

A Simple Long Memory Model of Realized Volatility

Fulvio Corsi ¹

November 2002

This version: 18th August 2004

Abstract

In the present work we propose a new realized volatility model to directly model and forecast the time series behavior of volatility. The purpose is to obtain a conditional volatility model based on realized volatility which is able to reproduce the memory persistence observed in the data but, at the same time, remains parsimonious and easy to estimate. Inspired by the Heterogeneous Market Hypothesis and the asymmetric propagation of volatility between long and short time horizons, we propose an additive cascade of different volatility components generated by the actions of different types of market participants. This additive volatility cascade leads to a simple AR-type model in the realized volatility with the feature of considering volatilities realized over different time horizons. We term this model, Heterogeneous Autoregressive model of the Realized Volatility (HAR-RV). In spite of the simplicity of its structure, simulation results seem to confirm that the HAR-RV model successfully achieves the purpose of reproducing the main empirical features of financial data (long memory, fat tail, self-similarity) in a very simple and parsimonious way. Preliminary results on the estimation and forecast of the HAR-RV model on USD/CHF data, show remarkably good out of sample forecasting performance which steadily and substantially outperforms those of standard models.

¹Institute of Finance, University of Lugano, Via Buffi 13, CH-6904 - Lugano, Switzerland.
National Centre of Competence in Research "Financial Valuation and Risk Management" (NCCR - FINRISK)
supported by the Swiss National Science Foundation.
E-mail: fulvio.corsi@lu.unisi.ch

The author would like to gratefully acknowledges Michel Dacorogna, Ulrich Müller, Gilles Zumbach, Paul Lynch, Giovanni Barone-Adesi, Patrick Gagliardini and Lorian Mancini for insightful discussions and the Olsen Group (www.olsen.ch) for providing the data.

1 Introduction

Despite volatility is one of the prevailing features of financial markets, it is still an ambiguous term for which there is no unique, universally accepted definition.

So far most of the studies have considered volatility as an unobservable variable and therefore used a fully specified conditional mean and conditional variance model to estimate and analyze that latent volatility. Modelling the unobserved conditional variance was one of the most prolific topics in the financial literature which led to all ARCH-GARCH developments and stochastic volatility models. In general this kind of models suffer a twofold weakness: first, they are not able to replicate main empirical features of financial data; second, the estimation procedure required are often non trivial (especially in the case of stochastic volatility models).

An alternative approach is to construct an observable proxy for the latent volatility by using intraday high frequency data. This proxy has recently been labelled *Realized Volatility* by Andersen, Bollerslev, Diebold and Labys (2001). In the present work we will employ the high frequency realized volatility estimators developed in Zumbach, Corsi and Trapletti (2002) to directly analyze, model and forecast the time series behavior of FX volatility.

The final purpose is to obtain a conditional volatility model based on realized volatility which is able to account for all the main empirical features observed in the data and, at the same time, which remains very parsimonious and easy to estimate.

Inspired by the Heterogeneous Market Hypothesis (Müller et al. 1993) which led to the HARCH model of Müller et al. (1997) and Dacorogna et al. (1998) and by the asymmetric propagation of volatility between long and short time horizons, we propose an additive cascade model of different volatility components each of which generated by the actions of different types of market participants. This additive volatility cascade leads to a simple AR-type model in the realized volatility with the feature of considering volatilities realized over different time horizons. We thus term this model, Heterogeneous Autoregressive model of Realized Volatility (HAR-RV). Surprisingly, in spite of its simplicity and the fact that it does not formally belong to the class of long memory models, the HAR-RV model is able to reproduce the same memory persistence observed in volatility as well as many of the other main stylized facts of financial data.

The rest of the paper is organized as follows. Section 2 briefly reviews the notion of realized volatility, introducing our notation and discussing the empirical issues related to its practical implementation. Section 3 describes the data set employed in the study and reviews the general stylized facts of FX data. Section 4 describes the motivations and derivation of the HAR-RV model. Section 5 shows the properties of the simulated HAR-RV series while section 6 describes the estimation and forecast results of the model for the twelve years USD/CHF series. Section 7 concludes.

2 Realized Volatility Measures

In this section, we introduce notation for instantaneous and integrated latent volatilities, as well as for realized volatilities aggregated over different horizons. We then briefly discuss the conditions for realized volatility measures to be a consistent and unbiased estimates of the latent volatility.

2.1 Notation

To illustrate the concept of latent *integrated volatility*¹ for day t $\sigma_t^{(d)}$, let's consider the following stochastic volatility process²

$$dp(t) = \mu(t)dt + \sigma(t)dW(t) \quad (1)$$

where $p(t)$ is the logarithm of instantaneous price, $\mu(t)$ is a continuous, finite variation process, $dW(t)$ is a standard Brownian motion, and $\sigma(t)$ is a stochastic process independent of $dW(t)$. For this diffusion process, the integrated volatility associated with day t , is the integral of the instantaneous volatility over the one day interval $(t - 1d; t)$, where a full 24 hours day is represented by the time interval $1d$,

$$\sigma_t^{(d)} = \left(\int_{t-1d}^t \sigma^2(\omega) d\omega \right)^{1/2} \quad (2)$$

Merton (1980) showed that the integrated volatility of a Brownian motion can be approximated to an arbitrary precision using the sum of intraday squared returns. More recently Andersen et al. (2001), applying the quadratic variation theory, generalized this result to the class of special (finite mean) semimartingales. This very general class encompasses processes used in standard arbitrage-free asset pricing applications, such as, Ito diffusions, jump processes, and mixed jump diffusions. In fact, under such conditions, the sum of intraday squared returns converges (as the maximal length of returns go to zero) to the integrated volatility of the prices allowing us, in principle, to construct an error free estimate of the actual volatility over a fixed-length time interval. This nonparametric estimator is called *realized volatility*. The standard definition (for an equally spaced returns series) of the realized volatility over a time interval of one day is

$$RV_t^{(d)} = \sqrt{\sum_{j=0}^{M-1} r_{t-j\Delta}^2} \quad (3)$$

where $\Delta = \frac{1d}{M}$ and $r_{t-j\Delta} = p(t - j\Delta) - p(t - (j + 1)\Delta)$ defines continuously compounded Δ -frequency returns, that is, intraday returns sampled at time interval Δ .

Under those assumptions, the ex-post realized volatility is an unbiased volatility estimator³. Moreover as the sampling frequency from a diffusion (even with non zero mean process) is increased, the realized volatility provides a consistent nonparametric measure of the integrated volatility over the fixed time interval⁴: $\text{plim}_{M \rightarrow \infty} RV_t^{(d)} = \sigma_t^{(d)}$.

Notice that the definition of realized volatility (as any other definition of historical volatility) involves two time parameters: the intraday return interval Δ and the aggregation period $1d$. In the following we will also consider latent integrated volatility and realized volatility viewed over different time horizons longer than one day. These multi-period volatilities will simply be normalized sums of the one-period realized volatilities (i.e. a simple average of the daily quantities). For example, in our notation, a weekly realized volatility at time t will be given by the average

$$RV_t^{(w)} = \frac{1}{5} \left(RV_{t-1d}^{(d)} + RV_{t-2d}^{(d)} + \dots + RV_{t-1w}^{(d)} \right) \quad (4)$$

¹Also called actual or notional volatility.

²We use (t) to denote instantaneous variables and subscripts t to denote discrete quantities.

³Formally the zero mean assumption should be made, but the results remains approximately true for stochastically evolving mean process.

⁴See Andersen, Bollerslev and Diebold (2002).

where $1w = 5d$ indicate a time interval of one week (5 working days). In particular we will make use of weekly and monthly aggregation periods. Indicating the aggregation period as an upper script, the notation for weekly integrated and realized volatility will be respectively $\sigma_t^{(w)}$ and $RV_t^{(w)}$ while a monthly aggregation will be denoted as $\sigma_t^{(m)}$ and $RV_t^{(m)}$. In the following, in respect of their actual frequency, all return and volatility quantities are intended to be annualized.

2.2 Measurement errors

In practice, however, empirical data at very short time scales differ in many ways from the arbitrage-free continuous time price process making this estimator strongly biased for small return interval. Because of market microstructure effects, the assumption that log asset prices evolve as a diffusion process becomes less realistic as the time scale reduces. Thus the volatility computed with very short time intervals is no longer an unbiased and consistent estimator of the daily integrated volatility. Harris (1990) Zhou (1996) and Corsi, Zumbach, Müller and Dacorogna (2001) found that for return intervals less than a few hours, such a definition is affected by a considerable systematic error. For the FX this deviation has a positive sign, i.e. the expectation of daily realized volatility computed with returns at frequencies higher than one hour is systematically larger than the standard deviation of daily returns. Such bias increases with the sampling frequency: at the 1-minute level, it ranges from 30% to about 80% (depending on the liquidity of the currency) while at the tick-by-tick frequency the estimator is two times larger.

Therefore, a trade-off arises: on one hand, statistical considerations would impose a very high number of return observations to reduce the stochastic error of the measurement, on the other hand, market microstructure comes into play, introducing a bias that grows as the sampling frequency increases. Given such a trade-off between measurement error and bias a simple way out is to choose, for each financial instruments, the shortest return interval at which the resulting volatility is still not significantly affected by the bias. This approach has been pursued by Andersen et al. (2001) who settle on a return interval of 30 minutes for the most highly liquid exchange rates leading to only 48 observations per day⁵.

A better solution to this trade-off which permits to fully exploit the information contained in high frequency data, is to have an explicit treatment of the bias. This alternative approach of directly removes the causes of the bias at the tick-by-tick level, has been first followed by Zhou (1996), Corsi et al. (2001), Zumbach et al. (2002) and recently by Curci and Corsi (2003). Given the characteristic of the data and the purpose to keep the approach simple, in this paper we will study a realized volatility measure obtained as an average of three simple high frequency volatility estimators proposed in Zumbach et al. (2002). All the three estimators are computed after having previously applied to the raw series of log prices the simple adaptive filter proposed in Corsi et al. (2001). This filter consists in an adaptive exponential moving average implemented with the inhomogeneous time series operators developed in Zumbach and Müller (2001) which permits a computationally efficient treatment of unevenly spaced data. The filtering is done in a causal way, namely the parameter of the filter is calibrated on the autocorrelation structure of past tick-by-tick returns estimated on a moving window. The reduction of the bias, although not perfect, is quite effective allowing the use of much higher returns frequency in the construction of the realized volatility estimator.

⁵This “unbiased return frequency” mainly depends on the interpolation scheme employed. If, as it seems more appropriate, previous tick would be used, significant bias would still be present at the 30 minutes frequency and it would disappear only at a frequency of 2-3 hours, leaving us with only 8-12 observation per day.

On the type of data employed in this study, each of the three estimators is unbiased and consistent, and gives essentially the same results. An average of them has been taken only to further reduce the stochastic error of the final estimator. The daily integrated volatility can then be estimated exploiting all the information contained in high frequency data, which provide a measure of superior quality (in terms of measurement error) compared to the one obtained with 30 minute returns. However, we will explicitly incorporate the fact that, even with those very high sampling frequency, realized volatility will not be an error-free measure of ex post latent volatility.

3 Empirical properties of the data

3.1 The data

Our data set consists in almost 12 years (from December '89 to July 2001) of tick-by-tick logarithmic middle prices of several FX rates. Log mid prices are computed as averages of the logarithmic bid and ask quotes obtained from the Reuters FXX screen. The whole data set amounts to millions of quotes kindly provided by Olsen&Associates. In the following univariate analysis we will concentrate on the USD/CHF exchange rate (as a proxy for the USD/EUR).

In order to avoid to explicitly model the seasonal behavior of trading activity induced by the weekend we exclude all the realized volatility taking place from Friday 21:00 GMT to Sunday 22:00 GMT. Moreover, a confounding influence comes from low trading days associated to fixed and moving holidays. Since the FX market is a world market, it is not easy to identify on the calendar the relevant holidays which affect such a global market. We then decided to use a more flexible approach by deleting those days presenting a number of ticks smaller than a certain threshold. Highly liquid rates such as USD/CHF have an average daily quotes number on the sample period of approximately 2,800. For this rate we choose a conservative threshold of 200 ticks per day. With this criteria 41 days (partially corresponding to the major US holidays) have been removed leaving us with a final sample of 3,001 full working days. The realized volatility estimates are aggregated at different scales in order to have realized volatility measures of the integrated volatility over different periods: daily, weekly and monthly.

Given that the true volatility is not observable there is no direct evaluation criteria of the quality of the volatility estimators. However, general benchmark criteria can be easily constructed under the hypothesis of an underlying continuous time diffusion process for the logarithm price. In fact if the log-price follows a stochastic volatility diffusion as in (1) with a negligible conditional mean dynamics, the model for daily returns could be written as $r_t^{(d)} = \sigma_t^{(d)} \epsilon_t$ where $\epsilon_t \sim iid N(0, 1)$. Hence the 1-day return is conditionally Gaussian with variance equal to the integrated variance. The normality of ϵ_t is justified by appealing to the Central Limit Theorem for mixing process to argue that the returns over a reasonable aggregation time (such as daily for highly traded assets) should tend towards normality. Then if $RV^{(d)}$ adequately estimates the integrated volatility $\sigma^{(d)}$, the RV-standardized returns should be normally distributed with a variance of unity. Table 1 shows that this is exactly the case for our realized volatility estimator.

3.2 Stylized facts

Our results for the empirical study of USD/CHF are in line with those already found for other rates. Summarizing the main characteristics of the exchange rate data are:

	Mean	Std. Dev	Kurtosis	Jarque-Bera	Probability
Raw returns	0.0005	0.1425	4.7262	382.89	0.0000
RV-std. returns	0.018	1.0191	2.9951	0.2464	0.8840

Table 1: Comparison of daily raw and RV-standardized return distributions

1. Fat tails: the kurtosis of the returns is much higher than that of a normal distribution at intraday frequency and tends to decrease as the return length increases. Thus return pdfs are leptokurtic with shapes depending on the time scale and presenting a very slow convergence of the Central Limit Theorem to the normal distribution. For USD/CHF the kurtosis of hourly returns is 15.58, for daily return is 4.72 and 3.78 at the weekly horizons.
2. Long memory in the volatility: although the autocorrelation of the returns is insignificant at all scales, the autocorrelation of the square and absolute returns shows very strong persistence which lasts for long time interval. This persistence reflects on the (hyperbolic) autocorrelation of realized volatilities where the long memory of the process becomes even more evident. The autocorrelation of USD/CHF realized volatility remain very significant for at least 6 months. This result holds true for realized volatilities aggregated at all frequencies (hourly, daily, weekly and monthly).
3. Distributional properties of realized volatility: the unconditional distributions of realized variances posses high level of skewness and kurtosis which decrease with temporal aggregation but remain far from normal even at monthly scale. Realized volatility and logarithmic realized volatility are instead much closer to normal distributions.
4. Scaling and multiscaling: scaling behavior is readily tested by computing the power spectrum of the logarithmic price. For a scaling process power law behavior of the spectrum is expected. Then an approximated straight line as that in figure (7) should appear in the relative log-log plot. Moreover, empirical data clearly shows strong evidences of multifractality (as we will later discuss).

4 The Model

4.1 Some desired properties of a volatility model

Standard GARCH and SV models are not able to reproduce the features described above. Observed data contains noticeable fluctuations in the size of price changes at all time scales while standard GARCH and SV short memory models appear like white noise once aggregated over longer time periods. For instance, in the GARCH(1,1) models there is a stringent trade-off between the possibility of having sharp changes in the short term volatility (high value of the parameter α) and the ability to capture the long memory behavior of volatility (through high values of β). Moreover, even with high value of $\beta < 1$ GARCH models are subject to exponential decline in the autocorrelation, which is at odds with the observed hyperbolic decline observed in the data. Hence the recent interest in long memory process.

Long memory volatility is usually obtained by employing fractional difference operators like in the FIGARCH models of returns or ARFIMA models on realized volatility. Fractional integration achieves long memory in a parsimonious way by imposing a set of infinite-dimensional restrictions on the infinite variable lags. Those restrictions are transmitted by the fractional difference operators. However fractionally integrated models also pose some problems. Fractional integration is a convenient mathematical trick but completely lacks a clear economic interpretation. The use of the fractional difference operator $(1 - L)^d$ may destroy some useful information on the process and may happen to be not flexible enough to capture the real structure of the data (especially if this structure is dynamically changing in time). Fractionally integrated models are often non trivial to estimate and not easily extendible to multivariate processes. These shortcomings are evident in the FIGARCH case. But also for ARFIMA models it has been shown that the heuristic method of estimating d separately (via a Geweke, Porter-Hudak method, for instance), gives notably biased and inefficient estimates especially in the presence of large AR or MA roots (which seems to be our case). Joint ML estimation of all the parameters in ARFIMA(p, d, q) models, would then be necessary, making the estimation procedure more complex and even more difficult to extend to the multivariate case. Moreover the application of the fractional difference operator requires a very long build up period which results in a loss of many observations. Finally these kind of models are able to reproduce only the unifractal (or monofractal) type of scaling but not the empirical multifractal behavior found in many recent works.

Formally a random process $X(t)$ is said to be *fractal* or self-similar if it satisfies the following scaling rule:

$$E[|X(t)|^q] = ct^{\zeta(q)} \quad (5)$$

where c is a constant, $q > 0$ is the order of the moment and $\zeta(q)$ the scaling function or structure function exponent which is linked to the Hölder exponent⁶ simply by $H(q) = \frac{\zeta(q)}{q}$. For *unifractal* processes $\zeta(q)$ is linear and then fully determined by its unique parameter $H(q) = H$, hence the terminology unifractal or monofractal. *Multifractal* processes, on the contrary are characterized by continuously changing $H(q)$ and this leads to a nonlinear (concave) $\zeta(q)$ function.

In practice the scaling function $\zeta(q)$ is estimated by studying the scaling behavior of the moments of returns computed at different scale, the so called empirical structure (or partition) function $S(\Delta t, q)$. If the empirical financial process is scaling, then:

$$S(\Delta t, q) = \sum_{t=1}^{int[T/\Delta t]} |p(t + \Delta t) - p(t)|^q \sim \Delta t^{\zeta(q)} \quad (6)$$

Estimation of $\zeta(q)$ is then obtained by regressing $S(\Delta t, q)$ on Δt in log-log plots for different values of q .

Structure function analysis can then be seen as a study of "generalized" average volatilities (since only moments of order 1 and 2 are usually employed to define volatility) computed at different scales Δt . For this reason structure function has been already implicitly studied in the financial literature without explicitly referring to the structure function and scaling formalism. In fact, variability of the scaling exponent H for various powers of the returns has been already found to be a pervasive feature of financial data: Ding, Granger and Engle (1993), Lux (1996), Mills (1996), Andersen and Bollerslev (1997), Lobato and Savin (1998).

⁶A generalization of the Hurst exponent.

Although the above authors did not refer to the concepts of multifractality in their papers, their findings identify the existence of multifractal processes in financial data. Their basic message is, therefore, the same as that of recent contributions from physicists (Schmitt, Schertzer and Lovejoy 1999, Vandewalle and Ausloos 1998b, Vandewalle and Ausloos 1998a, Pasquini and Serva 2000).

Formally, any additive process can be shown to have only linear $\zeta(q)$ or constant $H(q)$ ⁷. Hence, theoretically, only multiplicative processes can lead to multifractal behavior. It is in fact often stated, mainly by physicists, that only random multiplicative cascade models, as those encountered in turbulent flows analysis and fragmentation processes, are able to reproduce the long memory and multifractal properties found in the empirical financial data. Does this mean that we should refrain to continue to employ additive models and resign ourself to use multiplicative cascade processes which will be extremely difficult to identify and estimate?

The crucial point is that the long memory and multiscaling features observed in the data could also be only an *apparent behavior* generated from a process which is not really long memory or multiscaling. In fact, if the aggregation level is not large enough compared to the lowest frequency component of the model, truly asymptotic short memory and monoscaling models can be mistaken for long memory and multiscaling ones. In other words, the usual tests employed on the empirical data can indicate the presence of long memory and multiscaling even when none exists, just because the largest aggregation level that we are able to consider is actually not large enough. This means that the set of stochastic processes able to generate the stylized facts found in the data is much larger than commonly thought. In particular LeBaron (2001) shows that a very simple additive model defined as the sum of only three different linear processes (AR(1) processes) each operating on a different time frame can display hyperbolic decaying memory and multiscaling, provided that the longest component has a half life that is long relative to the tested aggregation ranges. The appearance of long memory as a combination of short memory processes is not surprising given the result of Granger (1980) which shows that the sum of an infinite number of short memory processes can give rise to long memory. However, what is surprising is that those results can be obtained with only three different time scales.

As a result, it would be empirically impossible to statistically discern between true multiplicative processes and simple additive models with more than one (but far from infinite) time scales. Since it would be desirable to have a volatility model which, in addition to replicate the main stylized facts, is also simple to estimate and which possibly possesses a clear economic justification and interpretation, it seems reasonable to go in the direction of simple additive volatility models with a small number of components rather than in that of complicated multiplicative systems.

4.2 The basic idea

In the light of the above considerations, we will propose a multi-component volatility model with an additive hierarchical structure which will lead to a very simple additive time series model of the realized volatility.

The basic idea stems from the so called "Heterogeneous Market Hypothesis" presented by Müller et al. 1993, which recognizes the presence of heterogeneity in the traders. This view

⁷For Brownian motion for instance we have $\zeta(q) = \frac{q}{2}$ which implies $H(q) = \frac{1}{2}$. More generally for fractional Brownian motion with an order of fractional integration of d , $\zeta(q) = q(d - \frac{1}{2}) = qH$. It has been shown numerically that ARCH-GARCH process quickly converge to giving $\zeta(q) = \frac{q}{2}$. Even in the case of more exotic Lévy flight (additive process with Lévy noise) and truncated Lévy flight the behavior of $\zeta(q)$ is still linear.

of financial markets can be easily related with the "Fractal Market Hypothesis" of Peters (1994) and the "Interacting Agent View" of Lux and Marchesi (1999). The idea of a presence of multiple components in the volatility process has been also suggested by Andersen and Bollerslev (1997) in their "mixture of distribution" hypothesis. Yet in this latter view the multi-component structure stems from the heterogeneous nature of the information arrivals rather than from the heterogeneity of the agents.

The Heterogeneous Market Hypothesis tries to explain the empirical observation of a strong positive correlation between volatility and market presence. In fact, in a homogeneous market framework where all the participants are identical, the more agents are presents, the faster the price should converge to its real market value on which all agents agreed. Thus, the volatility should be negatively correlated with market presence and activity. On the contrary in an heterogeneous market, different actors are likely to settle for different prices and decide to execute their transactions in different market situations, hence they create volatility.

The heterogeneity of the agents may arise from various reasons: differences in the endowment, degree of information, prior belief, institutional constraints, temporal horizons, geographical location, risk profile and so on. Here we concentrate on the heterogeneity which originates from the difference in the time horizon. Typically a financial market is composed by participants having a large spectrum of dealing frequency. On one side of the dealing spectrum we have dealers, market makers and intraday speculator, with very high intraday frequency, on the other side there are central banks, commercial organization and, for example, pension fund investors with their currency hedging. Each such participant has different reaction times to news, related to his time horizon and characteristic dealing frequency. The basic idea is that agents with different time horizons perceive, react and cause different types of volatility components. Simplifying a bit, we can identify three primary volatility components: the short-term with daily or higher dealing frequency, the medium-term typically made of portfolio manager who rebalance their positions weekly, and the long-term with a characteristic time of one or more months.

Although this categorization finds its justification in the simple observation of financial markets and has a clear and appealing economic interpretation, it has been mainly overlooked in financial modelling. A noteworthy exception is the HARCH model of Müller et al. (1997) and Dacorogna et al. (1998). The HARCH process belongs to the wide ARCH family but differs from all other ARCH-type processes in the unique property of considering squared returns aggregated over different intervals. The equation of the latent variance is then a linear combination of the squared returns aggregated over different time horizons. The heterogeneous set of return interval sizes leads to the name HARCH for "Heterogeneous interval ARCH" (but the first "H" may also stand for "Heterogeneous market"). Because of the long memory of volatility, the HARCH process in its initial formulation requires a large number of returns measured at different frequency, making the log-likelihood optimization very difficult and computationally demanding. To overcome these problems Dacorogna et al. (1998) propose a new formulation of the HARCH process in terms of exponential moving averages (EMA): the EMA-HARCH process. The idea is to keep in the variance equation only a handful of representative interval sizes instead of having all of them, and replace the influence of the neighboring interval sizes by an exponential moving average of the few representative returns. This introduces a sort of GARCH-type elements in the HARCH process. In fact, broadly speaking, the variance equation of the EMA-HARCH process can be seen as a combination of several IGARCH processes defined over square returns aggregated at different frequencies. Each IGARCH component can be regarded as the contribution of the corresponding market component to the the total market volatility and is hence termed

partial volatility .

Studying the interrelations of volatility measured over different time horizons, permits to reveal the dynamics of the different market components. It has been recently observed that volatility over longer time intervals has stronger influence on those over shorter time intervals than conversely. This asymmetric behavior of the volatility has been found with different statistical tools. Müller et al. (1997) employ a lead lag correlation analysis of "fine" and "coarse" volatility to investigate causal relation in the sense of Granger, while Arneodo, Muzy and Sornette (1998) perform a wavelets analysis. More recently Zumbach and Lynch (2001) clearly visualize the asymmetric propagation of volatility by plotting the level of the correlation between the volatility first difference and the realized volatility for a grid of many different frequencies. These correlations measure the response function (in terms of induced volatility) of a given market component to changes of volatilities at various time scales.

The overall pattern that emerges is a volatility cascade from low frequencies to high frequencies. This can be economically explained by noticing that for short-term traders the level of long term volatility matters because it determines the expected future size of trends and risk. Then, on the one hand, short term traders react to changes in long term volatility by revising their trading behavior and so causing short term volatility. On the other hand the level of short-term volatility does not affect the trading strategies of long-term traders. This hierarchical structure has induced some authors to propose formal analogy between FX dynamics and the motion of turbulent fluid where a energy cascade from large to small spatial scales is present. Then borrowing from the Kolmogorov model of hydrodynamic turbulence, multiplicative cascade processes for volatility have been proposed (Ghashghaie et al. 1999, Muzy et al. 2000 and Breymann et al. 2000). Although these types of models are able in theory to reproduce the main features of the financial data, their empirical estimation still remains an open question. Moreover Kolmogorov model refers to the so called homogeneous cascade where the energy is homogeneously dissipated over an infinite number of scales; while in financial markets, only a limited number of scales (corresponding to the predominant components of the market) are the carriers of the financial turbulence (Lynch 2000).

Motivated by previous consideration on the ability of simple additive stochastic models to replicate equally well in practice the empirical behavior of the data, and from the observation that heterogeneous market structure generate an heterogeneous cascade with only few relevant time scales, we propose a stochastic additive cascade model of the volatility with three components.

4.3 The HAR-RV model

Defining the *partial volatility* $\tilde{\sigma}_t^{(\cdot)}$ as the volatility generated by a certain market component, the proposed model can be described as an additive cascade of partial volatilities, each of them having a "sort of AR(1) structure"⁸. We assume a hierarchical process where at each level of the cascade the future partial volatility depends on the past volatility experienced at that time scale (the "AR(1)" component) and on the partial volatility at the next higher level of the cascade i.e. the next longer horizons volatility (the hierarchical component). To simplify, we consider a hierarchical model with only 3 volatility components corresponding to time horizons of one day ($1d$), one week ($1w$) and one month ($1m$) denoted respectively $\tilde{\sigma}_t^{(d)}$, $\tilde{\sigma}_t^{(w)}$ and $\tilde{\sigma}_t^{(m)}$.

⁸Since on the right hand side there won't be the lagged latent volatility itself but the corresponding realized volatility, strictly speaking the process is not a true AR(1), but the fact that the realized volatility is a close proxy for the latent one, makes this process similar to an AR(1).

We assume that the market dynamics is completely determined by the behavior of the dealers. Hence the high frequency return process is determined by the highest frequency volatility component in the cascade (the daily one in this simplified case) with $\tilde{\sigma}_t^{(d)} = \sigma_t^{(d)}$ the daily integrated volatility. Then the return process is

$$r_t = \sigma_t^{(d)} \epsilon_t \quad (7)$$

with $\epsilon_t \sim NID(0, 1)$

The model for the unobserved partial volatility processes $\tilde{\sigma}_t^{(\cdot)}$ at each level of the cascade (or time scale), is assumed to be a function of the past realized volatility experienced at the same time scale and, due to the asymmetric propagation of volatility, of the expectation of the next period values of the longer term partial volatilities. For the longest time scale (monthly) only the "AR(1)" structure remains. Then the model reads:

$$\tilde{\sigma}_{t+1m}^{(m)} = c^{(m)} + \phi^{(m)} RV_t^{(m)} + \tilde{\omega}_{t+1m}^{(m)} \quad (8)$$

$$\tilde{\sigma}_{t+1w}^{(w)} = c^{(w)} + \phi^{(w)} RV_t^{(w)} + \gamma^{(w)} E_t[\tilde{\sigma}_{t+1m}^{(m)}] + \tilde{\omega}_{t+1w}^{(w)} \quad (9)$$

$$\tilde{\sigma}_{t+1d}^{(d)} = c^{(d)} + \phi^{(d)} RV_t^{(d)} + \gamma^{(d)} E_t[\tilde{\sigma}_{t+1w}^{(w)}] + \tilde{\omega}_{t+1d}^{(d)} \quad (10)$$

Where $RV_t^{(d)}$, $RV_t^{(w)}$, and $RV_t^{(m)}$ are respectively the daily, weekly and monthly (ex post) observed Realized Volatilities as previously described, while the volatility innovations $\tilde{\omega}_{t+1m}^{(m)}$, $\tilde{\omega}_{t+1w}^{(w)}$ and $\tilde{\omega}_{t+1d}^{(d)}$ are contemporaneously and serially independent zero mean nuisance variates with appropriately truncated left tail to guarantee the positivity of partial volatilities⁹.

The economic interpretation is that to each volatility component in the cascade corresponds a market component which forms expectation for the next period volatility based on the observation of the current realized volatility and on the expectation for the longer horizon volatility (which is known to affect the future level of their relevant volatility).

By straightforward recursive substitutions of the partial volatilities, such cascade model can be simply written as

$$\sigma_{t+1d}^{(d)} = c + \beta^{(d)} RV_t^{(d)} + \beta^{(w)} RV_t^{(w)} + \beta^{(m)} RV_t^{(m)} + \tilde{\omega}_{t+1d}^{(d)} \quad (11)$$

Equation (11) can be seen as a three factor stochastic volatility model, where the factors are directly the past realized volatilities viewed at different frequency. From this process for the latent volatility it is easy to derive the functional form for a time series model in terms of realized volatilities by simply noticing that, ex-post, $\sigma_{t+1d}^{(d)}$ can be written as

$$\sigma_{t+1d}^{(d)} = RV_{t+1d}^{(d)} + \omega_{t+1d}^{(d)} \quad (12)$$

where $\omega_t^{(d)}$ represent latent daily volatility measurement as well as estimation errors. Equation (12) makes clear that we are not treating realized volatility as an error-free measure of latent volatility. Here the importance of a proper treatment of microstructure effect in the computation of the realized volatility measures (as discussed in section 2.2) becomes apparent. The consistency of the realized volatility (which is directly valid for a broad class of processes) is not enough to state that $\omega_t^{(d)}$ is a mean zero error term. Unbiased estimators

⁹An alternative way to ensure positiveness of the partial volatilities would be to write the model in terms of the log of RV.

of latent volatilities are needed. Equation (12) links our ex post volatility estimate $RV_{t+1d}^{(d)}$ to the contemporaneous measure of daily latent volatility $\sigma_{t+1d}^{(d)}$. Substituting equation (12) in equation (11) and recalling that measurement errors on the dependent variable can be absorbed into the disturbance term of the regression, we obtain a very simple time series representation of the proposed cascade model:

$$RV_{t+1d}^{(d)} = c + \beta^{(d)}RV_t^{(d)} + \beta^{(w)}RV_t^{(w)} + \beta^{(m)}RV_t^{(m)} + \omega_{t+1d} \quad (13)$$

with $\omega_{t+1d} = \tilde{\omega}_{t+1d}^{(d)} - \omega_{t+1d}^{(d)}$.

Equation (13) has a simple autoregressive structure in the realized volatility. In general, denoting l and h respectively the lowest and highest frequency in the cascade, eq.(13) is an $AR(\frac{l}{h})$ model reparametrized in a parsimonious way by imposing economically meaningful restrictions. In other words eq.(13) is an AR-type process but with the feature of considering volatilities realized over different interval sizes; it could than be labeled as an Heterogeneous Autoregressive model for the Realized Volatility (HAR-RV).

5 Simulation results

In spite of its simplicity the proposed model is able to produce rich dynamics for the returns and the volatility which closely resemble the empirical ones. This dynamic is generated by the heterogeneous reaction of the different market components to a given price change which in turns affect the future size of price changes. This causes a complex process by which the market reacts to its own price history with different reaction times. Thus market volatilities feed on themselves¹⁰.

To asses the ability of the model to replicate the main stylized facts of the empirical data, we compare the time series returns and volatilities produced by the simulation with those of twelve years of USD/CHF. In order to give the model the time to unfold its dynamics at daily level, the HAR-RV(3) process is simulated at the frequency of 2 hours ($2h$). The simulated model then reads:

$$r_t^{(2h)} = \sigma_t^{(d)} \epsilon_t \quad (14)$$

$$\sigma_{t+2h}^{(d)} = c + \beta^{(d)}RV_t^{(d)} + \beta^{(w)}RV_t^{(w)} + \beta^{(m)}RV_t^{(m)} + \omega_{t+2h}^{(d)} \quad (15)$$

The parameters of the model ($\beta^{(\cdot)}$) are just hand made calibrated to obtain realistic results.

The analysis begins with a simple visual inspection of the two time series for the returns (figure 1) and the realized volatilities (figure 2). In both figure 1 and 2, the upper panels show the empirical data for USD/CHF from December '89 to July '01, while the lower panels display a sample realization of the simulated process for a similar period. From the visual inspection alone is difficult to discern much difference.

Figure (3) summarizes the character of the simulated and actual return distribution for 1, 5, and 20 day interval. In these and the subsequent comparison figures, the number of observations for the real and simulated data is very different. The twelve years of USD/CHF gives 3001 daily observations, while the HAR-RV(3) process is simulated (at 2 hours frequency) for a period corresponding to approximately 600 years i.e. 150,000 daily observations.

Table 2 reports the values of the kurtosis of those distributions for the three aggregation interval. This table clearly shows how the simple HAR model for the realized volatility is

¹⁰This mechanism is sometimes called "price-driven volatility" in contrast to the "event-driven volatility" consistent with the EMH and the "error-driven volatility" due to over and under reaction of the market to incoming informations.




Figure 1: Comparison of actual (top) and simulated (bottom) daily returns series




Figure 2: Comparison of actual (top) and simulated (bottom) daily RV series




Figure 3: Comparison of actual (dotted) and simulated (solid) PDF of returns for different time horizons. Respectively from left to right: daily, weekly and monthly.

Kurtosis	daily returns	weekly returns	monthly returns
USD/CHF	4.72	3.78	3.04
HAR-RV(3)	4.89	3.90	3.50

Table 2: Comparison of actual and simulated kurtosis of returns for different time horizons

able to reproduce not only the excess of kurtosis of the daily returns, but also the empirical cross-over from fat tail to thin tail distributions as the aggregation interval increases.

But what we are mainly interested in, is the ability of the model to reproduce the volatility persistence of empirical data. Figure (4) shows the actual autocorrelation function of USD/CHF daily realized volatility together with the autocorrelation of HAR daily realized volatility simulated over a period corresponding to 600 years. This figure shows that the purpose of reproducing the long memory of empirical volatility seems to be very well fulfilled. It is important to remark that theoretically the HAR model for volatility is a short memory process which asymptotically should not exhibit hyperbolic decay of the autocorrelation. However, for the aggregation interval considered, the simulated model shows a volatility memory which is at least as long as that of actual data (actually it could be even much longer for different choices of the parameters). Also the partial autocorrelation functions show quite good agreement.

In figure (6) we also compare the distribution of the daily realized volatility, finding reasonable agreement between the real data and the simulated one.

Finally we investigate the scaling behavior of the real and simulated data. In figure (7) the periodogram of the daily returns for the two series is plotted in a log-log plane. Again real data cover a period of twelve years while the simulation is performed for a virtual period of more than 600 years. Both series display high degree of linearity as the one expected for true self-similar process.




Figure 4: Comparison of actual (dotted) and simulated (solid) autocorrelation of daily realized volatility




Figure 5: Comparison of actual (top) and simulated (bottom) partial autocorrelation of daily realized volatility

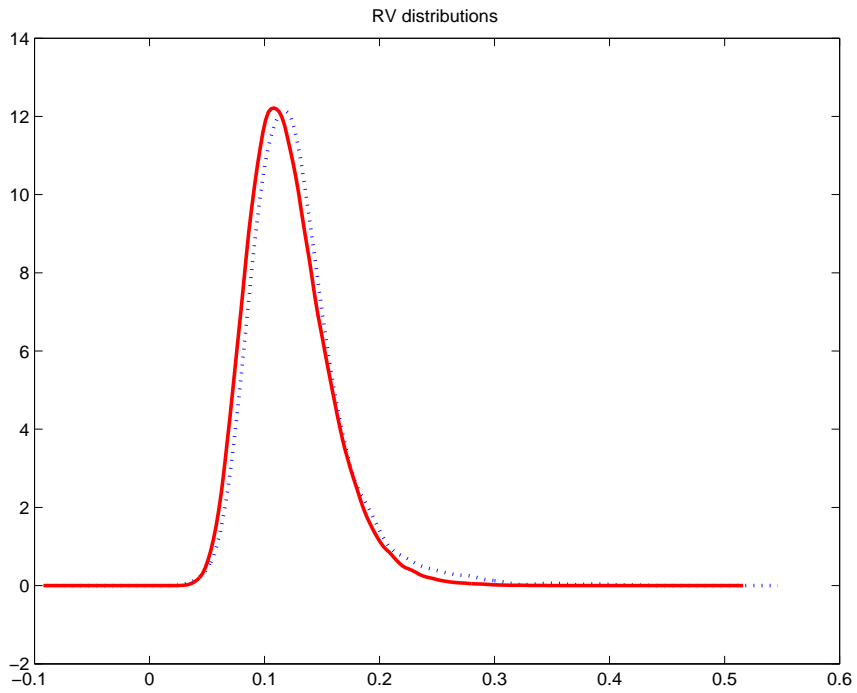


Figure 6: Comparison of actual (dotted) and simulated (solid) distribution of daily realized volatility




Figure 7: Comparison of actual (top) and simulated (bottom) periodogram of daily returns on log-log plane.

6 Estimation and Forecast

Following the recent literature on the realized volatility, we can consider all the terms in (13) as observed and then easily estimate its parameters $\beta^{(\cdot)}$ by applying simple linear regression. Standard OLS regression are consistent and normally distributed, but when multi-step ahead forecast are considered, the presence of regressors which overlap, makes the usual inference no longer appropriate. We then employ the Newey-West covariance correction with a conservative number of lags equal to 20.

Since the uses of intraday measures of realized volatility poses problems either of measurement accuracy and strong intraday seasonalities, we choose to estimate the variance equation (13) at daily frequency¹¹. Table 3 reports the results of the estimation of the HAR-RV model for twelve years of USD/CHF daily realized volatilities.

It is worth noticing that if we are ready to believe that realized volatilities aggregated over different horizons are reasonable proxies for volatilities generated by the corresponding market components, an interesting byproduct of this simple OLS regression is a direct estimate of the market components weights, that is, a readily evaluation of the contribution of each market component to the overall market activity. Moreover, if a moving window regression is performed, a time series evolution of such weights is easily attained as well.

As we have already seen, the HAR-RV process is an autoregressive model reparametrized in a parsimonious way by imposing economically meaningful restrictions. We can then evaluate if those restrictions are valid by comparing the restricted HAR model with the unrestricted AR one. Since the HAR model considered here employs monthly realized volatility (which corresponds to 20 working days) the corresponding unrestricted autoregressive model is an AR(20). A multiple hypothesis test based on the difference between restricted and unrestricted residual sums of squares is then computed. The result of this F-test is 2.48 which is significant. Looking at the information criteria, instead, gives less clear results: on the basis of the AIC, the unrestricted AR(20) model would be slightly preferred, while on the basis of the SIC (which imposes larger penalty for additional coefficients) the HAR-RV is preferred.

The in-sample 1 day ahead forecasts of the model are shown in table 4 and 5. These forecasts are obtained by first estimating the parameters of the models on the full sample and then performing a series of static one-step ahead forecasts. For comparison purposes other models are added: the standard GARCH(1,1) and J.P.Morgan's RiskMetrics, together with an AR(1) and AR(3) model of the realized volatility. Moreover a fractionally integrated model for the realized volatility as employed by Andersen et al. 2002 is considered. They propose a fractional differentiation of the realized volatility series with a fractional coefficient estimated on the full sample with the GPH algorithm (which gives $d = 0.401$) followed by an AR(5) fit. Hence the model is an ARFIMA(5,0.401,0) estimated with a two steps procedure.

In table 4 the forecasting performance are evaluated on the basis of Root Mean Square Errors (RMSE), Mean Absolut Error (MAE), Mean Absolute Percentege Error and Theil's Inequality coefficient. Following the analysis of Andersen and Bollerslev (1998) table 5 reports the results of the Mincer-Zarnowitz regressions of the realized volatility on a constant and the various model forecasts based on time $t - 1$ information. That is

$$RV_t^{(d)} = b_0 + b_1 E_{t-1} \left[(RV_t^{(d)}) \right] + error \quad (16)$$

¹¹A point of caution should be considered here. As for the GARCH, the $\beta^{(\cdot)}$ parameters of the variance equation (13) have a time frequency dimension; i.e. they are defined for a certain time frequency which is the frequency at which the model is estimated. The parameters will then have different values consistent with the time frequency employed and, in general, for the HAR-RV model time aggregation will tend to reduce the impact of shorter realized volatilities and increases that of longer horizons.




Figure 8: Comparison of actual (dotted) and in sample prediction (solid) of the HAR model for daily realized volatilities of USD/CHF exchange data from December '89 to July 2001.

HAR-RV(3) MODEL

Included observations: 3000 after adjusting endpoints
Newey-West Standard Errors and Covariance (lag=20)

Variable	Coefficient	Std. Error	t - Statistic	Probability
C	0.017845	0.002824	6.319572	0.00000
RVD(-1)	0.369542	0.028449	12.98979	0.00000
RVW(-1)	0.265822	0.041865	6.349472	0.00000
RVM(-1)	0.215011	0.037637	5.712704	0.00000

R-squared	0.45592	Mean dependent var	0.12764
Adjusted R-squared	0.45538	S.D. dependent var	0.04081
S.E. of regression	0.03012	Akaike info criter	-4.16567
Sum squared resid	2.71869	Schwarz criterion	-4.15766

Table 3: Estimation results of the least squares regression of HAR-RV(3) model for the USD/CHF exchange data from December '89 to July 2001.

In both tables the difference in forecasting performance between the standard models and the ones based on realized volatility is evident.

But what we are mainly interested in, is to compare the models on the basis of truly out of sample forecasts. Table 6 and 7 reports the results for out of sample forecast of the realized volatility in which the models are daily reestimated on a moving window of 1000 observations¹². An exception is made for the ARFIMA model for which the fractional difference operator requires a long build up period equal to the cut off of its Taylor expansion. We choose the standard cut off limits of 1000 which for a value of d of 0.401 induces a cut off error of 4.2%. After fractional differentiation, the optimal length of the moving window used in the estimation of the AR parameters turns out to be of about 250 days. The forecasting performance are compared over three different time horizons: 1 day, 1 week and 2 weeks. The multi-step ahead forecasts are evaluated considering the aggregated volatility realized and predicted over the multi-period horizon. For a h steps ahead forecast the target function is then $\sum_{j=0}^h RV_{t+j}^{(d)}$ and the Mincer-Zarnowitz regression becomes:

$$\sum_{j=0}^h RV_{t+j}^{(d)} = b_0 + b_1 E_{t-h} \left[\sum_{j=0}^h RV_{t+j}^{(d)} \right] + \text{error} \quad (17)$$

It turns out that, out of sample, the parsimonious HAR(3) model steadily outperforms the others at all the three time horizons considered (1 day, 1 week and 2 weeks). Moreover the HAR(3) model is the only one always presenting the values of 0 and 1 falling in the confidence interval of respectively b_0 and b_1 (the sufficient condition for unbiased forecasts).

It is noteworthy noticing that though the superior performance of the ARFIMA and HAR(3) were already apparent at daily horizon, it becomes striking at weekly and biweekly horizons. The reason is that the other models have a memory which is too short compared to the forecasting horizon (AR(1) and AR(3)) or they adjust too late to the movements of the realized volatility (RiskMetrics). This explanation is confirmed by figure 10 and 11 which compare the dynamic behavior of the forecasts of the different models for one week and two weeks periods ahead. For these time horizons the importance of long memory becomes manifest. What is surprising is the ability of the HAR-RV model to attain these results with only few parameters.

7 Conclusions

The additive volatility cascade inspired by the Heterogeneous Market Hypothesis leads to a simple AR-type model in the realized volatility which has the feature of considering volatilities realized over different interval sizes. We term this model, Heterogeneous Autoregressive model of the Realized Volatility (HAR-RV). The new HAR-RV model seems to successfully achieve the purpose of modelling the long memory behavior of volatility in a very simple and parsimonious way. In spite of the simplicity of its structure and estimation, the HAR-RV model shows remarkably good out of sample forecasting performance. These promising results together with its simple autoregressive structure suggest that a natural way to extend the model to the multivariate case would be to develop a Vector-HAR analogously to the standard VAR model.

¹²Hence these results refer only to the last 2000 observations of the sample and are, therefore, not directly comparable with those in table 4 and 5.

IN SAMPLE PERFORMANCE

	GARCH	RM	AR(1)	AR(3)	ARFIMA	HAR(3)
RMSE x 100	3.6579	3.6789	3.1600	3.0706	3.0787	3.0041
MAE x 100	2.8299	2.7742	2.1978	2.1214	2.0593	2.0577
MAPE %	25.24%	23.54%	17.87%	17.11%	15.80%	16.57%
Theil Inequality coefficient.x100	13.238	13.293	11.958	11.610	11.882	11.355

Table 4: Comparison of the in-sample performances of the 1 day ahead forecast of GARCH, RiskMetrics, AR(1), AR(3), ARFIMA(5,0.401,0) and HAR(3) RV models for 12 years of USD/CHF.

IN SAMPLE MINCER-ZARNOWITZ REGRESSION

	b_0	b_1	R^2
GARCH	-0.027631 (-0.0361, -0.0192)	1.101517 (1.0420, 1.1610)	0.3055
RM	0.032200 (0.0271, 0.0373)	0.688880 (0.6534, 0.7244)	0.3254
AR(1)	-0.000339 (-0.0061, 0.0054)	1.002469 (0.9586, 1.0464)	0.4007
AR(3)	-0.000553 (-0.0059, 0.0048)	1.004037 (0.9630, 1.0451)	0.4341
ARFIMA	0.005210 (0.0002, 0.0102)	1.002926 (0.9632, 1.0427)	0.4496
HAR(3)	-0.000861 (-0.0060, 0.0043)	1.006168 (0.9669, 1.0454)	0.4589

Table 5: In-sample Mincer-Zarnowitz regression for the GARCH, RiskMetrics, AR(1), AR(3), ARFIMA(5,0.401,0) and HAR(3) model for the 1 day ahead realized volatility of USD/CHF (95% confidence interval in parenthesis).




Figure 9: Top: comparison of actual (dotted) and out of sample prediction (solid) of the HAR(3) model for daily realized volatilities. Middle: residuals. Bottom: time evolution of the regression coefficients which according to the model represent market component weights.

OUT OF SAMPLE PERFORMANCE

		RM	AR(1)	AR(3)	ARFIMA	HAR(3)
1 D A Y	RMSE x 100	3.5945	2.9404	2.9088	2.8916	2.8472
	MAE x 100	2.6786	2.0520	2.0061	1.9842	1.9477
	MAPE %	24.01%	17.55%	16.91%	16.90%	16.27%
	Theil Inequality coefficient.x100	13.901	11.741	11.619	11.682	11.384
1 W E E K	RMSE x 100	3.0065	2.7788	2.4372	2.7864	2.2939
	MAE x 100	2.3426	2.1324	1.8089	2.0589	1.6403
	MAPE %	22.08%	19.19%	16.074%	18.05%	14.15%
	Theil Inequality coefficient.x100	11.601	11.124	9.774	11.258	9.205
2 W E E K S	RMSE x 100	2.9734	2.8111	2.4660	2.3743	2.1713
	MAE x 100	2.4254	2.3004	1.9254	1.7728	1.6339
	MAPE %	23.28%	21.46%	17.91%	15.76%	14.47%
	Theil Inequality coefficient.x100	11.421	11.212	9.880	9.593	8.722

Table 6: Comparison of the out of sample performances of the RiskMetrics, AR(1), AR(3), ARFIMA(5,0.401,0) and HAR(3) RV model of 12 years of USD/CHF for 1 day, 1 week and 2 weeks ahead aggregated realized volatility of USD/CHF.

OUT OF SAMPLE MINCER-ZARNOWITZ REGRESSION

		b_0	b_1	R^2
1 D A Y	RM	0.044168 (0.0384, 0.0500)	0.580141 (0.5367, 0.6236)	0.2552
	AR(1)	0.002169 (-0.0051, 0.0095)	0.977008 (0.9179, 1.0361)	0.3764
	AR(3)	0.004260 (-0.0027, 0.0113)	0.961717 (0.9052, 1.0182)	0.3896
	ARFIMA	0.010049 (0.0035, 0.0166)	0.916610 (0.8637, 0.9695)	0.3982
	HAR(3)	0.002030 (-0.0047, 0.0088)	0.982624 (0.9278, 1.0374)	0.4150
1 W E E K	RM	0.048970 (0.0410, 0.0570)	0.536633 (0.4705, 0.5929)	0.1333
	AR(1)	-0.021963 (-0.0448, 0.0009)	1.164388 (0.9786, 1.3502)	0.0801
	AR(3)	-0.055300 (-0.0675, -0.0431)	1.444399 (1.3452, 1.5436)	0.3196
	ARFIMA	0.007047 (-0.0002, 0.0143)	0.938308 (0.8790, 0.9976)	0.3569
	HAR(3)	0.000191 (-0.0073, 0.0077)	0.997471 (0.9361, 1.0588)	0.3692
2 W E E K S	RM	0.072274 (0.0614, 0.0831)	0.370246 (0.2901, 0.4504)	0.0371
	AR(1)	0.148271 (0.1323, 0.1642)	-0.229492 (-0.3562, -0.1028)	0.0063
	AR(3)	-0.027788 (-0.0474, -0.0082)	1.214038 (1.0544, 1.3737)	0.1142
	ARFIMA	0.006737 (-0.0016, 0.0151)	0.939015 (0.8708, 1.0072)	0.2969
	HAR(3)	0.002461 (-0.0060, 0.0109)	0.978118 (0.9089, 1.0474)	0.3079

Table 7: Out of sample Mincer-Zarnowitz regression for the RiskMetrics, AR(1), AR(3), ARFIMA(5,0.401,0) and HAR(3) model for the 1 day, 1 week and 2 weeks ahead aggregated realized volatility of USD/CHF (95% confidence interval in parenthesis).




Figure 10: Comparison of out of sample 1 week aggregated volatility predictions for respectively from top to bottom, GARCH, RiskMetrics, AR(3) and HAR(3) model. The continuous line is the prediction while the dotted line is the ex-post realized volatility over a 1 week period.




Figure 11: Comparison of out of sample 2 week aggregated volatility predictions for, respectively from top to bottom, GARCH, RiskMetrics, AR(3) and HAR(3) model. The continuous line is the prediction while the dotted line is the ex-post realized volatility over a 2 weeks period.

References

- Andersen, T. G. and Bollerslev, T.: 1997, Heterogeneous information arrivals and return volatility dynamics: Uncovering the long run in high frequency data, *Journal of Finance* **52**, 975–1005.
- Andersen, T. G. and Bollerslev, T.: 1998, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* **39**, 885–905.
- Andersen, T. G., Bollerslev, T. and Diebold, F. X.: 2002, Parametric and nonparametric volatility measurement.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P.: 2001, The distribution of exchange rate volatility, *Journal of the American Statistical Association* **96**, 42–55.
- Arneodo, A., Muzy, J. and Sornette, D.: 1998, Casual cascade in stock market from the "infrared" to the "ultraviolet", *European Physical Journal B* **2**, 277.
- Breymann, W., Ghashghaie, S. and Talkner, P.: 2000, A stochastic cascade model for fx dynamics, *International Journal of Theoretical and Applied Finance* **3**, 357–360.
- Corsi, F., Zumbach, G., Müller, U. A. and Dacorogna, M.: 2001, Consistent high-precision volatility from high-frequency data, *Economic Notes* **30**(2), 183–204.
- Curci, G. and Corsi, F.: 2003, A discrete sine transform approach for realized volatility measurement. NCCR FINRISK Working Paper No. 44.
- Dacorogna, M., Müller, U., Dav, R., Olsen, R. and Pictet, O.: 1998, Modelling short-term volatility with garch and harch models, In *"Nonlinear Modelling of High Frequency Financial Time Series"* pp. 161–76. edited by C. Dunis and B. Zhou, John Wiley, Chichester.
- Ding, Z., Granger, C. and Engle, R.: 1993, A long memory property of stock market returns and a new model, *Journal of Empirical Finance* pp. 83–106.
- Ghashghaie, S., Breymann, W., Peinke, J., Talkner, P. and Dodge, Y.: 1996, Turbulent cascades in foreign exchange markets, *Nature* **381**, 767–770.
- Granger, C.: 1980, Long memory relationships and the aggregation of dynamic models, *Journal of Econometrics* **14**, 227–38.
- Harris, L.: 1990, Estimation of stock price variances and serial covariances from discrete observations, *Journal of Financial and Quantitative Analysis* **25**(3), 291–306.
- LeBaron, B.: 2001, Stochastic volatility as a simple generator of financial power-laws and long memory, *Working Paper* **30**.
- Lobato, I. and Savin, N.: 1998, Real and spurious long-memory properties of stock market data, *Journal of Business and Economics Statistics* **16**, 261–283.
- Lux, T.: 1996, Long-term stochastic dependence in financial prices: evidence from the german stock market, *Applied Economics Letters* **3**, 701–6.
- Lux, T. and Marchesi, M.: 1999, Scaling and criticality in a stochastic multi-agent model of financial market, *Nature* **397**, 498–500.

- Lynch, P. E.: 2000, *Adaptive models for speculative turbulence*, PhD thesis, Department of Electrical and Electronic Engineering, University of Manchester Institute of Science and Technology. Manchester, M60. UK.
- Merton, R. C.: 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* **8**, 323–61.
- Mills, T.: 1996, Stylized facts of the temporal and distributional properties of daily ft-se returns, *Applied Financial Economics* **7**, 599–604.
- Müller, U., Dacorogna, M., Dav, R., Olsen, R., Pictet, O. and von Weizsacker, J.: 1997, Volatilities of different time resolutions - analysing the dynamics of market components, *Journal of Empirical Finance* **4**, 213–39.
- Müller, U., Dacorogna, M., Dav, R., Pictet, O., Olsen, R. and Ward, J.: 1993, Fractals and intrinsic time - a challenge to econometricians. XXXIXth International AEA Conference on Real Time Econometrics, 14-15 Oct 1993, Luxembourg.
- Muzy, J., Delour, J. and Bacry, E.: 2000, Modelling fluctuations of financial time series: from cascade process to stochastic volatility model, *Preprint cond-mat/0005400* .
- Pasquini, M. and Serva, M.: 2000, Clustering of volatility as a multiscale phenomenon, *European Physical Journal B* **16**, 195.
- Peters, E.: 1994, *Fractal Market Analysis*, A Wiley Finance Edition. John Wiley & Sons, New York.
- Schmitt, F., Schertzer, D. and Lovejoy, S.: 1999, Multifractal analysis of foreign exchange data, *Applied Stochastic Models and Data Analysis* **15**, 29–53.
- Vandewalle, N. and Ausloos, M.: 1998a, Multi-affine analysis of typical currency exchange rates, *European Physical Journal B* **4**, 257–61.
- Vandewalle, N. and Ausloos, M.: 1998b, Sparseness and roughness of foreign exchange rates, *International Journal of Modern Physics* **9**, 711–9.
- Zhou, B.: 1996, High-frequency data and volatility in foreign-exchange rates, *Journal of Business & Economic Statistics* **14**(1), 45–52.
- Zumbach, G. O., Corsi, F. and Trapletti, A.: 2002, Efficient estimation of volatility using high frequency data, *Working Paper* .
- Zumbach, G. O. and Lynch, P.: 2001, Heterogeneous volatility cascade in financial markets, *Forthcoming in Physica A* .
- Zumbach, G. O. and Müller, U. A.: 2001, Operators on inhomogeneous time series, *International Journal of Theoretical and Applied Finance* **4**(1), 147–178.