Modelling Short-Term Volatility with GARCH and HARCH Models

Michel M. Dacorogna, Ulrich A. Müller,
Richard B. Olsen, Olivier V. Pictet *

MMD.1997-01-08

March 23, 1997

A preprint by the O&A Research Group

This document will be a chapter of the book:
"Nonlinear Modelling of High Frequency Financial
Time Series" edited by Christian Dunis and Bin Zhou
and published by Wiley & Sons, Ltd.

Contents

*Olsen & Associates, Research Institute for Applied Economics, Seefeldstrasse 233, CH-8008 Zürich, Switzerland. e-mail: research@olsen.ch

Abstract

In this paper we present both a new formulation of the HARCH process and a study of the forecasting accuracy of ARCH-type models for predicting short-term volatility.

Using high frequency data, the market volatility is expressed in terms of partial volatilities which are formally exponential moving averages of squared returns measured at different frequencies. This new formulation is shown to produce more accurate fits to the data and, at the same time, to be easier to compute than the earlier version of the HARCH process. This is obtained without losing the nice property of the HARCH process to identify different market components.

In a second part, some performance measures of forecasting accuracy are discussed and the ARCH-type models are shown to be good predictors of the short-term hourly historical volatility with the new formulation of the HARCH process being the best predictor.

1 Introduction

One of the many challenges posed by the study of high frequency data in finance is to build models that can explain the empirical behavior of the data at any frequency from minutes to months at which they are measured. For instance, the well documented clustering of volatility of financial assets. The most popular model among researchers in the field for this behavior is undoubtedly the GARCH (Generalized Autoregressive Conditional Heteroskedasticity) model (Bollerslev et al., 1992). This model was originally developed to study data measured at daily or lower frequencies (Engle, 1982; Bollerslev, 1986). The persistence of volatility has, on the one hand, also been seen with high frequency data, and, on the other hand, the aggregation properties of GARCH models have been theoretically derived by two groups (Nelson, 1990; Drost and Nijman, 1993). Yet, the question remains whether the GARCH models are able to reproduce the heteroskedastic behavior under aggregation. Recent studies of this problem show the failure of simple GARCH models in this respect (Andersen and Bollerslev, 1994; Guillaume et al., 1994; Ghose and Kroner, 1995) even after a correct treatment of the intra-day seasonality of the volatility. The level of volatility clustering is relatively constant under aggregation. In other words, the volatility memory seems quite short-lived when measured with high-frequency data while it seems long-lived when measured with daily or lower frequency data. We attribute this, along with other authors (Andersen and Bollerslev, 1996), to the presence of many independent volatility components in the data. We identify these components to heterogeneous market agents following various investment strategies depending on their institutional constraints, geographical location and risk profile (Müller et al., 1993).

Moreover, in a recent paper (Müller et al., 1997), we have shown that there is asymmetry in the interaction between volatilities measured at different frequencies. A coarsely defined volatility predicts a fine volatility better than the other way around. This effect is not present in a simple GARCH model. All these reasons speak for the development of new and more complex type of ARCH models that would be able to account for the heterogeneity found in high frequency data. We propose to use for this the HARCH (Heterogeneous Autoregression Conditional Heteroskedasticity) model. We presented a first formulation of this model together with its stationarity properties in two papers (Müller et al., 1997; Dacorogna et al., 1996a). Because of the long memory detected in high frequency data (Dacorogna et al., 1993), this initial formulation of HARCH requires numerous sums of returns measured at different frequencies going from 30 minutes to few weeks. This makes the model optimization heavy and requires a lot of computational power when the model is evaluated on high frequency data. To overcome this dilemma, we propose here a reformulation of the model in terms of exponential moving averages, which both simplifies the numerical estimation of the model and preserves the stationarity condition derived for the original form of the process equation. This new formulation also preserves the idea of modeling the impacts of market components by defining "partial" volatilities originating from each component. We compare the new and the old formulation of the process in terms of their optimization results (impacts and likelihoods) and show that they give rise to similar impacts for the same market component.

The real challenge for a model is its ability to forecast the future behavior of the modeled quantity. The difficulty in volatility models is a good definition of the quantity to which the forecast should be compared. We develop in this paper a framework to test the forecasting accuracy of various models. This framework is used to analyze the performance of GARCH and HARCH models in predicting the hourly realized volatility out-of-sample.

In section 2, the new formulation of the HARCH process is presented and discussed. The estimation of the model parameters is explained in section 3 together with the results obtained, for both formulations, over a sample of 10 years of 30m returns. Section 4 deals with the forecasting performance of volatility models both in establishing the framework and presenting the results for various models. The conclusions are drawn in section 5. In a technical appendix, we give some additional results for the estimation of HARCH models on four different foreign-exchange (FX) rates and the respective forecasting performance.

2 A new formulation of the HARCH process equation

The original formulation of the HARCH process has a variance equation based on price changes over intervals of *different sizes*. The returns $r(t)$ of a HARCH(n) process are defined with the help of the random variable $\varepsilon(t)$ which is i. i. d. and follows a distribution function with zero expectation and unit variance (in this paper, we take a normal distribution).

$$\begin{aligned} r(t) &= \sigma(t)\,\varepsilon(t)\ , \\ \sigma^2(t) &= c_0 + \sum_{j=1}^{n} c_j \left(\sum_{i=1}^{j} r(t - i\Delta t) \right)^2\ , \end{aligned} \tag{2.1}$$

where

$$c_0 > 0\ ,\quad c_n > 0\ ,\quad c_j \geq 0 \ \text{ for } j = 1 \ldots n-1 \tag{2.2}$$

and Δt is the grid interval of the original time series. The returns are computed from the logarithmic price x as follows: $r(t) = x(t) - x(t - \Delta t)$ (Guillaume et al., 1997). The equation for the variance $\sigma^2(t)$ is a linear combination of the squares of *aggregated* returns. Aggregated price changes may extend over some long intervals from a time point in the distant past up to time $t - \Delta t$. The heterogeneous set of relevant interval sizes leads to the process name HARCH for "Heterogeneous interval, autoregressive, conditional heteroskedasticity". The first "H" may also stand for the heterogeneous market if we follow that hypothesis as proposed in (Müller et al., 1993). The HARCH process belongs to the wider ARCH family, but differs from all other ARCH-type processes in the unique property of considering the volatilities of price changes measured over different interval sizes. The Quadratic ARCH process (Sentana, 1991) is an exception; although it was not developed for treating different interval sizes, it can be regarded as a generalized form of HARCH.

In (Müller et al., 1997), the coefficients $c_1 \dots c_n$ are not regarded as free parameters of the model. The heterogeneous market approach leads to a low number of free model parameters which determine a much higher number n of dependent coefficients modeling the long memory of volatility.

The new idea is to keep in the equation only a handful of representative interval sizes instead of keeping all of them, and replace the influence of the neighboring interval sizes by an exponential moving average (EMA) of the returns measured on each interval. This has also the advantage of including a memory of the past intervals. Let us now introduce the concept of *partial* volatility σ_j^2, which can be regarded as the contribution of the jth component to the total market volatility σ^2. Here the volatility σ_j^2 is defined as the volatility observed over an interval of size $k_j \Delta t$. We can reformulate the HARCH equation in terms of σ_j as follows:

$$\begin{aligned} r(t) &= \sigma(t)\, \varepsilon(t) \;, \\ \sigma^2(t) &= c_0 + \sum_{j=1}^{n} C_j\, \sigma_j^2(t) \end{aligned} \tag{2.3}$$

where n is now the number of time components in the model (we choose here 7 as in (Müller et al., 1997)). It is also why the coefficients are termed C_j instead of c_j in the old formulation. Unlike the standard HARCH but similar to the generalized HARCH introduced in (Müller et al., 1997), the partial volatility σ_j^2 has a memory of the volatility of *past* intervals of size $k_j \Delta t$. The formal definition of σ_j^2 is

$$\sigma_j^2(t) = \mu_j\, \sigma_j^2(t - \Delta t) + (1 - \mu_j) \left(\sum_{i=1}^{k_j} r(t - i\Delta t) \right)^2 \tag{2.4}$$

where k_j is the aggregation factor of the returns and takes n possible values following the relation

$$k_j = p^{j-2} + 1 \quad \text{for} \quad j > 1 \quad \text{with} \quad k_1 \equiv 1 \;. \tag{2.5}$$

The same value $p = 4$ as in (Müller et al., 1997) is chosen here. Thus, k_j can only take the values $1, 2, 5, 17, 65, 257, 1025, \cdots, 4^{n-2}+1$. Eq.(2.4) is the iterative formula for an exponentially weighted moving average. The volatility memory is defined as a moving average of recent volatility. The depth of the volatility memory is determined by the constant μ_j:

$$\mu_j = \mathrm{e}^{-\frac{\Delta t}{M(k_j \Delta t)}} \tag{2.6}$$

where the memory decay time constant of the component is given as the function M of the component's volatility interval $k_j \Delta t$. Instead of introducing new parameters for the characterization of $M(k_j \Delta t)$, it is simply chosen as

$$M(k_j \Delta t) = \frac{(k_{j+1} - k_j)\; \Delta t}{2} \;. \tag{2.7}$$

The memory is defined by the start and the end point of the component interval k_j. In principle, a more complicated function M can be chosen with independent parameters.

It is easy to prove that a necessary stationarity condition for the new formulation is

$$\sum_{j=1}^{n} k_j C_j \quad < \quad 1 \; . \tag{2.8}$$

The proof relies on the fact that the expectation of the exponential moving average is the same as the expectation of the underlying time series and that the expectation of cross terms is zero. A similar proof as in (Dacorogna et al., 1996a) can be given for the sufficiency of this condition.

We can now define the impact I_j of each component:

$$I_j \quad = \quad k_j C_j \; . \tag{2.9}$$

There is no need for a summation here since each time component is represented by only one coefficient.

An iterative formula needs an initial value for σ_j^2 at the very beginning of the time series. A reasonable assumption of that initial value is the unconditional expectation of $\sigma_j^2(t)$ but the first value is computed here from a data sample prior to the first optimization point. We term this sample the "buildup" sample.

3 Optimization of HARCH – determining market components

We use time series homogeneous in ϑ-time (Dacorogna et al., 1993) to remove the seasonal pattern of intra-day volatility. In this section, the basic time interval is 30 minutes which means only some 7 minutes during the daily volatility peaks around 14:00 GMT, some 80 minutes during the Far Eastern lunch break, and even more during weekends and holidays with their very low volatility. Our optimization sample includes 10 years of data from 1.1.87 to 31.12.96. For getting a reasonable starting value for the iterations of eq.(2.4), some data before the first point in the optimization sample are used.

To achieve parsimony in the old HARCH formulation, we chose only seven market components. The choice was guided by the typical horizons of traders present in the market from intra-day market makers to long-term investors and central banks. We settled on seven components because the optimization did not show any significant improvement of the likelihood when adding an eighth one. For the computation of the new HARCH, we optimize the model with 7 components. This time, the component is built from only one time interval but includes, according to eq.(2.4), a moving average that extends over a certain range which should account for the neighboring time intervals. In fact, we now have two parameters controlling the component definition: the time interval size over which price changes are computed, $k_j \Delta t$, and the range of the moving average, $M(k_j \Delta t)$. We fix both of them and let the optimization find the C_j parameters.

The optimization is done by searching for the maximum of the log-likelihood function. The method we follow to find this maximum is a two-step method – first a genetic algorithm (GA) search (Pictet et al., 1995) and then the use of the Berndt, Hall, Hall and Hausman (BHHH) algorithm (Berndt et al., 1974).

USD-DEM	HARCH			EMA-HARCH		
coefficient	estimate	standard error	t-statistics	estimate	standard error	t-statistics
c_0	1.276×10^{-7}	0.03994×10^{-7}	31.94	0.529×10^{-7}	0.04399×10^{-7}	21.01
I_1	0.1309	0.007151	18.30	0.1476	0.008295	17.80
I_2	0.1930	0.010010	19.28	0.1875	0.012297	15.25
I_3	0.1618	0.009179	17.62	0.1829	0.012545	14.58
I_4	0.0703	0.007363	9.55	0.0507	0.010324	4.91
I_5	0.1003	0.006774	14.81	0.1434	0.010952	13.10
I_6	0.1014	0.006892	14.71	0.1120	0.011835	9.47
I_7	0.0990	0.006118	16.18	0.1145	0.010540	10.86
Log-likelihood	5.794741			5.801367		

Table 1: Comparison between the coefficients and impacts of the two HARCH processes, fitting a half-hourly USD-DEM series which is equally spaced in ϑ-time over 10 years. Instead of the coefficients C_i, the impacts I_i are given. These provide a direct measure of the impacts of the market components on the HARCH variance. The market components are those defined in (Müller et al., 1997) for HARCH and as in eqs.(2.4 and 2.6) for EMA-HARCH. The distribution of the random variable $\varepsilon(t)$ is normal with zero mean and unit variance.

We initialize a first generation of potential solutions for the parameters and store them in "genes" which will form an initial population. The log-likelihood of these solutions are evaluated and constitutes the "fitness" of the genes. Starting from this population, the genetic algorithm construct a new population using its selection and reproduction method (Pictet et al., 1995). The best solutions found by the genetic algorithm are then used as initial solutions for the BHHH algorithm. The BHHH algorithm is a variant of gradient descent which helps the convergence to the local maximum. Once convergence of the BHHH is achieved, the next generation of the GA is computed on the basis of the previous solutions obtained with the BHHH algorithm and the set of solutions of the previous generation. This iterative procedure continues until no improvement of the solution is found. The two-step procedure ensures that the optimization algorithm is not trapped in a local minimum.

The result of the optimization procedure is a set of C_j coefficients from which we can compute the component impact using eq.(2.9). The sum of impacts I_j must be below 1 for stationarity of the process (eq.(2.8)). In Table 1, the coefficients for both the HARCH and EMA-HARCH are shown with their t-statistics for USD-DEM. They are obtained on the exact same dataset. The likelihoods (here log-likelihoods) can be compared since both models have the same number of independent coefficients (the values displayed in Table 1 are per observation). Clearly, the log-likelihood is improved by going to EMA-HARCH. In both cases, all coefficients are highly significant according to the t-statistics and contribute to the variance equation. The stationarity property is fulfilled in both cases. The HARCH has a sum of impacts of of 0.8567 and the EMA-HARCH of 0.9386. The impacts of the different components are remarkably similar. Two small differences are worth noticing: the relative importance of the long-term components is slightly higher for EMA-HARCH (37% instead of 35%) and the minimum for the fourth component is more pronounced in EMA-HARCH. The t-statistics is also consistently smaller for EMA-HARCH than for HARCH but still highly significant in all cases. In the appendix, we present similar tables for four other FX-rates they show the same behavior: improved log-likelihood, slightly stronger long-term components, more pronounced minimum for the fourth component, stationarity condition fulfilled in all cases. The residuals in both formulations still present an excess kurtosis as was noticed in (Müller et al., 1997) for HARCH.

These results show that we have achieved the goal of redesigning the HARCH process in terms of moving averages. We are able to keep and even improve on the properties of the original HARCH and to considerably reduce the computational time to optimize the model. The new formulation of the process equation reduces this time by a factor 1,000, making the problem of computation of HARCH volatility much more tractable even with limited cpu power. In the next section, we will explore the forecasting ability of these models and compare it to a more traditional approach to volatility.

4 Forecasting performance of ARCH-type models

The true test of the veracity of a volatility model is its ability to forecast future movements. Since the seminal work of Meese and Rogoff in 1983 , the forecasting quality of a model of financial data is known to be best measured out-of-sample. This means that the data used to test the model are distinct from the data used to find the model parameters. All the analyses described in this section are performed out-of-sample.

There is some added complexity in the case of volatility models: the definition of the quantity against which the model should be tested. There is no unique definition of volatility. We choose here a path similar to that proposed in (Taylor and Xu, 1997). We construct a time series of realized hourly $v_h(t)$ from our time series of returns as follows,

$$v_h(t) \quad = \quad \sum_{i=1}^{a_h} r^2(t - i\delta t) \tag{4.1}$$

where a_h is the aggregation factor and δt the time interval size. In this case, we use data every $\delta t = 10$minutes in ϑ-time so the aggregation factor is $a_h = 6$. We do not need any factor in front of the summation if we assume Gaussian random walk aggregation properties for the variance.

We produce a forecast using different models that are compared to the realized volatility of eq.(4.1). In order to simply test the one-step ahead forecast, we consider models based on *hourly returns*, $\Delta t = 1h$ in ϑ-time to treat the seasonality. The advantage of using hourly returns instead of 30-minute returns as in the previous section is that hourly forecasts are compatible with the historical hourly volatility defined in eq.(4.1). Four models are studied here.

1. The first model, which is also used as a benchmark, is a naive historical model inspired by the effect described in (Müller et al., 1997): low frequency volatility predicts high frequency volatility. We compute the historical volatility over one day measured from returns computed over 1 hour (lower frequency than the volatility we want to predict). This quantity, properly normalized, is used as a predictor for the next hour volatility, $v_(t + \Delta t)$, as defined in eq.(4.1). Formally the forecasting model v_b is

$$v_b(t) \quad = \quad \frac{1}{24} \sum_{j=1}^{24} \left(\sum_{i=6(j-1)+1}^{6j} r(t - i\delta t) \right)^2 \tag{4.2}$$

where the factor in front of the summation is here to normalize v_b to hourly volatility.

2. GARCH(1,1):

$$v_{garch}(t) \;=\; h(t) \;=\; \alpha_0 \;+\; \alpha_1 \varepsilon^2(t - \Delta t) \;+\; \beta_1 h(t - \Delta t) \tag{4.3}$$

   where $\varepsilon(t)$ is i.i.d. and follows a normal distribution function with zero expectation and unit variance.

3. the old HARCH model following eq.(2.1) and the 7 components proposed in (Müller et al., 1997).

4. the new EMA-HARCH model following eqs.(2.3) and (2.4) with 7 components.

The three parameter-dependent models are optimized over a sample of 5 years of hourly data using the fitting procedure described in section 3. The forecast is then analyzed over the 5 remaining years. We term this procedure the static optimization. To account for possible changes in the model parameters, we also recompute them every year using a moving sample of 5 years. We term this procedure dynamic optimization. In this case, the performance is always tested outside of the gliding sample to ensure that the test is fully out-of-sample. In both studies, we use an out-of-sample period of 5 years of hourly data which represents more than 43,000 independent observations.

We compare the accuracy of the four forecasting models to the realized hourly volatility of eq.(4.1). The quantities of interest are the forecasting signal

$$s_f \;=\; \tilde{v}_f(t) \;-\; v_h(t) \tag{4.4}$$

where $\tilde{v}_f$ is any of the forecasting models, and the real signal

$$s_r \;=\; v_h(t + \Delta t) \;-\; v_h(t)\ . \tag{4.5}$$

The quantity $\tilde{v}_f(t)$ can be used as is or could be rescaled by the ratio of the averages $< v_h >$ and $< \tilde{v}_f >$ taken in the optimization sample. This makes the forecast values on average closer to the historical volatility and does not imply using any future information. In the rest of the paper, we call the quantity $\tilde{v}_f(t) \cdot < v_h > / < \tilde{v}_f >$ the rescaled forecast.

Formulated like this, performance measures proposed in (Dacorogna et al., 1996b) can be applied because the quantities defined in eqs.(4.4) and (4.5) can take positive and negative values contrary to the volatilities which are positive definite quantities. One of these measures is the direction quality:

$$Q_d \;=\; \frac{\mathcal{N}(\{\tilde{v}_f \mid s_f \cdot s_r > 0\})}{\mathcal{N}(\{\tilde{v}_f \mid s_f \cdot s_r \neq 0\})} \tag{4.6}$$

where $\mathcal{N}$ is a function that gives the number of elements of a particular set of variables. It should be noted that this definition does not test the cases where either the forecast is the same as the current volatility or when the volatility at time $t + \Delta t$ is the same as the current one. This occurrence is, of course, unlikely to occur in our particular case. A detailed statistical discussion of this measure can be found in (Pesaran and Timmerman, 1992).

USD-DEM	Q_d	Q_r	Q_f
	Static Optimization		
benchmark	67.7% (67.6%)	54.2% (54.3%)	0.000
GARCH(1,1)	67.8% (67.3%)	58.5% (59.7%)	0.085 (0.072)
HARCH(7c)	69.2% (68.7%)	58.3% (59.2%)	0.134 (0.129)
EMA-HARCH(7c)	69.4% (68.8%)	60.7% (62.5%)	0.140 (0.128)
	Dynamic Optimization		
benchmark	67.7% (67.4%)	54.2% (54.6%)	0.000
GARCH(1,1)	67.0% (66.0%)	59.5% (59.8%)	0.074 (0.057)
HARCH(7c)	67.7% (66.8%)	60.1% (60.8%)	0.113 (0.102)
EMA-HARCH(7c)	68.8% (67.7%)	62.4% (62.9%)	0.133 (0.117)

Table 2: The forecasting accuracy of various models in predicting the short-term market volatility. The performance is measured every hours over 5 years which means 43,230 independent observations. In parentheses, the accuracy of rescaled forecasts is shown.

In addition to this measure, we use a measure that combines the size of the movements and the direction quality. It is often called the *realized potential*

$$Q_r \quad = \quad \frac{\sum \mathrm{sign}(s_f \cdot s_r)\ |s_r|}{\sum |s_r|} \tag{4.7}$$

In fact, the measures Q_r and Q_d are not independent and Q_r is a weighted average of $\mathrm{sign}(s_f \cdot s_r)$ whereas $2Q_d - 1$ is the corresponding unweighted average. It is easy to show that if

$$Q_r \quad > \quad 2\,Q_d \; - \; 1\ , \tag{4.8}$$

the forecast of the sign of s_r for large $|s_r|$ values is better than average.

A more traditional measure is also used: the comparison of the absolute error of a model to a benchmark model. This benchmark model is chosen to be the historical volatility as defined in eq.(4.2), v_b. We compute the following quantity

$$Q_f \quad = \quad 1 \; - \; \frac{\sum |s_r - s_f^{arch}|}{\sum |s_r - s_f^{benchmark}|}\ . \tag{4.9}$$

This particular form is chosen to have a quality measure that increases with increasing performance of the model. If $Q_f > 0$, the model outperforms the benchmark. If $Q_f < 0$, the benchmark outperforms the model.

The summations (including $\mathcal{N}$) in eqs. (4.6), (4.7) and (4.9) are over all hours in the out-of-sample period. The number of independent observations is so large that all the statistical results presented in this study are highly significant. We did not use performance measures based on squares such as the signal correlation or squared errors since we are testing a forecast for essentially squared returns and the fourth moment of the distribution of returns does not converge (Dacorogna et al., 1994).

In Table 2, the results for the different performance measures are presented for the most traded FX-rate, USD-DEM, in the case of static and dynamic optimization. In parentheses, we give also the results for the scaled forecasts. For all measures, the three parameter-dependent models perform better than the benchmark and the new model, EMA-HARCH, performs the best. The forecast accuracy is remarkable for all ARCH-type models. The significance of the values shown on Table 2 is very high since the number of independent observations is 43,230 which means that the 95% significance level for a Gaussian random walk for Q_d is less than 0.5%. The significance levels for the two other measures (Q_r and Q_f) are more difficult to compute but the relative error, in Gaussian approximation, is $1/\sqrt{n}$ which, in our case, is about $5 \cdot 10^{-3}$. One of the advantages of working with high frequency data is to be able to achieve very high statistical significance.

In more than 2/3 of the cases, the forecast direction is correctly predicted and the mean absolute errors are smaller than the benchmark errors in for all the models. The realized potential measure shows that the forecast of volatility change is good not only for small $|s_r|$ but also for large ones. The condition expressed in eq.(4.8) is always satisfied for all models. Neither the scaled forecast nor the dynamic optimization seem to significantly improve the forecasting accuracy; the best results achieved so far are with the plain models. The realized potential Q_r is the only measure that consistently improves with dynamic optimization. Examining the model coefficients computed in moving samples shows that they oscillate around mean values. No structural changes in the coefficients were detected. The accuracy improvement in Q_r together with the loss in Q_f in the case of dynamic optimization shows that the prediction of large movements is improved at the cost of the prediction of direction and of small real movements. From the point of view of forecasting short-term volatility, the EMA-HARCH is the best of the models considered in this study and compares favorably to HARCH. Similar conclusions can be drawn from the results shown in the table of appendix B for four other FX rates. The cross rate JPY-DEM presents results slightly less good than the other currencies but it should be noted that the early half of the sample has been synthetically computed from USD-DEM and USD-JPY. This may lead to noise in the computation of hourly volatility and affect the forecast quality.

5 Conclusion

By introducing partial volatilities, the HARCH formalism can be significantly improved both from the computational point of view and its ability to describe the real market volatility both in-sample (higher maximum likelihood) and out-of-sample (more accurate forecasts). The partial volatility can be interpreted straightforwardly as the contribution of one market component to the market volatility. The optimization results allow us to assess the relative impacts of all components which are very close to those published in (Müller et al., 1997). Formulating σ_j^2 as a function of its past values introduces an element that was missing in the early formulation of HARCH and brings it slightly closer to a GARCH-type of model.

In general ARCH-type models are able to significantly predict the realized hourly short-term volatility out-of-sample with a limited optimization effort. Models including volatility measured at different temporal resolutions (as in HARCH and EMA-HARCH) outperforms those that do not consider this effect. This is further evidence of the market heterogeneity. It also emphasizes the need of high frequency data to properly analyze financial markets. The next research step will be to study how from the EMA-HARCH one can model volatilities measured at low temporal resolution such as daily or even monthly.

With EMA-HARCH, the use of volatility measured at different temporal resolutions becomes relatively cheap to implement as far as the computational time is concerned. It can be a good starting point to extend the formalism for predicting daily or even longer-term volatilities which are needed for option-pricing, risk management, and other portfolio management purposes. Another important use of HARCH models can be the study of market structures and possible changes in the influence of various market components over time.

References

Andersen T. G. and Bollerslev T., 1994, *Intraday seasonality and volatility persistence in foreign exhcnage and equity markets*, Kellogg Graduate School of Management, Northwestern University, **working paper 186**, 1–30.

Andersen T. G. and Bollerslev T., 1996, *Heterogeneous information arrivals and return volatility dynamics: uncovering the long-run in high frequency returns*, Kellogg Graduate School of Management, Northwestern University, **working paper No. 216**, 1–42.

Berndt E., Hall B., Hall R., and Hausman J., 1974, *Estimation and inference in nonlinear structural models*, Annals of Economic and Social Measurement, **3**, 653–665.

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

Bollerslev T., Chou R. Y., and Kroner K. F., 1992, *ARCH modeling in finance*, Journal of Econometrics, **52**, 5–59.

Dacorogna M. M., Embrechts P., Müller U. A., and Samorodnitsky G., 1996a, *How heavy are the tails of a stationary HARCH(k) process?*, to be published in a book in memory of Stamatis Cambanis edited by Yannis Karatzas, 1–31.

Dacorogna M. M., Gauvreau C. L., Müller U. A., Olsen R. B., and Pictet O. V., 1996b, *Changing time scale for short-term forecasting in financial markets*, Journal of Forecasting, **15**(3), 203–227.

Dacorogna M. M., Müller U. A., Nagler R. J., Olsen R. B., and 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**(4), 413–438.

Dacorogna M. M., Pictet O. V., Müller U. A., and de Vries C. G., 1994, *The distribution of extremal foreign exchange rate returns in extremely large data sets*, Internal document UAM.1992-10-22, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.

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

Engle R. F., 1982, *Autoregressive conditional heteroskedasticity with estimates of the variance of U. K. inflation*, Econometrica, **50**, 987–1008.

Ghose D. and Kroner K. F., 1995, *Temporal aggregation of high frequency financial data*, Proceedings of the HFDF-I conference, Zurich, Switzerland, March 29-31, 1995, **2**, 1–31.

Guillaume D. M., Dacorogna M. M., Davé R. D., Müller U. A., Olsen R. B., and Pictet O. V., 1997, *From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets*, To be published in Finance and Stochastics, **2**.

Guillaume D. M., Pictet O. V., and Dacorogna M. M., 1994, *On the intra-day performance of GARCH processes*, Internal document DMG.1994-07-30, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.

Meese R. A. and Rogoff J., 1983, *Empirical exchange rate models of the seventies, do they fit out of sample ?*, Journal of International Economics, **14**, 3–24.

Müller U. A., Dacorogna M. M., Davé R. D., Olsen R. B., Pictet O. V., and von Weizsäcker J. E., 1997, *Volatilities of different time resolutions – analyzing the dynamics of market components*, forthcoming in the Journal of Empirical Finance.

Müller U. A., Dacorogna M. M., Davé R. D., Pictet O. V., Olsen R. B., and Ward J. R., 1993, *Fractals and intrinsic time – a challenge to econometricians*, Invited presentation at the XXXIXth International AEA Conference on Real Time Econometrics, 14-15 Oct 1993 in Luxembourg, and the 4th International PASE Workshop, 22-26 Nov 1993 in Ascona (Switzerland); also in "Erfolgreiche Zinsprognose", ed. by B. Lüthje, Verband öffentlicher Banken, Bonn 1994, ISBN 3-927466-20-4; UAM.1993-08-16, Olsen & Associates, Seefeldstrasse 233, 8008 Zürich, Switzerland.

Nelson D., 1990, *Arch models as diffusion approximations*, Journal of Econometrics, **45**, 7–39.

Pesaran M. H. and Timmerman A., 1992, *A simple nonparametric test of predictive performance*, Journal of Business and Economic Statistics, **10**(4), 461–465.

Pictet O. V., Dacorogna M. M., Chopard B., Oussaidène M., Schirru R., and Tomassini M., 1995, *Using genetic algorithms for robust optimization in financial applications*, Neural Network World, **5**(4), 573–587.

Sentana E., 1991, *Quadratic ARCH models: a potential re-interpretation of ARCH models*, LSE Financial Markets Group Discussion Paper, **122**, 1–45.

Taylor S. J. and Xu X., 1997, *The incremental volatility information in one million foreign exchange quotations*, forthcoming in the Journal of Empirical Finance.

A Tables of comparative optimization results

USD-JPY	HARCH			EMA-HARCH		
coefficient	estimate	standard error	t-statistics	estimate	standard error	t-statistics
c_0	1.291×10^{-7}	0.03706×10^{-7}	34.83	0.640×10^{-7}	0.03962×10^{-7}	16.17
I_1	0.1342	0.006346	21.14	0.1520	0.007589	20.03
I_2	0.1979	0.009112	21.72	0.1829	0.011498	15.90
I_3	0.1815	0.009688	18.73	0.2292	0.013041	17.57
I_4	0.0942	0.008671	10.86	0.0868	0.011622	7.47
I_5	0.1360	0.007412	18.35	0.1642	0.011366	14.45
I_6	0.0932	0.006845	13.62	0.1063	0.011687	9.10
I_7	0.0486	0.004448	10.93	0.0376	0.008709	4.31
Log-likelihood	5.824818			5.833673		

Table 3: Comparison between the coefficients and impacts of the two HARCH processes, fitting a half-hourly USD-JPY series which is equally spaced in ϑ-time over 10 years. Instead of the coefficients C_i, the impacts I_i are given. These provide a direct measure of the impacts of the market components on the HARCH variance. The market components are those defined in (Müller et al., 1997) for HARCH and as in eqs.(2.4 and 2.6) for EMA-HARCH. The distribution of the random variable $\varepsilon(t)$ is normal with zero mean and unit variance.

GBP-USD	HARCH			EMA-HARCH		
coefficient	estimate	standard error	t-statistics	estimate	standard error	t-statistics
c_0	1.489×10^{-7}	0.03430×10^{-7}	43.42	0.876×10^{-7}	0.03198×10^{-7}	27.38
I_1	0.1455	0.007284	19.98	0.1672	0.008498	19.67
I_2	0.1809	0.009423	19.19	0.1621	0.011263	14.39
I_3	0.1469	0.008744	16.80	0.1609	0.011251	14.30
I_4	0.0732	0.007307	10.02	0.0491	0.009415	5.22
I_5	0.1077	0.007099	15.18	0.1420	0.010769	13.18
I_6	0.0655	0.006114	10.71	0.0823	0.011088	7.42
I_7	0.0502	0.004026	12.46	0.0759	0.007694	9.87
Log-likelihood	5.856335			5.864520		

Table 4: Comparison between the coefficients and impacts of the two HARCH processes, fitting a half-hourly GBP-USD series which is equally spaced in ϑ-time over 10 years. Instead of the coefficients C_i, the impacts I_i are given. These provide a direct measure of the impacts of the market components on the HARCH variance. The market components are those defined in (Müller et al., 1997) for HARCH and as in eqs.(2.4 and 2.6) for EMA-HARCH. The distribution of the random variable $\varepsilon(t)$ is normal with zero mean and unit variance.

USD-CHF	HARCH			EMA-HARCH		
coefficient	estimate	standard error	t-statistics	estimate	standard error	t-statistics
c_0	2.156×10^{-7}	0.05376×10^{-7}	40.10	1.000×10^{-7}	0.06646×10^{-7}	15.05
I_1	0.1342	0.006670	20.13	0.1530	0.007736	19.78
I_2	0.1822	0.009068	20.09	0.1659	0.011189	14.83
I_3	0.1436	0.008065	17.81	0.1781	0.011465	15.53
I_4	0.0556	0.006661	8.35	0.0446	0.009663	4.62
I_5	0.0905	0.006312	14.34	0.1293	0.010249	12.61
I_6	0.0968	0.007037	13.75	0.1350	0.011312	11.93
I_7	0.0761	0.005717	13.30	0.0915	0.010395	8.79
Log-likelihood	5.662343			5.668801		

Table 5: Comparison between the coefficients and impacts of the two HARCH processes, fitting a half-hourly USD-CHF series which is equally spaced in ϑ-time over 10 years. Instead of the coefficients C_i, the impacts I_i are given. These provide a direct measure of the impacts of the market components on the HARCH variance. The market components are those defined in (Müller et al., 1997) for HARCH and as in eqs.(2.4 and 2.6) for EMA-HARCH. The distribution of the random variable $\varepsilon(t)$ is normal with zero mean and unit variance.

DEM-JPY	HARCH			EMA-HARCH		
coefficient	estimate	standard error	t-statistics	estimate	standard error	t-statistics
c_0	1.110×10^{-7}	0.02466×10^{-7}	45.01	0.607×10^{-7}	0.02504×10^{-7}	24.23
I_1	0.1554	0.006465	24.04	0.1697	0.007544	22.49
I_2	0.1627	0.008697	18.67	0.1485	0.010434	14.24
I_3	0.1411	0.008057	17.52	0.1767	0.011027	16.03
I_4	0.0856	0.007121	12.02	0.0765	0.010345	7.40
I_5	0.0940	0.006382	14.73	0.1470	0.011048	13.30
I_6	0.0752	0.005814	12.85	0.1013	0.010310	9.82
I_7	0.0781	0.004442	17.60	0.0704	0.007451	9.44
Log-likelihood	5.958331			5.968624		

Table 6: Comparison between the coefficients and impacts of the two HARCH processes, fitting a half-hourly DEM-JPY series which is equally spaced in ϑ-time over 10 years. Instead of the coefficients C_i, the impacts I_i are given. These provide a direct measure of the impacts of the market components on the HARCH variance. The market components are those defined in (Müller et al., 1997) for HARCH and as in eqs.(2.4 and 2.6) for EMA-HARCH. The distribution of the random variable $\varepsilon(t)$ is normal with zero mean and unit variance.

B Tables of forecasting performance

We present here forecasting performance results for 4 other FX rates. They are computed on the same sample as the results in Table 2. Because of slight variations in the ϑ-time, the number of observations can also vary slightly. This number is reported in each table.

USD-JPY	Q_d	Q_r	Q_f
	Static Optimization		
benchmark	67.9% (67.4%)	49.0% (50.5%)	0.000
GARCH(1,1)	68.6% (67.9%)	53.1% (54.3%)	0.065 (0.070)
HARCH(7c)	69.8% (69.5%)	53.8% (54.9%)	0.123 (0.139)
EMA-HARCH(7c)	70.2% (69.7%)	54.5% (55.8%)	0.140 (0.148)
	Dynamic Optimization		
benchmark	67.9% (67.2%)	49.0% (50.9%)	0.000
GARCH(1,1)	68.6% (67.5%)	53.0% (54.8%)	0.070 (0.073)
HARCH(7c)	69.5% (68.4%)	56.1% (56.4%)	0.105 (0.115)
EMA-HARCH(7c)	69.9% (68.6%)	55.5% (57.1%)	0.126 (0.129)

Table 7: The forecasting accuracy of various models in predicting the short-term market volatility. The performance is measured every hours over 5 years which means 43,040 independent observations. In parentheses, the accuracy of rescaled forecasts is shown.

GBP-USD	Q_d	Q_r	Q_f
	Static Optimization		
benchmark	67.7% (67.8%)	52.7% (52.4%)	0.000
GARCH(1,1)	66.8% (66.6%)	58.3% (58.5%)	0.090 (0.069)
HARCH(7c)	68.1% (68.1%)	59.0% (59.1%)	0.120 (0.105)
EMA-HARCH(7c)	69.4% (69.2%)	59.0% (59.4%)	0.134 (0.113)
	Dynamic Optimization		
benchmark	67.7% (67.9%)	52.7% (52.4%)	0.000
GARCH(1,1)	67.4% (67.3%)	58.5% (58.7%)	0.089 (0.071)
HARCH(7c)	67.7% (68.0%)	58.9% (59.0%)	0.108 (0.099)
EMA-HARCH(7c)	68.8% (68.8%)	59.3% (59.5%)	0.125 (0.108)

Table 8: The forecasting accuracy of various models in predicting the short-term market volatility. The performance is measured every hours over 5 years which means 43,215 independent observations. In parentheses, the accuracy of rescaled forecasts is shown.

USD-CHF	Q_d	Q_r	Q_f
	Static Optimization		
benchmark	67.7% (67.6%)	53.7% (54.1%)	0.000
GARCH(1,1)	68.9% (68.5%)	58.9% (59.5%)	0.099 (0.090)
HARCH(7c)	69.2% (68.6%)	55.9% (55.9%)	0.115 (0.112)
EMA-HARCH(7c)	69.3% (68.7%)	57.9% (58.3%)	0.117 (0.107)
	Dynamic Optimization		
benchmark	67.7% (67.6%)	53.7% (54.3%)	0.000
GARCH(1,1)	68.3% (67.5%)	59.1% (59.5%)	0.088 (0.076)
HARCH(7c)	68.4% (67.8%)	57.1% (57.6%)	0.102 (0.097)
EMA-HARCH(7c)	68.9% (68.2%)	58.5% (59.1%)	0.112 (0.103)

Table 9: The forecasting accuracy of various models in predicting the short-term market volatility. The performance is measured every hours over 5 years which means 43,261 independent observations. In parentheses, the accuracy of rescaled forecasts is shown.

JPY-DEM	Q_d	Q_r	Q_f
	Static Optimization		
benchmark	66.2% (64.2%)	58.4% (58.3%)	0.000
GARCH(1,1)	66.2% (63.4%)	60.0% (60.4%)	0.063 (0.061)
HARCH(7c)	65.3% (62.7%)	59.1% (56.8%)	0.072 (0.092)
EMA-HARCH(7c)	65.3% (62.3%)	61.6% (58.6%)	0.072 (0.075)
	Dynamic Optimization		
benchmark	66.2% (65.7%)	58.4% (58.6%)	0.000
GARCH(1,1)	65.4% (64.5%)	59.8% (59.6%)	0.050 (0.045)
HARCH(7c)	64.2% (63.3%)	59.3% (58.8%)	0.044 (0.050)
EMA-HARCH(7c)	65.0% (63.9%)	62.1% (61.3%)	0.074 (0.064)

Table 10: The forecasting accuracy of various models in predicting the short-term market volatility. The performance is measured every hours over 5 years which means 43,292 independent observations. In parentheses, the accuracy of rescaled forecasts is shown.