A MODEL FOR INTRA-DAILY VOLATILITY WITH MULTIPLE INDICATORS

Robert F. Engle* and Giampiero M. Gallo**

PRELIMINARY and INCOMPLETE: submitted to ESEM 2002

Abstract
The literature on volatility forecasting is being enriched by models of intra-daily volatility. In principle, as the frequency of the data grows larger, the quality of forecasts should improve. Yet, there is no consensus about a "true" measure of volatility. In this paper we propose to use three such indicators of volatility and to analyze the dynamic interactions between them. We compare the outcome which can be obtained with two different model selection procedures and we show the performances of the models in terms of volatility forecasting over a month horizon by resorting to a market-based volatility measure such as VIX. The results show that the variables derived from the multiple indicators offer explanatory power both in and out of sample.

Keywords: volatility modelling, volatility forecasting, GARCH, VIX, high-low range, realized volatility.

JEL Codes: C22, C32, C53.

* Department of Finance, Stern School of Business, New York University and Department of Economics, University of California San Diego rengle@stern.nyu.edu.
** Dipartimento di Statistica. Università di Firenze, gallog@ds.unifi.it

1. Introduction

Volatility models abound. The usefulness of being able to forecast volatility is undisputed as a way to evaluate risk and being able to hedge against it, to price derivative products, and to produce measures of value at risk. Yet, the concept itself of volatility is somewhat elusive as many ways exist to measure it and hence to model it. In recent times, the work done on ultra-high frequency data has shed some light on the concept of volatility and on the dependence of the measure upon the frequency of observation of the data and the different performance of various models for different forecast horizons. In principle, the volatility measures based on tick-by tick data should provide a more accurate observation of volatility, hence resulting in forecast efficiency gains. Nevertheless, the measures that make more intensive of ultra-high frequency data are prone to all sorts of microstructure problems: bid-ask bounces (Roll, 1984), screen fighting (Zhou, 1996), price discreteness, irregular spacing of quotes and transactions can all bias volatility estimates.

Even if the growing literature on realized volatility has delivered promising results (cf. the survey by Andersen *et al.*, 2000), nevertheless what is of interest is the appropriate way to provide accurate forecasts in the medium-to-long run, and the problem remains open as to whether daily models or intra-daily models deliver the most successful answer. In the first category one can mention those models favoring the existence of long memory or high persistence in the process ruling volatility, such as the daily component model by Engle and Lee (1999), or Fractionally Integrated GARCH (FIGARCH – Baillie *et al.*, 1996; or FIEGARCH, Bollerslev and Mikkelsen, 1996) or Long Memory Stochastic Volatility Models (LMSV – Breidt *et al.*, 1998; Deo and Hurvich, 1999). Intra-daily models are more recent: Ghose and Kroner (1996) adopt a signal plus noise model towards the estimation of a persistent component as in Engle and Lee (1999), Andersen and Bollerslev (1998) show how accuracy of volatility forecasts can be improved if one moves to analysing five-minute returns; Engle (2000) derives a measure of volatility from transaction data.

The approach which we will pursue here is one in which several measures of volatility can be jointly used to see whether different features of observed time series can deliver an enrichment of volatility forecasting for the medium run.

2. Definitions and data description

We will use data on the Standard and Poor 500 index from January 4, 1988 to July 14, 1999 (2730 observations, leaving the last 217 for out-of-sample forecast comparison purposes). Let us start by introducing the notation and concepts behind the variables we will use hereafter. We will indicate the daily closing price as C_t, which is used to construct the daily returns as $\tilde{r}_t = \log(C_t / C_{t-1})$. The GARCH literature considers as one measure of volatility at time t, the (square root of the) return variance conditional on the information set. To construct the first measure of volatility, we use the absolute value of the daily return, $|\tilde{r}_t|$, for reasons that will be clearer later.

Other two variables which are observable and widely available (even without access to ultra-high frequency data) are the highest recorded daily price (Daily High) which we will denote by H_t and the lowest recorded daily price (Daily Low), L_t. It has long been recognized that the spread between daily high and daily low is a function of the volatility during the day and as such jointly used with closing or opening prices can lead to an improvement of the volatility estimates (Parkinson, 1980; Garman and Klass, 1980, Beckers, 1983; Gallant *et al.,* 1999). In what follows we will adopt for such a daily range the symbol $\tilde{h}l_t$ defined as the log $\left(\frac{H_t}{L_t}\right)$.

Finally, we will take on a measure of realized volatility, denoted by $d\tilde{v}_t$ which is the standard deviation of observed returns over *J* subperiods within the day between market opening time and market closing time. The merits of the choice of the appropriate J to balance estimation accuracy and microstructure pitfalls are discussed in Andersen *et al.* (2000): empirically, 5 minutes intervals "work" in a foreign exchange framework (Andersen *et al.*, 2001), 15 minutes intervals are adopted by Schwert (1998), but we still lack a theory of what is the optimal length to choose. In this context we choose *J*=78 and we adopted the method discussed by Zhou (1996) on raw data to purge bid-ask bounce effects.

To get the three volatility measures which we will be using we need to rescale variables so that all of them are expressed in terms of annualised percentage volatilities. We will therefore have:

$|r_t| = |\tilde{r}_t| * 100 * \sqrt{252}$ for absolute returns (variable ABSR)

$hl_t = \tilde{h}l_t \sqrt{\frac{\sum r_t^2}{\sum \tilde{h}l_t^2}}$ for daily range High-Low (variable HL), and

$dv_t = \tilde{d}v_t \sqrt{\frac{\sum r_t^2}{\sum \tilde{d}v_t^2}}$ for realized Daily Volatility (variable DV)

where the three variables share the same quadratic mean ($\sum r_t^2$). These are the three indicators used throughout, for which we report some descriptive statistics in Table 1 and their time series profile (on the same scale) in

Figure 1. It is clear from the figure that both the daily range and the realized volatility are persistent, and that the latter exhibits higher spikes than the other two variables. It may be worth noting that the three variables have different features which makes them different from one another: the main difference is that the daily return uses information about the closing price of the previous trading day, while the high-low spread and the realized volatility are measured on the basis of what is observed during the day. Thus, a zero return is not necessarily informative about what happened during the day, and, by the same token, a high return may wrongly signal high volatility during the day while being accompanied by an opening price very close to the closing price of the same trading day with a small high-low spread.

Also, note that the same value of realized volatility may correspond to different values of the range since we could both have recorded high and low fairly apart during the day with a smooth transition of price movements between the two, or some vivacious price swings concentrated in a short period of time.

Table 1. S&P500 Absolute returns, daily range and realized volatility. Summary statistics.

	ABSR	HL	DV
Mean	9.33	11.16	7.02
Median	6.76	9.56	4.25
Maximum	112.93	84.07	202.16
Minimum	0.000	0.000	0.004
Std. Dev.	9.21	6.87	11.07
Skewness	2.99	3.07	8.45
Kurtosis	22.76	23.16	117.63

All variables are expressed in terms of percentage annualised volatilities.

3. Some Modelling Choices

Working with absolute returns instead of returns to estimate a GARCH-type model is not an obstacle. In fact, we can set the intercept in the mean equation to zero and treat absolute returns just like returns. We could then write

$$
\begin{aligned}
| r_t | &= \sqrt{ha_t}\,\varepsilon_t \\
ha_t &= \omega + \alpha r_{t-1}^{\;2} + \beta ha_{t-1}
\end{aligned}
\tag{1}
$$

with ε_t a sequence of i.i.d. (positive valued) random variables. The behaviour of the absolute return is assumed to be time-varying and ruled by a variable (first indicator of volatility) the square of which behaves conditionally as a GARCH(1,1). The standard estimation techniques can be applied in this context, being aware that, as in Engle and Russell (1998), the estimator has the properties of a Quasi-Maximum Likelihood estimator.

This model can be expanded, first by defining the customary dummy variable for negative returns $d_t = I(r_t < 0)$, and writing threshold-type model where the expression is replaced by

$$
ha_t = \omega + \alpha\, | r_{t-1} |^2 + \beta \cdot ha_{t-1} + \gamma\, | r_{t-1} |^2 * d_{t-1}. \tag{2}
$$

More general model specifications can be suggested by inserting other variables in the information set I_{t-1}: for example, lagged returns r_{t-1} are a different way of accounting for asymmetry (as in the APARCH model) given that they maintain their sign; or the other two lagged indicators hl_{t-1}^2 and dv_{t-1}^2 to investigate whether they have explanatory power; or, finally, the product of these indicators by the dummy variable d_{t-1}, to capture some asymmetric impact of the volatility measures when returns in the previous day are negative. Of course, model selection can be performed to check for the significance of these effects.

With this in mind, the model can be written as

$$ha_t = \left(\omega_r + \alpha_r r_{t-1}^2 + \beta_r ha_{t-1}\right) + \delta_r r_{t-1} + \gamma_r r_{t-1}^2 d_{t-1} + \\ + \varphi_r hl_{t-1}^2 + \vartheta_r hl_{t-1}^2 d_{t-1} + \psi_r dv_{t-1}^2 + \lambda_r dv_{t-1}^2 d_{t-1}, \tag{3}$$

where a subscript "*r*" was added to the coefficients indicating that they refer to the specification for absolute returns. The augmentation of the GARCH(1,1) model in this case includes the six variables r_{t-1}, $r_{t-1}^2 d_{t-1}$, hl_{t-1}^2, $hl_{t-1}^2 d_{t-1}$, dv_{t-1}^2, $dv_{t-1}^2 d_{t-1}$ which are known at time *t*-1. A similar approach can be followed for the other two indicators of volatility, namely, hl_t and dv_t. We have

$$hl_t = \sqrt{hh_t}\,\eta_t \tag{4}$$

with i.i.d. η_t's, ruled by a time-varying variable hh_t for the daily range. As before, the general specification for hh_t is the result of an augmentation of a GARCH(1,1) model with the six variables r_{t-1}, $hl_{t-1}^2 d_{t-1}$, r_{t-1}^2, $r_{t-1}^2 d_{t-1}$, dv_{t-1}^2, $dv_{t-1}^2 d_{t-1}$, that is,

$$hh_t = \left(\omega_h + \alpha_h hl_{t-1}^2 + \beta_h hh_{t-1}\right) + \delta_h r_{t-1} + \gamma_h hl_{t-1}^2 d_{t-1} + \\ + \varphi_h r_{t-1}^2 + \vartheta_h r_{t-1}^2 d_{t-1} + \psi_h dv_{t-1}^2 + \lambda_h dv_{t-1}^2 d_{t-1}. \tag{5}$$

Finally, we can assume that for the realized daily volatility we have

$$dv_t = \sqrt{hd_t}\zeta_t , \tag{6}$$

with i.i.d. ζ_t's. Once again, we can adopt a general specification for the hd_t by adding six variables, r_{t-1}, $dv_{t-1}^2 d_{t-1}$, r_{t-1}^2, $r_{t-1}^2 d_{t-1}$, hl_{t-1}^2, $hl_{t-1}^2 d_{t-1}$, to the GARCH(1,1) in order to get

$$\begin{aligned} hd_t = & \left(\omega_d + \alpha_d dv_{t-1}^2 + \beta_d hd_{t-1}\right) + \delta_d r_{t-1} + \gamma_d dv_{t-1}^2 d_{t-1} + \\ & + \varphi_d r_{t-1}^2 + \vartheta_d r_{t-1}^2 d_{t-1} + \psi_d hl_{t-1}^2 + \lambda_d hl_{t-1}^2 d_{t-1}. \end{aligned} \tag{7}$$

For each of these models we envisage the introduction of six variables in addition to the standard specification, each of which may be irrelevant. Therefore, we are estimating 2^6=64 models ranging from the most general forms (3), (5) and (7) to the most specific ones which are the GARCH(1,1) models. The model selection strategies which we will adopt and compare are

1. a general-to-specific strategy whereby we start pruning the coefficients which appear to be statistically insignificant (using Bollerslev and Wooldridge robust standard errors) in the most general expressions and go on to search down to the level where all coefficients are significant;
2. the smallest value of the Schwartz Information Criteria (BIC) among the 64 models.

The chosen models from the general-to-specific selection procedure are reported in Table 2.

Table 2. S&P500 - General-to-specific Model Selection.
Equations for the square of the time-varying component in absolute returns, high-low range, and daily realized volatility.

$$ha_t = \underset{2.123}{3.389} - \underset{1.121}{0.031} \; |r_{t-1}|^2 + \underset{45.537}{0.901} \; ha_{t-1} - \underset{3.552}{0.747} \; r_{t-1} + \underset{2.535}{0.11} \; hl_{t-1}^2 - \underset{2.700}{0.010} \; dv_{t-1}^2$$

$$hh_t = \underset{4.885}{7.622} + \underset{5.407}{0.109} \; hl_{t-1}^{\;2} + \underset{32.713}{0.850} \; hh_{t-1} - \underset{3.608}{0.878} \; r_{t-1}$$

$$hd_t = \underset{4.005}{1.654} - \underset{1.292}{0.021} \; dv_{t-1}^{\;2} + \underset{80.909}{0.728} \; hd_{t-1} - \underset{14.932}{1.596} \; r_{t-1} + \underset{7.278}{0.211} \; dv_{t-1}d_{t-1} +$$
$$+ \underset{9.801}{0.094} \; r_{t-1}^2 + \underset{4.286}{0.098} \; hl_{t-1}^2 - \underset{4.127}{0.114} \; hl_{t-1}^2 d_{t-1}.$$

The results show that, when evaluated in terms of coefficient significance, the inclusion of other variables in the information set appears to add explanatory power in each of the expressions. The model for absolute returns includes an asymmetric effect captured by r_{t-1} (with the appropriate sign; interestingly, not by $r_{t-1}^2 d_{t-1}$), and both the squared daily range and the squared daily volatility. The model for the high-low range seems to be the most parsimonious with the presence of just an asymmetric response of hh_t to lagged values of the returns. Interestingly, the one model which attracts the highest number of significant variables is the model for the daily realized volatility in which there appear to be asymmetric effects from all variables, as well as lagged squared returns and lagged squared daily range. Some diagnostics (reported in Table 3) show that there are no major specification problems: for reference we give the values of the BIC and the estimated log-likelihood, as well as the results of an ARCH(2) test (5% critical value =

5.99) and the Ljung-Box test Q(12) for the squared residuals (5% critical value = 21.03).

Table 3: S&P500 - Diagnostics on the General-to-Specific Models

	BIC	ARCH(2)	Q(12)	LOGLIK
ha	7.8693	1.770	4.016	-9868.28
hh	7.8622	2.25	5.86	-9867.127
hd	7.2581	0.425	15.248	-9092.116

The specification search guided by the lowest value of the BIC gives the results we present in Table 4.

Table 4. S&P500 - Smallest BIC Model Selection.
Equations for the square of the time-varying component in absolute returns, high-low range, and daily realized volatility.

$$ha_t = \underset{2.805}{5.026} - \underset{1.068}{0.030} r_{t-1}^2 + \underset{43.432}{0.901}\ ha_{t-1} - \underset{3.293}{0.745}\ r_{t-1} + \underset{2.328}{0.101}\ hl_{t-1}^2$$

$$hh_t = \underset{4.885}{7.622} + \underset{5.407}{0.109}\ hl_{t-1}^2 + \underset{32.713}{0.850}\ hh_{t-1} - \underset{3.608}{0.878}\ r_{t-1}$$

$$hd_t = \underset{8.061}{2.123} + \underset{2.366}{0.035}\ dv_{t-1}^{\ 2} + \underset{91.479}{0.736}\ hd_{t-1} - \underset{28.350}{1.183}\ r_{t-1} + \underset{6.688}{0.122}\ dv_{t-1}d_{t-1} + + \underset{23.911}{0.123}\ r_{t-1}^2$$

The model selected for the daily range is the same as before. For absolute returns, the model selected here does not contain lagged square daily volatility; the values of the coefficients do not change much, so the profile of forecasts between the two models will differ mainly because of the impact of this variable. The model for the daily volatility here does not include the terms involving the daily range and the values of the coefficients are fairly different. Also here no major problems arise for the residuals, at least judging from the diagnostics reported in Table 5.

Table 5: S&P500 - Diagnostics on the Lowest BIC Models

	BIC	ARCH(2)	Q(12)	LOGLIK
ha	7.8682	1.272	4.174	-9870.81
hh	7.8622	2.25	5.86	-9867.127
hd	7.2561	0.470	12.136	-9097.375

4. Model Performance

4.1 The Model Used for Forecasting

The three models thus estimated could be used separately for one-step ahead predictions using the estimated coefficients and the actual value of the right-hand side variables: the performance of these single-equation models could be individually evaluated in relationship to the performance of the corresponding GARCH models. The interest of what is being done here, though, lies in the fact that the three equations together can be seen as a system which can be used as a tool for multi-step forecasting for medium horizons. Let us consider the left-hand side as a three dimensional vector h_t, and consider that at time T+1 we have

$$h_{T+1|T} = \begin{pmatrix} ha_{T+1|T} \\ hh_{T+1|T} \\ hd_{T+1|T} \end{pmatrix} = \begin{pmatrix} \omega_a \\ \omega_h \\ \omega_d \end{pmatrix} + \mathrm{A}^* \left(r_T, r_T^2, hl_T^2, hl_T^2 d_T, dv_T^2, dv_T^2 d_T, ha_T, hh_T, hd_T \right)' \tag{8}$$

where A^* is a 3 by 9 matrix which includes the coefficients on the variables the value of which is known at time T. To forecast the various future second-order moments conditional on information at time T for maturities greater than 1, we need to substitute the right-hand side variables with their conditional expectation as of time T. We will then have for a generic horizon k,

$$E_T(r_{T+k-1}) = 0,$$
$$E_T(r_{T+k-1}^2) = ha_{T+k-1|T},$$
$$E_T(hl_{T+k-1}^2) = hh_{T+k-1|T}$$
$$E_T(hl_{T+k-1}^2 d_{T+k-1}) = \frac{1}{2} hh_{T+k-1|T}$$
$$E_T(dv_{T+k-1}^2) = hd_{T+k-1|T}$$
$$E_T(dv_{T+k-1}^2 d_{T+k-1}) = \frac{1}{2} hd_{T+k-1|T}$$

and, therefore the expression (8) is substituted by

$$h_{T+k|T} = \begin{pmatrix} ha_{T+k|T} \\ hh_{T+k|T} \\ hd_{T+k|T} \end{pmatrix} = \begin{pmatrix} \omega_a \\ \omega_h \\ \omega_d \end{pmatrix} + \mathrm{A} \begin{pmatrix} ha_{T+k-1|T} \\ hh_{T+k-1|T} \\ hd_{T+k-1|T} \end{pmatrix} = \omega + A h_{T+k-1|T} \tag{9}$$

The dynamic properties of the estimated system can therefore be evaluated by examining the characteristic roots of the matrix A. In principle there could be stable complex conjugate roots which would give the system some dampened cyclicality in the forecasting. For the case at hand the roots for the two sets of estimates are given as in Table 6.

Table 6. Characteristic Roots of the Matrix A

Model selection			
General to specific	0.958	0.860	0.833
Smallest BIC	0.958	0.870	0.832

The roots are fairly similar so that the two multi-equation models will provide different forecasting profiles for short horizon, while they will tend to be very similar in the medium to long run. For this reason we will report the graphs just for the models selected according to the smallest BIC criterion.

4.2 Out-of-sample Forecasts

To evaluate the performance of the models let us first consider the profile of the out-of-sample forecasts generated by the system starting on Jan. 2, 1998, the first day after the estimation sample, and varying the starting date. In Figure 2 we report various profiles of the forecasts in intervals of 5 days, for a horizon of 132 periods ahead, taking the observed values as starting values for day T (corresponding to Dec. 30, 1997 and forecasting from Jan 2, 1998 to July 13, 1998), T+4 (corresponding to Jan 8, 1998), and so on until the last starting value which corresponds to Mar. 6, 1998 for a forecasting horizon going from Mar.9, 1998 to Sep. 14, 1998). The forecasts from the daily range indicator exhibit the typical monotonic profile of a GARCH(1,1) model. This is not the case for the curves corresponding to forecasts obtained from the other two indicators, indicating the possibility that the forecasts may overshoot or undershoot their long term (unconditional) value.

4.3 Term Structure of Volatility

Let us start by defining in this context the term structure of volatility for an asset with a given maturity *T+ k* as the square root of the cumulated sum of *j*-step ahead forecasts generated by expression (8), when *j*=1, or (9), when *j* is between 1 and k (cf. Engle and Patton, 2000). This corresponds to evaluating the square root of the cumulated sum of the expected values of each volatility indicator at any time between 1 and *k*. As *k* increases the terms of the sum will tend to repeat themselves We have

$$\begin{pmatrix} v^{a}_{T+k|T} \\ v^{h}_{T+k|T} \\ v^{d}_{T+k|T} \end{pmatrix} = \begin{pmatrix} \sqrt{\sum_{j=1}^{k} ha_{T+j|T}} \\ \sqrt{\sum_{j=1}^{k} hh_{T+j|T}} \\ \sqrt{\sum_{j=1}^{k} hd_{T+j|T}} \end{pmatrix} = \begin{pmatrix} \sqrt{\sum_{j=1}^{k} E_T r^2_{T+j}} \\ \sqrt{\sum_{j=1}^{k} E_T hl^2_{T+j}} \\ \sqrt{\sum_{j=1}^{k} E_T dv^2_{T+j}} \end{pmatrix} \qquad (10)$$

For reasons that will be clearer in the next section when we discuss the comparison of these forecasts with the VIX volatility index, we choose a horizon k equal to 22, that is one-month ahead forecasts. The in-sample calculations are shown in Figure 3, where we report the value of the cumulative volatility forecast as the (square root of the) average of 1-step, 2-steps, …, 22-steps ahead forecasts obtained by the general-to-specific three-equation system (8) and (9) and with a standard GARCH(1,1) model. As one would expect, the differences between the values obtained with the system and with our model are quite minimal and refer to the presence of lagged returns in the specification for the hh equation in the one-step ahead forecast from expression (8). For the other two indicators the result show that the estimates obtained with our model are generally higher and more persistent than the estimates obtained with a GARCH model when the absolute return indicator is used whereas the term structure of the volatility obtained from the daily realized volatility equation in our model are systematically lower than the GARCH counterpart.

A similar pattern arises when we examine the out-of-sample comparison between the system-based and the GARCH(1,1) based term structure of volatility as in Figure 4. We can detect quite a different response of the two models started at the same (observed) initial values and solved for multi-period forecasts. A remarkable difference is observed for estimates at or around August 31, 1998 for which the GARCH(1,1) model provides much higher estimates. Whether this signals excessive volatility forecasts by the latter set of models is an issue that requires a type of more specific evaluation. What we need is an overall comparison of the two sets of forecasts gauged in reference to a market-based volatility measure such as the volatility index VIX. We turn to this in the next section.

5. A Volatility Index as a Forecast Target

The Chicago Board Options Exchange (CBOE) is the world largest options exchange where standardized stock and index options are traded.[1] Among these options, CBOE offers an index (American-style) option on the S&P 100 index called OEX, the value of which is established as 100 times the current value of the index (e.g. if the index is 613.95 as on July 25, 2001, the dollar value of the index is $61,395). OEX options are the most actively traded on the CBOE and are very liquid. One interesting feature of OEX is that in 1993 the CBOE has created a volatility index called VIX, which is calculated as a

[1] For more edtailed information cf. the CBOE web site www.cboe.com.

weighted average of the implied volatilities from eight calls and puts on the OEX which have an average time to maturity of 30 days, and represent in-the, at-the, and out-of-the money options. The VIX index therefore measures the implied volatility of a hypothetical option that is at-the-money and has 30 days to expiration.[2] The behavior of the series from Jan. 4, 1988 to Dec. 14, 1998 is shown in Figure 5. Standard Augmented Dickey Fuller tests show that the unit root hypothesis is rejected, although the degree of persistence in the series is very high.

For the purposes of this paper, the analysis of VIX in reference to volatility estimates with a multiple indicator conditional model is relevant because we can use the value of VIX as a target for our volatility estimates. The 22 period ahead horizon for the term structure of volatilities was chosen to ensure compatibility with the horizon considered when VIX is constructed.[3]

We will therefore use a combination of forecast approach to examine two main questions:

- whether volatility forecasts provided by a GARCH-type model help in explaining the behavior of VIX both in and out of sample;
- whether the volatility forecasts provided by the multiple indicator model are complements or substitutes relative to GARCH(1,1)-based forecasts.

In order to answer these questions we ran two sets of regressions each over two separate periods (estimation sample and out-of-sample). In this context, we use both the estimates provided by the general to specific model

---

[2] The value of the index is reckoned to measure investors' fears. A very high value signals bearishness as downward risk is perceived as higher than upward. Chartists look at VIX for possible trend reversals.

[3] We are not overly concerned with the fact that VIX Is calculated on the S&P 100 index while we are calculating volatility measures on the S&P500. The two indices have a high correlation (in the sample period of interest here about 98% for the returns and 99% for the returns squared, with high cross-correlations). One should keep in mind that the S&P500 is an average over a larger number of stocks and hence has a potentially lower volatility.

selection procedure, and the ones obtained with the smallest BIC criterion since they provide partially different outcomes. Two models were estimated, both using the log of VIX as the dependent variable and the log of the term structure of volatility estimates according to the multiple indicator model and the standard GARCH-type model, as independent variables. The difference between the two is that the second includes an AR(1) term for the error term. The in-sample analysis is shown in Table 7 (the first panel refers to the general-to-specific models, the second panel to the smallest BIC models, in addition to GARCH(1,1) models). The results are encouraging, since almost all coefficients are statistically significant. One notices that in both specifications the constant term is highly significant and positive signalling that the volatility estimates together somewhat underestimate the volatility level. This could just be due to the fact that our volatility estimates are referred to a wider-based index taking the average over a larger number of stocks. The specification which seems to have better properties is one in which an AR(1) correction is introduced which interestingly, keeps the joint explanatory power of the variables we derived. Beside an obvious improvement in the R^2 and in the value of the likelihood function, the diagnostics show a remarkable improvement since the residual autocorrelation and ARCH effects are reduced. The joint significance of either set of coefficients from the multiple indicator volatility models and from the GARCH models is established by means of F-tests (labelled F-test on MIMVOL and on GARCH, respectively) the values of which are very high. There is no clear indication of a better in-sample model performance as a consequence of the selection criterion adopted. In both equation one notices that there appear some negative coefficients: concentrating on the high-low range, for example, other things being equal, an increase in the derived volatility value would imply that there is a reduction in the estimated value of VIX.

Table 7. Dependent variable VIX. In-sample analysis (Jan. 4, 1988 – Dec.30, 1997)

Lowest BIC		**Model 1**	**Model 2**
Multi-step Average Volatility (GARCH)	Constant	0.778**	0.412**
		(6.57)	(3.27)
	Absolute Return	0.724**	0.766**
		(29.91)	(11.89)
	High Low	0.099	-1.215**
		(1.30)	(-16.97)
	Daily Realized	-0.216**	-0.071**
		(-7.51)	(-4.96)
Multi-step Average Volatility (MIMVOL)	Absolute Return	0.723**	1.328**
		(4.15)	(5.95)
	High Low	-1.479**	-0.070
		(-6.71)	(-0.26)
	Daily Realized	1.032**	0.233**
		(10.39)	(5.51)
	AR(1)		0.950**
			(142.90)
	R-squared	0.82	0.97
	Log likelihood	1832.13	4332.75
Residuals	AR(2) - (TR^2)	2031.53**	23.14**
	ARCH(4) - (TR^2)	1519.63**	46.85**
F-test	GARCH terms	595.25**	197.94**
	MIMVOL terms	85.33**	490.63**

General-to Specific		**Model 1**	**Model 2**
Multi-step Average Volatility (GARCH)	Constant	1.178**	0.581**
		(10.89)	(5.68)
	Absolute Return	0.679**	0.836**
		(28.91)	(13.57)
	High Low	0.056	-1.282**
		(0.78)	(-18.89)
	Daily Realized	-0.068*	-0.023
		(-2.54)	(-1.68)
Multi-step Average Volatility (MIMVOL)	Absolute Return	1.199**	1.259**
		(9.39)	(7.26)
	High Low	-2.359**	-0.239
		(-12.65)	(-1.03)
	Daily Realized	1.237**	0.362**
		(13.64)	(8.88)
	AR(1)		0.946
			(139.52)
Diagnostics	R-squared	0.84	0.97
	Log likelihood	1917.66	4349.05
Residuals	AR(2) - (TR^2)	1985.01**	22.04**
	ARCH(4) - (TR^2)	1496.67**	46.86**
F-tests	GARCH terms	490.50**	516.97**
	MIMVOL terms	450.89**	507.51**

Robust t-values in parentheses. * = significant at 5%; ** = significant at 1%

The same kind of analysis may be carried over to an out-of-sample combination of forecasts exercise over the period January 2, 1998 to November 10, 1998, in which the same variable (log(VIX)) is regressed on a constant and on the (logs of) term structures of volatility obtained from the same models, but this time keeping the coefficients fixed at their estimated in-sample values. The results, again broken down between lowest BIC models and general-to-specific models are presented in Table 8. Once again, the AR(1) correction achieves a substantial reduction in the residual diagnostics (eliminating residual autocorrelation, for example), and, again, the significance of the GARCH-based and MIMVOL-based terms is maintained judging on the F-tests. The lowest BIC provides a slightly better performance: all parameters are significant: the daily high-low range provides volatility measures which have a significantly negative impact on the forecasted values. The only coefficient which is not significant is the one on daily realized volatility from the GARCH specification. Some features of this forecast combination can be assessed also graphically (Figure 6) since it seems that the highest degree of heteroskedasticity appears in correspondence to periods in which the VIX index has higher and more volatile values.

6. Conclusions

In this paper we have suggested a dynamic model where different indicators of volatility are jointly modelled. The results show that the realized volatility can be explained by other measures of volatility such as the daily volatility or the daily range. The forecasts that can be derived from our model prove themselves useful in forecasting a volatility index such as VIX.

Table 8. Dependent variable VIX. Out-of-sample analysis (Jan. 2, 1998 - Nov. 10, 1998)

Lowest BIC		Model 1	Model 2
	Constant	1.393^{**}	0.972^{**}
		(6.63)	(3.40)
Multi-step Average Volatility (GARCH)	Absolute Return	0.263^{**}	0.471^{**}
		(4.03)	(2.85)
	High Low	-0.960^{**}	-1.564^{**}
		(-4.17)	(-6.95)
	Daily Realized	0.048	-0.022
		(1.335)	(-0.60)
Multi-step Average Volatility (MIMVOL)	Absolute Return	1.681^{**}	2.512^{**}
		(5.30)	(11.85)
	High Low	-0.897^{**}	-0.789^{**}
		(-5.25)	(-8.40)
	Daily Realized	0.537^{**}	0.237^{**}
		(4.57)	(3.83)
	AR(1)		0.907^{**}
			(26.94)
Diagnostics	R-squared	0.91	0.97
	Log likelihood	253.64	374.64
Residuals	AR(2) test (TR^2)	136.79^{**}	4.18
	ARCH(4)	103.02^{**}	12.41*
F-tests	GARCH terms	6.55^{**}	28.67^{**}
	MIMVOL terms	23.90^{**}	85.01^{**}

General-to Specific		Model 1	Model 2
	Constant	1.234^{**}	1.202^{**}
		(6.67)	(3.72)
Multi-step Average Volatility (GARCH)	Absolute Return	0.193^{**}	0.557^{**}
		(3.01)	(2.94)
	High Low	-0.238	-1.293^{**}
		(-1.49)	(-5.97)
	Daily Realized	0.165^{**}	0.126^{**}
		(4.15)	(3.48)
Multi-step Average Volatility (MIMVOL)	Absolute Return	0.683^{**}	1.741^{**}
		(4.07)	(11.06)
	High Low	-0.510^{**}	-0.802^{**}
		(-3.71)	(-8.06)
	Daily Realized	0.429^{**}	0.430^{**}
		(5.41)	(11.61)
	AR(1)		0.944^{**}
			(37.32)
Diagnostics	R-squared	0.91	0.97
	Log likelihood	249.35	372.43
Residuals	AR(2) test (TR^2)	126.88^{**}	1.82
	ARCH(4)	98.22^{**}	24.87^{**}
F-tests	GARCH terms	12.24^{**}	14.74^{**}
	MIMVOL terms	20.26^{**}	85.39^{**}

Robust t-values in parentheses. * = significant at 5%; ** = significant at 1%

REFERENCES

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

Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys, 2000, "Great Realizations," *Risk*, 13, 105-108.

Bollerslev, T., and H.O. Mikkelsen, 1996, "Modeling and Pricing Long Memory in Stock Market Volatility", *Journal of Econometrics*, Vol.73, 151-184.

Baillie, R.T., T. Bollerslev and H.O. Mikkelsen, 1996, "Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics*, Vol.74, 3-30.

Beckers, S., 1983, Variances of Security Price Returns Based on High, Low, and Closing Prices », *Journal of Business*, 56, 97-112.

Breidt, F.J., N. Crato and P. de Lima, 1998, "The Detection and Estimation of Long Memory in Stochastic Volatility", *Journal of Econometrics*, 83, 325-48.

Deo, R.S. and C.M. Hurvich, 1999, Estimation of Long Memory in Volatility, mimeo, Stern Business School.

Engle, R.F., 2000, "The Econometrics of Ultra High Frequency Data", *Econometrica*, 68, 1-22.

Engle, R.F. and G.J. Lee, 1999, "A Permanent and Transitory Component Model of Stock Return Volatility", (with G.J. Lee), in R.F. Engle and H. White (eds.) *Cointegration, Causality, and Forecasting: A Festschrift in Honor of C.W.J. Granger*, Oxford: Oxford University Press, 475-497.

Engle, R.F. and A. Patton, 2000, "What Good is a Volatility Model?" mimeo, UCSD and NYU.

Engle, R.F. and J.R. Russell, 1998, "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data", *Econometrica*, 66, 1127-1162.

Gallant A.R., C.-T. Hsu, and G.Tauchen, 1999, "Using Daily Range Data to Calibrate Volatility Diffusions and Extract the Forward Integrated Variance'', *Review of Economics and Statistics*, 81, 617-631.

Ghose D. and K. Kroner, 1996, “Components of Volatility in Foreign Exchange Markets: An Empirical Analysis of High Frequency Data”, mimeo Dept. of Economics, University of Arizona.

Garman M.B. and M.J Klass, 1980, “On the Estimation of Security Price Volatilities from Historical Data”, *Journal of Business*, 53, 67-78.

Parkinson, M. 1980, “The Extreme Value Method for Estimating the Variance of the Rate of Return”, *Journal of Business*, 53, 61-66.

Roll, R., 1984, “A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market”, *Journal of Finance*, 39, 1127-1139.

Schwert, G.W., 1998, „Stock Market Volatility: Ten Years after the Crash“, *Brookings-Wharton Papers on Financial Services*, Washington D.C.: Brookings Institutions.

Zhou, B. 1996, “High-Frequency Data and Volatility in Foreign-Exchange Rates”, *Journal of Business and Economic Statistics*, 14, 45-52.

FIGURES

Figure 1. The S&P500 Index: Time Series Profile of Absolute returns (ABSR), Daily High-Low Range (HL) and Daily Realized Volatility (DV) from Jan, 4, 1988 to Dec. 30, 1997.

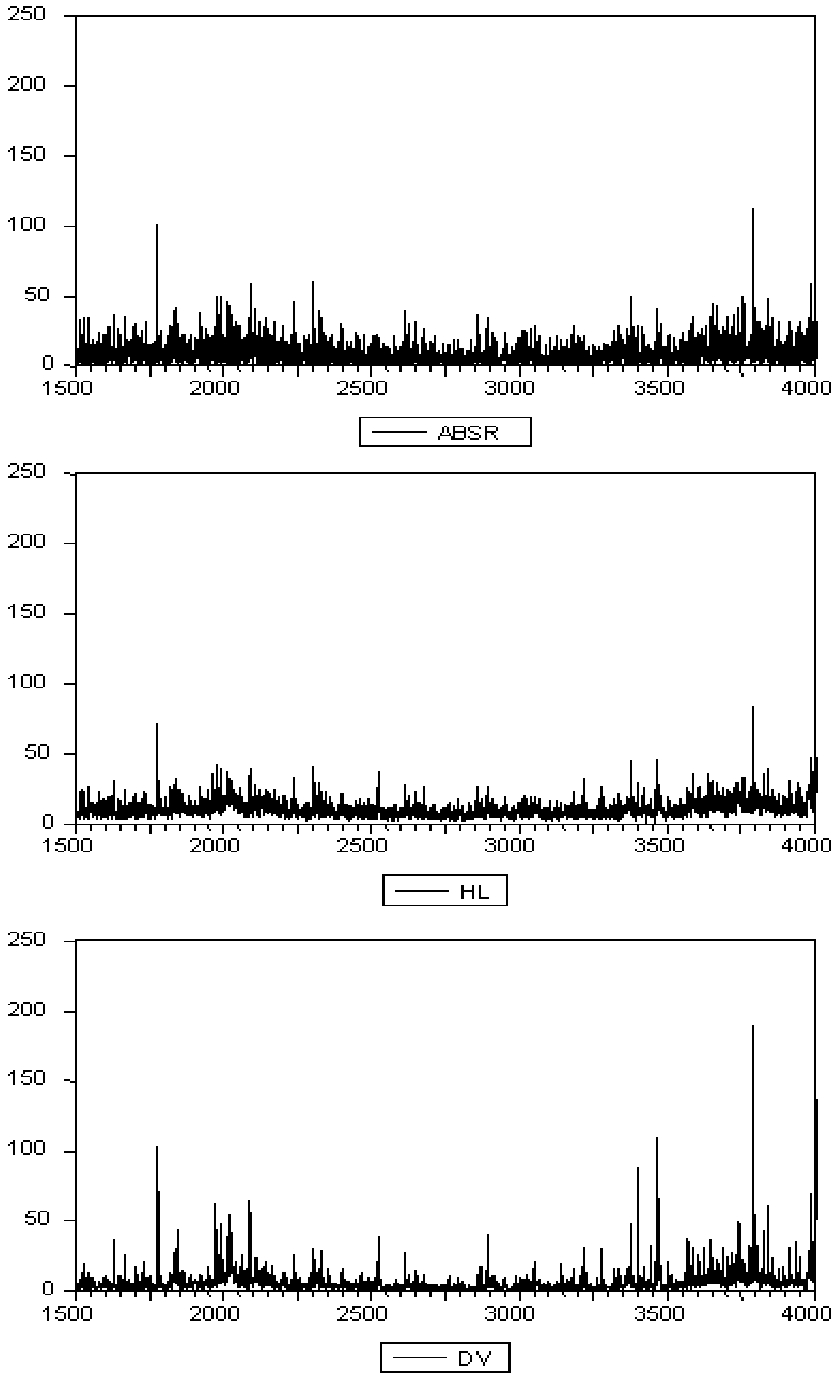

Figure 2. The S&P500 Index: Out-of-sample Volatility Forecasts: Abs. Returns, High-Low, Daily Volatility. Various starting dates.

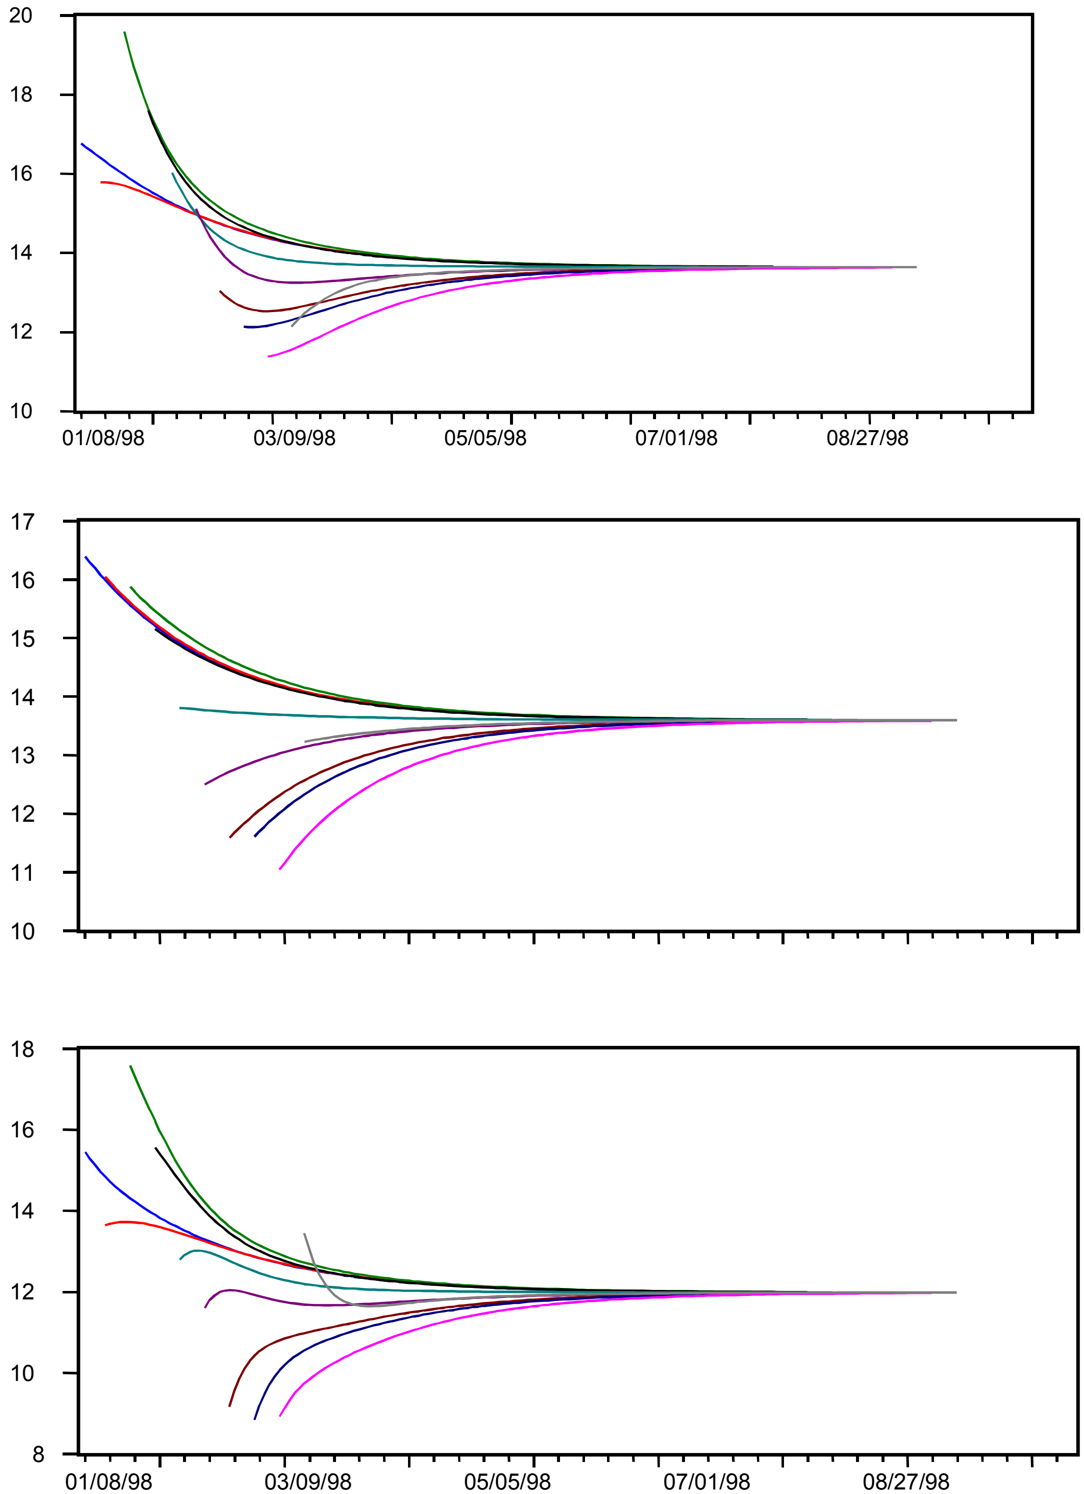

Figure 3. The S&P500 Index: One month ahead term structure of volatility: Absolute Returns, High-Low, Daily Volatility. IN-SAMPLE

Figure 3.a – One-month ahead term structure of volatility measured by absolute returns: System (HA-system) and GARCH (HA-GARCH) estimations.

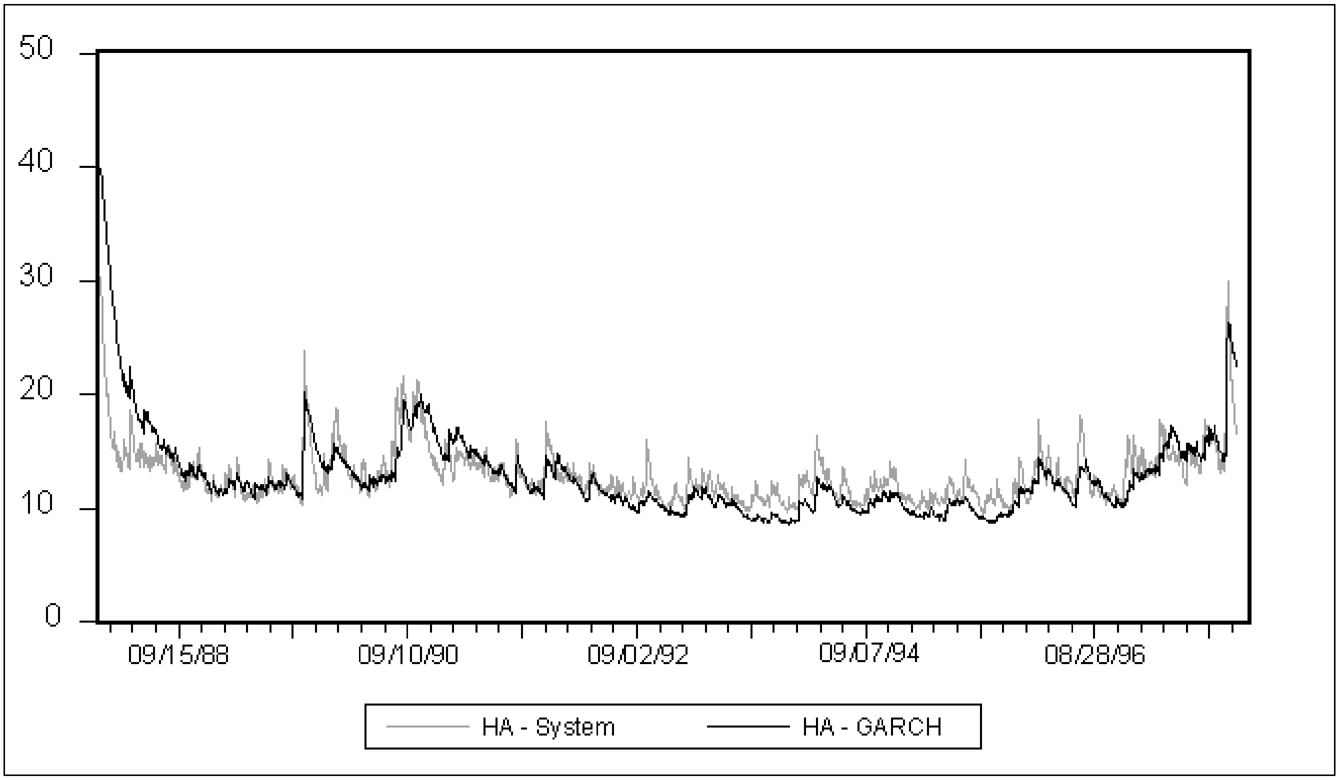

Figure 3.b – One-month ahead term structure of volatility measured by daily high-low range: System (HA-system) and GARCH (HA-GARCH) estimations.

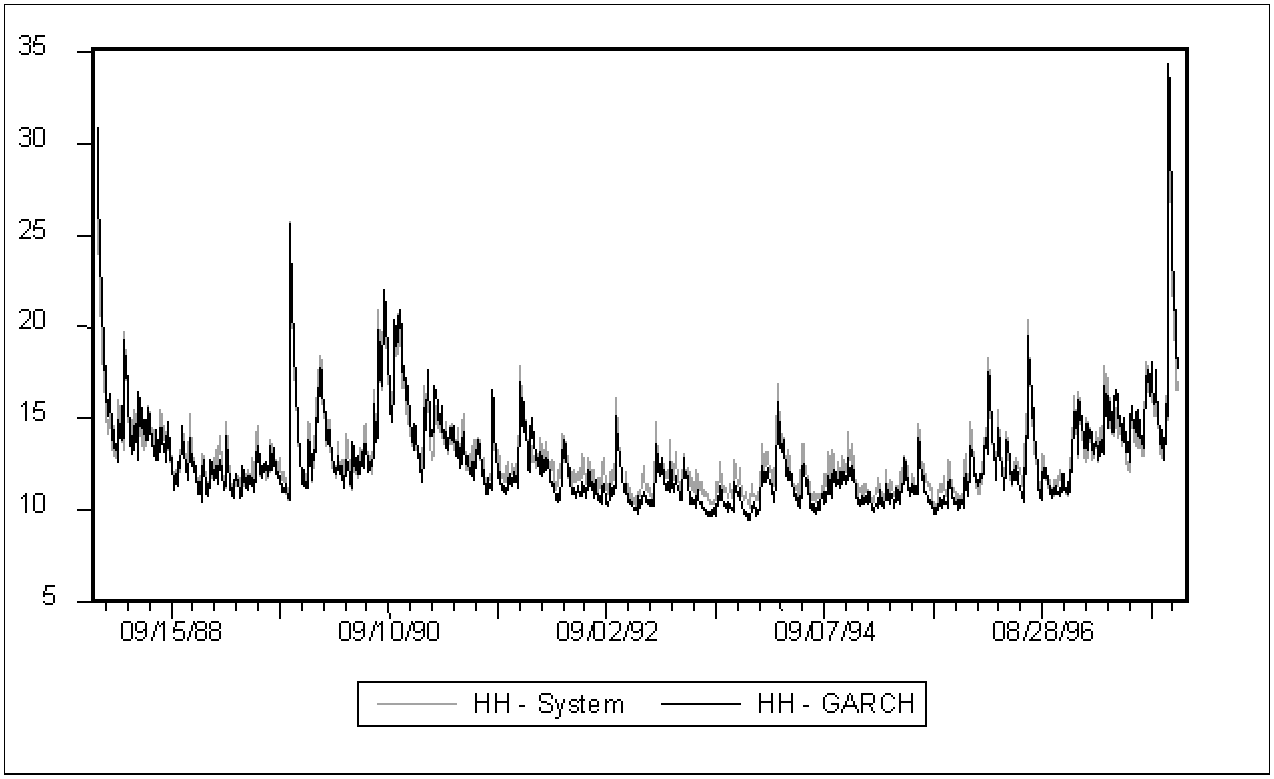

Figure 3.c – One-month ahead term structure of volatility measured by daily realized volatility: System (HA-system) and GARCH (HA-GARCH) estimations.

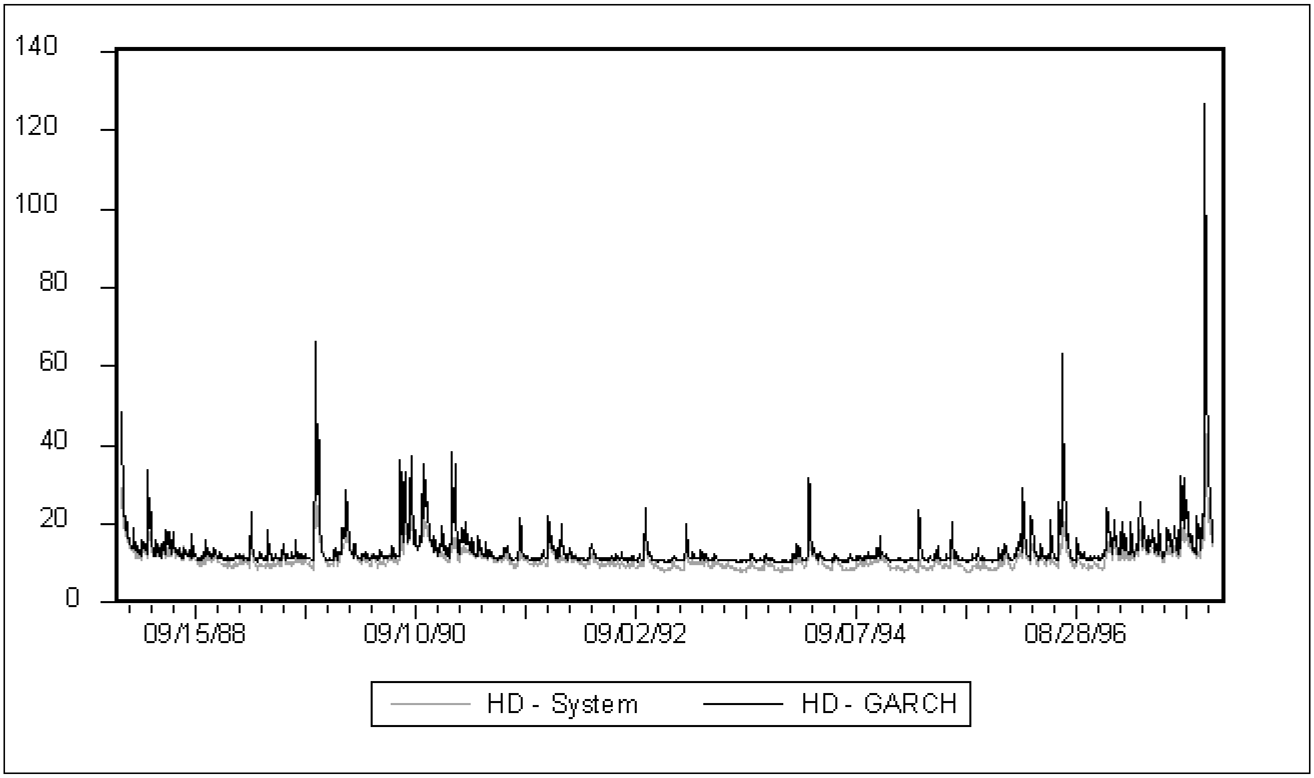

Figure 4. The S&P500 Index: One month ahead term structure of volatility: Abs. Returns, High-Low, Daily Volatility. OUT OF SAMPLE

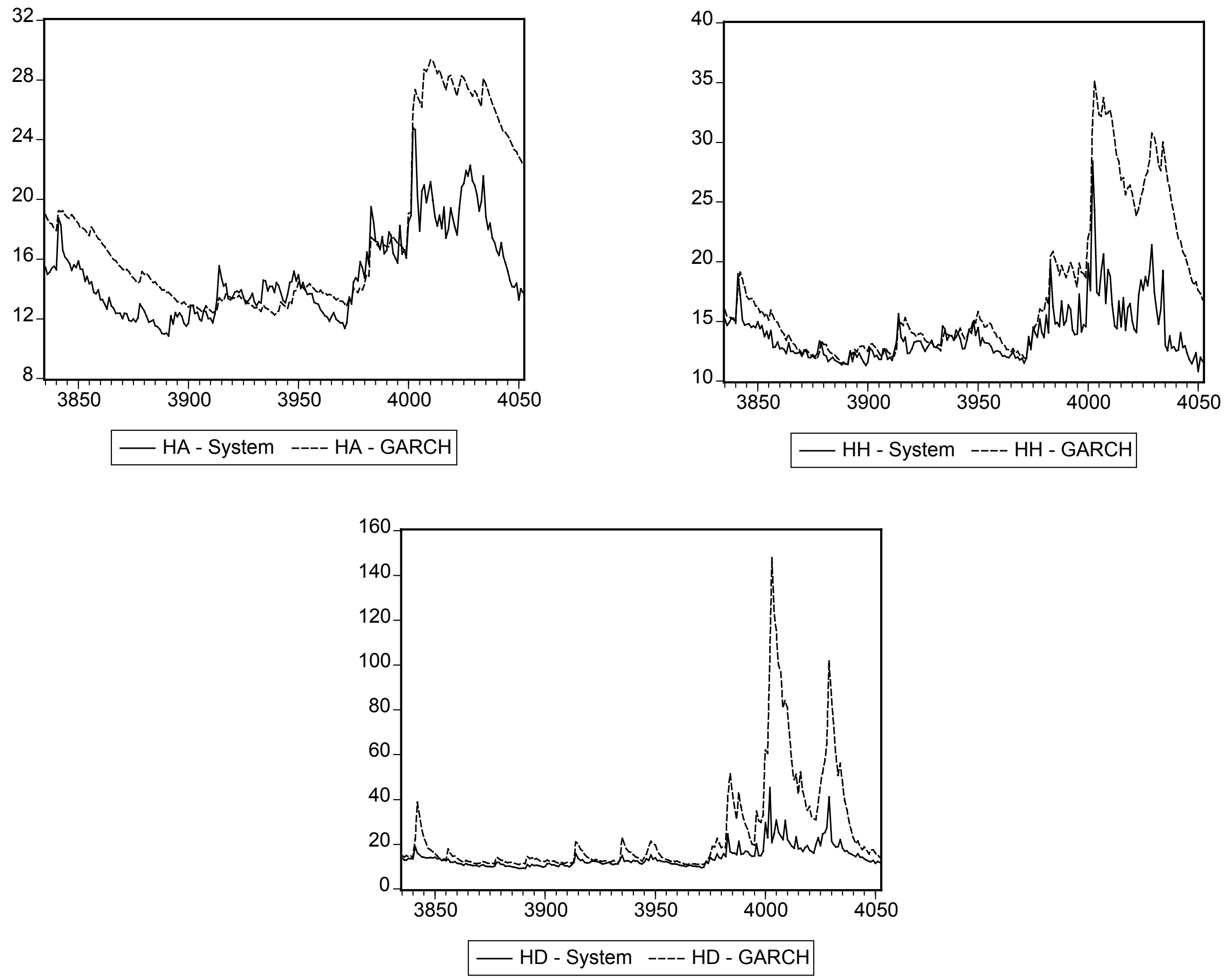

Figure 5. The CBOE VIX Index: Jan, 4, 1988 to Dec. 14, 1998.

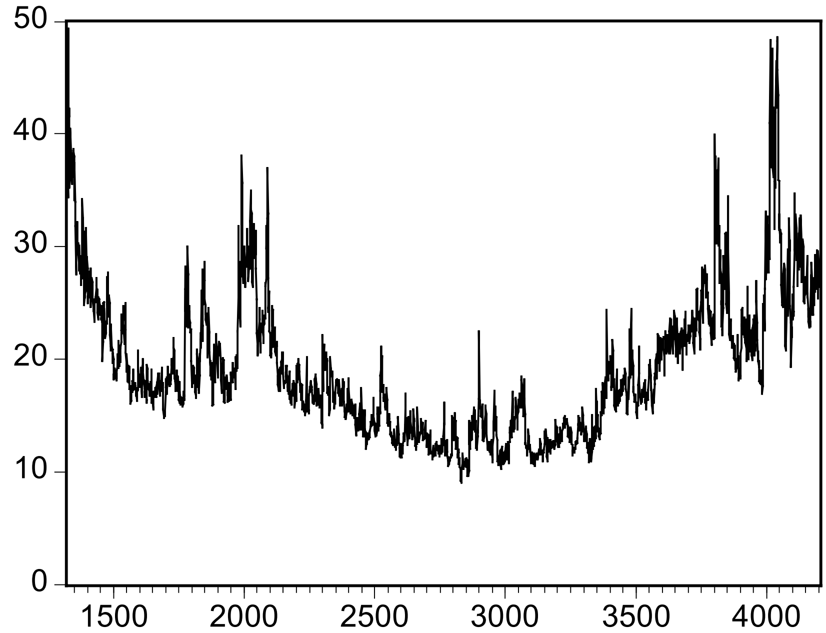

Figure 6. The CBOE VIX Index: Out of sample combination of forecasts - Jan, 2, 1998 to Nov. 10, 1998 – Smallest BIC Models with AR(1) correction (Table 8 – Model 2).