and daily recorded data. The estimated dummy variables give remarkably similar and distinct characteristics of the volatility dynamics embodied in the two currencies. Some discussion of the possible sources of the observed volatility dynamics in the two currencies is provided. A comparison between GARCH, FIGARCH and SV models is also provided. The estimated hourly and daily dummy variables suggests that euro is considerably more volatile when compared to British pound, a result with important implications for the economic policy making in the two regions.

JEL classification: F31; C22
Keywords: Conditional variance; Dummy variables; GARCH; Private information; Volatility

1. Introduction

As the debate intensifies on Britain's joining Europe's new single currency, the task of economists is manifold. A rational evaluation of the gains and losses to be made following the decision to join has to be produced. An economic perspective to the problem could definitely lead to a better resolution than relying on cultural or sentimental perspectives. Within this economic perspective lies many dimensions which

* Tel.: +44-7984-643453.
E-mail address: a.malik@postgrad.umist.ac.uk (A.K. Malik).

doi:10.1016/j.jempfin.2003.09.004

need to be addressed and one such dimension is the volatility of the British pound and the euro. Volatility is definitely an important financial phenomenon which needs to be addressed. Questions such as: which of the two currencies is more volatile, do the currencies follow the same intra- and inter-daily volatility patterns, does the volatility in the two currencies depict the underlying volatility in their economic fundamentals, etc., need to be answered. Answers to these questions are important in determining the relative trade and financial position of a country.

The foreign exchange market is the world's biggest financial market; in operation 24 h a day, 7 days a week, it is the closest analogue to the concept of a continuous time global market place. This particular study will specifically focus on just one aspect of trading in the foreign exchange market, i.e. the volatility of exchange rates and that too the focus would be on just two currencies: the euro and the British pound (BP). An enormous amount of research on the volatility of exchange rates has been done in the past decade or so. However, research along this dimension, i.e. exclusively providing an insight into the BP and euro argument is still pending. Considering the importance of this issue, some research on the volatility of BP and euro is conducted making use of the high frequency data. After identifying the hourly volatility patterns embodied in the two currencies the research is subsequently extended to a daily basis. The objective is to get a thorough analysis of the two currencies from a volatility dimension within the economic perspective.

Even though there is widespread agreement among the researchers that exchange rates are well described as martingales, the returns are not independently distributed over time because of 'volatility clustering'. It was this observation by Mandelbrot (1963) and Fama (1965) that periods of large absolute changes tend to cluster together, followed by periods of relatively small absolute changes that generated a huge amount of research on the volatility of financial market returns. The main econometric tool used in the paper is the well-known Generalised Autoregressive Conditional Heteroskedasticity (GARCH) model (Bollerslev, 1986) together with some of its recently developed extensions. The basic insight behind GARCH is that the second moment of the distribution of exchange rates is serially correlated. The rigorous treatment of the exchange rates using the fascinating tools of GARCH modelling would certainly make this research a valuable addition to the existing literature on the volatility of exchange rates and hopefully would pave the way for analysing Europe's new currency along this dimension. A comparison between GARCH and some of its recently developed versions and an alternative modelling technique is also provided using both the high and low frequency data, an exercise often found to be missing in the literature when comparing the superiority of different variance specifications.

The paper is organized as follows. Section 2 discusses the data and some preliminary testing. Section 3 estimates the hourly and daily volatility patterns embodied in the two currencies, i.e. the hourly and daily dummy variables in the conditional variance formulation of the two currencies and reviews some of the reasons documented in the literature on the observed volatility dynamics of exchange rates; their applicability to the volatility dynamics exhibited by euro and BP is also examined. Section 4 compares the GARCH with some of its recently developed extensions and an alternative modelling technique known as the stochastic volatility model. Section 5 concludes.

2. Data and preliminary testing

As mentioned above, the foreign exchange market is virtually continuously active around the clock with the same basic assets being traded in several different locations, permitting the examination of the behaviour of different currencies across different markets at fairly fine intervals of time. The spot exchange rates which we employ are for the British pound and the euro vis-à-vis the U.S. dollar, recorded on an hourly basis. The basic data set is supplied by Reuters Network and remains their property. The data set constitutes a total of 1892 trading hours for the BP and some slightly less trading hours for the euro, due to some missing observations. The period to which the data correspond starts in December 2001 and ends in March 2002 and is approximately the same for both the currencies. The data recorded is the last bid price coming in an hour or more specifically the closing bid price for the respective hour. We could be assured that these bid prices truly represent the exchange rates at which the deals were subsequently executed and thus are representative of the particular hour to which they correspond, since as argued by Bollerslev and Domowitz (1993) that while it would be possible to post quotes in the hope of influencing market movements and then subsequently refuse to deal at these prices, such behaviour is effectively precluded by reputation effects.[1] All observations from Friday 21:00 GMT through to Sunday 21:00 GMT have been excluded. Low activity over the weekend justifies this exclusion. This weekend definition is based on the analysis of the inter-bank foreign exchange market quotes by Bollerslev and Domowitz (1993).

The exchange rate for currency i at day t and hour τ is denoted by $s(i)_{t,\tau}$, where i = BP, euro; t = 1,2...112 and τ = 1,2...24.The rates are quoted as the amount of dollars per currency i.

To test for the non-stationarity of exchange rates, i.e. the presence of unit root in the exchange rate series, the Augmented Dickey–Fuller test (Dickey and Fuller, 1979, 1981) and Phillips and Perron (1988) test (which allows for weakly dependent and heterogeneously distributed innovations) were applied to the logarithm of exchange rates. Both tests however gave the unanimous conclusion that the null hypothesis of a unit root in the logarithm of the exchange rate series could not be rejected for any of the currency.[2] The conclusion that the null hypothesis of a unit root cannot be rejected for the hourly spot exchange rate series is consistent with the findings of Baillie and Bollerslev (1991) and Goodhart and Giugale (1993) for the hourly exchange rates. In light of the preliminary analysis, the research which follows will only make use of the first differences of the two exchange rates:[3]

$$y(i)_{t,\tau} = 100 \cdot [log(s(i)_{t,\tau}) - log(s(i)_{t,\tau-1})] \tag{1}$$

[1] This may however be only partially true, since as pointed out by the referee that in the real world the quoted bid/ask spreads are probably wider than the true bid/ask spreads.

[2] The ADF test statistics were computed with various number of autoregressive lags which made no difference to the result stated above. Similarly, the Phillips and Perron test statistics were also computed with different truncation lags without any alteration to the result stated above. The results of the unit root tests are not presented due to limitations of space but are available from author on request.

[3] Natural logarithms of the exchange rates have been taken for all the analysis conducted.

where s is the spot exchange rate. The series $y(i)_{t,\tau}$ corresponds to the approximate percentage nominal return on currency i obtained from time $t,\tau-1$ to t,τ.

Initially a simple regression equation with a MA(1) term in the conditional mean return equation, under the auxiliary assumptions of homoskedasticity and conditional normality is estimated for both the return series. The inclusion of MA(1) term is in accordance with literature on the analysis of hourly exchange rates by Baillie and Bollerslev (1991), Goodhart and Giugale (1993) and is in accordance with nearly all the research done on the analysis of the exchange rate returns using high frequency data. The argument forwarded by these authors justifies the inclusion of a MA(1) term on the basis of averaging the last five bid rates coming in an hour and also due to the non-synchroneity of those last five bid rates, which would naturally generate a negative first-order serial correlation in the return series. However, the argument presented corresponds to the data set used in the above mentioned studies but not for the return series utilised in this study, since the exchange rates used simply correspond to the last bid price coming in an hour.[4] To find the best model a detailed analysis of the correlogram of the exchange rate return series of both the currencies was performed. For the BP the autocorrelation at the first lag was marginally significant at the 5% level, while for the euro none of the autocorrelations or partial autocorrelations at any lags were significant. No other seasonal or cyclical patterns in the correlogram were visible for any of the currency.[5] Based on this analysis it was decided to include a MA(1) term in the conditional mean equation of both the currencies, so as not only to make the models comparable with each other but also to make them comparable with the existing literature. Even though more complicated parametric formulations for the conditional mean are possible, research has shown that the degree of predictability in the mean is very minor and inconsequential for the conditional variance formulations. The model initially estimated for the two return series is:

$$y(i)_{t,\tau} = u_0(i) + \theta_1(i)\epsilon(i)_{t,\tau-1} + \epsilon(i)_{t,\tau} \tag{2}$$

$$\epsilon(i)_{t,\tau} \mid \Omega_{t,\tau-1} \sim N(0, \omega_0(i)) \tag{3}$$

where $\Omega_{t,\tau}$ is the set of all relevant and available information at time t,τ. Table 1 reports the Quasi Maximum Likelihood Estimates (QMLE) obtained for this simple MA(1) model under the auxiliary assumptions of homoskedasticity and conditional normality. The standard errors reported in parentheses are the White heteroskedasticity-consistent standard errors. As noted by Baillie and Bollerslev (1991) that in the presence of severe kurtosis and heteroskedasticity (evident from the statistics reported in the table), the conventional standard errors reported are unreliable. In their analysis of the hourly exchange rates they report the more robust standard errors developed by Bollerslev and Wooldridge (1992), but subsequently also report the finding that under

[4] The negative serial correlation could also be induced from fluctuations across the bid/ask spread if the returns are calculated from both the bid and ask prices. Since the returns series used are calculated exclusively from the bid prices such an argument is not valid over here.

[5] The correlograms are available from the author on request.

Table 1
Quasi-maximum likelihood estimates

	BP	Euro
$u_0(i)$	− 0.00012 (0.00196)	− 0.00128 (0.00273)
$\theta_1(i)$	− 0.04763 (0.03368)	− 0.01081 (0.03491)
$\omega_0(i)$	0.08847	0.11446
$m_3(i)$	− 0.01709	− 0.43250
$m_4(i)$	6.08219	8.40103
$Q_{20}(i)$	15.299	21.037
$Q^2_{20}(i)$	188.89	135.88

White heteroskedasticity-consistent standard errors are in parentheses. The statistics $m_3(i)$ and $m_4(i)$ denote the standardized residuals skewness and kurtosis. $Q_{20}(i)$ and $Q^2_{20}(i)$ give the Ljung–Box portmanteau test for up to the 20th-order serial correlation in the standardized residuals and the squared standardized residuals, respectively.

the assumption of conditional normality (homoskedastic errors) the robust standard errors are equivalent to the standard errors developed by White (1982).

Looking at the results from the estimated models in Table 1, the MA(1) term is not significant for any of the return series using the heteroskedastic-consistent standard errors. Using the conventional standard errors (not reported here) the MA(1) term is marginally significant for the BP at the 5% level, which replicates the warning of Bollerslev and Wooldridge (1992) and Baillie and Bollerslev (1991) that when working with high frequency data use of the traditional standard errors could lead to wrong inference being made. This particular result is also in direct contradiction to the finding of Baillie and Bollerslev (1991), who find that other than for DM, the MA(1) term is significant for all the other three currencies analysed in their study (BP, SF and YEN). This result also reinforce the earlier point that when using a single bid price corresponding to an hour, the inclusion of a MA(1) term in the conditional mean return equation is completely an arbitrary decision. No doubts on the efficiency of the foreign exchange market could be made; the returns appear to be martingales, since no term in the conditional mean return equation for the two currencies is significant.[6]

However, before making any further conclusions it is important to check the efficiency of the estimated models itself. The statistics $m_3(i)$ and $m_4(i)$ are the usual sample estimates of skewness and kurtosis in the standardized residuals.Under the null hypothesis of normality, $\sqrt{6/Tm_3(i)}$ and $\sqrt{24/Tm_4(i)}$ both have asymptotic standard normal distributions. Severe excess kurtosis and negative skewness is a feature of the residual series for both the currencies. The $Q_{20}(i)$ and $Q^2_{20}(i)$ are the standard Ljung and Box (1978) portmanteau test statistics for up to 20th-order serial correlation in $\epsilon\hat{}(i)_{t,\tau}$ (standardized residuals) and $\epsilon\hat{}(i)^2_{t,\tau}$ (squared standardized residuals), respectively. Even though the $Q_{20}(i)$ statistics for the two currencies are insignificant, the Ljung–Box tests are indicative of a misspecification. The extremely high and significant values of $Q^2_{20}(i)$ (in the corresponding asymptotic chi-square distribution) strongly suggest the presence of

[6] As noted by Andersen and Bollerslev (1998) the large differences between the heteroskedasticity and autocorrelation consistent standard errors and the OLS standard errors signify the importance of accounting for the effects of strong volatility clustering and outlying observations when conducting the inference.

Table 2
LM diagnostic tests for the models in Table 1

	BP	Euro
ARCH (1)	82.0604 (0.0000)	37.0777 (0.0000)
ARCH (4)	107.9342 (0.0000)	81.0446 (0.0000)

p-values are reported in parentheses.

conditional heteroskedasticity.[7] Theses results are qualitatively similar to the ones reported by Baillie and Bollerslev (1991) and Goodhart and Giugale (1993) for the hourly exchange rate series and of Bollerslev and Domowitz (1993) for the 5-min return series. In fact, excess kurtosis is a feature of virtually all the work done on intra- and interday foreign exchange bid/ask returns. The presence of such excess kurtosis could be due to the presence of serial correlation in the returns variance process.

To further check the model specification and more importantly the presence of conditional heteroskedasticity in the two residual series, the LM test for the presence of ARCH(1) and ARCH(4) effects in the residuals of the estimated models is conducted. The results of the LM diagnostic tests together with the p-values are reported in Table 2. The LM test statistics for the presence of ARCH(1) and ARCH(4) effects are all highly significant at virtually any level, lending support to the presence of volatility clustering in the neighbouring hourly observations. Some other misspecification tests for the presence of higher order moving average and autoregressive terms in the conditional mean equation were also conducted but none of them was found to be significant. The presence of conditional heteroskedasticity seems to be the most important feature of the estimated models.

3. A model for hourly exchange rate movements

Based on the preliminary data analysis and the estimated models in the preceding section, we now move on to estimate the hourly dummy variables (representing the hourly volatility patterns) in the conditional variance specification of the two return series. For each of the hourly first differences of the logarithmic exchange rates the following model is estimated,

$$y(i)_{t,\tau} = u_o(i) + \theta_1(i)\epsilon(i)_{t,\tau-1} + \epsilon(i)_{t,\tau} \tag{4}$$

$$\epsilon(i)_{t,\tau} \mid \Omega_{t,\tau-1} \sim N(0, \sigma(i)^2_{t,\tau}) \tag{5}$$

$$\sigma(i)^2_{t,\tau} = \omega_\tau(i) - [\alpha_1(i) + \beta_1(i)]\omega_{\tau-1}(i) + \delta(i)V_{t,\tau} + \alpha(i)\epsilon(i)^2_{t,\tau-1} + \beta_1(i)\sigma(i)^2_{t,\tau-1} \tag{6}$$

[7] McLeod and Li (1994) have shown that the asymptotic distribution of the squared-residual autocorrelations does not depend on the parameter estimates for the conditional mean under the null hypothesis of homoskedasticity.

$V_{t,\tau}$ is the vacation dummy which takes the value of one for the hour immediately following the market being closed and zero otherwise. ω_1 to ω_{24} are the hourly dummy variables, one for their respective hour and zero otherwise. The conditional variance has a GARCH(1,1) formulation in the fourth and fifth terms, as in Bollerslev (1986) to describe the time-dependent volatility. The model specification is virtually similar to the specification in Baillie and Bollerslev (1991) for the hourly exchange rates.[8] The advantage of above specification as documented by Baillie and Bollerslev is that the initial two terms in the conditional variance equation use a multiplicative parameterization in $\alpha_1(i)$ and $\beta_1(i)$, so that, apart from any deterministic vacation effects, the unconditional variance for hour τ is equal to $\omega_\tau(i)$. The parameterization also implies that the effect of the vacation dummy in the conditional variance is geometrically declining.[9]

Table 3 presents the details of the estimated models. In the conditional mean equation, using the robust standard errors, the unconditional expected return term $u\hat{}_0(i)$ is insignificant for both the currencies. The MA(1) term remains insignificantly different from zero for the BP, while is now significant for the euro at the 5% level. Most of the parameters in the conditional variance equation are significant for both the currencies. Characteristics of the hourly dummy variables included in the variance specification are more or less similar for the two currencies. The estimated coefficients of the hourly dummies are small but most of them are significant. The estimated GARCH(1,1) parameters are highly significant for the two currencies, significantly different from zero at the conventional levels of significance.

The standard errors reported are the robust standard errors, conventional ones are not reported.[10] The standardized residuals for the models in Table 3 still possess a large degree of leptokurtosis, evident from the $m_4(i)$ statistics reported in the table. Although the degree of leptokurtosis is considerably less than the models estimated with conditional homoskedasticity, it is nonetheless significant.[11] This finding is in accordance with some of the previous findings in the literature for daily exchange rates in Baillie and Bollerslev (1989), Hsieh (1988, 1989) and for the intraday exchange rates in Engle et al. (1990), Baillie and Bollerslev (1991) and Goodhart and Giugale (1993). The $Q_{20}(i)$ and $Q^2_{20}(i)$ test statistics for the presence of serial correlation in the standardized and squared standardized residuals are not significant at the usual levels of significance. This result is in accordance with

[8] For the sake of simplicity, and to some extent clarity, we have chosen not to render our model even more complex by the inclusion of daily dummies. This paper follows previous practise by researchers in this area of including hourly but not daily dummies, in the hourly exchange rate models. See, for example, Baillie and Bollerslev (1991).

[9] As noted by a referee, the conditional variance specification might be improved by the inclusion of dummy variables for macroeconomic news releases. This is easier to do with the UK and U.S., but it is a formidable task to achieve for the EU which spans many countries. The author intends to focus on this issue in a separate study.

[10] Bollerslev and Wooldridge (1992) in a simulation study show that the conventional OPG standard errors tend to under-estimate the true parameter estimator uncertainty. Similarly Baillie and Bollerslev (1991) in an analysis of the hourly exchange rate series find that the conventional standard errors differ from the robust standard errors by a factor of two for the variance parameters, when conditional leptokurtosis is present.

[11] The lower values of the $m_4(i)$ statistics compared to values obtained for the models with conditional homoskedasticity indicates that the models are correctly specified. Since as documented by Hsieh (1989) that if the models are correctly specified then by Jensen's inequality the coefficients of kurtosis for the standardized residuals should be less than the kurtosis for the raw residuals.

Table 3
Quasi-maximum likelihood estimates

	BP	Euro
$u_0(i)$	− 0.00035 (0.00138)	− 0.00051 (0.00183)
$\theta_1(i)$	− 0.05280 (0.03162)	− 0.05883 (0.02701)
$\omega_1(i)$	0.00752 (0.00121)	0.00614 (0.00337)
$\omega_2(i)$	0.00348 (0.00059)	0.00134 (0.00273)
$\omega_3(i)$	0.00271 (0.00037)	0.00147 (0.00164)
$\omega_4(i)$	0.00419 (0.00080)	0.00646 (0.00461)
$\omega_5(i)$	0.00345 (0.00078)	0.00191 (0.00291)
$\omega_6(i)$	0.00241 (0.00031)	0.00226 (0.00143)
$\omega_7(i)$	0.00460 (0.00064)	0.01442 (0.00478)
$\omega_8(i)$	0.01356 (0.00199)	0.03317 (0.0110)
$\omega_9(i)$	0.01366 (0.00210)	0.01999 (0.0100)
$\omega_{10}(i)$	0.02002 (0.00264)	0.00670 (0.00921)
$\omega_{11}(i)$	0.02293 (0.00293)	0.02515 (0.0168)
$\omega_{12}(i)$	0.02569 (0.00395)	0.03920 (0.0297)
$\omega_{13}(i)$	0.01734 (0.00218)	0.00203 (0.0184)
$\omega_{14}(i)$	0.02193 (0.00252)	0.04459 (0.0153)
$\omega_{15}(i)$	0.02508 (0.00300)	0.02886 (0.0149)
$\omega_{16}(i)$	0.01506 (0.00178)	0.02744 (0.0139)
$\omega_{17}(i)$	0.01934 (0.00239)	0.00095 (0.0124)
$\omega_{18}(i)$	0.02318 (0.00363)	0.03469 (0.0203)
$\omega_{19}(i)$	0.00949 (0.00185)	0.00891 (0.0159)
$\omega_{20}(i)$	0.00849 (0.00123)	0.01055 (0.00528)
$\omega_{21}(i)$	0.00651 (0.00102)	0.00226 (0.00453)
$\omega_{22}(i)$	0.00411 (0.00106)	0.00405 (0.00407)
$\omega_{23}(i)$	0.00281 (0.00074)	0.00157 (0.00329)
$\omega_{24}(i)$	0.00274 (0.00034)	0.00702 (0.00278)
$\delta(i)$	0.00805 (0.00415)	0.00595 (0.00545)
$\alpha_1(i)$	0.09395 (0.03108)	0.09588 (0.03271)
$\beta_1(i)$	0.30669 (0.11283)	0.50993 (0.11074)
$m_3(i)$	0.13747	− 0.10188
$m_4(i)$	4.5171	5.0562
$Q_{20}(i)$	22.7759	29.2286
$Q^2_{20}(i)$	27.331	12.754

See footnote to Table 1. Asymptotic robust standard errors are in parentheses.

those obtained by Baillie and Bollerslev (1991) and Bollerslev and Domowitz (1993) for the intraday exchange rates.[12] The degree of skewness and kurtosis drops with the GARCH specification but is still significant. Lastrapes (1989) has shown that the inclusion of dummy variables in the conditional variance reduces the degree of leptokurtosis in the standardized residuals. The estimates of the hourly dummy variables $\hat{\omega}_1(i),\ldots,\hat{\omega}_{24}(i)$ display a very distinct and similar pattern for the two currencies. A more detailed analysis of the hourly volatility patterns using the graphical depiction follows shortly. As will be

[12] Box and Pierce (1970) have shown that when testing the residuals from an estimated ARMA model, the portmanteau test Q_k is asymptotically chi-square distributed with $K-k$ degrees of freedom, where k is the number of estimated ARMA parameters. Following this logic, Bollerslev and Mikkelsen (1996) recommend adjusting the degrees of freedom for the Q_k^2 test statistics by the number of estimated ARCH parameters.

seen the graphical analysis is probably the best way to understand the hourly volatility patterns exhibited by the two currencies. Similarities between the hourly dummy variables of the two currencies will also be evident from the graphs. Judged from the diagnostic statistics, GARCH(1,1) model does a reasonably good job of tracking the time series behaviour in the second moments of the two return series.

The significant estimates for $\alpha_1(i)$ and $\beta_1(i)$ could indicate the apparent autocorrelation of news as documented by Baillie and Bollerslev (1991), with the euro being the more persistent of the two. For none of the currencies there is any sign of Integrated GARCH or IGARCH characteristics, the sum of $\alpha_1+\beta_1$ being considerably less than 1 for both the currencies. This autocorrelation in the conditional second moments of the currencies, as documented by Baillie and Bollerslev (1991), is consistent with both the so-called "meteor shower" and "heat wave" hypothesis in the terminology of Engle et al. (1990). The GARCH specification of conditional variance for both the currencies was also estimated with an additional term $\alpha_{24}(i)\epsilon(i)^2_{t-1,\tau}$. The coefficient $\alpha_{24}(i)$ (i.e. whether high volatility yesterday at hour τ implies a higher volatility today at hour τ) was found insignificant for both the currencies, lending support to "meteor shower" type hypothesis of Engle et al. (1990), in the arrival of news to the market. The result is inferred using the robust QMLE standard errors.

The vacation dummy estimate is not significant for any of the currency at the 5% level. If the amount of new information that accumulates is proportional to the length of time that the market is closed, then volatility should be approximately 49 times greater following vacations when compared to the average volatility across the day, 49 being the average hourly length of a vacation in the sample. The effects as measured by the estimates for $\delta(i)$ are much smaller, which Baillie and Bollerslev (1991) interpret as being consistent with the notion of information as essentially a private phenomenon, with market participants acting on their own acquisition of new information. This result confirms the previous empirical finding of French and Roll (1986) for the U.S. stock market.

To assess the general specification validity of the models, some misspecification tests were performed. To check for the presence of any further ARCH effects in the residuals of the estimated models, the Engel's ARCH LM test was computed. The LM test statistics for the presence of any further ARCH effects were not found significant for any of the currency and the null of no ARCH effects could not be rejected. LM test statistics for the GARCH in mean effects, i.e. whether or not the conditional mean return is linearly related to its own conditional variance, were also computed. However, the LM statistics were found to be insignificant and as such no evidence for the presence of a time-varying risk premium was evident. In brief further testing indicates a reasonably good fit for the two univariate time-series models conditional on the own past history of each of the exchange rates. The result is in accordance with the finding of Baillie and Bollerslev (1991) that GARCH(1,1) model does a good job of tracking the short run intra-daily volatility dependencies.

3.1. Intraday volatility behaviour

As already mentioned, the estimated hourly dummy variables in the conditional variance equations possess a very characteristic and remarkably similar pattern across both the currencies. Looking at the graphs of the estimated dummies in Figs. 1 and 2, some

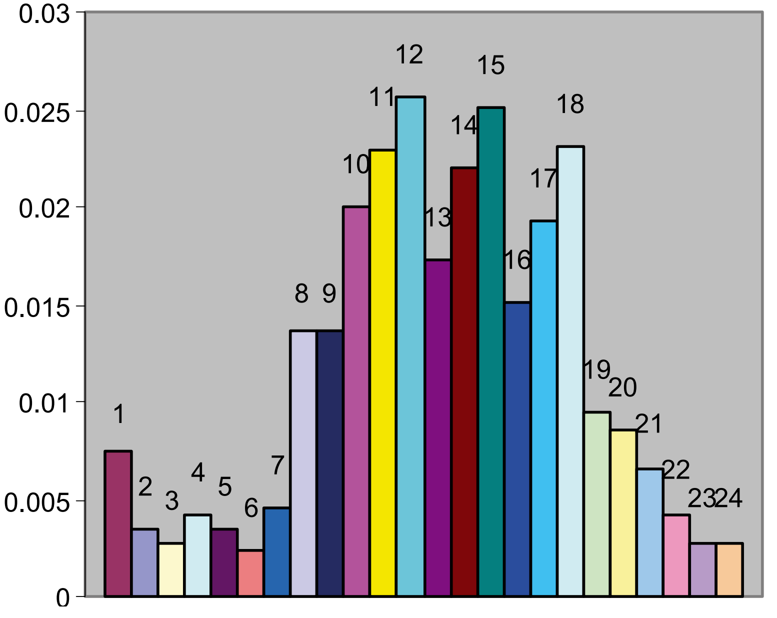

Fig. 1. Estimates of the hourly dummy variables in the conditional variance function of BP.

interesting results could be inferred. There is a slight increase in the volatility for both the currencies at hour 1 with the opening of the Tokyo market.[13] The magnitude of this volatility is slightly higher for the BP compared to euro. During the next 5 h or so, volatility remains low, since it is the time when trade predominantly takes place in the Asian markets, with Tokyo being the lead market. At hour 7 the volatility picks up a little in the euro but not in the BP, this could be due to the pre-market opening of the Paris foreign exchange (France and Germany are probably the strongest of the economies embracing euro). Finally at hour 8, the leading European foreign exchange markets open officially, with London, Paris and Frankfurt foreign exchange markets in full swing the volatility in both the currencies increase dramatically. The increase in volatility could be due to the release of public information related to the economies of the two currencies respectively, the regularly scheduled macroeconomic announcements. It would not be fair to attribute the increase in volatility solely to public information at this stage, since investors could also be trading on the basis of their accumulated private information (estimates of vacation dummy in the previous section were also supportive of the private information hypothesis). A discussion of the probable reasons for this observed increase in volatility during the trading hours of the leading European foreign exchange markets is provided shortly.

The volatility continues to increase and remains high during the next 10–11 h. Other than a slight anomaly for the euro at the hour 13, at which time volatility drops off sharply,

[13] All times are quoted in GMT.

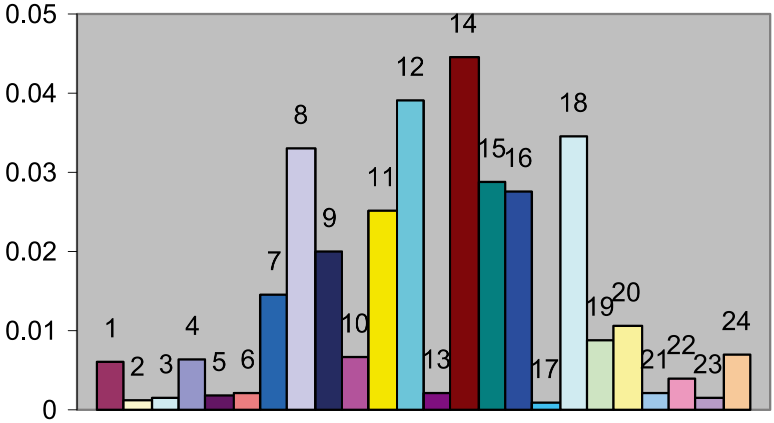

Fig. 2. Estimates of the hourly dummy variables in the conditional variance function of euro.

the volatility generally remains high during the opening and the following trading hours of the European foreign exchange markets.[14] At 14 GMT there is a sudden burst of volatility highly pronounced for the euro. This increase in volatility coincides with the opening of the New York foreign exchange market. The jump in volatility in euro is significantly higher than the increase in volatility of BP. This could be due to perhaps the reason that the EU as a whole is much more affected by the happenings in American economy or it may be due to more dissemination of the private information when the trading starts in U.S., which affects euro considerably more. Whatever the source, there is evidence to believe that euro is considerably much more responsive to the trading in U.S. compared to the BP. Recently this fact has become even more emergent, the dramatic moves in dollar mostly in the form of depreciation, due to the strains in the U.S. economy has produced much more volatility (mostly in the form of appreciation) in the euro and the Yen. It is certainly possible that movements in the currencies of these three biggest economies of the world, directly competing with each other, produce significant volatility spillover effects for each other to absorb.

Keeping aside this discussion, the increase in volatility at 14 GMT, establishes the importance of the U.S. for the rest of the world. The volatility remains high in the next few trading hours with the European and American markets fully opened. For BP the volatility continues to be high till the closure of the New York foreign exchange by hour 20/21 (approximately). Even though the volatility generally continues to be high till the closure of the New York foreign exchange market, there is a sharp drop in volatility after hour 18: the hour immediately following the last hour of trading in the European foreign exchanges. The London and Paris markets approximately close by hour 17. The volatility finally drops of almost completely by the end of hour 21, after which only the Asian markets are available for trading. High volatility at the closing hour of the European foreign exchange

[14] The sharp drop off in the volatility of Euro, at hour 13 is considered to be an anomaly since the European foreign exchange markets trade continuously. There is no official break in the trading at the lunch time.

markets is also observed for BP, even though the peak volatilities are observed during the middle of the trading day.

For the euro the volatility pattern is more or less similar. The volatility remains high during the simultaneous trading of the European and American foreign exchange markets. There is a sudden sharp drop in the volatility at hour 17 (in fact the volatility at this hour is virtually non-existent).The source of this sharp drop off in volatility is not exactly known but I believe a plausible reason for this sharp drop in volatility could be the closure of Paris and London foreign exchange markets by the end of this hour. This drop off in the volatility of euro is in sharp contrast to the behaviour of BP for which the volatility remains high even at the closing of European foreign exchange markets. The volatility then rises sharply in the next hour, even after the closure of the leading European foreign exchange markets. The New York market is open during this hour and the volatility in euro remains high too. Finally the volatility drops sharply compared to the previous hours but generally remains pronounced till the end of hour 20, the time at which Frankfurt and New York markets are approximately closed. Even though the early trade in the Paris foreign exchange and some what late closing of the foreign exchange market in Germany produces volatility in euro, the main volatile periods happen to be those at which the three leading European exchanges are trading simultaneously. The same conclusion is drawn for the BP as well. Additionally it also lends support to prove London as the main financial district of Europe.[15]

Other than the pronounced volatility for the two currencies during the European trading hours, another surprise finding is the observed spike in volatility for hour 18, the hour immediately following the closure of leading European foreign exchange markets, including the London foreign exchange (the world's biggest foreign exchange market). This high volatility at hour 18, the time at which New York foreign exchange market is still active and remains active for some further 3 h or so could be due to the release of private information, i.e. traders dealing in BP and euro disseminating important information related to the European currencies immediately in the following hour by trading in the U.S. markets, even though the European markets are closed. To summarise, the volatilities displayed by the two European currencies over a particular trading day in the European markets could well be described by an approximate inverse U-shaped pattern (peak volatilities are observed during the middle of the day).

Baillie and Bollerslev (1991) divide the trading hours in the foreign exchange markets across the world into three time-zones; the Asian (A) market which is defined as being open between hour 23 and hour 7, the European (E) market which is open from hour 8 through 16, and the United States (U) which is open between hours 14 and 22.[16]

[15] Obviously the above analysis assumes that the opening and closing timings of the foreign exchange markets used in the study are correct. The timings were provided by Reuters and were subsequently converted to GMT by adjusting the differences. Broadly, we believe that these timings do correspond to the opening and closing of the respective markets, since some verification is provided by the existing literature.

[16] No account is made of the fact that Europe goes on daylight saving time from the last weekend of March. This causes a small misalignment in the definition of the time zones around these dates. However, this is of no concern to this study, since the data used does not extend over to the time, when daylight saving change is made.

Using these three time-zones Baillie and Bollerslev compare the volatility of the three market segments for the four currencies considered in their study. The same analysis is conducted here with the objective to compare the volatility of the markets across the globe for the two currencies under study. A summary of the average estimates of the unconditional variances for the two currencies across the three different market locations is provided in Table 4.

For the two currencies considered the European market appears to be on average the most volatile market with the U.S. being the second most. The Asian market is in general the least volatile market. This phenomenon seems surprising at least when considered in the foreign exchange market. The foreign exchange market is virtually a continuously active market thorough out the day, with the traders having the flexibility to trade around the clock. This finding reinforce the result from the previous findings that relevant markets need to open to absorb the news related to their respective currencies, even though markets in other regions of the world have opened earlier [See Goodhart and Giugale (1993)]. This result contradicts the previous finding of Baillie and Bollerslev (1991) who find U.S. market to be the most volatile even for the European currencies considered in their study and surprisingly even for YEN, the Asian market was the least volatile, with U.S. being the most volatile. This particular result for YEN in Baillie and Bollerslev and the overall finding that the Asian markets are the least volatile for the currencies considered in their study and for the currencies considered in this study seems to reinforce the widely documented importance of the world's two biggest foreign exchange markets, London and New York in absorbing news related to the most frequently traded currencies of the world. The second point to note from estimates of the average volatilities in Table 4 is that on average euro is more volatile than BP. In all three markets it exhibits more volatility than BP. The average volatility in euro presented in the first row of Table 4 is also of greater magnitude than that of the BP as well.

3.2. A model for daily exchange rate movements

Before discussing the reasons behind the observed volatility behaviour in the two currencies, documented in the literature on foreign exchange market, we consider it useful to extend the analysis to a lower frequency, i.e. at the daily level. The objective is to study the behaviour of volatility embodied in the two currencies at the daily level and thus providing an even clearer picture of the volatility dynamics displayed by the two currencies. Since the euro only started trading from the beginning of January 1999, the data set consists of a total of 920 observations for both the BP and the euro; starting

Table 4
Estimated average hourly variances

	BP	Euro
$\hat{\omega}_0(i)=(1/24)(\hat{\omega}_1(i)+\cdots+\hat{\omega}_{24}(i))$	0.01168	0.01380
$\hat{\omega}_A(i)=(1/9)(\hat{\omega}_{23}(i)+\cdots+\hat{\omega}_7(i))$	0.00377	0.00473
$\hat{\omega}_E(i)=(1/9)(\hat{\omega}_8(i)+\cdots+\hat{\omega}_{16}(i))$	0.01948	0.02524
$\hat{\omega}_U(i)=(1/9)(\hat{\omega}_{14}(i)+\cdots+\hat{\omega}_{22}(i))$	0.01480	0.01803

from 5th January 1999 and ending on 12th July 2002.The data is extracted from the DataStream International.[17] The BP is quoted as the amount of U.S. dollars per British pound whereas euro is quoted as the number of euros per U.S. dollar. The results of the initial preliminary data testing were similar to that conducted for the hourly exchange rates. The null of the unit root could not be rejected for the logarithm of the exchange rates. The result is accordance with those reported in literature on the analysis of daily exchange rates by Baillie and Bollerslev (1989), Hsieh (1988, 1989) and Baillie et al. (1996).[18]

Based on the preliminary testing, only the first differences of the two exchange rate series will be considered in the subsequent analysis, i.e.

$$y(i)_t = 100 \cdot [log(s(i)_t) - log(s(i)_{t-1})] \tag{7}$$

where s is the spot exchange rate. The series $y(i)_t$ corresponds to the approximate percentage nominal return on currency i obtained from day $t-1$ to t.

As widely documented in the literature skewness and excess kurtosis are key features of the financial market returns, the following model (similar to the one estimated for the hourly returns) is estimated for the two daily return series:

$$y(i)_t = u_o(i) + \epsilon(i)_{t,\tau} \tag{8}$$

$$\epsilon(i)_t \mid \Omega_t \sim N(0, \sigma(i)_t^2) \tag{9}$$

$$\begin{aligned}\sigma(i)_t^2 = {} & \omega_0(i) + \omega_{\mathrm{M}}(i)D_{\mathrm{M}t} + \omega_{\mathrm{T}}(i)D_{\mathrm{T}t} + \omega_{\mathrm{W}}(i)D_{\mathrm{W}t} + \omega_{\mathrm{R}}(i)D_{\mathrm{R}t} + \alpha_1(i)\epsilon(i)_t^2 \\ & + \beta_1(i)\sigma(i)_t^2\end{aligned} \tag{10}$$

where $D_{\mathrm{M}t}$, $D_{\mathrm{T}t}$, $D_{\mathrm{W}t}$ and $D_{\mathrm{R}t}$ are the dummy variables for Monday, Tuesday, Wednesday and Thursday, respectively. The model has a GARCH(1,1) formulation in the sixth and seventh terms. The model specified has some similarities with the model specifications in Baillie and Bollerslev (1989) and Hsieh (1988). The results from the estimated models are presented in Table 5. In estimation of the above model, the constant term was replaced by another (fifth) dummy variable for Friday, to make the interpretation simpler.

Most of the estimated parameters for the two models in the conditional variance function are significant at the 5% level. The GARCH(1,1) parameter estimates for the two currencies are highly significant and sum close to 1, indicating the presence of integrated GARCH or IGARCH behaviour at the daily frequency. Hsieh (1989) and

[17] As a useful exercise we checked the degree of consistency between the two data sets [hourly (Reuters) and daily (DataStream)] by looking at the observations from the two data sets when they coincided (at a particular time). On average, the degree of consistency was quite high.

[18] The results are not presented due to limitations of space.

Table 5
Quasi-maximum likelihood estimates

	BP	Euro
$u_0(i)$	0.00357 (0.01547)	0.01990 (0.02144)
$\omega_M(i)$	0.12449 (0.03533)	0.22019 (0.08195)
$\omega_T(i)$	0.11896 (0.02906)	0.13523 (0.08459)
$\omega_W(i)$	0.01936 (0.03763)	0.07134 (0.09438)
$\omega_R(i)$	0.01112 (0.03832)	0.09093 (0.07253)
$\omega_F(i)$	0.05463 (0.04026)	0.12421 (0.07652)
$\alpha_1(i)$	0.04853 (0.01679)	0.00694 (0.00429)
$\beta_1(i)$	0.93660 (0.01714)	0.98372 (0.00488)
$m_3(i)$	0.16712	− 0.25438
$m_4(i)$	3.93084	4.3535
$Q_{20}(i)$	15.0807	13.8862
$Q^2_{20}(i)$	26.9183	19.2351

See footnote to Table 1. Asymptotic robust standard errors are in parentheses.

Baillie and Bollerslev (1989) also find the sum of GARCH parameters to be close to 1 on analysing the daily bid prices. All other diagnostic test statistics reported in the lower part of the table indicate that GARCH(1,1) model provide a reasonably good fit. The Ljung–Box portmanteau test statistics for the presence of up to 20th-order serial correlation in the standardized and squared standardized residuals of the two currencies are not significant at the 5% level. The LM test statistics for the presence of ARCH effects when applied to the residuals of the estimated models gave no indication of any additional conditional heteroskedasticity (the statistics were insignificant at the usual levels of significance) and as such there seems to be no need of using any higher order GARCH models. The value of kurtosis is greater than 3 for the residuals of the two currencies indicating non-normality but this is something in accordance with the literature on daily exchange rates. The graphs of the estimated daily dummy variables are presented in Figs. 3 and 4.[19]

The estimated daily dummy variables display a broadly similar volatility pattern for the two currencies. Volatility for the two currencies is at peak on Monday and stays high on Tuesday as well (even though it is slightly lesser than on Monday). For the next 2 days, Wednesday and Thursday, the volatility drops off for both the currencies. Low volatility on these 2 days is clearly evident from the graphs. The volatility then picks up on Friday, the last trading day of the week. Even though the magnitude of the volatility on Friday is not comparable to that on Monday and Tuesday, the pronounced increase in volatility is clearly visible. The increase in volatility on Friday is much higher and more pronounced for the euro compared to the BP.

[19] Please note that through out the analysis no discussion on the sign of price changes, i.e. whether the volatility on a particular hour or day is positive or negative is provided. The direction of volatility would be of interest if the purpose is to forecast the volatility but since the main objective is to compare the volatility across the hours and days, i.e. whether the volatility at some specific hour or day is high or lower compared to some other hour or day, we need not to be bothered by the direction of the volatility. The estimates of the dummy variables could be considered as the absolute values of the true estimates.

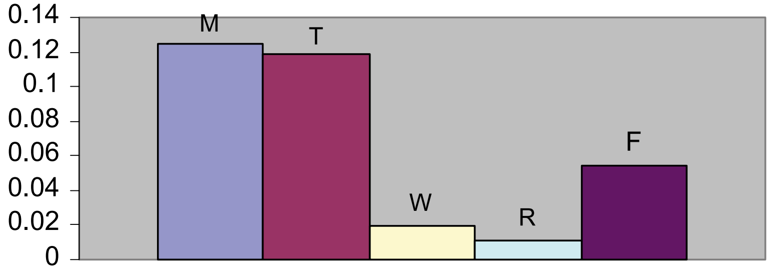

Fig. 3. Estimates of the daily dummy variables in the conditional variance function of BP.

Of the two currencies, the estimates of the daily dummy variables reinforce the previous conclusion of euro being more volatile of the two currencies. Nearly all the estimated daily dummies are of greater magnitude for the euro than for BP. The magnitude of the pronounced volatility on Monday is considerably higher for the euro compared to that for the BP. The high volatility on Monday for the two currencies could be due to the accumulated news (public or private) over the weekend. The volatility for the BP follows a monotonically decreasing pattern over the first four weekdays before picking up on Friday, whereas for the more volatile euro it starts picking up from Thursday onwards. The volatility in euro fluctuates in a much wider interval than BP, over the weekdays.

Since volatility in financial markets is considered to be of utmost importance, it is considered imperative to go through the possible sources of volatility which could have led to the observed volatility dynamics at the hourly and daily frequency. Provided below is a discussion of some of the most important studies which have investigated the underlying source of inter- and intra-day volatility. As we briefly go through these studies, to check the applicability of their conclusions to the observed volatility dynamics for BP and euro, a thorough comparison is made of the results obtained by these studies using various currencies and the one's obtained in this study.

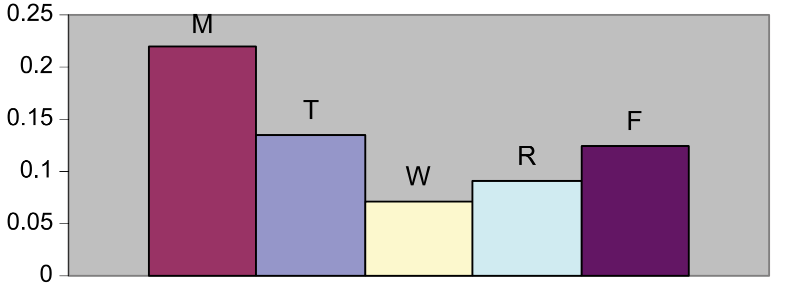

Fig. 4. Estimates of the daily dummy variables in the conditional variance function of euro.

3.3. Sources of the observed volatility dynamics

French and Roll (1986) on analysing the equity returns find them to be more volatile during exchange trading hours than during non-trading hours.[20] French and Roll on differentiating between the public and private information as the source of trading time volatility find evidence in favour of the private information hypothesis.[21] French and Roll (1986) conclude that the large difference between trading and non-trading return variances is caused by differences in the arrival and incorporation of private information during trading and non-trading hours.

Harvey and Huang (1991) examine the volatility implications of continuous foreign exchange trading using transactions data on foreign exchange futures contracts for DM, YEN, SF, BP and CD (Canadian dollar) vis-à-vis the U.S. dollar. Harvey and Huang on analysing the currencies at the inter- and intra-daily basis, find higher U.S.–European and U.S.–Japanese exchange rate volatilities during U.S. trading hours and higher European cross-rate volatilities during European trading hours. In specific the U.S.–European foreign exchange futures are found to be on average twice as volatile during U.S. trading hours than during European trading hours. This particular result of Harvey and Huang directly conflicts with the finding of higher BP and euro volatilities during the European trading hours.[22] In discriminating between the private or public information as the source of the observed volatility, Harvey and Huang discount the probability that higher volatility during the U.S. trading hours could be induced by trading that acts on private information. They argue that with the nature of the foreign exchange market and the possibility to trade at any time of the day, private information cannot be considered as the underlying source of volatility. The ability of U.S. investors to trade in many markets around the clock and the high liquidity in non-U.S. markets suggest that the concentration of volatility for the U.S. exchange rates during U.S. trading hours is driven by public rather than private information-based trading.

Further Harvey and Huang argue that the effect of U.S. macroeconomic news released during U.S. trading hours dominates the effect of news released in the European markets in determining the volatility of U.S.–European rates. The estimated hourly dummy variables in Section 3 also suggest the importance of U.S. news announcements, since volatility generally remains high during the U.S. trading hours. However, if Harvey and Huang interpretation of public news announcements as the main source of underlying volatility is given more weight as the determinant of the observed volatility dynamics, then for the research conducted above the European news announcements appear to be more important in determination of the two European exchange rates considered. More information released during the European trading hours could be the main source of the observed

[20] This phenomenon is widely documented in the literature. Asset returns display a puzzling difference in volatility between exchange trading hours and non-trading hours.

[21] French and Roll also consider a third hypothesis; trading noise hypothesis, i.e. the high daily return variances are caused due to pricing errors but find them to have a trivial effect on the difference between trading and non-trading time volatility.

[22] Baillie and Bollerslev (1991) find that even though the European currencies are more volatile during the U.S. trading hours, the phenomenon is not much pronounced.

volatility patterns. Volatility might be concentrated during the times when the most relevant macroeconomic news is released.

Subsequent analysis of the inter-daily data by Harvey and Huang (1991) suggest that the concentration of U.S. macroeconomic announcements on Thursdays and Fridays appears to manifest itself in higher volatilities during the opening hours of the U.S. trading on these days; Friday in particular appears to be the most volatile day of the week. Harvey and Huang attribute the higher volatilities on Friday to the large number of U.S. macroeconomic news announcements made on Friday. For the two currencies considered in this study, volatility for the two currencies certainly increase on Friday, supportive of the fact that U.S. news does play an important role in determining European exchange rates as well, but the most volatile days are the Mondays and Tuesdays, with Monday being the most volatile. The major macroeconomic announcements for the United Kingdom and the European Union are released on Wednesdays and Thursdays and the low volatilities observed on these two weekdays for these two regions currencies, respectively, jeopardize the importance of public information hypothesis. The volatility in euro is considerably greater on Friday compared to BP suggesting U.S. news to be more important in determining the euro exchange rate. Harvey and Huang also find that the currency volatilities when reported by day of the week display a distinct U-shaped pattern characterising the volatility during the week. The daily volatility patterns for the euro and the BP does correspond roughly to a U-shaped pattern with heightened volatility on Monday, decreasing some what on Tuesday (but still keeping high), further decreasing on Wednesday and remaining low on Thursday as well and finally rising on Friday. For BP, Thursday is the least volatile day, whereas for euro Wednesday is the least volatile; volatility starts increasing from Thursday. Undoubtedly for both the currencies Monday is the most volatile day followed by Tuesday.

Andersen and Bollerslev (1998) on analysing the Deutsche mark (DM)–U.S. dollar (USD) five min returns over a 1-year period show that the largest returns appear linked to the release of public information and in particular, to certain macroeconomic announcements. However, they also find that public information arrivals induce abrupt price changes but the average price move is typically attained within minutes, whereas the finding of Kim and Verrecchia (1991) that the volatility and trading volume remain elevated for several hours is observed in reality.[23] Andersen and Bollerslev therefore conclude in favour of models with heterogeneously informed agents, since if agents have identical information sets and interpret news similarly, the observed protracted response pattern is hard to explain.

Some additional results of Andersen and Bollerslev (1998) include the finding of a significant jump in volatility at the time of opening of the Tokyo foreign exchange interbank market (the Tokyo opening effect) and the increased volatility during the overlap in the Asian and European and the European and North American business hours. Volatility is also found to be notably higher during European trading hours which could be due to the relevant economic news hitting the market during that part of the daily cycle. After the

[23] A significant finding of Andersen and Bollerslev in this regard is that some unscheduled news compared to the regularly scheduled news could have more prolonged impact on the volatility.

closure of the London market, volatility displays a monotonic decline until it reaches the plateau associated with the pacific segment and further no signs of elevated volatility when trading closes down in New York are observed. These findings overall seems consistent with the hourly volatility patterns observed for euro and BP, with slight volatility observed at the opening of Tokyo market and higher volatilities observed for the two European currencies during the European trading hours. The two currencies remain volatile during U.S. trading hours but the magnitude of the hourly dummies is considerably less towards the closure of the U.S. markets. In terms of the days of the week Andersen and Bollerslev find that Mondays appear the least volatile, while Thursdays and Fridays are the most volatile, with the large Thursday and Friday dummies reflecting the clustering of scheduled news releases on these weekdays; a finding consistent with that of Harvey and Huang (1991) but not with the findings in Section 3.2 for the two European currencies.

Andersen et al. (2001) test for a change in the intraday volatility pattern of the USD–YEN foreign exchange rate, following the removal of trading restrictions over the lunch period for the Japanese-based banks; on December 22, 1994. Andersen et al. examine the relevance of the private information for the increase in volatility (if any found) over the lunch period. In other words they argue that if the volatility is caused solely by public information and public information flow is unaltered, volatility should not be affected by the introduction of lunch-period trading. On the other hand, the presence of private information or mispricing could increase the volatility over the lunch period following deregulation. Andersen et al. find significant evidence of volatility increasing over the Tokyo lunch period and therefore strongly reject the null hypothesis of no change versus the alternative of higher volatility following the relaxation of the Tokyo trading restrictions. No evidence for any other qualitative shifts in volatility pattern outside the lunch period was observed. Andersen, Bollerslev and Das conclude that the relevance of private information for the spot foreign exchange inter-bank market cannot be denied. Similar results have been obtained by Ito et al. (1998), who also find private information as the main contributing force behind the post-period increase in volatility. Covrig and Melvin (2002) identify a period in the foreign exchange market when there is a high concentration of informed YEN–USD traders active in Tokyo foreign exchange market. Comparing the period of informed trader clustering to a similar period without informed traders, they show that exchange rate quotes adjust to full information levels much more quickly when informed traders are active in the market than when they are not.

The research on the possible sources of the observed volatility dynamics is by no means conclusive. The importance of public news in the determination of the exchange rates cannot be denied; however as documented above, the support for private information hypothesis is increasing at least when analysing the exchange rates over hourly or daily frequencies. Even though it is hard to digest the relevance of private information in the foreign exchange market, the foreign exchange market is virtually a continuously active market with the traders having the flexibility to trade anywhere in the world. However, as documented above the empirical evidence in support of private information is substantial. Researchers have found that the proportion of return variability that can be ascribed to public news announcements is low. Proxies used by economists for the flow of public information, e.g. the number of news items released by Reuter's News Service, etc. explain little of the overall variation. Goodhart and Giugale (1993) find that volatility is smaller during intervals when trading volume is known to be smaller (for example, weekends and lunch hours) and is larger during the first trading hour on Monday for each currency in the domestic country, regardless of the fact that other markets have opened earlier. Goodhart and Giugale interprets the heightened volatility on Monday as being consistent with the fact that accumulated (weekend) news is not fully incorporated into prices until the main relevant markets are open for business, even though subsidiary markets may have opened earlier. Trading in key markets is necessary to complete the assimilation of news into prices. A similar Monday market opening effect is also reported in Baillie and Bollerslev (1989) and Hsieh (1988). This evidence is considered to be supportive of the private information hypothesis. Foster and Viswanathan (1990) in a theoretical model show that due to the weekend closing of the foreign exchange market the greatest information advantage takes place on Mondays and it is also beneficial for the informed trader to carry some of his/her information from Monday to Tuesday. Their theoretical model provides evidence in favour of the French and Roll (1986) private information hypothesis. Private information however does not necessarily have to be the only explanation for the observed volatility dynamics. Hsieh and Kleidon (1996) argue that the volatility in the foreign exchange market associated with (local) market openings is probably due to the traders trying to gather order-flow information and due to inventory management. Bollerslev and Domowitz (1993) find that the behaviour of bid/ask spread plays an important role in determining the conditional volatility of the returns process. A strong contemporaneous correlation between trading volume and volatility is also reported in the literature. To figure out a unique culprit as the source of volatility in the foreign exchange market is perhaps nearly impossible. A significant proportion of the trading volume is conducted mainly among the market makers. More extensive research is required to solve this mystery, perhaps using continuous-time GARCH models.

To summarize, for the two European currencies considered, the main finding to have emerged is higher return volatility during the trading hours of European foreign exchange markets and Monday being the most volatile day of the week for the two currencies followed by Tuesday. Keeping in view these facts together with the fact that London is the biggest foreign exchange market in the world, the applicability of the French and Roll (1986) conclusion that high trading time volatility is caused due to the incorporation of private information into prices through trading cannot be discounted (at least superficially) even when applied to the foreign exchange market. Clearly the evidence on the whole makes a strong case in favour of the private information hypothesis, which is incorporated into the exchange rates during the trading hours of the European foreign exchange markets. Further the appearance of Monday as the most volatile day of the week lends additional support to the hypothesis. However, more research is warranted before presenting any solid linkage/hypothesis of causality between the two variables, volatility and private information.

4. GARCH versus FIGARCH versus SV model

The models estimated in Section 3 are robust to the misspecification tests and judged by the standard diagnostic statistics provide a reasonable good fit. The estimates of the hourly and daily dummy variables are mostly significant and provide some interesting results. Further there are no signs of any additional ARCH effects both at the hourly and daily frequencies in the residuals of the estimated models of the two currencies. The GARCH parameters in the four models are highly significant, together with most of the other estimated parameters. The sum of the GARCH parameters for both the currencies using the hourly data is considerably less than 1 and as such no evidence for the presence of integrated variance or IGARCH volatility process is forthcoming. Compared to this the sum of estimated parameters for the GARCH(1,1) models using the daily data is very close to 1, indicating the presence of integrated variance for the two currencies. However, Baillie et al. (1996) documented the argument that the frequent finding of IGARCH behaviour at the daily frequencies is actually due to the presence of fractionally integrated variance, i.e. fractionally integrated GARCH (FIGARCH) behaviour.

The purpose of this section then is to assure whether or not the GARCH model is best for describing the dynamic behaviour of the two currencies at the hourly and daily frequencies. What follows below is a comparison between the simple GARCH model and some of its recently developed extensions, more importantly with the FIGARCH model. For the FIGARCH process even though a shock to the volatility process eventually does die out in a forecasting sense, the decay occurs at a slow hyperbolic rate. The data set employed in this comparison is the same as used at the hourly and daily frequencies; the analysis is conducted for both the currencies. The comparison which follows thus not only allows us to check the rigourness of the results obtained in the previous section but also helps to check the appropriateness of GARCH in modelling the dynamics of data recorded at high and low frequencies.

Some care need to be exercised in estimating the FIGARCH models. As discussed in Baillie et al. (1996) that in the presence of long memory and relatively slow decay of a response to a lagged squared innovation, the effect of pre-sample values might be expected to have a bigger impact than with stationary GARCH processes. Thus truncating at too low a lag may destroy important long-run dependencies. To keep this estimation problem aside the truncation lag was set at 1000, similar to the truncation lag used by Baillie et al. in their analysis of the daily exchange rates.[24] Additionally in a simulation study Baillie et al. (1996) find that results indicate that the proposed QMLE procedure performs very well for the FIGARCH models with the sample sizes typically encountered with high frequency financial data. And more importantly, for comparisons of the GARCH versus FIGARCH or IGARCH versus FIGARCH, the actual size of the standard *t*-test is very reliable. Bollerslev and Mikkelsen (1996) have shown that the AIC and SIC information criterion could also be effectively used in discriminating between the GARCH, IGARCH and FIGARCH models. Therefore the AIC and SIC model selection criterion and the portmanteau test statistics for residual autocorrelation could be used effectively in deciding on the correct conditional variance specification. There is however a word of caution over here, the data sets which we utilise are relatively short, with the data sets spanning over approximately 4 months for the hourly

[24] For the daily data set, it was set equal to 900.

Table 6
Quasi-maximum likelihood estimates (the estimated parameters are for the BP, using hourly returns)

(p,d,q)	(1,0,1)	(1,d,1)	(1,0,0)	(1,d,0)
u_0	0.00061 (0.00179)	0.00061 (0.00179)	− 0.00079 (0.00191)	− 0.00034 (0.00226)
u_1	− 0.00547 (0.03515)	− 0.00551 (0.03527)	− 0.01022 (0.00852)	− 0.04125 (0.03548)
ω	0.00231 (0.00035)	0.00233 (0.00064)	− 4.85855 (0.07247)	− 4.67976 (0.10908)
β_1	0.45673 (0.05254)	0.45623 (0.05440)	0.63458 (0.05002)	0.57074 (0.09185)
α_1	0.26992 (0.05355)	0.72736 (0.05011)	–	–
θ	–	–	− 0.13332 (0.04519)	− 0.13430 (0.07826)
γ	–	–	0.47892 (0.05733)	0.51173 (0.08885)
d	–	− 0.00147 (0.04082)	–	0.15971 (0.08480)
AIC	− 2.107050	− 2.105960	− 2.120056	− 2.065158
SIC	− 2.092016	− 2.087920	− 2.105022	− 2.044112
m_3	0.12490	0.12469	0.18837	0.17787
m_4	5.3851	5.3851	5.2494	5.5088
Q_{10}	12.131	12.1426	12.4189	9.45086
Q^2_{10}	27.4744	27.4488	28.4265	83.5829

Asymptotic robust standard errors are in parentheses. The statistics $m_3(i)$ and $m_4(i)$ denote the standardized residuals skewness and kurtosis. Q_{10} and Q^2_{10} give the Ljung–Box portmanteau test for up to 10th-order serial correlation in the standardized residuals and the squared standardized residuals, respectively. AIC and SIC refer to the Akaike and the Schwarz Information Criterion, respectively.

data and 3.5 years for the daily data. Short data sets may not contain much information on long memory volatility processes.

Four models (GARCH, FIGARCH, EGARCH and FIEGARCH) for each of the two currencies (returns) are estimated and explicitly compared, using both the hourly and daily returns data. The fractionally integrated EGARCH model was developed by Bollerslev and Mikkelsen (1996) and similar to FIGARCH the conditional variance is modelled as a mean-reverting fractionally integrated process. The estimates for the models estimated using the hourly frequency are presented in Table 6 for the BP and in Table 7

Table 7
Quasi-maximum likelihood estimates (the estimated parameters are for the euro, using hourly returns)

(p,d,q)	(1,0,1)	(1,d,1)	(1,0,0)	(1,d,0)
u_0	− 0.00124 (0.00224)	− 0.00122 (0.00226)	− 0.00336 (0.00244)	− 0.00259 (0.00299)
u_1	− 0.00821 (0.03473)	− 0.00882 (0.03470)	− 0.00736 (0.03248)	0.00055 (0.03400)
ω	0.00372 (0.00061)	0.00415 (0.00143)	− 4.32035 (0.09736)	− 4.17103 (0.14627)
β_1	0.46362 (0.05398)	0.45399 (0.05956)	0.71884 (0.04095)	0.69361 (0.09851)
α_1	0.28369 (0.06443)	0.75245 (0.06378)	–	–
θ	–	–	− 0.06001 (0.04505)	0.08265 (0.10004)
γ	–	–	0.42111 (0.05499)	0.40848 (0.06672)
d	–	− 0.01957 (0.04851)	–	0.13035 (0.11784)
AIC	− 1.597095	− 1.596096	− 1.605192	− 1.534854
SIC	− 1.581214	− 1.577039	− 1.586135	− 1.512620
m_3	− 0.33741	− 0.34286	− 0.28693	− 0.42041
m_4	7.4560	7.4332	7.0325	8.0540
Q_{10}	13.4558	13.4588	13.3919	12.534
Q^2_{10}	7.25607	6.88955	7.04777	56.908

See footnote to Table 6.

for the euro. The first column in both the tables present the estimated parameters from the simple GARCH(1,1) model. The sum of the GARCH(1,1) parameters for both the currencies is considerably less than 1, signalling no presence of the integrated variance at the hourly frequency. This result is consistent with the results obtained in the previous section with hourly dummies included in the variance specification. There is simply no evidence of IGARCH behaviour in the high frequency data for the two currencies. The next column in the tables present the results obtained for the FIGARCH(1,d,1) model. The model explicitly considers whether there is any indication (presence) of long memory volatility (fractionally integrated variance) process in the data. To test the null hypothesis of short against long memory alternative, the usual t-test for the hypothesis $d=0$ against $d\neq 0$ is performed. The hypothesis of $d=0$, i.e. the estimated fractional differencing parameter, is not significantly different from zero could not be rejected at the 5% level for any of the currency. Since d is not significantly different from zero it is considered unnecessary to make any further comparisons between these two models.[25] The results from the EGARCH(1,0,0) models appear in the third column of the tables and those for the FIEGARCH(1,d,0) models in the final column. For the FIEGARCH, the result is similar to that for the FIGARCH models, the estimated fractional differencing parameter is not significant at the conventional significance levels, and thus we could discard the FIEGARCH specification as well. The EGARCH model provides a reasonably good fit. Selecting between the EGARCH and GARCH models could be a difficult scenario. Both models provide good estimates of the parameters, the diagnostic test statistics from the two models are more or less similar for the two currencies; in fact there is slight preference for the EGARCH model based on the diagnostic test statistics for the two currencies (the AIC and SIC information criterion both selects the EGARCH specification, the value of m_4 statistics are also lower for the EGARCH processes). The asymmetric relation between the past returns and changes in volatility represented by the θ parameter in the EGARCH specifications is significant for the BP but insignificant for the euro, at the 5% level. Therefore the choice between the GARCH and the EGARCH process seems to be an arbitrary one.

The results for the daily frequencies are presented in Tables 8 and 9 for the BP and the euro, respectively. The results for the daily returns indicate the presence of fractionally integrated variance for the two currencies. The sum of the GARCH(1,1) parameters is close to 1, cautioning us against the inference of integrated variance. The estimates of the fractional differencing parameter (d) are significant at the usual significance levels for the two currencies under both the FIGARCH and FIEGARCH specifications. In selecting the appropriate model for the daily returns, a selection needs to be made between the FIGARCH(1,d,1) and FIEGARCH(1,d,0). The test diagnostics prefer the FIGARCH models with lower kurtosis, skewness and slightly lower values for the Ljung–Box test statistics as well. The AIC and

[25] Beside, the conditions presented in Baillie et al. (1996)to ensure that the conditional variance of the FIGARCH(1,d,1) model is positive are not fulfilled.

Table 8
Quasi-maximum likelihood estimates (the estimated parameters are for the BP, using daily returns)

(p,d,q)	(1,0,1)	(1,d,1)	(1,0,0)	(1,d,0)
u_0	− 0.00091 (0.01506)	0.00043 (0.01525)	− 0.00674 (0.01503)	− 0.01600 (0.01997)
ω	0.00338 (0.00475)	0.01015 (0.00824)	− 1.51409 (0.08167)	− 1.64109 (0.26113)
β_1	0.95930 (0.03681)	0.69759 (0.06944)	0.93640 (0.02330)	0.75825 (0.16048)
α_1	0.02510 (0.01775)	0.35425 (0.10017)	–	–
θ	–	–	− 0.04013 (0.02094)	0.02575 (0.06577)
γ	–	–	0.07373 (0.02926)	0.20617 (0.12406)
d	–	0.33758 (0.11611)	–	0.52962 (0.15095)
AIC	1.290312	1.286408	1.301353	1.308334
SIC	1.311324	1.312673	1.327617	1.339851
m_3	0.088745	0.078807	0.062911	0.093958
m_4	3.84934	3.85141	3.84212	3.88443
Q_{10}	7.63067	8.27574	8.14286	7.71839
Q_{10}^2	18.4915	12.5675	18.7425	25.3942

See footnote to Table 6.

SIC information criterion also selects the FIGARCH specifications over the FIEGARCH (the estimate of θ parameter for the two currencies is not significant at the 5% level as well). This finding of fractionally integrated variance at the daily frequency for the two currencies is consistent with the finding of Baillie et al. (1996), who also report significant evidence for the presence of fractionally integrated behaviour in the conditional variance of the USD−DM exchange rate at the daily frequency.

The evidence so for seems consistent with the exiting literature. Researchers have found that when working with high frequency data the dynamic impact of identifiable intraday volatility shocks seems rather low, with any noticeable effects on the overall volatility gone in a matter of hours or less. Short lived volatility dynamics

Table 9
Quasi-maximum likelihood estimates (the estimated parameters are for the euro, using daily returns)

(p,d,q)	(1,0,1)	(1,d,1)	(1,0,0)	(1,d,0)
u_0	0.01538 (0.02087)	0.01110 (0.02267)	0.01727 (0.02151)	− 0.01110 (0.02382)
ω	0.00277 (0.00192)	0.03364 (0.03596)	− 0.84918 (0.09477)	− 1.09945 (0.45708)
β_1	0.98094 (0.00542)	0.68018 (0.14446)	0.96145 (0.02870)	0.60188 (0.50311)
α_1	0.01237 (0.00444)	0.38072 (0.06613)	–	–
θ	–	–	0.00893 (0.01709)	0.32291 (0.15691)
γ	–	–	0.04795 (0.01711)	0.97105 (0.38602)
d	–	0.26201 (0.14476)	–	0.36473 (0.10219)
AIC	1.945896	1.944267	1.959714	1.952858
SIC	1.966908	1.970531	1.985978	1.984376
m_3	− 0.19821	− 0.14051	− 0.15597	− 0.19363
m_4	4.2347	4.1618	4.2342	4.2192
Q_{10}	3.68062	3.67172	4.1897	5.14093
Q_{10}^2	9.18246	6.05157	8.85045	8.68123

See footnote to Table 6.

seems to be the norm observed with high frequency intraday returns.[26] We could be assured that our estimation of the high frequency data in the previous section was definitely based on the best specification. The GARCH(1,1) model does an excellent job in estimating the short-run volatility dependencies. For the data recorded at daily frequency even though the results in this section indicate that fractionally integrated GARCH models are more appropriate (with shocks to the conditional variance of the two currencies dieing at a slow hyperbolic rate of decay), however this conclusion does not invalidate the authenticity of the results obtained in previous section. The sole purpose of estimating the GARCH(1,1) variance specification in the previous section was to estimate the daily dummy variables. GARCH or FIGARCH processes could have important implications for the dissipation of shocks to conditional variance when forecasting the conditional variance itself, but no significant alterations on the behaviour of estimated daily dummy variables are expected. Even on re-estimating the models for the daily data with a FIGARCH specification, the estimated dummy variables for the two currencies remained approximately the same.[27]

So far the simple GARCH model has stood up remarkably well against some of its newly developed extensions (even if not for the daily data). One final check on the robustness of GARCH model comes from the comparison of GARCH with an alternative conditional variance modelling technique called the stochastic variance model. Recently huge literature on the application of stochastic variance models to financial market returns have emerged. In the stochastic variance model the logarithm of σ_t^2 (variance) is modelled as a stochastic process. The results on applying the SV model to the two currencies both at the hourly and daily frequencies are presented in the Tables 10 and 11, respectively.

The results on applying the SV models at both the frequencies to the two currencies are qualitatively similar to those obtained from the GARCH(1,1) models. For the two currencies, using the hourly data the estimate of ϕ (the AR(1) coefficient in the variance process) is considerably less than 1, similar to the sum of the estimated GARCH parameters and thus reinforcing the earlier result of short-run volatility dependencies in the high frequency data. The estimate of ϕ is significant for both the currencies at the hourly frequency. In terms of other diagnostic statistics the remarkably lower statistics for the skewness and kurtosis for the two currencies on applying the stochastic variance models favours the SV model over the GARCH(1,1) model. The degree of kurtosis in the standardized residuals from the SV model for the two currencies is remarkably lower compared to

[26] Andersen and Bollerslev (1997) have recently explained the finding of insignificant fractional differencing parameter obtained with high frequency data. They argue that while a sudden burst of volatility typically possess both short- and long-run components, the short run decay stands out most clearly over the intra-daily frequencies, whereas the highly persistent processes are only noticeable over longer horizons. Verification of their argument would require longer data sets, which are currently not available for euro. See Baillie et al. (2000) for a more recent discussion.

[27] Only some slight change in the magnitudes of the daily dummy variables was observed.

Table 10
Stochastic variance model (quasi-maximum likelihood estimates for the hourly frequency)

	BP	Euro
ϕ	0.62748 (0.7307)	0.72012 (0.6976)
Log-likelihood	− 1153.76	− 1075.35
AIC	1.258043	1.253948
BIC	1.267063	1.263476
Q_{10}	19.8365	17.1392
$m_3(i)$	− 0.03074	− 0.05424
$m_4(i)$	2.24289	2.25876

See footnote to Table 6. p-values are reported in parentheses.

the GARCH model, a result considered to be ideal when analysing the high frequency data.[28]

Comparison of the two models at the daily frequency is not much meaningful, since we know that the data is best characterised by fractionally integrated variance or long memory volatility process. Keeping aside this fact for a moment and interpreting the estimates obtained from the application of SV models, the estimated value of the ϕ is very close to 1 for the two currencies. This particular result is similar to that obtained from the application of GARCH. The coefficient of kurtosis for the two currencies is considerably less for the stochastic variance model, a very strong result once again displaying the superiority of the SV models. Other diagnostic statistics however give no clear picture on the relative efficiency of the SV models.[29] A similar comparison of the SV model and the GARCH(1,1) model is provided by Kim et al. (1998) using the daily sterling/dollar exchange rate return series. Based on the diagnostic statistics they find SV model to fit the data better than the Gaussian GARCH model, however, not better than the Student's t distributed GARCH model.

The stochastic variance model could definitely be used as an alternative to the GARCH specification. The estimation results indicate that SV models may perform better than GARCH when applied to the high frequency data. Beside the SV model many other alternative models to GARCH have also been applied to financial market returns. Such models include the application of ARFIMA(p,d,q) models with GARCH type innovations to the exchange rates and other macro-economic variables. However, the success which the researchers have obtained with the GARCH models and the relative ease of computability of the GARCH models is

[28] The values of Ljung–Box statistics for the presence of up to 10th-order serial correlation in the standardised residuals of SV models are of greater magnitude than from the GARCH models, implying that the models could not fully account for the serial correlation present in the data. However, perhaps by increasing the number of autoregressive terms in the variance specification of the SV models an even better fit could be obtained.

[29] Application and estimation of a long memory stochastic volatility (LMSV) model is provided in Breidt et al. (1998). The authors formulate the LMSV by incorporating an ARFIMA process in a standard stochastic volatility scheme.

Table 11
Stochastic variance model (quasi-maximum likelihood estimates for the daily frequency)

	BP	Euro
ϕ	0.97237 (0.6150)	0.99391 (0.6084)
Log-likelihood	−592.451	−595.494
AIC	1.291941	1.297002
BIC	1.307700	1.312760
Q_{10}	5.5787	24.3404
$m_3(i)$	−0.28744	−0.33634
$m_4(i)$	2.20944	2.22549

See footnote to Table 6. p-values are reported in parentheses.

probably the reason why the GARCH still remains to be a valuable tool in the field of finance.

5. Conclusion

As already mentioned, volatility is just one dimension of the many, which need to be considered before making any conclusive argument as to whether Britain should replace pound with euro or not. This particular study makes an important contribution along the volatility dimension and significant results were obtained in the previous chapter. Even though the estimated hourly and daily dummy variables for the two currencies possess identical characteristics, they differ substantially when compared in terms of their magnitudes. Euro is considerably more volatile than BP both at the hourly and daily frequencies. Whether the volatility of the two currencies reflects the underlying volatility in the economic fundamentals of two regions is definitely an important avenue for future research. The higher volatility of euro than BP has obviously important implications for many other financial markets and this particular finding definitely requires more research along this dimension before economists converge to any specific decision on this highly debated issue. Short data sets does appear to be a potential weakness, but then to obtain more authenticated results we may have to wait for another 5 years. Further horizons will definitely be explored with the passage of time, as larger data sets would allow even a much more rigorous and deeper study of the volatility dynamics of euro. Hopefully, this study would prove to be a cornerstone in generating future research on the Europe's new currency and on the importance of volatility dynamics in policy making.

Acknowledgements

I am grateful to Martyn Duffy (my supervisor) for the help provided during the course of this study. And I would like to thank the anonymous referee of this journal for providing some very helpful comments and suggestions. Any remaining errors are the sole responsibility of the author.

References

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

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

Andersen, T.G., Bollerslev, T., Das, A., 2001. Variance-ratio statistics and high-frequency data: testing for changes in intraday volatility patterns. Journal of Finance 56, 305–327.

Baillie, R.T., Bollerslev, T., 1989. The message in daily exchange rates: a conditional-variance tale. Journal of Business and Economic Statistics 7, 297–305.

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

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

Baillie, R.T., Cecen, A.A., Han, Y.W., 2000. High frequency Deutsche mark–US dollar returns: FIGARCH representations and non linearities. Multinational Finance Journal 4, 247–267.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307–327.

Bollerslev, T., Domowitz, I., 1993. Trading patterns and prices in the inter-bank foreign exchange market. Journal of Finance 48, 1421–1443.

Bollerslev, T., Mikkelsen, H.O., 1996. Modelling and pricing long memory in stock market volatility. Journal of Econometrics 73, 151–184.

Bollerslev, T., Wooldridge, J.M., 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. Econometric Reviews 11, 143–172.

Box, G.E.P., Pierce, D.A., 1970. Distribution of residual autocorrelations in autoregressive-integrated moving average time series models. Journal of the American Statistical Association 65, 1509–1526.

Breidt, F.J., Crato, N., Lima, P.D., 1998. The detection and estimation of long memory in stochastic volatility. Journal of Econometrics 83, 325–348.

Covrig, V., Melvin, M.T., 2002. Asymmetric information and price discovery in the FX market: does Tokyo know more about the Yen? Journal of Empirical Finance 9, 271–285.

Dickey, D.A., Fuller, W.A., 1979. Distribution of the estimators for autoregressive time series with a unit root. Journal of the American Statistical Association 74, 427–431.

Dickey, D.A., Fuller, W.A., 1981. Likelihood ratio statistics for autoregressive time series with a unit root. Econometrica 49, 1057–1072.

Engle, R.F., Ito, T., Lin, W.-L., 1990. Meteor showers or heat waves? Heteroskedastic intra-daily volatility in the foreign exchange market. Econometrica 58, 525–542.

Fama, E.F., 1965. The behaviour of stock market prices. Journal of Business 38, 34–105.

Foster, F.D., Viswanathan, S., 1990. A theory of the interday variations in volume, variance, and trading costs in securities markets. Review of Financial Studies 3, 593–624.

French, K.R., Roll, R., 1986. Stock return variances: the arrival of information and the reaction of traders. Journal of Financial Economics 17, 5–26.

Goodhart, C.A.E., Giugale, M., 1993. From hour to hour in the foreign exchange market. Manchester School of Economic and Social Studies 61, 1–34.

Harvey, C.R., Huang, R.D., 1991. Volatility in the foreign currency futures market. Review of Financial Studies 4, 543–569.

Hsieh, D.A., 1988. The statistical properties of daily foreign exchange rates: 1974–1983. Journal of International Economics 24, 129–145.

Hsieh, D.A., 1989. Modelling heteroskedasticity in daily foreign exchange rates. Journal of Business and Economic Statistics 7, 307–317.

Hsieh, D.A., Kleidon, A.W., 1996. Bid-Ask spreads in foreign exchange markets: implications for models of asymmetric information. In: Frankel, J.A., Galli, G., Giovannini, A. (Eds.), The Microstructure of Foreign Exchange Markets. NBER, University of Chicago Press, pp. 41–65.

Ito, T., Lyons, R., Melvin, M.T., 1998. Is there private information in the FX market: the Tokyo experiment. Journal of Finance 53, 1111–1130.

Kim, O., Verrecchia, R.E., 1991. Trading volume and price reactions to public announcements. Journal of Accounting Research 29, 302–321.
Kim, S., Shephard, N., Chib, S., 1998. Stochastic volatility: likelihood inference and comparison with ARCH models. Review of Economic Studies 65, 361–393.
Lastrapes, W.D., 1989. Exchange rate volatility and U.S. monetary policy: an ARCH application. Journal of Money, Credit and Banking 21, 66–77.
Ljung, G., Box, G., 1978. On a measure of lack of fit in time series models. Biometrika 65, 297–303.
Mandelbrot, B., 1963. The variation of certain speculative prices. Journal of Business 36, 394–419.
McLeod, A.I., Li, W.K., 1994. Diagnostic checking ARMA time Series models using squared-residual autocorrelations. In: Harvey, A. (Ed.), Time Series, vol. 2, Elgar Reference Collection, International Library of critical writings in Econometrics, vol. 5, Elgar Publishing UK, pp. 301–305.
Phillips, P.C.B., Perron, P., 1988. Testing for a unit root in time series regression. Biometrika 75, 335–346.
White, H., 1982. Maximum likelihood estimation of misspecified models. Econometrica 50, 1–25.