Report n. 194

A Note on volatility estimate-forecast with GARCH models

Emilio Barucci, Roberto Renò

Pisa, Novembre 2000

A Note on Volatility Estimate-Forecast with GARCH models

Emilio Barucci
Dipartimento di Statistica e Matematica
Applicata all'Economia, Università di Pisa

Roberto Renò
Scuola Normale Superiore, Pisa
and INFM, Sezione di Palermo, Italy

November 8, 2000

Keywords: Volatility, Forecasting, High frequency data, GARCH models

Abstract

We apply a methodology based on Fourier analysis to compute the volatility of a diffusion process. In a GARCH setting we show that this method performs well in computing the integrated volatility and that the forecasting performance of the GARCH model is better than the one established with classical methods. These results are confirmed by the analysis of exchange rate high frequency time series. The method allows us to identify microstructure effects on the asset price dynamics. We applied it to the evaluation of the forecast performanc of a GARCH(1,1) model when the sampling frequency of returns is increased at the intraday level

1 Introduction

Estimate and forecast of the volatility for a financial time series is an outstanding topic in the finance literature. Indeed, volatility estimate-forecast plays a crucial role in many different fields, e.g., risk management, time series forecasting, contingent claim pricing, etc..

Starting out from empirical investigations showing that volatility in financial time series is highly persistent with clustering phenomena, in the last twenty years many models have been proposed to describe the volatility evolution. The literature is now quite large with many different specifications of GARCH models.

Empirical analyses have shown a high degree of intertemporal volatility persistence, but in many papers it has also been observed that forecasting with GARCH models can be extremely unsatisfactory when the squared (daily) return is taken as an ex post measure of the volatility in a day, e.g. see [4, 8, 15, 23]. In [4], it is shown that the forecasting performance of a GARCH(1,1) is improved when the daily volatility (*integrated volatility*) is measured by means of the cumulative squared intraday returns. Monte Carlo experiments for the processes estimated on exchange rate time series ($DM-\$$ and $Yen-\$$) show that the high frequency estimator of the volatility leads to an improvement in measuring volatility. The noise of the cumulative intraday squared returns is much lower than that of the daily squared returns. Also the forecasting performance of the GARCH(1,1) model turns out to be higher when the cumulative squared intraday returns estimator is employed.

In this note we address estimate-forecast volatility in a GARCH setting with high frequency data by applying a new methodology developed in [19] to compute the volatility. This method is based on Fourier analysis techniques. Volatility for diffusion processes is defined as the limit of the quadratic variation of the process. This definition motivates standard estimation methods of the volatility of a process based on a *differentiation procedure*: the quadratic variation of a process with a given frequency (day, week, month) is taken as an estimate of the volatility. Estimating volatility by this method using high frequency data presents some drawbacks. Tick by tick data are not equally spaced, in the papers cited above an equally spaced time series for intraday returns is constructed from the linearly interpolated logarithmic midpoint of the bid-ask quotes or by taking the last quote before the reference time. This procedure induces some distortions in the analysis, e.g. it may be the origin of returns autocorrelation, see [17], and reduces the number of observations. The procedure employed below avoids these problems. The proceudre is based on the *integration* of the time series, it employs all the observations without any manipulation, data can be irregularly spaced, no interpolation of the data is needed. Using all the data, the computation of the volatility is more precise. We prove this fact through a Monte Carlo simulation of a GARCH(1,1) model with the parameters estimated in [4]. We show that the variance of our estimator of the integrated volatility is smaller than the one of the cumulative squared intraday returns. Moreover, by measuring the volatility according to our method also the forecasting performance of the GARCH(1,1) model is improved.

This paper is organized as follows. In Section 2 we present the methodology to compute the volatility. In Section 3 we address some problems arising in implementing the methodology. In Section 4 we compute the volatility when the asset price process is described by a GARCH(1,1) model, through Monte Carlo experiments we compare it to the estimate obtained through the cumulative squared intraday rateurns. In Section 5 we evaluate the GARCH(1,1) volatility forecasting performance by measuring volatility with the new method. In Section 6 we evaluate the forecasting performance of the GARCH model for exchange rates high frequency time series.

2 Measuring volatility via Fourier analysis

The methodology developed in [19] allows to compute volatility of diffusion processes through Fourier analysis. In what follows we present the method, for a detailed analysis we refer the reader to the paper.

Given a multivariate process $p_i(t)$, $i = 1, \ldots, D$, representing economic data, we only require its quadratic variation be limited. Among the processses satisfying this assumption we have diffusion processes:

$$dp_i(t) = \sum_{j=1}^{D} \sigma_{ij}(t)dW_j(t) + \mu_i(t)dt \qquad i = 1, \ldots, D, \tag{1}$$

where σ and μ are time dependent random functions and $W_1, \ldots, W_D$ are indipendent Brownian motions.

The volatility of any process with well defined quadratic variation at time t can be defined as follows:

$$\Sigma^2_{ij}(t) := \lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \boldsymbol{E}^{N_t}\Big[(p_i(t+\epsilon) - p_i(t)) \cdot (p_j(t+\epsilon) - p_j(t))\Big] \ \ i, j = 1, \ldots, D, \tag{2}$$

where $\boldsymbol{E}^{N_t}[\cdot]$ denotes the expectation operator and N_t is the $\sigma-$field generated by the full observation of the economic data until time t. The relation between $\Sigma(t)$ and the model (1) is given by

$$\Sigma^2_{ij}(t) = \sum_{k=1}^{D} \sigma_{ik}(t)\sigma_{jk}(t).$$

For the computation of the volatility we normalize the time window [0,T] in which the time series is recorded to $[0, 2\pi]$. In [19, Theorem 1.2] it is shown that the Fourier coefficients of Σ^2 can be computed by means of the Fourier coefficients of dp, then classical results of Fourier theory allows us to reconstruct $\Sigma^2(t)$ $\forall t \in [0, 2\pi]$.

First of all, we consider a univariate setting. The Fourier coefficients of dp are

$$\begin{array}{c} a_0(dp) = \frac{1}{2\pi}\int_0^{2\pi} dp(t) \\ a_k(dp) = \frac{1}{\pi}\int_0^{2\pi} \cos(kt)dp(t) \\ b_k(dp) = \frac{1}{\pi}\int_0^{2\pi} \sin(kt)dp(t) \quad k \geq 1. \end{array} \tag{3}$$

Then we obtain the Fourier coefficients of Σ^2 through the formulae:

$$a_0(\Sigma^2) = \lim_{n\to\infty} \frac{\pi}{n+1-n_0} \sum_{s=n_0}^{n} \frac{1}{2}(a_s^2(dp) + b_s^2(dp)) \tag{4}$$

$$a_k(\Sigma^2) = \lim_{n\to\infty} \frac{2\pi}{n+1-n_0} \sum_{s=n_0}^{n} a_s(dp)a_{s+k}(dp) \tag{5}$$

$$b_k(\Sigma^2) = \lim_{n\to\infty} \frac{2\pi}{n+1-n_0} \sum_{s=n_0}^{n} b_s(dp)b_{s+k}(dp). \tag{6}$$

Note that we are left with the choice of omitting the first n_0 coefficients, since they could be too sensitive to the drift term in equation (1).

By the classical Fourier-Féjer inversion formula, we can reconstruct $\Sigma^2(t)$:

$$\Sigma^2(t) = \lim_{n\to\infty} \sum_{k=0}^{n} (1 - \frac{k}{n}) \cdot (a_k(\Sigma^2)\cos(kt) + b_k(\Sigma^2)\sin(kt)). \tag{7}$$

The generalization to multivariate volatility is straightforward. The Fourier coefficients of Σ^2_{ij} $(i, j = 1, \ldots, D)$ are

$$a_k(\Sigma^2_{ij}) = \lim_{n\to\infty} \frac{\pi}{n+1-n_0} \sum_{s=n_0}^{n} \frac{1}{2}(a_s(dp_i)a_{s+k}(dp_j) + a_s(dp_j)a_{s+k}(dp_i)). \tag{8}$$

Similar formulas hold for (4) and (6).

If $p(t)$ could be observed in continuous time, this method would allow us to reconstruct the volatility at any time inside the interval $[0, 2\pi]$. Note that the Fourier coefficients for cross-volatilities are obtained through the coefficients computed for a single realization of the time series.

3 Implementation

Given a time series of N observations $(t_i, p(t_i))$, $i = 1, \ldots N$, we will compress the data in the interval $[0, 2\pi]$ and compute the integrals in (3) through *integration by parts*:

$$a_k(du) = \frac{1}{\pi}\int_0^{2\pi} \cos(kt)dp(t) = \frac{p(2\pi) - p(0)}{\pi} - \frac{k}{\pi}\int_0^{2\pi} \sin(kt)p(t)dt. \tag{9}$$

As observations are finite, to implement the method and in particular the integration we need an assumption on how data are connected. Our choice is $p(t)$ be equal to $p(t_i)$ in the interval $[t_i, t_{i+1}]$ (piecewise constant). Then, the integral in equation (9) in the interval $[t_i, t_{i+1}]$ becomes

$$\frac{k}{\pi}\int_{t_i}^{t_{i+1}} \sin(kt)p(t)dt = p(t_i)\frac{k}{\pi}\int_{t_i}^{t_{i+1}} \sin(kt)dt = p(t_i)\frac{1}{\pi}(cos(kt_i) - cos(kt_{i+1})) \tag{10}$$

thus avoiding the multiplication by k which amplifies cancellation errors when k becomes large.

The smallest wavelength that can be evaluated is twice the smallest distance between two consecutive prices; in the case of equally spaced data, it will correspond to $k = N/2$ (Nyquist frequency): this frequency is the largest that can be evaluated. Then the Fourier coefficients (4),(5),(6) will be evaluated for $0 \leq k \leq J$ and $n = M$ such that:

$$M + J = \frac{N}{2}. \tag{11}$$

In the integration by parts formula (9), the constant term $(p(2\pi) - p(0))/2\pi$ appears in every coefficient; this term can be set to zero by adding in equation (1) the drift term $-\frac{p(2\pi)-p(0)}{2\pi}dt$. This change of variables will not affect the volatility estimate and will remove a possible source of bias, since we compute the volatility on a single realization.

In [9] we tested the Fourier method. Monte Carlo experiments with a diffusion process with constant volatility have shown that the method is able to estimate volatility in a univariate setting and cross-volatility in a multivariate setting. The precision of the estimate is similar to the one of classical methods. When applied to the daily time series of the Dow Jones Industrial and Transportation index the method replicates the volatility estimates obtained with the classsical method.

4 Volatility Computation

From now on, we will focus on price return volatility, so we will choose $p(t) = \log S(t)$ where $S(t)$ is a generic asset price. Following a large literature, we model the volatility of the asset price through a GARCH(1,1) model:

$$\begin{aligned} dp(t) &= \sigma(t)dW_1(t) \\ d\sigma^2(t) &= \theta(\omega - \sigma^2(t))dt + \sqrt{2\lambda\theta}\sigma^2(t)dW_2(t), \end{aligned} \tag{12}$$

where θ, ω, λ are constant and W_1, W_2 are two indipendent Brownian motions.

Given a time window $[0,1]$ (a day, week, month), we want to compute the integrated volatility of the process, i.e., $\int_0^1 \sigma^2(t+\tau)d\tau$, see [4]. An unbiased estimator of this quantity is given by $(p(1) - p(0))^2$. However this estimator is very noisy: for example, it can be zero also if the volatility is very high. In [4] for exchange rates and in [20] for a stock index, it is shown that an estimator with less noise is provided by the cumulative sum of squared intra-day returns: $\sum_{i=2}^{N}(p(\frac{i}{N}) - p(\frac{i-1}{N}))^2$. Using as evaluation criterion the difference $\int_0^1 \sigma^2(t)dt - \sum_{i=2}^{N}(p(\frac{i}{N}) - p(\frac{i-1}{N}))^2$, both authors find better results using the highest N they analyze, corresponding for both to 5 minutes returns.

The analysis in [19] gives us an estimator of the integrated volatility. Integrating $\Sigma^2(t)$ between 0 and 2π in fact we have

$$\int_0^{2\pi} \sigma^2(t)dt = 2\pi a_0(\Sigma^2), \qquad (13)$$

where $a_0(\Sigma^2)$ is given by (4).

The computation of $a_0(\Sigma^2)$ provides an estimate of the integrated volatility avoiding any manipulation of high frequency data; moreover, all the data are used to compute the Fourier coefficients of dp and of Σ.

To illustrate the validity of the Fourier approach, we simulate the diffusion process (12) by an Euler discretization scheme, see for example [16]. We use the parameters $\theta = 0.035, \omega = 0.636, \lambda = 0.296$ estimated in [4] on the daily returns time series of the Deutsche Mark-U.S. Dollar exchange rate.

Taking a day as a reference measure, we simulate 24 hours of trading by taking $dt = 1/86400$ when discretizing (12), which correspond to an observation every second. In order to simulate high-frequency unevenly sampled times, we select a subset of $[1, 86400]$ extracting the time differences from an exponential distribution of mean $\tau = 14$ seconds. This choice is motivated by the fact that the empirical distributon of $t_i - t_{i-1}$ can be approximated with an exponential shape and 14 seconds is the mean distance between quotes in the DM-$ exchange rate time series which will be analyzed in the following sections. As a result, we will have a data set $(t_k, p(t_k), k = 1, \dots, N)$ with t_k unevenly sampled, and $\sigma(t)$ recorded for every t. Then we compute three estimators of the integrated volatility: the squared daily return, the cumulative intra-day five minutes returns[1] and the Fourier-based estimator (13). The result is illustrated in figure 1, where the distribution of the difference between the integrated volatility and the estimator is shown. As expected, the squared daily return is a very noisy estimator. As argued in [4], estimating the volatility through the sum of the squared intraday returns we have a smaller variance. However, measuring the volatility according to the Fourier method we have a further reduction of the variance.

[1]Five minutes returns are obtained from the unevenly sampled data stream by the procedure illustrated in [21].

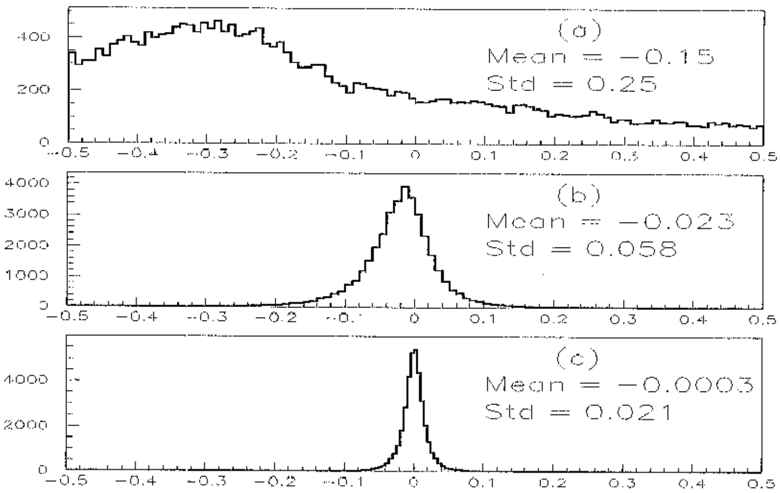

FIGURE 1: Distribution of $\int_0^1 \sigma^2(t)dt - \hat{\sigma}$ where $\hat{\sigma}$ are three different estimators of the integrated volatility: (a) $\hat{\sigma} = (p(1) - p(0))^2$; (b) $\hat{\sigma} = \sum_{i=2}^{288}(p(\frac{i}{288}) - p(\frac{i-1}{288}))^2$; (c) $\hat{\sigma} = 2\pi a_0(\Sigma^2)$.

5 Volatility forecast evaluation

Following [12, 13], the GARCH continuos time diffusion (12) can be discretized, obtaining the weak GARCH process:

$$\zeta_{t+1}^2 = \psi + \alpha \cdot r_t^2 + \beta \cdot \zeta_t^2, \tag{14}$$

where $r_t = p_t - p_{t-1}$ and ζ_t is the best linear predictor of r_t^2 expressed as a linear combination of the lagged squared returns. In [13] the exact relation bewteen ψ, α, β and θ, ω, λ is provided, so that one can calibrate the process (14) on a given time series to obtain the coefficients of the continuous diffusion (12).

If the continuous time GARCH model holds true for the time series at hand, then ζ_{t+1}^2 provides us with a forecast of $\int_0^1 \sigma^2(t+\tau)d\tau$. While there is a strong support in favour of a high persistence in the volatility dynamics, the one day ahead forecasting performance of the above model has been evaluated in the literature to be very poor. However, as outlined in [4] the reason for this apparently disappointing result hinges on the fact that the measure of the volatility used to evaluate ex post the forecasting performance was the squared daily return. The authors point out that a careful measure of the realized integrated volatility is needed in order to correctly evaluate the performance of the model.

Estimator	R^2 (DM-\$)	R^2 (Y-\$)
$(p(1)-p(0))^2$	0.061	0.092
$\sum_{i=2}^{288}(p(\frac{i}{288})-p(\frac{i-1}{288}))^2$	0.446	0.498
Fourier	0.464	0.507
R^2_∞	0.465	0.508

TABLE 1: The R^2 obtained on simulated data with different estimators of the integrated volatility, compared with the ideal R^2_∞.

We use the Fourier method to measure the realized integrated volatility and we evaluate the GARCH(1,1) forecasting performance according to it. The performance with respect to this estimator is compared to the one associated to the cumulative squared intraday returns. The comparison can be done through the R^2 of the linear regression:

$$\hat{\sigma}_t^2 = a + b \cdot \zeta_t^2 + \epsilon_t, \tag{15}$$

which is given by

$$R^2 = 1 - \frac{Var[\hat{\sigma}_t^2 - \zeta_t^2]}{Var[\hat{\sigma}_t^2]} \tag{16}$$

where $\hat{\sigma}^2$ is the estimator of the ex post volatility.

As before, we used the simulation according to an Euler scheme of the process (12), using again $\Delta t = 1/86400$ and 86400 time steps per day (24 hours of trading) with the same parameters of Section 4. To simulate a high frequency unevenly sampled time series we used the procedure described in Section 4. The forecasting model is given by equation (14) with the parameters $\psi = 0.022, \alpha = 0.068, \beta = 0.898$ corresponding to those employed in the above Section according to [13].

We point out that the R^2 obtained for an estimator must be compared to the R^2 obtained when the ex post volatility measure is perfectly known; its value is given by:

$$R^2_\infty = 1 - \frac{Var[\int_0^1 \sigma^2(t)dt - \zeta_t^2]}{Var[\int_0^1 \sigma^2(t)dt]}. \tag{17}$$

This value can be measured in our simulation setting.

The result is given in table 1: the Fourier estimator gives an R^2 which is very close R^2_∞. Employing the Fourier estimator, the GARCH forecasting performamce turns out to be better than the one obtained by measuring the volatility through the cumulative squared intraday returns.

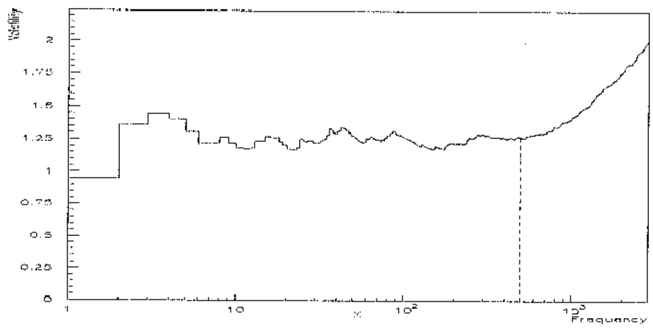

FIGURE 2: Volatility Estimate $\Sigma^{1/2}$ as a function of n as in (4), obtained October 1^{st}, 1992 for the DM-$ exchange rate. The dashed line indicates the cut we perform to calculate volatility.

6 Exchange Rates Time Series

The data set we analyzed is the collection of all high frequency tick by tick quotes (bid and ask) of the foreign exchange rate for Deutsche Mark-U.S. Dollar and Japanese Yen-U.S. Dollar, rounded to the near even second. The data was collected by Olsen & Associates, which delivered the data set to us, from October, 1^{st} 1992 to September, 30^{th} 1993. Such a data set has been extensively studied, e.g. see [1, 3, 4, 5, 8, 21, 25] among others.

We define the price to be the midprice between bid and ask quote. The trading day is chosen to begin and end at 21:00 GMT. We excluded weekends and, in general trading days with less than 1000 quotes for the DM-$ and less than 400 quotes for the Y-$. Moreover we applied the filter described in [11] which removes roughly 0.35 % of the quotes. We end up with 1·466·944 quotes for the DM-$ and 567·758 quotes for Y-$, distributed over 258 trading days.

In the previous sections we showed that the estimate through the Fourier analysis can improve the measurement and forecasting evaluation of volatility with a GARCH model; all the analysis was based on simulated time series. When we apply the Fourier estimator to real high frequency data, we have to calculate $a_0(\Sigma^2)$ through the expansion (4) stopped at the frequency $N/2$. Here we encounter a severe difficulty: the diffusion model (12) cannot hold at time steps at which microstructure effects, such us price discreteness or bid-ask bounce effect, are present. This affects the calculation of the Fourier coefficients at the highest frequencies; indeed, such a calculation appears to be distorted. This is illustrated in figure 2 which shows the calculation of the volatility as in (4), as a function of the highest frequency n employed in the computation, for n ranging from 0 to $N/2$. The plot in figure 2 can

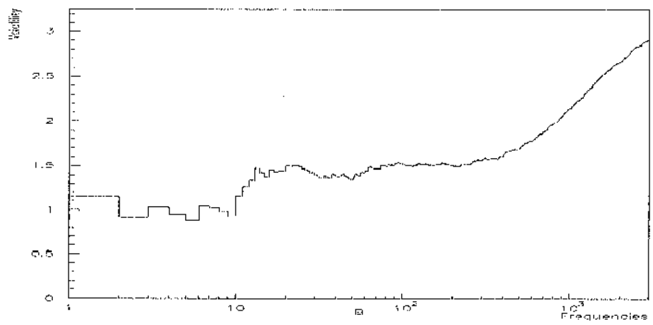

FIGURE 3: Volatility estimate $\Sigma^{1/2}$ for the simulated process (18) as a function of the frequency n.

also be interpreted as the mean of the power spectrum of the exchange rate taken from frequency 0 to frequency n. If $du(t)$ is Normally distributed, as in a model like (12), then the spectrum of u would be flat and so would be its mean as a function of n. Figure 2 shows that the power spectrum of u is not flat; for a frequency larger than a certain value N_{cut}, to be determined empirically, the Fourier coefficients are higher than expected. For the time series we considered, it turns out that $N_{cut} \simeq 500$ for the DM-\$ exchange rate and $N_{cut} \simeq 200$ for the Y-\$. These frequencies correspond to a time step, calculated as $\frac{86400}{2 \cdot N_{cut}}$ seconds, of roughly 1.5 and 3.5 minutes respectively. We conclude that the price process cannot be modelled by equation (12) for time steps smaller than two or three minutes.

This behaviour of the exchange rate spectrum can be motivated by the fact that high frequency returns are negatively correlated, a phenomenon that has been documented in [1, 10] among others. We can show this by simulating the process:

$$\begin{cases} \sigma_{t+1}^2 = \psi + \alpha \cdot p_t^2 + \beta \cdot \sigma_t^2 \\ \epsilon_{t+1} = \rho\epsilon_t + \sqrt{1-\rho^2}\eta_t \\ p_{t+1} = \sigma_{t+1}\epsilon_t \end{cases} \tag{18}$$

where ρ is the first-order serial correlation coefficient and η_t is a normal distributed random variable with mean 0 and variance 1. Figure 3 shows the analogous of figure 2 for this simulated time series with $\rho = -0.985$ and a time step of one second in (18), and the same phenomenon occurs. Smaller values of $|\rho|$ would lead to a larger cut frequency N_{cut}, as we checked in our simulation setting. Such a negative correlation can be linked to non-synchronous trading (see [17]) or to the way market makers manage their inventory positions (see [1]). The analysis of microstructure effects on high

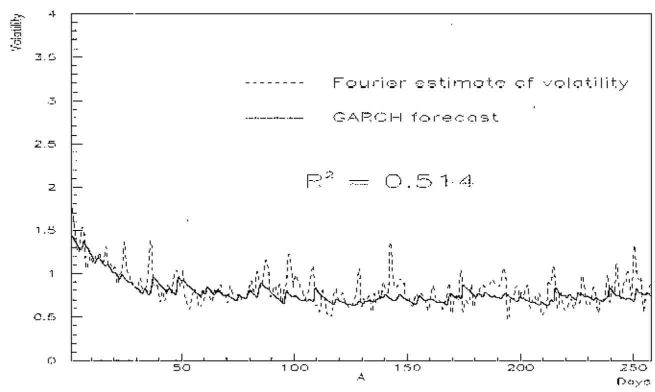

FIGURE 4: Comparison between the GARCH forecast and the realized volatility for the DM-$ exchange rate, from October, 2^{th} 1992 to September, 30^{th} 1993. The realized volatility is measured with the Fourier estimator.

Estimator	DM-$		Y-$	
	R^2	MSE	R^2	MSE
$\sum_{i=2}^{288}(p(\frac{i}{288}) - p(\frac{i-1}{288}))^2$	0.444	0.082	0.134	0.260
Fourier	0.514	0.071	0.158	0.245

TABLE 2: R^2 and MSE obtained on the two time series considered.

frequency time series is beyond the scopes of this paper, we refer the interested reader to a forthcoming paper.

For the subsequent analysis, we cut the highest frequencies in the computation of volatility, i.e. we will evaluate the Fourier coefficient (4) for $n = \min(N/2, N_{cut})$. In the above cited papers, to prevent microstructure effects distorting the estimate of the volatility, the authors have manipulated the data by building (through interpolation) five minutes returns time series. This procedure has the drawbacks of any data manipulation, drawbacks that can be completely avoided with our method. As a matter of fact, we do not manipulate the original time series and we use all the data in order to compute the Fourier coefficients of dp, we only stop the expansion (7) properly.

We evaluate the GARCH(1,1) model forecasting performance when the volatility is computed according to the Fourier method. In Figure 4, we show the forecast of the GARCH model together with the actual volatility measurement with the Fourier method. Table 2 shows the corresponding R^2 and MSE ($MSE = Var[\hat{\sigma_t^2} - \zeta_t^2]$).

We observe that the GARCH model performs well in forecasting when the Fourier method is employed to compute the volatility. Its performance is higher than the one associated with the one obtained when the cumulative squared intraday returns is employed as a measure of the volatility.

The R^2 on the Y-\$ time series is disappointingly low; this is due to few events with pronounced unexpected volatility; for example, if we exclude from the sample the realization of August 19^{th} 1993, when the realized squared volatility is measured to be 5.0 (the exchange rate went up of 3.35 % that day), then R^2 becomes 0.270.

7 Forecasting exchange rate volatility increasing the sampling frequency

As an illustrative example of the validity of the Fourier approach, we use it to evaluate the forecast performance of GARCH(1,1) when the sampling frequency of returns is increased, following [8].

Since, as discussed in [12, 13], for the model (12) temporal aggregation of the weak GARCH processes holds, then we can discretize the process (14) at finer frequencies in a straightforward manner. If we denote :

$$r_m(t) = p(t) - p(t - \tfrac{1}{m})$$

then we can write:

$$\zeta_m^2(t) = \psi_m + \alpha_m \cdot r_m^2(t - \tfrac{1}{m}) + \beta_m \cdot \zeta_m^2(t - \tfrac{1}{m}) \tag{19}$$

where $\zeta_m^2(t)$ is the best linear predictor of $r_m^2(t)$ from the space spanned by the lagged squared returns. The relation between $(\psi_m, \alpha_m, \beta_m)$ in equation (19) and $(\omega, \theta, \lambda)$ in equation (12) can be obtained for every m in a closed form following [13]. Equation (14) corresponds to (19) with $m = 1$. Table 3 reports the GARCH(1,1) coefficients at different frequencies for the two time-series considered.

Following [8], the optimal forecast for the h period hence integrated volatility, $\int_t^{t+h} \sigma^2(s)ds$, using returns spaced by $1/m$, is given by:

$$F_{m,h}(t) = mh\sigma_m^2 + \frac{\alpha_m + \beta_m}{1 - \alpha_m - \beta_m}[1 - (\alpha_m + \beta_m)^{m \cdot h}](\zeta_m^2(t) - \sigma_m^2), \tag{20}$$

where $\sigma_m^2 = \psi_m \cdot (1 - \alpha_m - \beta_m)^{-1}$. In this analysis, we will concentrate only on daily forecast evaluation, or $h = 1$ in equation (20), since our data sample covers only one year of data. The realized ex-post integrated volatility will be indicated as $\hat{\sigma}_h^2(t)$.

In order to make this analysis comparable to the ones of other authors, we choose to evaluate the forecast performance of the GARCH(1,1) model with the following statistical indicators:

$$RMSE = \boldsymbol{E}\left[(\hat{\sigma}_h^2(t) - F_{m,h}(t))^2\right]^{\frac{1}{2}}$$

DM-$				Y-$			
m	σ_m^2	α_m	β_m	m	σ_m^2	α_m	β_m
1	0.6365	0.0679	0.8978	1	0.4760	0.1043	0.8431
2	0.318	0.0541	0.9285	2	0.2380	0.0840	0.8893
3	0.2122	0.0466	0.9418	3	0.1587	0.0726	0.9095
6	0.1061	0.0353	0.9589	6	0.0793	0.0553	0.9358
12	0.0530	0.0262	0.9709	12	0.0397	0.0411	0.9544
24	0.0265	0.0192	0.9794	24	0.0198	0.0302	0.9676
48	0.0133	0.0139	0.9854	48	0.0099	0.0219	0.9770
96	0.0066	0.0100	0.9896	96	0.0050	0.0157	0.9837
144	0.0044	0.0082	0.9915	144	0.0033	0.0130	0.9867
288	0.0022	0.0059	0.9940	288	0.0016	0.0093	0.9906

TABLE 3: Coefficients of GARCH(1,1) at different frequencies, obtained (according to [13]) from the continuous-time coefficients $\theta = 0.035, \omega = 0.636, \lambda = 0.296$ for the DM-$ exchange rate, and $\theta = 0.054, \omega = 0.476, \lambda = 0.480$ for the Y-$ exchange rate.

$$MAE = \boldsymbol{E}\,|\hat{\sigma}_h^2(t) - F_{m,h}(t)|$$

plus the following heteroskedasticity adjusted statistics:

$$HRMSE = \boldsymbol{E}\,[(1 - F_{m,h}(t)/\hat{\sigma}_h^2(t))^2]^{\frac{1}{2}}$$

$$HMAE = \boldsymbol{E}\,|1 - F_{m,h}(t)/\hat{\sigma}_h^2(t)|$$

plus a logarithmic loss function which exaggerates the influence of low volatility observations:

$$LL = \boldsymbol{E}\,[\log(F_{m,h}(t)/\hat{\sigma}_h^2(t))]$$

These estimators are reported in table 4 for the DM-$ time series and in table 5 for the Y-$ time series, with $\hat{\sigma}_h^2(t)$ calculated as the sum of 5-minutes intraday returns and with the Fourier estimator.

We stress that, for sake of simplicity, we are completely neglecting intraday patterns and macro-economic announcement effects, which have been documented to be important at the intra-dey level; for a summary of the extant literature on this topic, see [2]. Moreover we are neglecting the fact that temporal aggregation of the continuous GARCH process (12) has not been confirmed empirically, see for example [2, 25].

Our results on the forecast performance of a GARCH model as a function of the sampling frequency are substantially in agreement with the ones reported in [8]. Volatility forecasts improve if higher frequency returns are used, but this procedure has an intrinsic limit due to the fact that beyond a certain time scale intra-day features and microstructure effects become

5 minutes returns					
m	RMSE	MAE	HRMSE	HMAE	LL
1	0.300	0.213	0.621	0.418	0.108
2	0.297	0.208	0.546	0.381	0.039
3	0.308	0.211	0.517	0.363	0.002
6	0.309	0.213	0.523	0.366	-0.007
12	0.311	0.202	0.487	0.338	-0.051
24	0.314	0.214	0.550	0.366	-0.048
48	0.324	0.214	0.566	0.373	-0.019
96	0.331	0.221	0.655	0.401	0.041
144	0.337	0.228	0.679	0.426	0.102
288	0.381	0.251	0.800	0.475	0.154
Fourier Estimator					
m	RMSE	MAE	HRMSE	HMAE	LL
1	0.293	0.203	0.378	0.282	-0.081
2	0.300	0.209	0.363	0.279	-0.151
3	0.312	0.219	0.363	0.286	-0.189
6	0.314	0.227	0.385	0.302	-0.198
12	0.324	0.231	0.390	0.304	-0.242
24	0.327	0.237	0.431	0.317	-0.240
48	0.330	0.233	0.431	0.316	-0.210
96	0.326	0.228	0.474	0.320	-0.151
144	0.321	0.222	0.470	0.316	-0.090
288	0.357	0.232	0.559	0.336	-0.038

TABLE 4: Statistics for the forecast evaluation of the GARCH(1,1) model on the DM-$ time series when returns are spaced by $1/m$ days. The results are reported for both the estimators considered.

5 minutes returns					
m	RMSE	MAE	HRMSE	HMAE	LL
1	0.504	0.282	0.589	0.421	-0.083
2	0.522	0.296	0.536	0.404	-0.204
3	0.537	0.296	0.664	0.422	-0.170
6	0.521	0.284	0.560	0.393	-0.175
12	0.480	0.273	0.485	0.377	-0.216
24	0.475	0.272	0.508	0.393	-0.162
48	0.467	0.270	0.514	0.398	-0.142
96	0.479	0.289	0.582	0.448	-0.075
144	0.512	0.307	0.665	0.490	-0.009
288	0.521	0.321	0.767	0.546	0.066
Fourier Estimator					
m	RMSE	MAE	HRMSE	HMAE	LL
1	0.495	0.300	0.523	0.405	-0.152
2	0.517	0.312	0.508	0.396	-0.274
3	0.536	0.315	0.616	0.416	-0.241
6	0.518	0.306	0.532	0.394	-0.247
12	0.480	0.294	0.467	0.373	-0.289
24	0.472	0.288	0.477	0.379	-0.235
48	0.465	0.289	0.484	0.389	-0.216
96	0.474	0.301	0.548	0.424	-0.149
144	0.499	0.312	0.603	0.449	-0.084
288	0.506	0.323	0.686	0.494	-0.009

TABLE 5: Statistics for the forecast evaluation of the GARCH(1,1) model on the Y-$ time series when returns are spaced by $1/m$ days. The results are reported for both the estimators considered.

prominent. We confirm the observation that such a time scale is around few hours, this turns out to be true also for the Y-$ time series, which has not yet been analyzein this context. However we cannot appreciate this effect very clearly, especially with the RMSE and MAE indicators, because of the fact that we are using a time series which is quite short (one year), while in [8] a ten years long time series has been used.

However, our aim is to show that the Fourier estimator performs better than the cumulative sum of intra-day returns, and this turns out to be the case; this is particulary clear looking at the heteroskedasticity adjusted statistics. Moreover we can see that the Fourier estimator does a better job on the DM-$ than on the Y-$ in decreasing the measurement error, since the DM-$ exchange rate is more liquid.

An exception to these encouraging results is given by the LL indicator, which always performs worst than the 5-minutes estimator. We explain this result by the fact that such an estimator gives a larger weight to observations when the volatility is low, i.e. when the number of quotes is low, so that the Fourier method is less efficient.

8 Conclusions

In the recent years, a voluminose literature has been devoted to the study of volatility measuring and forecast. In this field, the importance of high-frequency data has been stressed, expecially in the forecasting evaluation of the popular GARCH models.

In this paper we introduced a different method to estimate volatility; the main feature of this method is that it is based upon integration instead of differentiation, so that it naturally exploits the time-structure of high frequency data and all quotes are included in the computation with no need of data manipulation. Using simulations we proved that, at least theoretically, this method performs better than the other used in the literature in both measuring volatility and evaluating the forecast performance of a GARCH model.

We used this method on two time series of foreign exchange rates: on real data one has to deal with microstructure effects, which become dominant when the sampling frequency of returns becomes comparable to the frequency of tick-by-tick quotes. We gave a precise estimation of the time step above which these effects can be neglected in the classical diffusion scheme, and we showed how to eliminate the microstructure distortions: since the Fourier estimator is given by an expansion of the Fourier coefficients as a function of the frequency, it is enough to cut the highest frequency in a suitable way.

We used the Fourier method in the evaluation of the forecast performance of the GARCH(1,1) model when the returns sampling frequency is increased;

the results in recent literature are confirmed, moreover the Fourier estimator reduces substantially the error on the forecast evaluation.

References

[1] Andersen, T. Bollerslev, (1997) Heterogeneus Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Return, *Journal of Finance*, 52, 975-1005

[2] Andersen, T., Bollerslev, T., (1998) Towars a unified framework for high and low frequency return volatility modeling, *Statistica Neerlandica*, 52, 273-302

[3] Andersen, T. Bollerslev, (1998) Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies. *Journal of Finance*, 53: 219-265.

[4] Andersen, T. and Bollerslev, T. (1998) Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts. *International Economic Review*, 39: 885-905.

[5] Andersen, T. Bollerslev, T., Diebold, F. and Labys, P. (1999) The Distribution of Realized Exchange Rate Volatility, *Journal of American Statistical Association*, forthcoming.

[6] Andersen, T. Bollerslev, T., Diebold, F. and Labys, P. (1999), (Understanding, Optimizing, Using and Forecasting) Realized Volatility and Correlation, working Paper FIN-99-061, Department of Finance, Stern School of Business, New York University .

[7] Andersen, T. Bollerslev, T., Diebold, F. and Labys, P. (2000) The Distribution of Stock Return Volatility, *Journal of Financial Economics*, forthcoming.

[8] Andersen, T., Bollerslev, T. and Lange, S. (1998) Forecasting Financial Market Volatility: Sample Frequency vis-à-vis Forecast Horizon. Forthcoming *Journal of Empirical Finance.*

[9] Barucci, E., Mancino, M., Renò, R. (2000) Volatility Estimate via Fourier Analysis, manuscript

[10] Bollerslev, T., Domowitz, I., Trading patterns and prices in the interbank foreign exchange market, *Journal of Finance*, 48, 1421-1443

[11] Dacorogna , M.M., Muller, U., Nagler, R.J., Olsen, R., Pictet, O.V. (1993) A geographical model for the daily and weekly seasonal volatility in the FX market, *Journal of International Money and Finance*, 12, 413-438

[12] Drost, F. and Nijman, T. (1993) Temporal aggregation of GARCH processes, Econometrica, 61, 909-927

[13] Drost, F. and Werker, B. (1996) Closing the GARCH Gap: Continuous time GARCH Models. *Journal of Econometrics*, 74: 31-57.

[14] Ebens, H. (1999), Realized Stock Volatility, manuscript.

[15] Figlewski, S. (1997) Forecasting Volatility. *Financial Markets, Institutions and Instruments*, 6: 1-88.

[16] Kloeden, P., Platen, E., Numerical Solution of Stochastic Differential Equations, Springer

[17] Lo, A. and MacKinlay, C. (1990) An Econometric Analysis of Nonsynchronous Trading, *Journal of Econometrics*, 45: 181-211.

[18] Lundin, M., Dacorogna, M.M., Muller, U.A., (1999), Correlation of High-Frequency Financial Time Series, in Financial Markets Tick by Tick, ed. Pierre Lequeux, Wiley

[19] P. Malliavin, M. Mancino, (1999) Fourier series method in Statistical Determination of Multivariate Volatilities, Infinitesimal Principal Component Analysis and a geometric vision of Market evolution. Dimadefas, Firenze, December 1999, 99-12.

[20] Martens, M. (2000), Measuring and forecasting stock market volatility using high-frequency data, manuscript, University of New South Wales

[21] P.Muller, et al. (1990), Statistical Study of Foreign Exchange Rates, Empirical Evidence of a Price Change Scaling Law and Intraday Analysis, Journal of Banking and Finance, 14, 1189-1208

[22] Nelson, D. (1990) Arch Models as Diffusion Approximations. *Journal of Econometrics*, 45: 7-38.

[23] A. Pagan and G. Schwert (1990) Alternative Models for Conditional Stock Volatility; *Journal of Econometrics*, 45: 267-290.

[24] W.H.Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery (1992), *Numerical Recipes*, Cambridge University Press, Cambridge

[25] G.Zumbach (2000), The Pitfalls in Fitting GARCH Processes, in *Advances in Quantitative Asset Management*, ed. Christian L. Dunis, Kluwer Academic Publishers