Working Paper 06-25(09)
Statistics and Econometrics Series
Abril 2006

Departamento de Estadística
Universidad Carlos III de Madrid
Calle Madrid, 126
28903 Getafe (Spain)
Fax (34) 91 624-98-49

VOLATILITY FORECASTS: A CONTINUOUS TIME MODEL VERSUS DISCRETE TIME MODELS[1]

Helena Veiga[2]

Abstract

This paper compares empirically the forecasting performance of a continuous time stochastic volatility model with two volatility factors (SV2F) to a set of alternative models (GARCH, FIGARCH, HYGARCH, FIEGARCH and Component GARCH). We use two loss functions and two out-of-sample periods in the forecasting evaluation. The two out-of-sample periods are characterized by different patterns of volatility. The volatility is rather low and constant over the first period but shows a significant increase over the second out-of-sample period. The empirical results evidence that the performance of the alternative models depends on the characteristics of the out-of-sample periods and on the forecasting horizons. Contrarily, the SV2F forecasting performance seems to be unaffected by these two facts, since the model provides the most accurate volatility forecasts according to the loss functions we consider.

Keywords: Efficient Method of Moments, Reprojection, Volatility Forecasting.

[1] I thank Michael Creel for introducing me to the idea of Efficient Method of Moments and for his advice and encouragement. I also thank George Tauchen and Tim Bollerslev for ther helpful remarks during my stay at Duke University. This research has been made possible thanks to the financial support of Fundaçao para a Ciência e Tecnologia. The usual disclaimer applies.

[2] Departamento de Estadística, Universidad Carlos III de Madrid.

Volatility Forecasts: A Continuous Time Model versus Discrete Time Models

Helena Veiga*

Statistics Department

Universidad Carlos III de Madrid

C/ Madrid 126, 28903 Getafe (Madrid), Spain

Email: mhveiga@est-econ.uc3m.es

Abstract

This paper compares empirically the forecasting performance of a continuous time stochastic volatility model with two volatility factors (SV2F) to a set of alternative models (GARCH, FIGARCH, HYGARCH, FIEGARCH and Component GARCH). We use two loss functions and two out-of-sample periods in the forecasting evaluation. The two out-of-sample periods are characterized by different patterns of volatility. The volatility is rather low and constant over the first period but shows a significant increase over the second out-of-sample period. The empirical results evidence that the performance of the alternative models depends on the characteristics of the out-of-sample periods and on the forecasting horizons. Contrarily, the SV2F forecasting performance seems to be unaffected by these two facts, since the model provides the most accurate volatility forecasts according to the loss functions we consider.

JEL classification: C10; G13

Keywords: Efficient Method of Moments, Reprojection, Volatility Forecasting.

*I thank Michael Creel for introducing me to the idea of Efficient Method of Moments and for his advice and encouragement. I also thank George Tauchen and Tim Bollerslev for ther helpful remarks during my stay at Duke University. This research has been made possible thanks to the financial support of Fundaçao para a Ciência e Tecnologia. The usual disclaimer applies.

1 Introduction

Having accurate predictions of the future volatility of financial assets is a necessary ingredient for markets participants to be able to price assets correctly. For some time, the *implied volatility* obtained from the Black-Scholes model was the most popular measure of future volatility. However, the predictive power of the Black-Scholes model diminished severely after the October 1987 stock market crash due to the increasing changes in volatility and, consequently, the implied volatility became a less attractive estimator of future volatility, [see, Canina and Figlewski (1993)].

The availability of high-frequency data and the computational advances allow us to provide today better estimates of future volatility. In the literature, one finds two main methods of how to predict the future volatility. The first method uses intra-period data to calculate the *realized volatility* (the sum of the intra-period squared returns) and then fits to it models that incorporate its main features. In this case, volatility is treated as observed. Since much of the theoretical literature assumes that the logarithm of the asset price follows a continuous time model like a diffusion, one advantage of this procedure is the possibility of getting unbiased and efficient estimators of the underlying integrated volatility. Contrarily, the second method treats volatility as latent; that is, it is possible to filter the volatility after estimating the structural model. It is rather obvious that the success of this method depends on the specification, because the volatility estimates are model dependent.

In this paper, we follow the second approach and fit the continuous time model with two volatility factors of Gallant and Tauchen (2001) to the returns of Microsoft. We choose a two factor volatility specification, because one recent finding shows that stochastic volatility models with one volatility factor are not able to characterize all the moments of the return distributions, [see, Andersen et al. (2002), Chernov and Ghysels (2000), Eraker et al. (2003), Jones (2003), Pan (2002), Gallant and Tauchen (2001), and Chernov et al. (2003)]. The main reason for this result is that these models are not able to fit the fat tails of return distributions. The introduction of two stochastic volatility factors can solve this problem, because one factor is going to deal with the persistence while the other one tries to accommodate the kurtosis, [see, Chernov et al. (2003)].

In the first step of our analysis we compare the forecasting performance of a set of benchmark models for two different out-of-sample periods and for two different forecasting horizons (10-days-ahead and 1-day-ahead forecasts). As benchmark models we consider the GARCH, the HYGARCH, the FIEGARCH, the FIGARCH, and the Component GARCH model. The last model is the most direct alternative in discrete time to the SV2F since it allows for two volatility components.

The task of selecting the benchmark models is not an easy one. Often, the best model depends upon the objectives of the researcher. González-Rivera et al. (2004) provided evidence that the preferred models depend sharply upon the loss function being used. For instance, in the context when an individual intends to maximize her/his expected utility by chosing a portfolio that consists of a risk-free and a risky asset, then the asymmetric GARCH-type models perform best, while stochastic volatility models clearly dominate other specifications when the objective is to calculate value-at-risk. Andersen and Bollerslev (1998), Hansen and Lunde (2005a), Pagan and Schwert (1990), and West and Cho (1995) also provided evidence that ARCH-type models yield accurate volatility forecasts and Davidson (2004) reports encouraging empirical results for the hyperbolic GARCH (HYGARCH) with respect to Asian exchange rates. Our choice is based mainly on the previous findings and include models such as the GARCH, the FIGARCH, the HYGARCH, the FIEGARCH and the component GARCH model (CModel). The FIGARCH, the HYGARCH and the FIEGARCH model have in common that the volatility processes include fractional integrated roots whose purpose is to capture long memory.[1]

In the second step of our analysis, we select the two best benchmark models at each out-of-sample period and compare them to the SV2F by calculating the *"rolling" volatility forecasts*. To do so, we proceed as follows: We estimate the SV2F model with the help of the Efficient Method of Moments (EMM) of Gallant and Tauchen (1996) and filter the underlying volatility using the reprojection technique of Gallant and Tauchen (1998). Under the assumption that the model is correctly specified, we obtain a consistent estimator of the integrated volatility. Finally, we regress a function of the realized volatility on a constant and on a function of the volatility forecasts and evaluate the predictive power of the volatility forecasts via the corrected R^2 of the OLS regression and the corresponding mean squared forecast error (MSFE), [see, Andersen et al. (2005)].

One common problem is that the realized volatility is only a consistent estimator of the true volatility when prices are observed continuously and without measurement errors, [see, Merton (1980)]. Unfortunately, these hypotheses are not true in general, and, as a consequence, the realized volatility is often biased due to market microstructure noises. Moreover, its bias tends to get worse as the sampling frequency of intra-day returns increases, [see, Andreou and Ghysels (2002), Oomen (2002) and Bai et al. (2004)]. One way to minimize this problem is to compute the realized volatility from intra-day returns that are sampled

[1] According to Parzen (1981), a stationary process $\{y_t\}$ with an autocovariance γ_y is called a *long memory process in the covariance sense*, if $\sum_{\tau=-n}^{n} \gamma_y(\tau) \to +\infty$ as n tends to $+\infty$. Granger and Joyeux (1980) provided a different definition of long memory. According to them, $\{y_t\}$ is a *long memory process in the covariance sense with a speed of convergence of order* $2d$, $0 < d < 1/2$, whenever $\gamma_y(\tau) = C(d)\tau^{2d-1}$, as $\tau \to \infty$ (here, $C(d)$ is a function that depends on d).

at a moderate frequency. We use 15 minutes intra-day data. According to Andersen et al. (2005a) this frequency is effective in reducing the bias of the realized volatility. An alternative solution to decrease the biases is the application of kernel-based estimators [see, Zhou (1996), Hansen and Lunde (2005b), Barndorff-Nielsen et al. (2004b) and Hansen and Lunde (2005c)] or sub-sample based estimators [see, Zhou (1996), Zhang et al. (2005) and Zhang (2005)]. In this paper, we also apply the simplest kernel-estimator of Hansen and Lunde (2005c).

Our empirical results evidence the superiority of the SV2F model in forecasting volatility in both out-of-sample periods. Moreover, the performance of the benchmark models depends mainly on the characteristics of the out-of-sample periods and on the forecasting horizons.

The paper is organized as follows: In the next Section, we introduce all models and estimate them. In Section 3, we explain formally how the volatility forecasts are calculated from the data. Afterwards, we evaluate the forecasting performance of all specifications. Finally, we conclude. Figures and Tables are relegated to the Appendix.

2 Volatility Specifications

2.1 Continuous Time Stochastic Volatility

It is one of our main objectives to forecast the future volatility of the Microsoft share using the stochastic volatility model (SV2F) of Gallant and Tauchen (2001). Formally, let P_t be the value of one share of Microsoft at instant t (we reserve the notation U_{1t} for the logarithm of P_t) and assume that the instantaneous return of the asset at t, $\frac{dP_t}{P_t}$, is given by

$$\frac{dP_t}{P_t} = \alpha_{10}dt + \exp(\beta_{10} + \beta_{12}U_{2t} + \beta_{13}U_{3t})(\psi_{11}dW_{1t} + \psi_{12}dW_{2t} + \psi_{13}dW_{3t}) \qquad (1)$$

and

$$dU_{it} = (\alpha_{i0} + \alpha_{ii}U_{it})dt + dW_{it}, \text{ for } i = 2, 3. \qquad (2)$$

In equation (1), α_{10} denotes the instantaneous expected return, $\exp(\beta_{10} + \beta_{12}U_{2t} + \beta_{13}U_{3t})$ is the instantaneous standard deviation (or instantaneous volatility), W_i, $i = 1, 2, 3$, are Wiener processes, and ψ_{1i}, $i = 1, 2, 3$, are correlation coefficients that satisfy the restriction $\psi_{11} = \sqrt{1 - \psi_{12}^2 - \psi_{13}^2}$. A consequence of equation (1) and (2) is that the instantaneous correlation between returns and changes in variance (the leverage effect) is given by

$$corr\,(dU_{1t}, \beta_{12}dU_{2t} + \beta_{13}dU_{3t}) = \frac{\beta_{12}\psi_{12} + \beta_{13}\psi_{13}}{\sqrt{\beta_{12}^2 + \beta_{13}^2}}dt. \qquad (3)$$

In the SV2F model (equations (1) - (3)), U_2 and U_3 are volatility factors whose drifts allow for mean reversion (this is the case when α_{22} and α_{33} are negative). Moreover, if the absolute values of α_{22} and α_{33} are both smaller than one, then shocks to volatility take time to dissipate. In this case, the volatility factors are said to be *slow mean reverting*. Observe finally that the parameter β_{10} takes care of the long-run mean of volatility.

The SV2F model is not fully identified, but it is possible to deal with this problem by imposing additional restrictions on some selected parameters. In particular, we set $\alpha_{20} = \alpha_{30} = 0$. These restrictions are common in the literature on systems of differential equations, because they provide flexibility and numerical stability in the estimation phase, [see, Gallant and Tauchen (2001)].

In order to be able to forecast the volatility, we have to estimate the SV2F model first. To do so, we use the *Efficient Method of Moments* (EMM) of Gallant and Tauchen (1996), an estimation technique that is based on two compulsory phases. The first phase (Projection) consists of projecting the observed data onto a transition density that is a good approximation of the distribution implicit in the true data generating process. The simulated density is called the *auxiliary model* and its score is said to be *the score generator for EMM*. The advantage of EMM is that the score has an analytical expression. In the projection step, we proceed carefully along an expansion path with tree structure and the auxiliary model comes out to be a semi-parametric GARCH, as in Gallant and Tauchen (2001). In the second phase (Estimation), the parameters of the models are estimated with the help of the score generator. The score enters the moment conditions in which we replace the parameters of the auxiliary model by their quasi-MLEs obtained in the projection step. The estimates are finally obtained by minimizing the GMM criterion function. EMM includes a post-estimation simulation (Reprojection) as an optional step. This step becomes crucial in our forecasting analysis, because it allows us to filter the volatility implicit in the model.

The SV2F model and all alternative benchmark models presented later on are estimated using data adjusted for stock splits from March 13, 1986 until February 23, 2001. In total, we have 3.778 observations [for the time series of the Microsoft share over this period and the corresponding daily returns, see Figure 2]. The SV2F model is estimated using the package EMM which is available online at Duke University.

Table 1 reports the values of the diagnostic tests. The test statistic follows an asymptotic chi-square distribution with $p_\theta - p_\rho$ degrees of freedom, where $p_\theta - p_\rho$ denotes the difference between the number of parameters of the auxiliary model obtained in the projection step and the structural SV2F model. We observe that the two factor volatility models (with and without leverage effect) pass the specification test and that all coefficients are statistically significant. Moreover, the first volatility factor is very slow mean reverting while the second is extremely fast mean reverting as in Gallant and Tauchen (2001), [see, again Table 2]. With respect to the coefficient estimates of leverage effect, ψ_{12} and ψ_{13}, we see that they are both negative and statistically significant.

2.2 Discrete Time: ARCH-Type Models

2.2.1 HYGARCH and FIEGARCH Models

Davidson (2004) proposed the HYGARCH model as an alternative to the FIGARCH since it is able to generate long memory without behaving oddly when d, the parameter of fractional integration, approximates 1. Formally, let the prediction error ε_t satisfy

$$\varepsilon_t = \sigma_t \epsilon_t, \tag{4}$$

where σ_t^2 is the conditional variance of ε_t given information at time $t-1$, $\sigma_t > 0$, and $\epsilon_t \sim NID(0,1)$. Additionally, it is assumed that σ_t^2 is such that

$$\sigma_t^2 = \omega + \theta(L)\varepsilon_t^2, \tag{5}$$

where

$$\theta(L) = 1 - \frac{\delta(L)}{\beta(L)}(1 + \alpha((1-L)^d - 1)). \tag{6}$$

In equation (6), $\theta(L)$, $\delta(L)$ and $\beta(L)$ are polynomials in the lag operator L. Moreover, $\omega > 0$, $\alpha \geq 0$ and $d \geq 0$. The HYGARCH model (equations (4)-(6)) simplifies to a GARCH(p,q) and to a FIGARCH(p,d,q) if $\alpha = 0$ and $\alpha = 1$, respectively. For $0 < \alpha < 1$, we have a nested model that behaves normal in the sense that increases in the parameter of fractional integration d leads to more persistence.

If, on the other hand, ε_t follows a FIEGARCH(p,d,q), then the volatility process is given by

$$\ln \sigma_t^2 = \omega + \phi(L)^{-1}(1-L)^{-d}[1 + \psi(L)]g(\epsilon_{t-1}), \qquad -1 \leqslant d \leqslant 1. \tag{7}$$

In equation (7), $\phi(L) = 1 - \phi_1 L - ... - \phi_p L^p$ and $\psi(L) = 1 + \theta_1 L + ... + \theta_q L^q$ are an autoregressive polynomial and a moving average polynomial in the lag operator L, respectively. It is assumed that the roots of $\phi(L)$ lie outside the unit circle and that both polynomials do not have common roots. Note that the objective of the function $g(\epsilon_{t-1}) = \gamma_1 \epsilon_{t-1} + \gamma_2[|\epsilon_{t-1}| - E(|\epsilon_{t-1}|)]$ is to introduce asymmetry between returns and changes in the variance, [see, Nelson (1991)].

2.2.2 The Component Model

Engle and Lee (1993) formulated a model with two components in the volatility specification. The first one deals with the long-run features that could affect volatility while the second tries to accommodate the short-run dynamics. Formally, let y_t be the returns of a financial asset with expected value m_t and conditional variance σ_t^2. It is assumed in the component model (CModel) that y_t follows a process such that

$$y_t = m_t + \varepsilon_t, \tag{8}$$

where ε_t is given by equation (4) and ϵ_t and σ_t satisfy the assumptions imposed on equation (4). However, the conditional variance is now assumed to be equal to

$$\sigma_t^2 = q_t + \alpha(\varepsilon_{t-1}^2 - q_{t-1}) + \beta(\sigma_{t-1}^2 - q_{t-1}). \tag{9}$$

In the former equation, q_t is the permanent component of the conditional variance and it is specified by

$$q_t = \varpi + \rho q_{t-1} + \phi(\varepsilon_{t-1}^2 - \sigma_{t-1}^2), \quad \rho > \alpha + \beta. \tag{10}$$

In the CModel (equations (4) and (8)-(10)), the short-run component of the conditional variance is given by the difference between the conditional variance σ_t^2 and q_t. The error $\varepsilon_{t-1}^2 - \sigma_{t-1}^2$, whose expected value is zero, drives the time-dependent movement of the permanent component, and therefore, it might be seen as a trend.

2.2.3 Estimation Results

The benchmark models are estimated with the Ox package Garch 4.0 of Laurent and Peters (2005) or with Eviews 5. We report our results in Tables 3, 3.1 and 4. With respect to the HYGARCH model we observe that the hyperbolic parameter $ln(\alpha)$ is not statistically different from 0. We have already commented before that the HYGARCH reduces to a FIGARCH whenever α is equal to one. We also see that the estimate of the persistence of the GARCH model is around 0.96 and that the asymmetric relation between returns and volatility in the FIEGARCH model is negative but not significant. For this reason, we do not consider models with leverage effect in the forecasting step. Finally, with respect to the CModel, we observe that the autoregressive parameter ρ is close to one. This means that the permanent component of the conditional variance has a high degree of persistence. Analogously, the persistence level of the transitory component is given by $\alpha + \beta$ in equation (10). Since this sum is equal to 0.649, deviations of the conditional variance from its trend seem to be temporary. Our estimation results for the CModel are similar to those of Engle and Lee (1993).

3 Forecasting Volatility

SV2F: Concerning the SV2F model, our principal objective is to obtain an estimator $\hat{\sigma}_{t+1}^2$ of the one-step-ahead conditional variance forecast σ_{t+1}^2. To do so, we proceed as follows: As a by-product of the estimation step we obtain a long simulation of y_t, $\{\hat{y}_\tau\}_{\tau=1}^N$ with $N = 100000$, at the estimated parameter vector of the structural model. Next, we impose the auxiliary model found in the projection step on the simulated values $\{\hat{y}_\tau\}_{\tau=1}^N$ in order to obtain a good representation of the conditional variance.[2] From the estimation of the auxiliary model we can now calculate the conditional variance $\tilde{\sigma}_\tau^2$ in the semi-parametric GARCH. Then, we regress $\tilde{\sigma}_\tau^2$ on its own lags and the lags of $\hat{y}_\tau$ and $|\hat{y}_\tau|$. In particular, the expression

$$\tilde{\sigma}_\tau^2 = \alpha_0 + \alpha_1 \tilde{\sigma}_{\tau-1}^2 + ... + \alpha_p \tilde{\sigma}_{\tau-p}^2 + \theta_1 \hat{y}_{\tau-1} + ... + \theta_q \hat{y}_{\tau-q} + \pi_1 |\hat{y}_{\tau-1}| + ... + \pi_r |\hat{y}_{\tau-r}| + u_t, \quad (11)$$

gives us a calibrated function inside the simulation. We obtain the *reprojected volatility* $\hat{\sigma}_t$ by replacing the simulated values on the right hand side of equation (11) by the true data; that is,

$$\hat{\sigma}_t^2 = \hat{\alpha}_0 + \hat{\alpha}_1 \sigma_{t-1}^2 + ... + \hat{\alpha}_p \sigma_{t-p}^2 + \hat{\theta}_1 y_{t-1} + ... + \hat{\theta}_q y_{t-q} + \hat{\pi}_1 |y_{t-1}| + ... + \hat{\pi}_r |y_{t-r}|. \quad (12)$$

Finally, we calculate $\hat{\sigma}_{t+1}^2$ by evaluating a filter equation similar to equation (12) on the observed data series. We ignore the conditional mean variation because it is negligible for short forecasting horizons, [see, Andersen et al. (2005b)].

GARCH$(1,1)$**:** Using recursive substitutions, the GARCH$(1,1)$ model can be written as an ARCH(∞); that is,

$$\sigma_t^2 = \omega(1-\beta)^{-1} + \alpha \sum_{i=1}^{+\infty} \beta^{i-1} \varepsilon_{t-i}^2. \quad (13)$$

Since the unconditional variance of the process is equal to $\sigma^2 = \omega(1-\alpha-\beta)^{-1}$, the multi-step-ahead forecast of the conditional variance based upon the available information at t is given by

$$\sigma_{t+k}^2 = \sigma^2 + (\alpha+\beta)^{k-1} \cdot (\sigma_{t+1}^2 - \sigma^2). \quad (14)$$

Observe that we need to assume that $(\alpha + \beta) < 1$ in order to guarantee that σ^2 exists. Moreover, the multi-step-ahead forecast of the conditional variance converges to the unconditional variance at an exponential rate fixed by $\alpha + \beta$, [see, Andersen et al. (2005b)].

FIGARCH$(1,d,1)$**:** If we consider a FIGARCH$(1,d,1)$, then the actual conditional variance forecasts are given by

$$\sigma_{t+k|t+k-1}^2 = \omega(1-\beta)^{-1} + \lambda(L)\sigma_{t+k-1|t+k-2}^2, \quad (15)$$

where $\sigma_{t+k|t+k-1}^2 \equiv \varepsilon_t^2$ for $k < 0$ and the coefficients of $\lambda(L) \equiv 1 - (1-\beta L)^{-1}(1-\alpha L - \beta L)(1-L)^d$ are computed from the expressions $\lambda_1 = \alpha + d$ and for all $j = 2, 3, ...$,

$$\lambda_j = \beta\lambda_{j-1} + [(j-1-d)j^{-1} - (\alpha+\beta)]\delta_{j-1}, \text{ with } \delta_j \equiv \delta_{j-1}(j-1-d)j^{-1}. \quad (16)$$

[2] Given the simulation length, these regressions are as Gallant and Tauchen (2001) say, analytic projections.

Note that the δ_j's are the coefficients in the Maclaurin series expansion of $(1-L)^d$, [see, Andersen et al. (2005b)].

CModel: Finally, forecasting volatility using the component model requires to elicit for all t, given the available information at $t-1$, the multi-step-ahead forecasts of the conditional trend q_{t+k} and the conditional variance σ^2_{t+k}. According to Engle and Lee (1993) the conditional trend and the conditional variance, respectively, are equal to

$$q_{t+k} = \left[(1-\rho^k)/(1-\rho)\right]\omega + \rho^k q_t \tag{17}$$

and

$$\sigma^2_{t+k} - q_{t+k} = (\alpha+\beta)^k \left(\sigma^2_t - q_t\right). \tag{18}$$

In the former equations, $\rho < 1$ and $(\alpha+\beta) < 1$. If $\rho > (\alpha+\beta)$, then the transitory component (equation (18)) will decay faster than the trend component (equation (17)), and therefore, the trend component will dominate the conditional variance forecasts as k increases. In this case, the conditional variance will converge to $\sigma^2_{t+k} = q_{t+k} = \omega/(1-\rho)$ as k tends to infinity, because the trend component is itself stationary.

4 Evaluating and Comparing Volatility Forecasts

4.1 Evaluation Procedure: Realized Volatility

Suppose for a second that we have obtained the volatility forecasts for every model. To address the question which model performs best in terms of volatility forecasts, we compare the volatility forecasts with the *realized volatility.* To do so, let $r_{j,t}$, $0 \leq j \leq n$, represent a set of $n+1$ intra-day returns for day t ($j=0$ refers to the last price at day $t-1$, $j=1$ to the first observation after the market has opened on day t, and $j=n$ is the last price at day t). It can then be shown, under innocuous regularity conditions, that the realized volatility $RV_t \equiv \sum_{j=0}^{n} r^2_{j,t}$ converges to the *integrated volatility* (the time integral of the instantaneous volatility), [see, e.g., Andersen and Bollerslev (1998), Andersen et al. (2001), Barndorff-Nielsen and Shephard (2001, 2002a, 2002b, 2004a), Comte and Renault (1998), Andersen et al. (2003) and Andersen et al. (2005a)]. In fact, for a given sample period, the higher the frequency of the data and the larger the number of observations, the better the approximation of the realized volatility estimator to the integrated volatility. In the cases where the logarithm of asset price is not a pure diffusion (for instance, if it follows a jump-diffusion process), Andersen et al. (2003) proved that the realized volatility converges to the total variation of the asset return.

In this paper, we use the intra-day 15-minutes return observations of Microsoft. Hence, we have for every day 26 observations.[3] Using this data set, we use two different estimators $\widehat{RV}_t$ of realized volatility: (1) We sum up the intra-day squared returns and (2) we implement the simplest kernel-estimator of Hansen and Lunde (2005c). In the latter case, $\widehat{RV}_t = \sum_{i=1}^{m} y_i^2 + 2\frac{m}{m-1}\sum_{i=1}^{m} y_i y_{i+1}$, where m denotes the number of observations at day t. Figure 1 reports the graphs for these two high-frequency volatility measures.

Finally, we regress, as it has been proposed for example by Andersen and Bollerslev (1998) and Andersen et al. (2003), a function of the realized volatility on a constant and on a function of the forecasts of the different models using the OLS estimation method. In particular, the two loss functions we consider are such that

$$(\widehat{RV}_{t+1})^{0.5} = \beta_0 + \beta_1 \cdot (\hat{\sigma}^2_{t+1|\text{model}})^{0.5} + u_{t+1} \tag{19}$$

and

$$\ln(\widehat{RV}_{t+1}) = \beta_0 + \beta_1 \cdot \ln(\hat{\sigma}^2_{t+1|\text{model}}) + u_{t+1}. \tag{20}$$

Note that the variable *model* can take the values GARCH, FIGARCH, CModel or SV2F. In order to decide which model performs best, we take into account *mean squared forecast error* (MSFE), the corrected R^2, and the t-statistics corresponding to the hypotheses $\beta_0 = 0$ and/or $\beta_1 = 1$.[4] We only calculate the corrected R^2, which we denote R^{*2}, when the dependent variable of regressions (19) and (20) is a function of the realized volatility. When the dependent variable is a function of the Kernel based estimator of Hansen and Lunde (2005c), we present the values of the normal R^2 because this measure is, supposedly, not affected by microstructure noises.

4.2 Empirical Results

We compare the forecasting performance of the SV2F model to the forecasting performance of the alternative models at two different out-of-sample periods. The first one ranges from January 4, 1999 until December 31, 1999 (252 observations) whereas the second one ranges from January 4, 2000 until January 23, 2001 (288 observations). We choose these two out-of-sample periods, because they allow to test the models in two different environments; one in a relatively constant volatility pattern (the first) and one in an increasing volatility pattern (the second), [see, Figure 1].

[3] The data was obtained from Price-data.com.

[4] In our analysis, we also account for possible measurement errors in the empirical realized volatility that will often result in a downward bias in any measure of predictability. To solve this problem we follow Andersen et al. (2005a) and compute the corrected R^2 by scaling the original R^2 by a multiplicative adjustment factor.

The comparison of the forecasting performance takes place in two steps: First, we evaluate the benchmark models and select the two best models at each of the two forecasting horizons (10 days ahead for the first out-of-sample and 1 day ahead for the second out-of-sample period). Then, in the second step, we compare the performance of the selected models to the forecasting performance of the SV2F model. Proceeding like this allows us to perform a thorough comparison of the alternative models, including several forecasting techniques, from which the analysis may benefit.

We focus now on the first step of the two step comparison and consider the first out-of-sample period. We compute the 10 days ahead volatility forecasts by re-estimating the alternative models every 10 days. The results are reported in Tables 5, 5.1 and 6. We observe in Tables 5 and 5.1 that both hypotheses $\beta_0 = 0$ and $\beta_1 = 1$ are not rejected at a 5% significance level for the GARCH and the FIGARCH model. Hence, the volatility forecasts are unbiased estimators of the two measures of realized volatility presented in this paper (intra-day squared returns and kernel estimator). Moreover, we see with respect to the CModel that its forecasting performance is the worst of the three alternative models, because (a) its volatility forecasts in equations (19) and (20) are not statistical significant, (b) it presents the smallest corrected R^2 for both loss functions and its mean squared forecasting error is bigger than the one of the GARCH model. Our results confirm the findings of Ederington and Guan (2004) who have shown that more complex and flexible models forecast worse out-of-sample because adding more parameters into the models increase the scope for estimation error. Although the ratios of the MSFE in Table 5 and 5.1 are different from 1, we do not get statistical evidence that their differences are statistical significant according to the S_1 statistic of Diebold and Mariano (2002).[5] Therefore, for the first out of sample period, the FIGARCH and the GARCH model are selected to proceed to the second comparison stage.

At the second out-of-sample period that ranges from January 4, we compute the 1 day ahead forecasts by re-estimating the alternative models every day. In Tables 7 and 7.1 we see, once more, that the FIGARCH is the best model in terms of forecasting performance. It presents the highest R^2's and the smallest MSFE. Moreover, it is the only model for which we do not reject the null hypotheses of $\beta_0 = 0$ and $\beta_1 = 1$ at the 5% significance level. With respect to GARCH and CModel, we observe that they behave quite similar under both measures of ex-post volatility. In fact, for both models the null hypotheses of $\beta_0 = 0$ and $\beta_1 = 1$ are rejected for any conventional significance level. Finally, the results of this first step evidence that the forecasting performance of the CModel gets better (comparatively to GARCH) at

[5] The null hypothesis is that the loss differential series between the GARCH and the FIGARCH (or the CModel, respectively), $SFE_{t,10}^{GARCH} - SFE_{t,10}^{model}$, is equal to zero. Observe that $SFE_{t,10} = \left(f\left(RV_{t,10}\right) - f\left(\sigma_{t,10}^2\right)\right)^2$. For alternative tests check West (1996) and Harvey et al. (1997).

out-of-sample periods where the volatility pattern is not constant. The introduction of two components into the volatility specification makes the model more flexible and allows the specialization of the components. One components accommodates the kurtosis while the other deals with the persistence of data. Hence, for the second out-of-sample period, the FIGARCH and the CModel are selected to proceed to the second comparison stage.

In the second step, we compare the forecasting performance of the two best alternatives models at each out-of-sample period to the forecasting performance of SV2F. We use the same forecasting horizons and the same measures of realized volatility as before, but, nevertheless, we proceed in a slightly different way. The difference stems from the fact that it would be too much time consuming to estimate the SV2F model every 10 days (or even every day depending on the forecasting horizon), to filter the volatility and to compute finally the volatility forecasts. Instead, we estimate the model the day before the beginning of the out-of-sample period and compute afterwards the 10 days (or the 1 day) *"rolling" volatility forecasts.*

Tables 9 and 9.1 and Figure 3 report the main results for the first out of sample period. In terms of predictability, we mean R^2 measures, the SV2F performs much better than the selected alternative models. Moreover, we observe that the hypotheses of $\beta_0 = 0$ and/or $\beta_1 = 1$ are sharply rejected for all regressions. These biases would disappear if we estimated the models every 10 days, as we had seen before. It is revealed in Figure 3, where the forecasts of the GARCH, the FIGARCH and the SV2F model are presented, that the GARCH type models seem to produce forecasts that overestimate the kernel-based estimator of realized volatility. With respect to the SV2F model, panel (a) of Figure 3 shows that the model overestimates the kernel-based estimator for the first part of the out-of-sample period but behaves quite well in the second part of the out-of-sample period.

The results for the second out-of-sample period confirm our previous finding that the SV2F is the best model in terms of volatility forecasting: It has the highest R^2 for the kernel-based estimator of realized volatility and it is shown in panel (a) of Figure 4 that its volatility forecasts track quite well the pattern of our realized volatility estimator (kernel-based estimator). This allows us to conclude that stochastic volatility models designed in continuous time seem to be more powerful in terms of volatility forecasting.

5 Conclusion

In this paper, we compare the forecasting performance of a continuous time stochastic volatility model with two factors of volatility (SV2F) to the one of the GARCH, the FIGARCH and the component model. As a proxy of ex-post volatility, we choose the realized volatility calculated from the intra-daily returns and the kernel-based estimator of Hansen and Lunde (2005c).

The main contributions of this paper include the calculation of volatility forecasts using the reprojection technique proposed by Gallant and Tauchen (1998), the evaluation and the comparison of the forecasting performance of a continuous time model to the forecasting performances of alternative models designed in discrete time, and the use of two different measures of realized volatility (the simple sum of squared intra-daily returns and the kernel-based estimator) at the Mincer-Zarnowitz style regressions of the ex-post realized volatility on the model forecasts.

Our empirical analysis reveals that, at the first out-of-sample period, the forecasting performance of the SV2F model is significatively better than the one of the GARCH and the FIGARCH model and, at the second out-of sample period, the SV2F model is the best model under the kernel-based estimator. Therefore, continuous time models with two volatility factors seem to predict future volatility better than other possible specifications at relevant time horizons and out-of-sample periods.

References

[1] Andersen, T.G., Benzoni, L. and J. Lund (2002): An Empirical Investigation of Continuous-Time Equity Return Models, *Journal of Finance* **57**, 1239-1284.

[2] Andersen, T.G. and T. Bollerslev (1998): Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts.

[3] Andersen, T. G., Bollerslev, T., Diebold, F.X. and P. Labys (2001): The Distribution of Exchange Rate Volatility, *Journal of the American Statistical Association* **96**, 42-55.

[4] Andersen, T.G., Bollerslev, T., Diebold, F.X. and P. Labys (2003): Modelling and Forecasting Realized Volatility, *Econometrica* **71**, 579-625.

[5] Andersen, T.G., Bollerslev, T. and N. Meddahi (2005): Correcting the Errors: Volatility Forescast Evaluation Using High-Frequency Data and Realized Volatilities, *Econometrica* **73**, 279-296.

[6] Andersen, T.G., Bollerslev, T., Christoffersen, P.F. and F.X. Diebold (2005b): Volatility Forecasting, *NBER Working Paper* no 11188.

[7] Andreou, E. and E. Ghysels (2002): Rolling-Sample Volatility Estimators: Some New Theoretical, Simulation, and Empirical Results, *Journal of Business and Economic Statistics* **20**, 363-376.

[8] Bai, C., Russel, J.R. and G.C. Tiao (2004): Effects of Non-Normality and Dependence on the Precision of Variance Estimates Using High-Frequency Financial Data, *Working Paper*, University of Chicago .

[9] Barndorff-Nielsen, O.E. and N. Shephard (2001): Non-Gaussian Ornstein-Uhlenbeck Based Models and Some of Their Uses in Financial Economics, *Journal of the Royal Statistical Society, Ser. B* **63**, 167-207.

[10] Barndorff-Nielsen, O.E. and N. Shephard (2002a): Econometric Analysis of Realized Volatility and Its Use in Estimating Stochastic Volatility Models, *Journal of the Royal Statistical Society, Ser. B* **64**, 253-280.

[11] Barndorff-Nielsen, O.E. and N. Shephard (2002b): Estimating Quadratic Variation Using Realized Variance, *Journal of Applied Econometrics* **17**, 457-477.

[12] Barndorff-Nielsen, O.E. and N. Shephard (2004a): Econometric Analysis of Realized Covariation: High Frequency Based Covariance, Regression, and Correlation in Financial Economics, *Econometrica* **72**, 885-925.

[13] Barndorff-Nielsen, O.E. and N. Shephard (2004b): Power and Bipower Variation with Stochastic Volatility and Jumps, *Journal of Financial Econometrics* **2**, 1-48.

[14] Canina, L. and S. Figlewski (1993): The Informational Content of Implied Volatility, *Review of Financial Studies* **6**, 659-681.

[15] Chernov, M., Gallant, A. R., Ghysels, E. and G. Tauchen (2003): Alternative Models for Stock Price Dynamics, *Journal of Econometrics* **116**, 225-257.

[16] Chernov, M. and E. Ghysels (2000): A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation, *Journal of Financial Economics* **56**, 407-458.

[17] Comte, F. and E. Renault (1998): Long Memory in Continuous Time Stochastic Volatility Models, *Mathematical Finance* **8**, 291-323.

[18] Davidson, J. (2004): Moment and Memory Properties of Linear Conditional Heteroscedasticity Models, and a New Model, *Journal of Business and Economic Statistics* **22**, 16-29.

[19] Diebold, F.X. and R.S. Mariano (2002): Comparing Predictive Accuracy, *Journal of Business and Economic Statistics* **20**, 134-144.

[20] Ederington, L.H. and W. Guan (2004): Forecasting Volatility, *Working Paper*, SSRN.

[21] Engle, R.F. and Gary G.J. Lee (1993): A Permament and Transitory Component Model of Stock Return Volatility, *Working Paper* **92-44R**, UCSD.

[22] Eraker, B., Johannes, M. and N. Polson (2003): The Impact of Jumps in Returns and Volatility, *Journal of Finance* **53**, 1269-1300.

[23] Gallant, A. R. and G. Tauchen (1996): Which Moments to Match?, *Econometric Theory* **12**, 657- 681.

[24] Gallant, A. R. and G. Tauchen (1998): Reprojecting Partially Observed Systems with Application to Interest Rate Diffusions, *Journal of American Statistical Association* **93**, 10-24.

[25] Gallant, A.R. and G. Tauchen (2001): Efficient Method of Moments, *Working Paper*, University of North Carolina at Chapel Hill.

[26] González-Rivera, G., Lee, T.H. and S. Mishra (2004): Forecasting Volatility: A Reality Check Based on Option Pricing, Utility Function, Value-at-Risk, and Predictive Likelihood, *International Journal of Forecasting* **20**, 624-645.

[27] Granger, C.W.J. and R. Joyeux (1980): An Introduction to Long Memory Time Series Models and Fractional Differencing, *Journal of Time Series Analysis* **1**, 15-29.

[28] Hansen, P.R. and A. Lunde (2005a): A Forecast Comparison of Volatility Models: Does Anything Beat a GARCH(1,1)?, *Journal of Applied Econometrics*, forthcoming.

[29] Hansen, P.R. and A. Lunde (2005b): A Realized Variance for the Whole Day Based on Intermittent High-Frequency Data, *Journal of Financial Econometrics*, forthcoming.

[30] Hansen, P.R. and A. Lunde (2005c): Realized Variance and Market Microstructure Noise, *Working Paper*, SSRN.

[31] Harvey, D., Leybourne, S. P. Newbold (1997): Testing the Equality of Prediction Mean Squared Errors, *International Journal of Forecasting* **13**, 281-291.

[32] Jones, C. (2003): The Dynamics of Stochastic Volatility, *Journal of Econometrics* **116**, 181-224.

[33] Laurent, S. and J.P. Peters (2005): G@RCH 4.0, Estimating and Forecasting ARCH Models, Timberlake Consultants.

[34] Merton, R.C. (1980): On Estimating the Expected Return on the Market: An Exploratory Investigation, *Journal of Financial Economics* **8**, 323-361.

[35] Nelson, D. B. (1991): Conditional Heteroskedasticity in Asset Returns: A New Approach, *Econometrica* **59**, 347–370.

[36] Oomen, R.A.C. (2002): Modelling Realized Variance When Returns Are Serially Correlated, *Working Paper*, Warwick Business School, The University of Warwick.

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

[38] Pan, J. (2002): The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study, *Journal of Financial Economics* **63**, 3-50.

[39] Parzen, E. (1981): Time Series Model Identification and Prediction Variance Horizon, in Applied Time Series Analysis II, ed. Findley, Academic Press, 415-447, New York.

[40] West, K. D. (1996): Asymptotic Inference About Predictive Ability, *Econometrica* **64**, 1067-1084.

[41] West, K. and D. Cho (1995): The Predictive Ability of Several Models of Exchange Rate Volatility, *Journal of Econometrics* **69**, 367-391.

[42] Zhang, L. (2005): Efficient Estimation of Stochastic Volatility Using Noisy Observations: A Multi-Scale Approach, *Working Paper*, Carnegie Mellon University.

[43] Zhang, L., Mykland, P. and Y. Aït-Sahalia (2005): A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High Frequency Data, *Journal of the American Statistical Association* **100**, 1394-1411.

[44] Zhou, B. (1996): High-Frequency Data and Volatility in Foreign-Exchange Rates, *Journal of Business and Economic Statistics* **14**, 45-52.

Tables and Figures

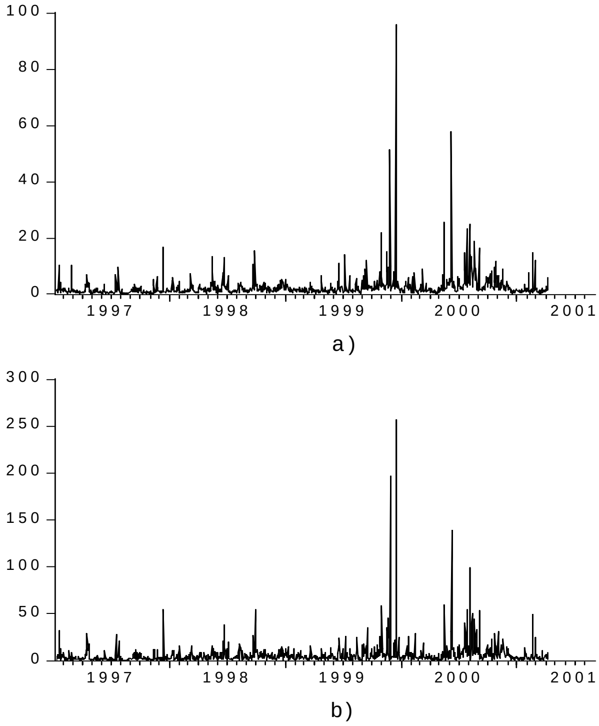

Figure 1. High-Frequency Volatilities (%): a) Realized Volatility obtained by summing up the intra-day squared returns and b) Kernel-estimator of Hansen and Lunde (2005c).

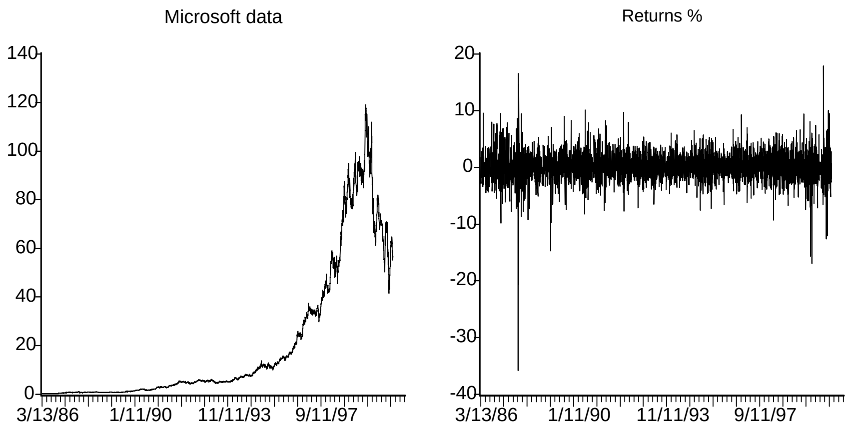

Figure 2. Evolution of the Microsoft stock price and returns.

Model	α_{10}	α_{22}	α_{33}	β_{10}	β_{12}	β_{13}	ψ_{12}	ψ_{13}	N	χ^2	df	p-val
SV2F	*	*	*	*	*	*			100k	6.70	5	0.24
Asym. SV2F	*	*	*	*	*	*	*	*	100k	6.83	3	0.08

Table 1: *is used for free parameters. 100k refers to a simulation of length 100 000 at step size $\Delta = 1/6048$, corresponding to 24 steps per day and 252 trading days per year.

SV2F	α_{10}	α_{22}	α_{33}	β_{10}	β_{12}	β_{13}	ψ_{12}	ψ_{13}
Estimate	0.424	-0.00028	-89.21	-0.110	0.006	-4.628		
Std. Dev.	0.074	0.00015	3.933	0.009	0.001	0.076		
95%Lower	0.269	-0.00049	-97.15	-0.123	0.004	-4.778		
95% Upper	0.579	-0.00008	-81.43	-0.097	0.008	-4.480		
Asym. SV2F								
Estimate	0.344	-0.179	-88.63	-0.054	0.160	-5.545	-0.217	-0.137
Std. Dev.	0.088	0.104	5.409	0.007	0.030	0.231	0.021	0.062
95%Lower	0.337	-0.180	-88.66	-0.054	0.159	-5.545	-0.2174	-0.1371
95% Upper	0.346	-0.179	-88.63	-0.054	0.160	-5.543	-0.2173	-0.1369

Table 2: Estimates, standard deviations and confidence intervals.

Models	*const* (*mean*)	*const* (*var*)	α	β	d	γ_1	γ_2	$hy = \ln(\alpha)$
GARCH	0.186	0.267	0.096	0.864				
Std. Dev.	0.037	0.123	0.034	0.046				
FIGARCH	0.181	7.687	0.098	0.226	0.275			
Std. Dev	0.035	1.415	0.180	0.195	0.036			
HYGARCH	0.180	0.922	0.115	0.252	0.300			- 0.060
Std. Dev.	0.037	0.473	0.275	0.313	0.138			0.169
FIEGARCH	0.142	2.344	0.113	0.360	0.485	- 0.051	0.259	
Std. Dev	0.041	0.261	0.486	0.242	0.066	0.037	0.103	

Table 3: Estimates and standard deviations.

Models	$const$ $(mean)$	$const$ (var)	α	β	d	γ_1	γ_2	$hy = \ln(\alpha)$
GARCH	0.186	0.267	0.096	0.864				
Std. Dev.	0.037	0.123	0.034	0.046				
FIGARCH	0.181	7.582		0.121	0.267			
Std. Dev	0.035	1.339		0.038	0.031			
HYGARCH	0.180	0.922			0.243			
Std. Dev.	0.037	0.242			0.075			
FIEGARCH	0.175	2.410		0.372	0.515		0.265	
Std. Dev	0.038	0.304		0.199	0.067		0.094	

Table 3.1: Final estimates and standard deviations.

Component Model	ϖ	ρ	ϕ	α	β
Estimates	5.676	0.976	0.033	0.129	0.520
Std. Dev.	0.300	0.005	0.007	0.013	0.068

Table 4: Estimates and standard deviations.

Dependent Variable $(RV)^{1/2}$ T=252	Est.	Std. Error	T-statistic	Prob	R^2	R^{*2}	$\frac{MSFE}{MSFE_{Garch}}$
GARCH					0.030	0.041	
β_0	0.739	0.724	1.020	0.308			
β_1	0.625	0.312	2.006	0.046			
FIGARCH					0.055	0.076	0.97
β_0	-0.213	0.959	-0.222	0.825			
β_1	1.025	0.407	2.252	0.012			
CModel					0.005	0.007	1.06
β_0	1.524	1.014	1.503	0.134			
β_1	0.286	0.432	0.662	0.508			
Dependent Variable $(Kernel)^{1/2}$ T=247							
GARCH					0.041		
β_0	0.257	0.809	0.318	0.751			
β_1	0.819	0.350	2.342	0.020			
FIGARCH					0.068		0.98
β_0	-0.855	1.114	-0.716	0.444			
β_1	1.285	0.474	2.711	0.007			
CModel					0.004		1.10
β_0	1.549	1.433	1.081	0.281			
β_1	0.264	0.612	0.431	0.667			

Table 5: First step forecasting evaluation (10 days-ahead forecasts). We report OLS variance estimates robust to autocorrelation (Newey-West HAC Standard Errors & Covariance).

Dependent Variable $ln\,(RV)$ T=252	Est.	Std. Error	T-statistic	Prob	R^2	R^{*2}	$\frac{MSFE}{MSFE_{Garch}}$
GARCH					0.033	0.043	
β_0	0.384	0.373	1.030	0.304			
β_1	0.637	0.219	2.912	0.004			
FIGARCH					0.059	0.077	1.02
β_0	-0.327	0.452	-0.722	0.471			
β_1	1.048	0.264	3.974	0.000			
CModel					0.005	0.007	1.06
β_0	0.982	0.447	2.197	0.029			
β_1	0.280	0.258	1.084	0.280			
Dependent Variable $ln\,(Kernel)$ T=247							
GARCH					0.048		
β_0	-0.390	0.506	-0.772	0.441			
β_1	1.039	0.297	3.499	0.001			
FIGARCH					0.068		1.00
β_0	-1.223	0.756	-1.619	0.107			
β_1	1.517	0.432	3.512	0.001			
CModel					0.0004		1.09
β_0	1.183	1.112	1.064	0.288			
β_1	0.108	0.656	0.164	0.870			

Table 5.1: First step forecasting evaluation (10 days-ahead forecasts). We report OLS variance estimates robust to autocorrelation (Newey-West HAC Standard Errors & Covariance).

Dependent Variable	$(RV)^{1/2}$	$ln(RV)$
$MSFE_{Figarch}/MSFE_{Garch}$	0.06	0.28
$MSFE_{CModel}/MSFE_{Garch}$	0.22	0.20
Dependent Variable	$(Kernel)^{1/2}$	$ln(Kernel)$
$MSFE_{Figarch}/MSFE_{Garch}$	0.14	0.50
$MSFE_{CModel}/MSFE_{Garch}$	0.18	0.11

Table 6: This Table provides the p-values from testing the null hypothese that FIGARCH and CModel have similar MSFE than the GARCH. A low p-value indicates that forecasts from the corresponding models would not be rejected in favor of GARCH forecasts. The test statistic is the $S_1$of Diebold and Mariano (2002).

Dependent Variable $(RV)^{1/2}$ T=288	Est.	Std. Error	T-statistic	Prob	R^2	R^{*2}	$\frac{MSFE}{MSFE_{Garch}}$
GARCH					0.060	0.069	
β_0	1.697	0.381	4.452	0.000			
β_1	0.414	0.119	3.490	0.001			
FIGARCH					0.160	0.185	0.822
β_0	0.768	0.482	1.592	0.113			
β_1	0.705	0.153	4.597	0.000			
CModel					0.068	0.079	0.962
β_0	1.567	0.387	4.051	0.000			
β_1	0.462	0.118	3.911	0.000			
Dependent Variable $(Kernel)^{1/2}$ T=279							
GARCH					0.043		
β_0	1.621	0.350	4.629	0.000			
β_1	0.378	0.107	3.526	0.001			
FIGARCH					0.124		0.867
β_0	0.706	0.463	1.523	0.129			
β_1	0.665	0.146	4.551	0.000			
CModel					0.050		0.961
β_0	1.507	0.361	4.180	0.000			
β_1	0.420	0.108	3.888	0.000			

Table 7: First step forecasting evaluation (1 day-ahead forecasts). we report the OLS variance estimates robust to autocorrelation (Newey-West HAC Standard Errors & Co-variance).

Dependent Variable $ln\,(RV)$ T=288	Est.	Std. Error	T-statistic	Prob	R^2	R^{*2}	$\frac{MSFE}{MSFE_{Garch}}$
GARCH					0.123	0.139	
β_0	0.844	0.298	2.835	0.005			
β_1	0.517	0.129	4.020	0.000			
FIGARCH					0.210	0.237	0.854
β_0	0.313	0.311	1.006	0.315			
β_1	0.750	0.132	5.668	0.000			
CModel					0.142	0.161	0.932
β_0	0.690	0.298	2.320	0.021			
β_1	0.594	0.128	4.631	0.000			
Dependent Variable $ln\,(Kernel)$ T=279							
GARCH					0.081		
β_0	0.705	0.288	2.451	0.015			
β_1	0.490	0.126	3.886	0.000			
FIGARCH					0.155		0.912
β_0	0.123	0.307	0.400	0.690			
β_1	0.746	0.132	5.647	0.000			
CModel					0.099		0.943
β_0	0.535	0.286	1.873	0.062			
β_1	0.574	0.124	4.610	0.000			

Table 7.1: First step forecasting evaluation (1 day-ahead forecasts). We report OLS variance estimates robust to autocorrelation (Newey-West HAC Standard Errors & Covariance).

Dependent Variable	$(RV)^{1/2}$	$ln(RV)$
$MSFE_{Figarch}/MSFE_{Garch}$	0.023	0.003
$MSFE_{CModel}/MSFE_{Garch}$	0.024	0.001
Dependent Variable	$(Kernel)^{1/2}$	$ln(Kernel)$
$MSFE_{Figarch}/MSFE_{Garch}$	0.024	0.009
$MSFE_{CModel}/MSFE_{Garch}$	0.020	0.001

Table 8: This table provides the p-values from testing the null hypothese that FIGARCH and CModel have similar MSFE than the GARCH. A low p-value indicates that forecasts from GARCH would be rejected in favor of the alternative models forecasts. The test statistic is the $S_1$of Diebold and Mariano (2002).

Dependent Variable $(RV)^{1/2}$ n=243	Est.	Std. Error	T-statistic	Prob	R^2	R^{*2}
GARCH					0.0003	0.0004
β_0	2.223	0.150	14.85	0.000		
β_1	-0.017	0.061	-0.271	0.787		
FIGARCH					0.0003	0.0004
β_0	2.223	0.150	14.85	0.000		
β_1	-0.012	0.046	-0.267	0.791		
SV2F					0.011	0.015
β_0	1.934	0.218	8.862	0.000		
β_1	0.088	0.073	1.207	0.229		
Dependent Variable $(Kernel)^{1/2}$ n=243						
GARCH					0.0014	
β_0	2.066	0.159	13.00	0.000		
β_1	0.041	0.069	0.594	0.223		
FIGARCH					0.0014	
β_0	2.065	0.163	12.691	0.000		
β_1	0.031	0.055	0.566	0.572		
SV2F					0.034	
β_0	1.660	0.238	6.967	0.000		
β_1	0.174	0.081	2.151	0.032		

Table 9: Second step forecasting evaluation (10 days rolling forecasts). We report OLS variance estimates robust to autocorrelation (Newey-West HAC Standard Errors & Covariance).

Dependent Variable $ln(RV)$ n=243	Est.	Std. Error	T-statistic	Prob	R^2	R^{*2}
GARCH					0.000	0.000
β_0	1.449	0.128	11.32	0.000		
β_1	-0.001	0.077	-0.009	0.993		
FIGARCH					0.000	0.000
β_0	1.449	0.167	8.664	0.000		
β_1	-0.0001	0.077	-0.002	0.999		
SV2F					0.021	0.027
β_0	1.187	0.165	7.174	0.000		
β_1	0.132	0.076	1.733	0.084		
Dependent Variable $ln(Kernel)$ n=238						
GARCH					0.002	
β_0	1.238	0.146	8.461	0.000		
β_1	0.075	0.086	0.871	0.385		
FIGARCH					0.003	
β_0	1.194	0.215	5.547	0.000		
β_1	0.076	0.098	0.775	0.439		
SV2F					0.050	
β_0	0.816	0.178	4.008	0.000		
β_1	0.272	0.091	2.987	0.003		

Table 9.1: Second step forecasting evaluation (10 days rolling forecasts). We report OLS variance estimates robust to autocorrelation (Newey-West HAC Standard Errors & Covariance).

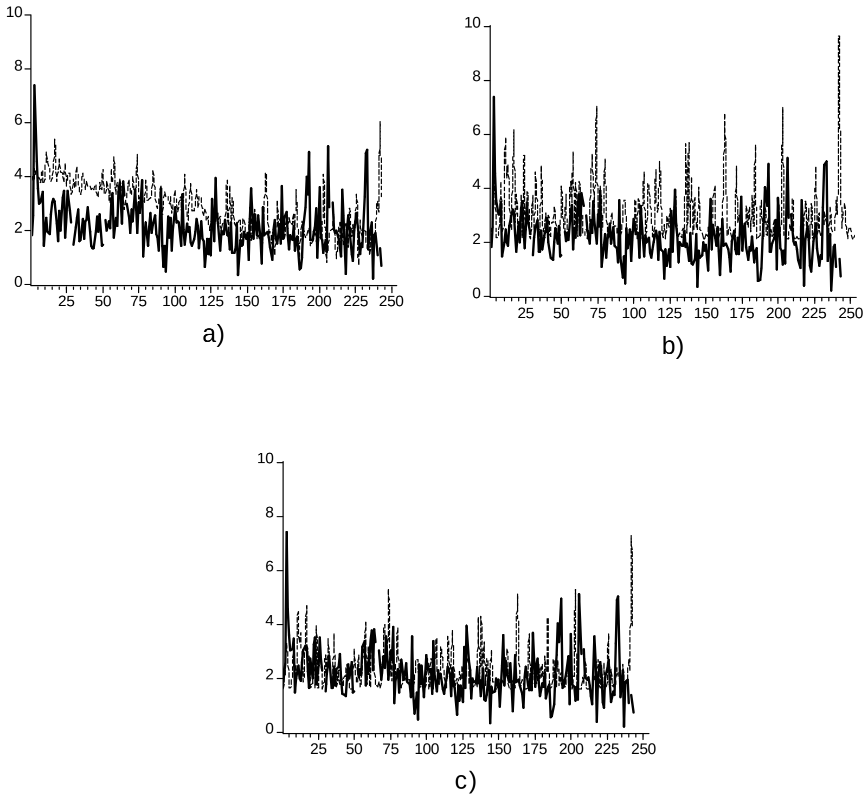

Figure 3: Comparison between the square roots of Kernel-estimator and volatility forecasts of the alternative models. Panel a) $\sqrt{Kernel}$ and $\sqrt{SV2F}$, Panel b) $\sqrt{Kernel}$ and $\sqrt{FIGARCH}$ and Panel c) $\sqrt{Kernel}$ and $\sqrt{GARCH}$. The continuous line corresponds to the Kernel and the dotted lines to volatility forecasts of the alternative models.

Dependent Variable $(RV)^{1/2}$ n=288	Est.	Std. Error	T-statistic	Prob	R^2	R^{*2}
FIGARCH					0.027	0.031
β_0	2.499	0.230	10.87	0.000		
β_1	0.128	0.04	3.019	0.003		
CModel					0.016	0.018
β_0	2.667	0.236	11.291	0.000		
β_1	0.120	0.061	1.971	1.971		
SV2F					0.023	0.027
β_0	2.727	0.203	13.41	0.000		
β_1	0.115	0.038	3.00	0.003		
Dependent Variable $(Kestimator)^{1/2}$ n=279						
FIGARCH					0.018	
β_0	2.376	0.227	10.45	0.000		
β_1	0.114	0.042	2.700	0.007		
CModel					0.014	
β_0	2.467	0.215	11.48	0.000		
β_1	0.125	0.055	2.274	0.024		
SV2F					0.020	
β_0	2.524	0.212	11.91	0.000		
β_1	0.115	0.042	2.743	0.007		

Table 10: Second step forecasting evaluation (1 day rolling forecasts). We report OLS variance estimates robust to autocorrelation (Newey-West HAC Standard Errors & Covariance).

Dependent Variable $ln(RV)$ n=288	Estimates	Std. Error	T-statistic	Prob	R^2	R^{*2}
FIGARCH					0.038	0.043
β_0	1.481	0.191	7.763	0.000		
β_1	0.200	0.065	3.076	0.002		
CModel					0.052	0.058
β_0	1.603	0.162	9.911	0.000		
β_1	0.218	0.066	3.292	0.001		
SV2F					0.040	0.045
β_0	1.790	0.113	15.81	0.000		
β_1	0.133	0.039	3.378	0.001		
Dependent Variable $ln(Kernel)$ n=279						
FIGARCH					0.028	
β_0	1.298	0.211	6.149	0.000		
β_1	0.196	0.075	2.606	0.010		
CModel					0.038	
β_0	1.407	0.162	8.702	0.000		
β_1	0.217	0.070	3.092	0.002		
SV2F					0.050	
β_0	1.498	0.117	12.76	0.000		
β_1	0.175	0.044	3.995	0.000		

Table 10.1: Second step forecasting evaluation (1 day rolling forecasts). We report OLS variance estimates robust to autocorrelation (Newey-West HAC Standard Errors & Covariance).

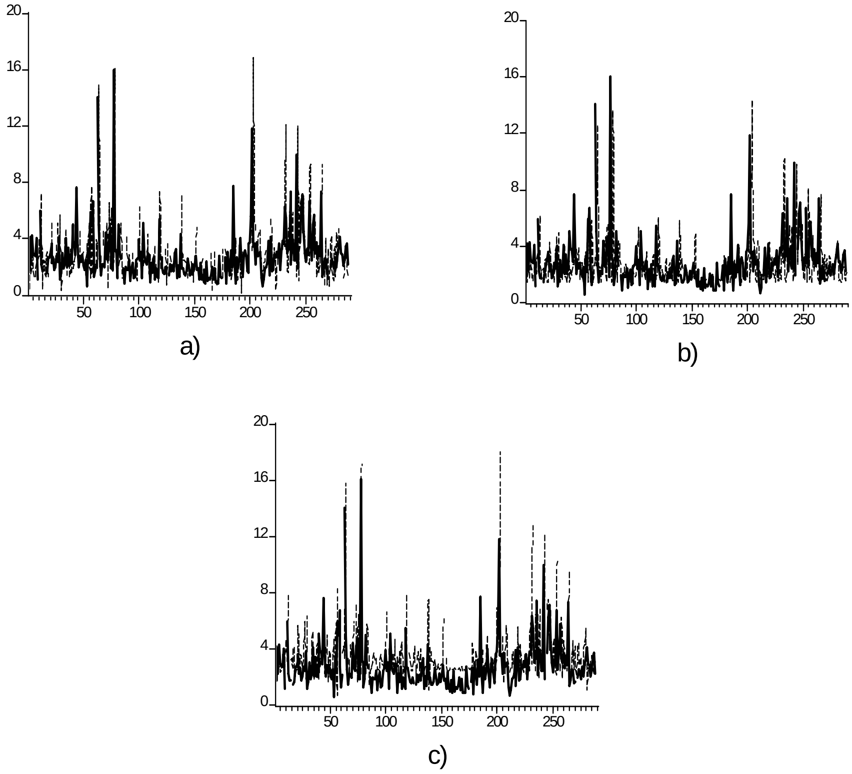

Figure 4: Comparison between the square roots of Kernel-estimator and volatility forecasts of the alternative models. Panel a) $\sqrt{Kernel}$ and $\sqrt{SV2F}$, Panel b) $\sqrt{Kernel}$ and $\sqrt{CModel}$ and Panel c) $\sqrt{Kernel}$ and $\sqrt{FIGARCH}$. The continuous line corresponds to the Kernel and the dotted lines to volatility forecasts of the alternative models.