On Measuring Volatility and the GARCH Forecasting Performance ⋆

Emilio Barucci [a,∗] , Roberto Renò [b]

[a] *Università di Pisa, Dipartimento di Statistica e Matematica Applicata all'Economia, via Cosimo Ridolfi 10, 56124, Pisa, ITALY*

[b] *Scuola Normale Superiore, Piazza dei Cavalieri 7, 56121, Pisa, ITALY*

Forthcoming

Journal of International Financial Markets, Institutions and Money

Abstract

We apply a new algorithm based on Fourier analysis to compute the volatility of a diffusion process. By using simulations of the continuous-time GARCH model, we show that our method performs well in computing integrated volatility. We show that linear interpolation of high frequency observations induces a downward bias in estimating integrated volatility. By measuring ex post volatility with our method, we find that the forecasting performance of the GARCH model is improved with respect to what is established when classical methods are employed. These results are confirmed by the analysis of exchange rate high frequency time series.

Key words: Volatility, High frequency data, GARCH models
JEL Classification: C22, C53, F31

⋆ We thank Carlo Bianchi, Paul Malliavin, Mavira Mancino, Rosario Mantegna and an anonymous referee for useful comments. We acknowledge Olsen & Associates for the data. The usual disclaimers apply.

∗ Corresponding author. Tel.: +39 050 945235; fax: +39 050 945375
Email addresses: ebarucci@ec.unipi.it (Emilio Barucci), reno@cibs.sns.it (Roberto Renò).

1 Introduction

Volatility estimation and forecasting is an outstanding topic in the financial literature. Indeed, it plays a crucial role in many different fields, e.g., risk management, time series forecasting and contingent claim pricing.

In the last twenty years, starting out from empirical investigations showing that volatility in financial time series is highly persistent with clustering phenomena, many models have been proposed to describe volatility evolution. The literature is now quite large with several specifications of auto-regressive models. Empirical analyses have shown a high degree of inter-temporal volatility persistence, but in many papers it has also been observed that forecasting with GARCH models can be extremely unsatisfactory when the daily volatility is measured ex post by the squared (daily) return, see Andersen and Bollerslev (1998c); Andersen, Bollerslev and Lange (1999); Figlewski (1997); Pagan and Schwert (1990). In Andersen and Bollerslev (1998c) it is shown that the forecasting performance of a GARCH(1,1) model is improved when the daily volatility (*integrated volatility*) is measured by means of the cumulative squared intraday returns. Monte Carlo experiments of processes, whose parameters have been estimated on exchange rate time series (DM-$ and Yen-$), show that the noise of the high frequency volatility estimator is smaller than that of the daily squared return. On this topic see also Andersen, Bollerslev, Diebold and Labys (1999b); Andersen, Bollerslev, Diebold and Ebens (2000); Barndorff-Nielsen and Shephard (2000a,b).

In this paper we address volatility estimation and forecasting in a GARCH setting with high frequency data by applying a new algorithm, which has been developed in Malliavin and Mancino (2000), to compute the volatility of a diffusion process. This method is based on Fourier analysis techniques (hereafter *Fourier method*). The volatility of a diffusion process is defined as the limit of its quadratic variation. This definition motivates standard volatility estimation methods based on a *differentiation procedure*: the quadratic variation of a process with a given frequency (day, week, month) is taken as a volatility estimate. Estimating volatility by this method presents some drawbacks when high frequency data are used. Tick-by-tick data are not equally spaced, in the above cited papers an equally spaced time series for intraday returns is constructed by linearly interpolating logarithmic midpoints of bid-ask adjacent quotes or by taking the last quote before a given reference time (henceforth called imputation method). This procedure induces some distortions in the analysis, e.g. it may be at the origin of returns autocorrelation, and it reduces the number of observations. The method proposed in this paper avoids these problems; it is based on *integration* of the time series, and it employs all the (irregularly spaced) observations. We assume the price to be piecewise constant, i.e. the price is constant between two subsequent observations.

Volatility computation by using all the data with the Fourier method should be more precise. We illustrate this fact through Monte Carlo simulations of a continuous-time GARCH(1,1) model with the parameters estimated in Andersen and Bollerslev (1998c). We show that the variance of our integrated volatility estimator is smaller than that of the cumulative squared intraday returns. The precision of the cumulative squared intraday returns in measuring volatility depends on the procedure employed to build an equally spaced time series. Linear interpolation causes a downward bias which increases with sampling frequency. The imputation method is immune from these drawbacks. Linear interpolation seems to be at the origin of a downward bias, and this conjecture is confirmed by implementing the Fourier method with linearly interpolated observations instead of assuming the price to be piecewise constant. In this case, a strong downward bias arises.

Through Monte Carlo simulations, we show that, by measuring integrated volatility according to the Fourier method, the forecasting performance of the GARCH(1,1) model is better than that obtained by computing volatility according to the cumulative squared intraday returns.

These results are confirmed when the method is applied to compute volatility of exchange rate high frequency time series. We apply the Fourier method to the evaluation of the forecasting performance of the daily GARCH model and of the intraday GARCH model, as in Andersen, Bollerslev and Lange (1999). For both the time series considered, the GARCH model forecasts are evaluated to be better if the Fourier method is employed as a volatility estimate instead of the cumulative squared intraday returns.

Our paper is organized as follows. In Section 2 we present the algorithm to compute volatility. In Section 3 we address the algorithm implementation. In Section 4 we compute volatility when the asset price process is described by a GARCH(1,1) model, and by using Monte Carlo experiments we compare it to the estimate obtained with the sum of squared intraday returns. In Section 5 we evaluate through simulations the GARCH(1,1) model volatility forecasting performance by measuring volatility with the new method. In Section 6 we evaluate the daily GARCH model performance in forecasting daily volatility of exchange rate high frequency time series. In Section 7 we consider the GARCH model forecasting performance when it is extended to intraday returns. Section 8 concludes.

2 Measuring volatility via Fourier analysis

The algorithm developed in Malliavin and Mancino (2000) allows us to compute the diffusion process volatility by Fourier analysis. In what follows we present the method, for a detailed analysis we refer the reader to the paper.

Given a process $p(t)$, we only require its quadratic variation to be bounded. Among the processes satisfying this assumption, we have the diffusion process:

$$dp(t) = \sigma(t)dW(t) + \mu(t)dt, \tag{1}$$

where σ and μ are time dependent random functions and $W(t)$ is a Brownian motion.

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

$$\Sigma^2(t) := \lim_{\epsilon \downarrow 0} \frac{1}{\epsilon} \boldsymbol{E} \left[(p(t+\epsilon) - p(t))^2 \, | \mathcal{F}_t \right], \tag{2}$$

where $\boldsymbol{E}[\cdot|\mathcal{F}_t]$ denotes the conditional expectation operator with respect to the $\sigma-$field $\mathcal{F}_t$ generated by the full observation of the process until time t; for model (1) we have $\Sigma^2(t) = \sigma^2(t)$. We normalize the time window [0,T] in which the time series is recorded to $[0, 2\pi]$. In Malliavin and Mancino (2000, 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 $\sigma^2(t)$ $\forall t \in [0, 2\pi]$. The Fourier coefficients of dp are

$$\begin{aligned} a_0(dp) &= \tfrac{1}{2\pi} \textstyle\int_0^{2\pi} dp(t) \\ a_k(dp) &= \tfrac{1}{\pi} \textstyle\int_0^{2\pi} \cos(kt) dp(t) \\ b_k(dp) &= \tfrac{1}{\pi} \textstyle\int_0^{2\pi} \sin(kt) dp(t); \; k \geq 1. \end{aligned} \tag{3}$$

Following Malliavin and Mancino (2000), we obtain the Fourier coefficients of σ^2 through the formulae:

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

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

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

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) [1]. By the classical Fourier-Féjer inversion formula, we can reconstruct $\sigma^2(t)$:

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

The generalization to multivariate volatility is straightforward, see Malliavin and Mancino (2000). If $p(t)$ is observed in continuous time, this method would allow us to reconstruct the instantaneous volatility at any time inside the interval $[0, 2\pi]$.

[1] The fact that the lowest frequency coefficients could be affected by a trend component has been argued by Malliavin and Mancino (2000). In our analysis, we will use $n_0 = 1$, since exchange rate high frequency time series do not display systematic trends.

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*:

$$a_k(dp) = \frac{1}{\pi}\int_0^{2\pi} \cos(kt)dp(t) = \frac{p(2\pi) - p(0)}{\pi} - \frac{k}{\pi}\int_0^{2\pi} \sin(kt)p(t)dt. \tag{8}$$

To implement the method and in particular the integration, we need an assumption on how the data are connected. Our choice is to set $p(t)$ equal to $p(t_i)$ in the interval $[t_i, t_{i+1}]$ (piecewise constant). Then, the integral in (8) in the interval $[t_i, t_{i+1}]$ becomes

$$\frac{k}{\pi}\int_{t_i}^{t_{i+1}} \sin(kt)p(t)dt = p(t_i)\frac{k}{\pi}\int_{t_i}^{t_{i+1}} \sin(kt)dt = p(t_i)\frac{1}{\pi}\left[\cos(kt_i) - \cos(kt_{i+1})\right]. \tag{9}$$

The smallest wavelength that can be evaluated in (3) to avoid aliasing effects is twice the smallest distance between two consecutive prices; in the case of equally spaced data, it will correspond to the frequency $N/2$ (Nyquist frequency). This frequency is the largest that can be evaluated. If the Fourier coefficients of σ^2 (4),(5),(6) have to be evaluated for $0 \leq k \leq J$, then the largest frequency n in (7) is

$$n = \frac{N}{2} - J, \tag{10}$$

because of the terms a_{s+k}, b_{s+k} in (5),(6). In the integration by parts formula (8), the constant term $\frac{p(2\pi)-p(0)}{2\pi}$ appears in each coefficient; this term can be set at 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.

In Barucci, Mancino and Renò (2000) we tested this method on equally spaced data. Monte Carlo experiments simulating a diffusion process with constant volatility showed that the method allows to estimate volatility in a univariate setting and cross-volatilities in a multivariate setting. The precision of the estimate is similar to that of classical methods. When applied to the daily time series of the Dow Jones Industrial and Dow Jones Transportation Index, the Fourier method replicates the volatility estimates obtained by the classical method.

4 Volatility Computation

Let $p(t) = \log S(t)$, where $S(t)$ is a generic asset price. Following a large literature, we model the asset price through the continuous-time GARCH model:

$$
\begin{aligned}
dp(t) &= \sigma(t)dW_1(t) \\
d\sigma^2(t) &= \theta\left[\omega - \sigma^2(t)\right]dt + \sqrt{2\lambda\theta}\sigma^2(t)dW_2(t),
\end{aligned}
\tag{11}
$$

where θ, ω, λ are constants and W_1, W_2 are two independent Brownian motions. Provided $\int \sigma(t)dW(t)$ is a continuous martingale, our method allows for jumps in the process $\sigma(t)$. Jumps inserted directly in the differential equation driving the price evolution are not allowed.

Given a time window $[0,1]$ (a day, week, month), we wish to compute the integrated volatility of the process, i.e. $\int_0^1 \sigma^2(t)dt$. An unbiased estimator of this quantity is provided by $[p(1)-p(0)]^2$. However, this estimator is very noisy. In Andersen and Bollerslev (1998c) for exchange rates and in Martens (2000) for a stock index, it is shown that an estimator with smaller noise is provided by the sum of squared intraday returns: $\sum_{i=2}^N [p(\frac{i}{N}) - p(\frac{i-1}{N})]^2$. Using as evaluation criterion the difference $\int_0^1 \sigma^2(t)dt - \sum_{i=2}^N [p(\frac{i}{N}) - p(\frac{i-1}{N})]^2$, the authors find better results on simulated time series with the highest frequency considered, corresponding to 5 minute returns.

The method proposed in Malliavin and Mancino (2000) gives us an estimator of integrated volatility. In fact, integrating $\sigma^2(t)$ between 0 and 2π we get

$$
\int_0^{2\pi} \sigma^2(t)dt = 2\pi a_0(\sigma^2),
\tag{12}
$$

where $a_0(\sigma^2)$ is given by (4). The computation of $a_0(\sigma^2)$ provides an estimate of the integrated volatility using all the observations.

To illustrate the validity of the Fourier approach, we simulate the diffusion process (11) by an Euler discretization scheme, see for example Kloeden and Platen (1992). We use the parameters $\theta = 0.035, \omega = 0.636, \lambda = 0.296$ estimated in Andersen and Bollerslev (1998c) on the daily return time series of the Deutsch Mark-U.S. Dollar exchange rate. Existence of the exact discretization of the process (11) is guaranteed by Drost and Werker (1996) in a weak sense. On this topic, see also the comprehensive results on temporal aggregation in Meddahi and Renault (2000).

Taking a day as a reference measure, we simulate 24 hours of trading with $dt = 1/86400$ when discretizing (11), which corresponds to an observation every second. In order to simulate high frequency unevenly sampled observations, we select a subset of $[1, 86400]$ by extracting the time differences from an exponential distribution with the mean equal to 14 seconds. This choice is motivated by the fact that the empirical distribution of $t_i - t_{i-1}$ can be approximated with an exponential shape and 14 seconds is the mean duration between quotes in the DM-\$ exchange rate time series. As a result, we will have a data set $(t_k, p(t_k), k = 1, \ldots, N)$ with t_k unevenly sampled,

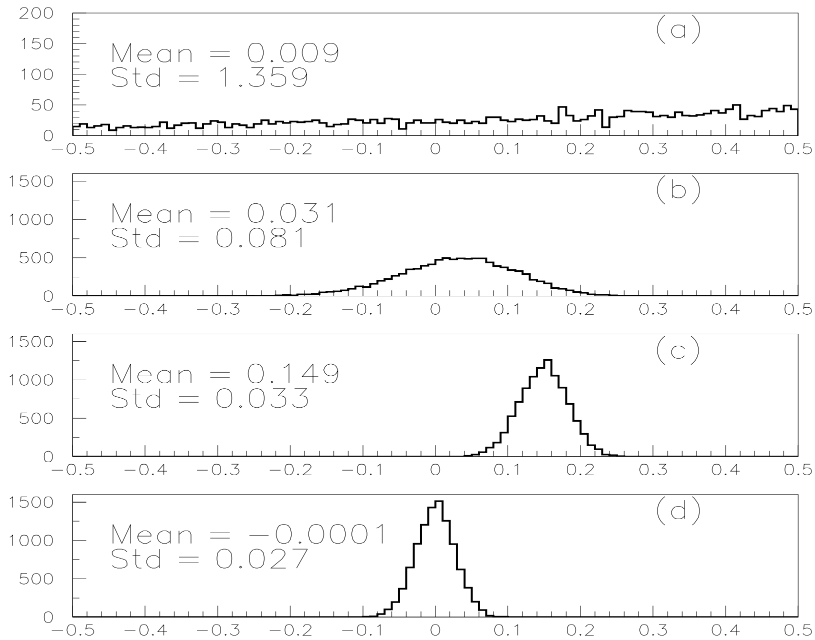

Fig. 1. Distribution of $\frac{\int_0^1 \sigma^2(t)dt - \hat{\sigma}^2}{\int_0^1 \sigma^2(t)dt}$ where $\hat{\sigma}^2$ is obtained with four different estimators of the integrated volatility: (a) $\hat{\sigma}^2 = [p(1) - p(0)]^2$; (b) $\hat{\sigma}^2 = \sum_{i=2}^{288}[p(\frac{i}{288}) - p(\frac{i-1}{288})]^2$ (five minute estimator); (c) $\hat{\sigma}^2 = \sum_{i=2}^{1440}[p(\frac{i}{1440}) - p(\frac{i-1}{1440})]^2$ (one minute estimator); (d) $\hat{\sigma}^2 = 2\pi a_0(\sigma^2)$ (Fourier estimator). In (b) and (c) returns are linearly interpolated. For each distribution, we indicate its mean and its standard deviation (*Std*). The distributions are computed with 10,000 "daily" replications.

and $\sigma(t)$ recorded every second, so that the generated value of the integrated volatility can be computed. Then we compute the integrated volatility according to three estimators: the squared daily return, the cumulative five minute squared intraday returns with linear interpolation of adjacent observations as implemented in Muller et al. (1990), and the Fourier estimator (12). Results are illustrated in Figure 1, where the distribution of the normalized difference between the integrated volatility and its estimate is shown. As expected, the squared daily return is a very noisy estimator. As argued in Andersen and Bollerslev (1998c), when estimating the volatility with the cumulative squared intraday returns, we get a smaller variance. However, measuring volatility according to the Fourier method we have a further reduction of the variance, as well as of the measurement bias.

Increasing the sampling frequency of a financial time series we encounter problems related to microstructure effects. In our setting (simulated time series) such effects are not present; as a consequence, by increasing the return sampling frequency from five minutes to, say, one minute the cumulative squared intraday return estimator should perform better. This is definitely not the case. Figure 1 shows that increasing the frequency from five minutes to one minute the variance of the cumulative squared intraday return estimator is reduced but a downward bias comes in. Considering diffusion processes belonging to the SR-SARV(1) class, in Barucci and Renò (2000) we show that, as the sampling frequency is increased, the cumulative squared intraday return estimator with linearly interpolated returns is more and more downward biased.

This downward bias does not depend on the estimation method. If we implement the Fourier method by computing the integrals in (3) interpolating linearly the prices in the interval $[t_i, t_{i+1}]$, instead of assuming the price to be constant, we get again a downward biased estimator; on the simulation sample of Figure 1, we get a downward bias in the mean of 0.426. Linear interpolation induces spurious autocorrelation; in some sense, a straight line is the "minimum variance" path between two observations.

The downward bias effect of linear interpolation was conjectured in Corsi et al. (2001). In that paper the authors suggest to use an imputation algorithm which coincides with our piecewise constant assumption. As a matter of fact, taking the last observation before time t as $p(t)$ is equivalent to assume $p(t) = p(t_i)$ in the interval $[t_i, t_{i+1}]$. With this imputation scheme, the cumulative squared intraday return estimator is unbiased, as shown in Figure 2 where the cumulative squared intraday returns with the adoption of the imputation scheme is computed for a sampling frequency of two minutes, one minute and 14 seconds. As suggested by the theory, increasing the sampling frequency the variance of this estimator decreases. In the limit, it converges from above to the variance of the Fourier estimator.

The Fourier estimator is characterized by the smallest variance for any sampling frequency, regardless of the adoption of the imputation or of the linear interpolation scheme for the cumulative squared intraday return estimator. In what follows, when volatility is computed according to the cumulative squared intraday returns, we will aggregate the data through linear interpolation, as it is done in large part of literature.

5 Volatility forecasting evaluation

Following Drost and Nijman (1993); Drost and Werker (1996), the GARCH continuous time diffusion (11) can be discretized, obtaining the weak GARCH(1,1) process:

$$\zeta_{t+1}^2 = \psi + \alpha \cdot r_t^2 + \beta \cdot \zeta_t^2, \tag{13}$$

where $r_t = p(t) - p(t-1)$ and ζ_t^2 is the best linear predictor of r_t^2 expressed as a linear combination of lagged squared returns. In Drost and Werker (1996) the exact relation between ψ, α, β and θ, ω, λ is provided, so that one can calibrate (13) on a given time series to obtain the coefficients of the continuous time

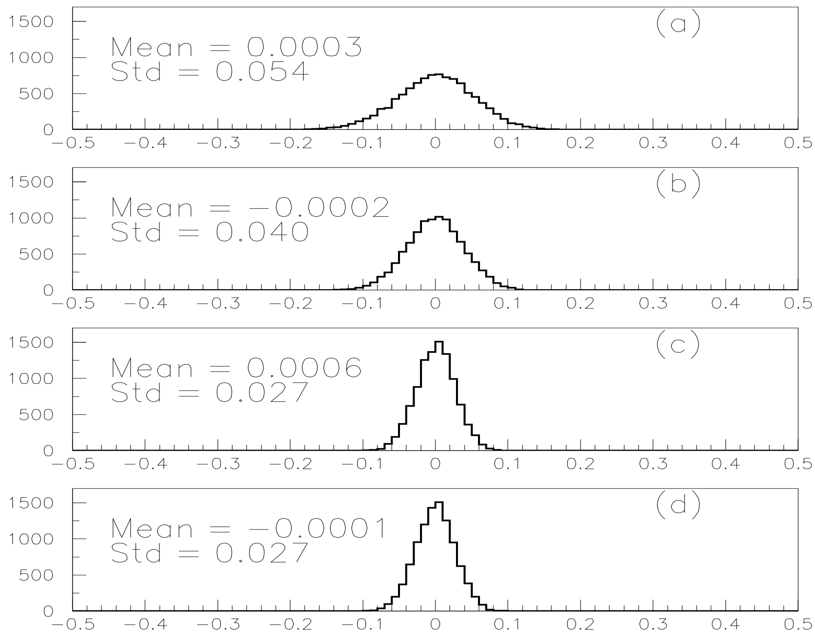

Fig. 2. Distribution of $\frac{\int_0^1 \sigma^2(t)dt - \hat{\sigma}^2}{\int_0^1 \sigma^2(t)dt}$ where $\hat{\sigma}^2$ is obtained with four different estimators of the integrated volatility: (a) 2-minute cumulative squared intraday returns; (b) 1-minute cumulative squared intraday returns; (c) 14-second cumulative squared intraday returns; (d) Fourier estimator. In (a),(b),(c) returns are obtained with an imputation scheme. For each distribution, we indicate its mean and its standard deviation (*Std*). The distributions are computed with 10,000 "daily" replications.

diffusion (11).

In this setting, ζ_{t+1}^2 provides us with an unbiased forecast of $\int_0^1 \sigma^2(t + \tau)d\tau$. While there is a strong support in favor of a high persistence in the volatility dynamics, the one day ahead forecasting performance of the above model has been evaluated to be very poor in the literature. However, as outlined in Andersen and Bollerslev (1998c), the reason for this apparently disappointing result hinges on the fact that the volatility measure which was used to evaluate ex post the forecasting performance was the squared daily return. These authors point out that a careful measure of the realized integrated volatility is needed in order to correctly evaluate the model forecasting performance.

Using simulated time series, we employ the Fourier method to measure the realized integrated volatility, and we evaluate the GARCH(1,1) model forecasting performance according to it. The GARCH model performance with respect to this estimator is compared to that associated with the cumulative squared intraday returns with linearly interpolated observations. Following Andersen and Bollerslev (1998c), the forecasting performance of the GARCH

Table 1
R^2 obtained on simulated data with different estimators of integrated volatility. The values are computed through 50,000 "daily" replications.

Estimator	R^2 (DM-\$)	R^2 (Y-\$)
$(p(1)-p(0))^2$	0.062	0.092
$\sum_{i=2}^{288}(p(\frac{i}{288})-p(\frac{i-1}{288}))^2$	0.476	0.488
Fourier	0.489	0.501
R^2_∞	0.491	0.505

model can be evaluated according to the R^2 of the linear regression:

$$\hat{\sigma}_t^2 = a + b \cdot \zeta_t^2 + \epsilon_t, \tag{14}$$

which is given by

$$R^2 = \left[corr\left(\hat{\sigma}_t^2, \zeta_t^2\right)\right]^2, \tag{15}$$

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

We use the simulation setting described in Section 4. For the Y-\$ time series we set the mean duration equal to 52. The forecasting model is given by (13) with the parameters as in Table 3 for $m = 1$, corresponding to those employed in the above Section according to Drost and Werker (1996).

We point out that the R^2 obtained in (15) must be compared to the R^2 obtained when the ex post integrated volatility measure is perfectly known; its value is given by

$$R^2_\infty = \left[corr\left(\int_{t-1}^{t} \sigma^2(s)ds, \zeta_t^2\right)\right]^2, \tag{16}$$

a value that can be computed in our simulation setting.

Results are shown in Table 1: the Fourier estimator gives an R^2 which is very close to R^2_∞. Employing the Fourier estimator, the GARCH forecasting performance is better than that obtained by measuring integrated volatility through the sum of squared intraday returns.

6 Exchange Rate Time Series analysis

The data set at hand consists of the one year (from October, 1^{st} 1992 to September 30^{th} 1993) collection of tick-by-tick quotes (bid and ask) of the Deutsch Mark-U.S. Dollar exchange rate and of Japanese Yen-U.S. Dollar exchange rate, with time stamps rounded to the nearest even second. The data set was collected and delivered to us by Olsen & Associates. This data set has been extensively studied, e.g. see Andersen and Bollerslev (1997, 1998b,c);

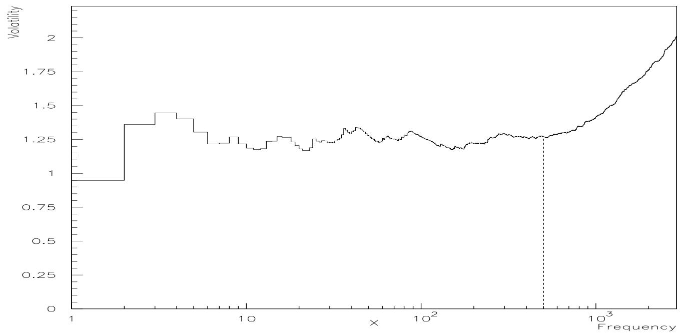

Fig. 3. $\sqrt{2\pi a_0(\sigma^2)}$ computed according to (4) as a function of n for the DM-$ exchange rate (October 1^{st}, 1992). The dashed line indicates the cut we perform to compute integrated volatility.

Andersen, Bollerslev, Diebold and Labys (1999a); Andersen, Bollerslev and Lange (1999); Muller et al. (1990); Zumbach (2000).

We define the price to be the mid-price between bid and ask quotes. The trading day is chosen to begin and to end at 21:00 GMT. We excluded weekends and trading days with less than 1000 quotes for the DM-$ and less than 320 quotes for the Y-$. Moreover we applied the filter described in Dacorogna et al. (1993) which removes roughly 0.36 % of the quotes. We end up with 1·466·944 quotes for the DM-$ and 567·758 quotes for Y-$, distributed over 258 trading days.

When applying the Fourier estimator to a high frequency time series, we compute $a_0(\sigma^2)$ through the expansion (4) stopped at the frequency $N/2$. Here we encounter a severe difficulty: the diffusion model (11) does not hold for small time steps because of microstructure effects such as price discreteness or bid-ask bounce effect, see Madhavan (2000) for a survey. Microstructure effects jeopardize the computation of the Fourier coefficients at high frequencies. This fact is shown in Figure 3 where the square root of the integrated volatility, computed according to (12), is plotted as a function of the highest frequency n employed in the sum (4), with n ranging from $n_0 = 1$ to $N/2$. The plot in Figure 3 can be interpreted as the mean of the power spectrum of the exchange rate from frequency 0 to n. If $dp(t)$ is Normally distributed, as in a model like (11), then the spectrum of p should be flat. Figure 3 shows that the power spectrum of p is not flat [2]; for a frequency larger than a certain value (denoted by N_{cut}) the Fourier coefficients become considerably higher than the lower frequency coefficients. In our setting, it turns out that $N_{cut} \simeq 500$ for the DM-$ exchange rate and $N_{cut} \simeq 160$ for the Y-$. These frequencies correspond roughly to a time step, computed as $\frac{86400}{2 \cdot N_{cut}}$ seconds, of 1.5 and 4.5 minutes respectively. We conclude that the price process cannot be modeled by (11) for time steps smaller than two (five) minutes for the DM-$ (Y-$)

[2] A plot analogous to Figure 3 can be found in Andersen, Bollerslev, Diebold and Labys (1999b).

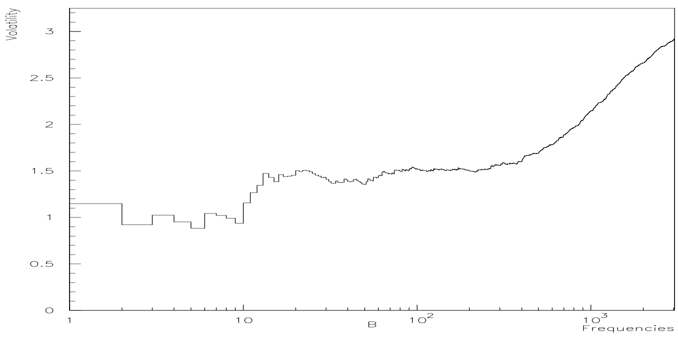

Fig. 4. $\sqrt{2\pi a_0(\sigma^2)}$ for the simulated process (17) as a function of the frequency n in (4).

exchange rate.

This behavior of the exchange rate spectrum can be motivated by the fact that high frequency returns are negatively correlated, a phenomenon that has been documented by Andersen and Bollerslev (1997); Bollerslev and Domowitz (1993) among others. We can show this by simulating the process

$$\begin{cases} \sigma_{t+1}^2 = \psi + \alpha \cdot p_t^2 + \beta \cdot \sigma_t^2 \\ \epsilon_{t+1} = \rho\epsilon_t + \sqrt{1-\rho^2}\eta_t \\ p_{t+1} = \sigma_{t+1}\epsilon_t, \end{cases} \tag{17}$$

where ρ is the first-order serial correlation coefficient and η_t is a sequence of i.i.d. Normal distributed random variables with mean 0 and variance 1. Figure 4 shows the analogous of Figure 3 for a time series simulated according to (17) with $\rho = -0.985$ and a time step of one second: the phenomenon illustrated in Figure 3 occurs. According to our simulations, smaller values of $|\rho|$ would lead to a larger cut frequency. Such a negative correlation can be linked to non-synchronous trading (Lo and MacKinlay, 1990), to the management of inventory positions by market makers (Andersen and Bollerslev, 1997) and to the bid-ask bounce effect (Madhavan, 2000).

In what follows, we cut the highest frequencies in the computation of the integrated volatility, i.e. we will evaluate the Fourier coefficient (4) for $n = \min(N/2, N_{cut})$. In the literature on volatility computation with high frequency data, microstructure effects are attenuated by aggregating the data through linear interpolation or an imputation scheme, building a five minute return time series. With our method we do not interpolate nor aggregate the original time series and we use all the data in order to compute the Fourier coefficients of dp, we only stop expansion (4) properly.

We evaluate the GARCH(1,1) model forecasting performance, by looking at the one-step-ahead daily forecasts, when the integrated volatility is computed according to the Fourier method. The parameters of the model are those estimated in Andersen and Bollerslev (1998c). In Figure 5, we show

Table 2
R^2 for the two foreign exchange rate time series.

Estimator	R^2 (DM-\$)	R^2 (Y-\$)
$\sum_{i=2}^{288}(p(\frac{i}{288}) - p(\frac{i-1}{288}))^2$	0.400	0.128
Fourier	0.470	0.143

the forecasts of the GARCH model together with the integrated volatility of the DM-\$ exchange rate computed according to the Fourier method. Table 2 shows the corresponding R^2. We observe that the GARCH model is evaluated to perform quite well in forecasting when the Fourier method is employed to compute the integrated volatility. Its performance is better than that associated with the cumulative squared intraday returns (with linear interpolation of observations) as an integrated volatility measure. The poor performance on the Y-\$ time series, when compared to that obtained in simulated data in Table 1, can be explained by the presence of few days with very high volatility, see also Andersen and Bollerslev (1998c).

7 Forecasting daily exchange rate volatility using intraday returns

As an illustrative example of the validity of the Fourier approach, we use it to evaluate the GARCH(1,1) model forecasting performance when it is extended to intraday returns, as in Andersen, Bollerslev and Lange (1999); Martens (2001).

Since temporal aggregation of the continuous-time GARCH process (11) holds, we can discretize it at any frequency in a straightforward manner. Denote

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

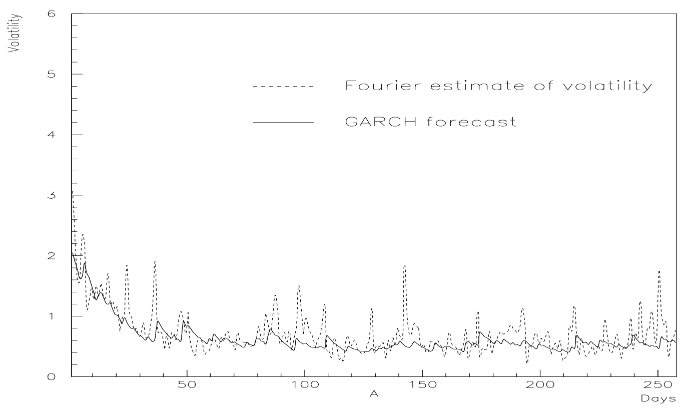

Fig. 5. Comparison between GARCH model forecasts and realized volatilities for the DM-\$ exchange rate, from October, 2^{th} 1992 to September, 30^{th} 1993. Realized volatility is measured with the Fourier estimator.

Table 3
Coefficients of the GARCH(1,1) model at different frequencies, obtained (according to Drost and Werker (1996)) 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 as estimated in Andersen and Bollerslev (1998c), $\psi_m = \sigma_m^2(1 - \alpha_m - \beta_m)$.

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	0.3180	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
4	0.1590	0.0417	0.9496	4	0.1190	0.0651	0.9215
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

then we can write:

$$\zeta_m^2(t) = \psi_m + \alpha_m \cdot r_m^2(t - \tfrac{1}{m}) + \beta_m \cdot \zeta_m^2(t - \tfrac{1}{m}), \tag{18}$$

where $\zeta_m^2(t)$ is the best linear predictor of $r_m^2(t)$ expressed as a linear combination of lagged squared intraday returns. The relation between $(\psi_m, \alpha_m, \beta_m)$ in equation (18) and $(\omega, \theta, \lambda)$ in equation (11) can be obtained for every m in closed form following Drost and Werker (1996); (13) corresponds to (18) with $m = 1$. Table 3 reports the GARCH(1,1) coefficients at different frequencies for the DM-\$ and Y-\$ exchange rate time series.

Following Andersen, Bollerslev and Lange (1999), the forecast for the integrated volatility ($\int_t^{t+h} \sigma^2(s)ds$) using returns spaced by $1/m$ is given by

$$F_{m,h}(t) = mh\sigma_m^2 + \frac{\alpha_m + \beta_m}{1 - \alpha_m - \beta_m}[1 - (\alpha_m + \beta_m)^{m \cdot h}](\zeta_m^2(t) - \sigma_m^2), \tag{19}$$

where $\sigma_m^2 = \psi_m \cdot (1 - \alpha_m - \beta_m)^{-1}$. The realized ex-post integrated volatility will be denoted by $\hat{\sigma}_h^2(t)$. In what follows, we will concentrate only on daily forecast evaluation ($h = 1$).

We choose to evaluate the GARCH(1,1) model forecasting performance with R^2 and the following statistics:

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

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

We analyze the GARCH(1,1) model forecasting performance using intraday returns both on simulated time series and on the DM-\$ and Y-\$ exchange rate time series.

Using the technique described in Section 5, we simulate the time series with the parameters of the DM-\$ exchange rate, setting the mean duration to 14 and $N_{cut} = 500$. The results, presented in Table 4, are in agreement with those in Andersen, Bollerslev and Lange (1999); Martens (2001): increasing the sampling of intraday returns, an improvement of the forecasting performance is observed. We remark that the forecasting performance associated with the Fourier estimator is better than that associated with the 5-minute estimator, with the exception of $RMSE$ which is slightly larger for $m = 1, 2, 3$. However, at frequencies higher than $m = 3$ the GARCH model is evaluated to perform better when volatility is estimated with the Fourier method, than when it is estimated with the cumulative squared intraday returns. These results are confirmed for all the frequencies by adjusting the RMSE for heteroskedasticity. In Table 4 we also report the results when the largest frequencies are included in the computation. As expected, on simulated data this inclusion leads to a performance improvement.

The above statistics are reported in Table 5 for the DM-\$, and in Table 6 for the Y-\$ time series, with $\hat{\sigma}_h^2(t)$ computed as the sum of 5-minute squared intraday returns and with the Fourier method. We stress that we are completely neglecting intraday patterns and macro-economic announcement effects, which have been documented to be important at the intraday level, see Andersen and Bollerslev (1998b); Martens (2001). We also stress that we are neglecting the fact that temporal aggregation of the continuous time GARCH process (11) has not been confirmed empirically, see for example Andersen and Bollerslev (1998a); Zumbach (2000). Our results on the forecasting performance of the GARCH model as a function of the sampling frequency are substantially in agreement with those reported in Andersen, Bollerslev and Lange (1999); Martens (2001). Volatility forecasting improves when the GARCH model is discretized at intraday frequencies, but this effect has an intrinsic limit due to the fact that beyond a certain time scale intraday features and microstructure effects become prominent. As in previous studies, 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 analyzed from this perspective. For the DM-\$ time series the best R^2 and $HRMSE$ are obtained at $m = 4$ (six hour returns) for both estimators, while on the Y-\$ time series the best R^2 is obtained at $m = 48$ (two hours returns) and the best $HRMSE$ is obtained at $m = 4$ (six hours returns) for both estimators.

In general, the GARCH model forecasting performance associated with

Table 4
Summary statistics of the GARCH(1,1) model forecasts for the simulated time series of daily volatility when returns are spaced by $1/m$ days. Between parenthesis we report the values when frequencies larger than N_{cut} are included in the computation. Results are computed with 10,000 "daily" replications.

5 minute returns				Fourier Estimator			
m	R^2	RMSE	HRMSE	m	R^2	RMSE	HRMSE
1	0.374	0.293	0.542	1	0.376 (0.378)	0.296 (0.296)	0.371 (0.357)
2	0.540	0.250	0.461	2	0.544 (0.545)	0.251 (0.251)	0.423 (0.405)
3	0.584	0.238	0.425	3	0.589 (0.590)	0.239 (0.239)	0.390 (0.372)
4	0.640	0.222	0.391	4	0.646 (0.647)	0.222 (0.221)	0.357 (0.341)
6	0.696	0.203	0.356	6	0.704 (0.706)	0.202 (0.202)	0.323 (0.308)
12	0.766	0.179	0.304	12	0.777 (0.780)	0.176 (0.175)	0.272 (0.258)
24	0.815	0.158	0.263	24	0.827 (0.830)	0.154 (0.153)	0.233 (0.221)
48	0.853	0.141	0.228	48	0.866 (0.868)	0.136 (0.135)	0.199 (0.190)
96	0.877	0.128	0.201	96	0.891 (0.894)	0.122 (0.121)	0.173 (0.165)
144	0.893	0.120	0.186	144	0.906 (0.909)	0.113 (0.112)	0.159 (0.153)
288	0.906	0.111	0.165	288	0.921 (0.924)	0.105 (0.105)	0.141 (0.137)

the Fourier estimator is better than that associated with the cumulative squared intraday returns. For the DM-$ time series, the R^2 associated with the Fourier estimator is higher than that of the five minute returns at any m, while the $HRMSE$ is largely lower. The $RMSE$ reports similar results for the two methods, this confirms what has been shown in Table 4 with simulated time series. For the Y-$ exchange rate and $m = 1, 2, 3, 6, 144$, we find that the Fourier estimator performs better than the cumulative squared intraday returns with all the statistics. For $m = 4, 12, 24$ the results are similar; with $m = 96$ the 5-minute estimator performs better. At high frequency the results are not clear cut; this can be due to intraday patterns, low liquidity of the time series and the breakdown of the GARCH temporal aggregation properties.

8 Conclusions

Recently, a large literature has been devoted to compute-forecast volatility for financial time series. In this field, the importance of high frequency data has been stressed, in particular to evaluate the forecasting performance of GARCH models.

In this paper we introduced a new method to compute volatility; the main feature of this method is that it is based upon integration instead of differentiation of the time series, so that it naturally exploits the time struc-

Table 5
Summary statistics of the GARCH(1,1) forecasts for the DM-$ daily volatility when returns are spaced by $1/m$ days.

5 minute returns				Fourier Estimator			
m	R^2	RMSE	HRMSE	m	R^2	RMSE	HRMSE
1	0.400	0.299	0.619	1	0.470	0.292	0.377
2	0.413	0.296	0.545	2	0.491	0.300	0.362
3	0.414	0.307	0.516	3	0.490	0.311	0.362
4	0.445	0.293	0.469	4	0.513	0.308	0.355
6	0.434	0.308	0.522	6	0.502	0.313	0.384
12	0.424	0.311	0.486	12	0.494	0.323	0.389
24	0.422	0.314	0.549	24	0.484	0.326	0.430
48	0.415	0.324	0.565	48	0.473	0.329	0.431
96	0.403	0.331	0.653	96	0.453	0.326	0.473
144	0.406	0.336	0.678	144	0.449	0.320	0.469
288	0.404	0.380	0.798	288	0.436	0.356	0.558

Table 6
Summary statistics of the GARCH(1,1) forecasts for the Y-$ daily volatility when returns are spaced by $1/m$ days.

5 minute returns				Fourier Estimator			
m	R^2	RMSE	HRMSE	m	R^2	RMSE	HRMSE
1	0.128	0.503	0.588	1	0.143	0.493	0.562
2	0.129	0.521	0.535	2	0.138	0.514	0.531
3	0.169	0.536	0.662	3	0.171	0.532	0.632
4	0.237	0.479	0.454	4	0.235	0.478	0.461
6	0.219	0.520	0.559	6	0.221	0.517	0.548
12	0.263	0.479	0.484	12	0.261	0.477	0.477
24	0.275	0.474	0.507	24	0.270	0.473	0.504
48	0.290	0.466	0.513	48	0.283	0.466	0.513
96	0.266	0.478	0.581	96	0.255	0.480	0.595
144	0.233	0.511	0.664	144	0.236	0.507	0.660
288	0.220	0.520	0.765	288	0.219	0.517	0.758

ture of high frequency data by including all the observations in the volatility computation. Using simulated time series, we illustrated that this method performs better than the cumulative squared intraday returns in measuring integrated volatility and that, according to it, the forecasting performance of the GARCH model is improved. We showed that linear interpolation of the time series induces a downward bias in the volatility estimate, and this effect is avoided by assuming the price process to be piecewise constant.

We applied this method to two exchange rate time series. On real data one has to deal with microstructure effects, which become dominant when the return sampling frequency becomes comparable to the frequency of tick-by-tick quotes. We gave a precise estimate of the time step above which these effects can be neglected, and we showed how to remove microstructure distortions: since the Fourier estimator is given by an expansion of the Fourier coefficients, it is enough to cut the highest frequencies in a suitable way. When employing the Fourier method, GARCH forecasts turn out to be more accurate than those associated with the sum of squared intraday returns.

We used the Fourier method to evaluate the forecasting performance of the GARCH(1,1) model when it is discretized at intraday frequencies; the results obtained in the recent literature are confirmed, moreover the forecasting properties of the GARCH model are evaluated to be better if the Fourier estimator is employed, instead of the cumulative squared intraday returns, to measure integrated volatility.

References

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

Andersen, T. and Bollerslev, T., 1998a. Towards a unified framework for high and low frequency return volatility modeling. *Statistica Neerlandica*, 52, 273-302.

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

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

Andersen, T., Bollerslev, T., Diebold, F. and Labys, P., 1999a. The Distribution of Realized Exchange Rate Volatility. *Journal of American Statistical Association*, 96, 42-55.

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

Andersen, T., Bollerslev, T., Diebold, F. and Ebens, H., 2000. The Distribution of Stock Return Volatility. *Journal of Financial Economics*, forthcoming.

Andersen, T., Bollerslev, T., and Lange, S., 1999. Forecasting Financial Market Volatility: Sample Frequency vis-à-vis Forecast Horizon. *Journal of Empirical Finance*, 6, 457-477.

Barndorff-Nielsen, O.E. and Shephard, N., 2000a. Non-Gaussian OU based models and some of their uses in financial economics. CAF Working Paper n.37, Aarhus University.

Barndorff-Nielsen, O.E. and Shephard, N., 2000b. Econometric Analysis of Realised Volatility and it Use in Estimating Levy Based Non-Gaussian OU Type Stochastic Volatility Models. CAF Working Paper n.72, Aarhus University.

Barucci, E., Mancino, M., Renò, R., 2000. Volatility Estimate via Fourier Analysis. In *Finanza Computazionale, Atti della scuola estiva 2000*, Università Ca' Foscari, Venezia.

Barucci, E. and Renò, R., 2000. On Measuring Volatility of Diffusion Processes with High Frequency Data. Manuscript, Università di Pisa and Scuola Normale Superiore, Pisa.

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

Corsi, F., Zumbach, G., Muller, U. and Dacorogna, M., 2001. Consistent High-Precision Volatility from High-Frequency Data. Manuscript, Olsen & Associates.

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.

Drost, F. and Nijman, T., 1993. Temporal aggregation of GARCH processes. *Econometrica*, 61, 909-927.

Drost, F. and Werker, B., 1996. Closing the GARCH Gap: Continuous Time GARCH Models. *Journal of Econometrics*, 74: 31-57.

Figlewski, S., 1997. Forecasting Volatility. *Financial Markets, Institutions and Instruments*, 6: 1-88.

Kloeden, P., Platen, E., 1992. *Numerical Solution of Stochastic Differential Equations*. Springer.

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

Madhavan, A., 2000. Market Microstructure: A Survey. *Journal of Financial Markets*, (3) 3: 205-258.

Malliavin, P. and Mancino, M., 2000. Fourier Series Method for Measurement of Multivariate Volatilities. Forthcoming *Finance & Stochastics*.

Martens, M., 2000. Measuring and forecasting stock market volatility using high-frequency data. Manuscript, University of New South Wales.

Martens, M., 2001. Forecasting daily exchange rate volatility using intraday returns. *Journal of International Money and Finance* (20)1: 1-23

Meddahi, N. and Renault, E., 2000. Temporal Aggregation of Volatility Models. Scientific Series 2000s-22, CIRANO, Montreal, Canada.

Muller, U., Dacorogna, M., Olsen, R., Pictet, O., Schwarz, M. and Morgenegg, C., 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.

Nelson, D., 1990. Arch Models as Diffusion Approximations. *Journal of Econometrics*, 45: 7-38.

Pagan, A. and Schwert G., 1990. Alternative Models for Conditional Stock Volatility. *Journal of Econometrics*, 45: 267-290.

Zumbach, G., 2000. The Pitfalls in Fitting GARCH Processes. In *Advances in Quantitative Asset Management*, ed. Christian L. Dunis, Kluwer Academic Publishers.