Dipartimento di Statistica
"Giuseppe Parenti"

WORKING PAPER 2007/01

On the Interaction between Ultra–high Frequency Measures of Volatility

Giampiero M. Gallo,
Margherita Velucchi

Università degli Studi
di Firenze

Dipartimento di Statistica "G. Parenti" – Viale Morgagni 59 – 50134 Firenze - www.ds.unifi.it

On the Interaction between Ultra–high Frequency Measures of Volatility

Giampiero M. Gallo Margherita Velucchi*

This version: May 2007

Abstract

We analyze several measures of volatility (realized variance, bipower variation and squared daily returns) as estimators of integrated variance of a continuous time stochastic process for an asset price. We use a Multiplicative Error Model to describe the evolution of each measure as the product of its conditional expectation and a positive valued iid innovation. By inserting past values of each measure and asymmetric effects based on the sign of the return in the specification of the conditional expectation, one can investigate the information content of each indicator relative to the others. The results show that there is a directed dynamic relationship among measures, with squared returns and bipower variance interdependent with one another, and affecting realized variance without any feed–back from the latter.

Keywords: Volatility, Multiplicative Error Models, Realized Variance, Bi-power Variance, Squared Returns, Jumps.

*Dipartimento di Statistica "G.Parenti", Università di Firenze, Viale G.B. Morgagni, 59, 50134 Firenze – Italy. gallog@ds.unifi.it, velucchi@ds.unifi.it. Thanks are due to Christian T. Brownlees for the massive effort he profused in programming an amazingly efficient utility for extracting data from TAQ databases and to Marco J. Lombardi for useful discussions on realized volatility. Comments by participants in the 2005 S.Co. Conference in Bressanone, especially the discussant Cosimo Vitale, and by Loriano Mancini helped improve the motivation of the paper. The lectures by Nour Meddahi at the 2006 CIDE Summer School of Econometrics provided a clear assessment of the existing literature on realized volatility. All remaining errors are ours.

1 Introduction

In the past few years, the availability of financial data at a very high frequency has allowed innovative measuring and modeling of many market phenomena, among which volatility is one of the most investigated. "In financial economics, volatility is often defined as the (instantaneous) standard deviation (or "sigma") of the random Wiener-driven component in a continuous-time diffusion model" (Andersen *et al.*, 2005): many authors have addressed the issue of consistently estimating a discrete time transformation of this (unobservable) continuous–time object. The most notable set of contributions is a series of papers by Anderson, Bollerslev, Diebold and coauthors (e.g. Anderson *et al.*, 2003) who exploit tick–by–tick (irregularly spaced) data to reconstruct intra-daily series of transaction prices at regular intervals and then derive a *realized* variance series as an estimator of integrated variance. Five minutes seems to be the frequency of choice, as a rule of thumb trade–off between getting close to continuous time and incurring in so-called market microstructure problems (liquidity, bid–ask bounce, etc.) which would distort the overall picture (cf. Hansen and Lunde, 2006, for a discussion of the treatment of microstructure noise in this context).

More recent research (Barndorff–Nielsen and Shephard, 2002, 2004, 2006) pointed out that in the presence of jumps the realized variance is affected by a non–vanishing bias. Squared returns have been long recognized to be a noisy measure of integrated variance, although they seem to possess some specific information (possibly connected with overnight behavior) which is not contained in intra-daily data derived measures (Gallo, 2001).

In this paper, we adopt a Multiplicative Error Model as employed by Engle and Gallo (2006) in order to analyze the interaction between three different indicators, namely squared returns, realized variance and bi-power variance. The goal is to detect which significant lagged relationships there exists, if any, in the specification for the conditional expectation of each indicator. The results show that the lagged squared returns enter significantly and with approximately the same coefficient size each specification when they are associated with the sign of past returns. Bi-power variance is also significant everywhere but in a way that is not differentiated according to whether past returns are positive or negative. The role of jumps is present only in the model for realized variance and in a separate form from bi–power variance.

The paper is organized as follows: section 2 summarizes the main results of the theoretical debate on high-frequency measures of volatility, establishing some notation as well. Section 3 is devoted to the specification of the Multiplicative Error Model when the conditional expectation term is augmented with predetermined variables. Section 4 contains the main estimation results and a comparison exercise based on Mincer–Zarnowitz regressions of the indicators on the forecasts produced by the MEMs. Concluding remarks follow.

2 Some Theoretical Background

Various arguments can justify the interest in the high–frequency based measures of volatility[1]. Andersen and Bollerslev (1998) pointed out that squared daily returns is a noisy measure of variation: with simulation arguments they show that Mincer–Zarnowitz type regressions of squared returns on any conditional variance forecast would produce a very low R^2. Given that volatility or variance of returns is not observed, it has to be substituted with a proxy whose measurement error should vanish under certain conditions. One solution suggested is to refer to the availability of ultra–high frequency data on returns and to compute a variable called *realized variance*, constructed as

$$rv_t(\tau) = \sum_{i=1}^{1/\tau} r^2_{t-1+i\tau} \tag{1}$$

where the generic term $r_{t-1+i\tau}$ is the return measured intra–daily as the price variation of an asset over a (very small) period τ, that is,

$$r_{t-1+i\tau} = \log(P_{t-1+i\tau}) - \log(P_{t-1+(i-1)\tau}) \equiv p_{t-1+i\tau} - p_{t-1+(i-1)\tau}.$$

The interval τ can be chosen as to have its reciprocal be an integer value, representing the number of intradaily time intervals considered during the day. When $\tau = 1$ we get squared returns back (but we will need to discuss the role of overnight returns); common choices are fractions of the trading day corresponding to five minutes or thirty minutes intervals. Using this variable as the left-hand side variable to evaluate forecasts, Andersen and Bollerslev showed that the R^2 of the Mincer–Zarnowitz regression would increase as τ decreases.

[1]This section is based on various survey papers by Andersen, Bollerslev and Diebold included in the references.

The theoretical support for such a result stems from the following considerations. Assuming that the (log of) asset prices evolve according to a stochastic differential equation (cf. Andersen, Bollerslev and Diebold, 2005)

$$dp(t) = \mu(t)dt + \sigma(t)dW(t) \tag{2}$$

where $\mu(t)$ denotes the drift term, $\sigma(t)$ is the instantaneous (or spot) volatility, $dW(t)$ is a standard Brownian motion.[2]

Over an infinitesimal discrete time interval τ, returns can be written as

$$r(t,\tau) \equiv p(t) - p(t-\tau) \approx \mu(t-\tau)\tau + \sigma(t-\tau)\Delta W(t)$$

where $\Delta W(t) = W(t) - W(t-\tau) \sim N(0,\tau)$. Over a day, we would then have

$$r(t) = p(t) - p(t-1) = \int_{t-1}^{t} \mu(s)ds + \int_{t-1}^{t} \sigma(s)dW(s);$$

if the volatility process $\sigma(t)$ is independent of the Wiener process, $W(t)$, then the distribution of the returns at time t, conditional on the path of the drift process $\mu(s)$ and the volatility process $\sigma(s)$ between $t-1$ and t, is actually Gaussian with variance

$$\int_{t-1}^{t} \sigma^2(s)ds \equiv iv_t \tag{3}$$

which is called the *integrated variance*. It has an interest in itself since it is at the basis of the most important derivative pricing formulas with stochastic variance. The link with existing discrete time counterparts is straightforward: the variance of the returns r_{t+1} measured in discrete time conditional on the information set at time t, I_t is actually

$$Var(r_{t+1}|I_t) = E\left[\int_{t-1}^{t} \sigma^2(s)ds\,|I_t\right] = E\left[iv_{t+1}|I_t\right].$$

This justifies using the integrated variance on the left–hand side of a Mincer–Zarnowitz regression where the regressor is the estimated conditional variance from some model. By the same token, since iv_{t+1} is a latent variable, an observable counterpart is needed. The theory states that, in the absence of jumps, the realized variance $rv_{t+1}(\tau)$ converges uniformly to the integrated variance as τ tends to zero.

[2]Some continuity conditions are required on $\mu(t)$ and $\sigma(t)$ in order to exploit the martingale properties of the resulting process, but they are well documented elsewhere – Andersen, Bollerslev and Diebold (2007) and will not be recalled here.

Empirical illustration of some of the data features can be provided in reference to a highly capitalized stock, General Electric (GE), whose transaction data have been extracted from the NYSE TAQ database between January 4, 1995 and December 29, 2000, following a procedure described in Brownlees and Gallo (2005). The tick–by–tick (irregularly spaced) transaction price data for a randomly chosen day in the sample of GE data are depicted in the top–left panel of Figure 1: note the highly irregular profile with frequent spikes in the data. The other panels show (left to right; top to bottom) prices sampled at regular intervals τ every 1 minute, every 5 minutes and every 10 minutes; the corresponding returns are reported in Figure 2.

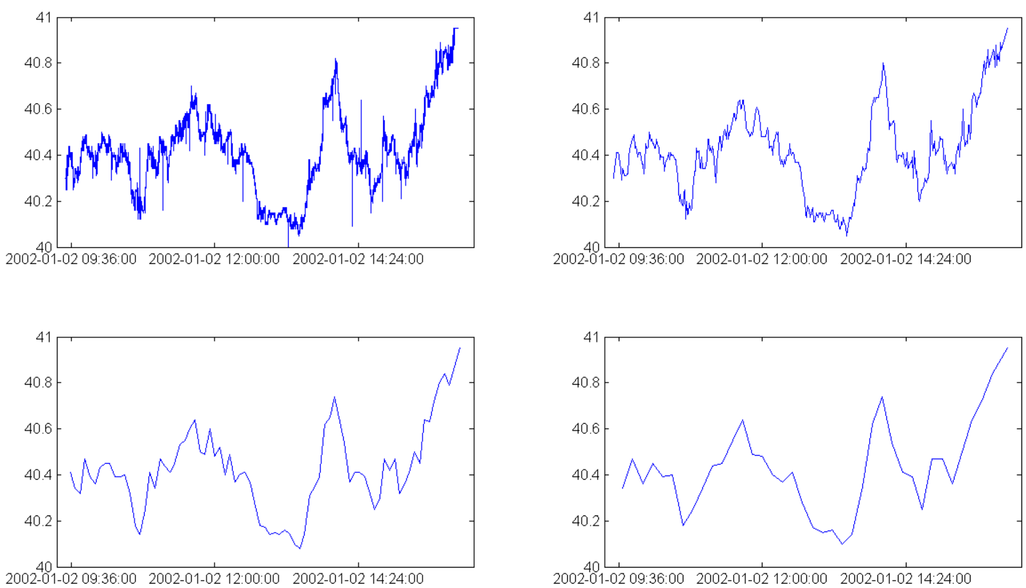

Figure 1: Time series profile of transaction prices during one randomly chosen day for GE – different sampling frequency. Tick–by–tick (top left); 1 minute (top right); 5 minutes (bottom left); 10 minutes (bottom right).

The presence of jumps in the process $dp(t)$ somewhat modifies the picture. Equation 2 is extended to include tractable jumps, namely,

$$dp(t) = \mu(t)dt + \sigma(t)dW(t) + \kappa(t)dq(t) \tag{4}$$

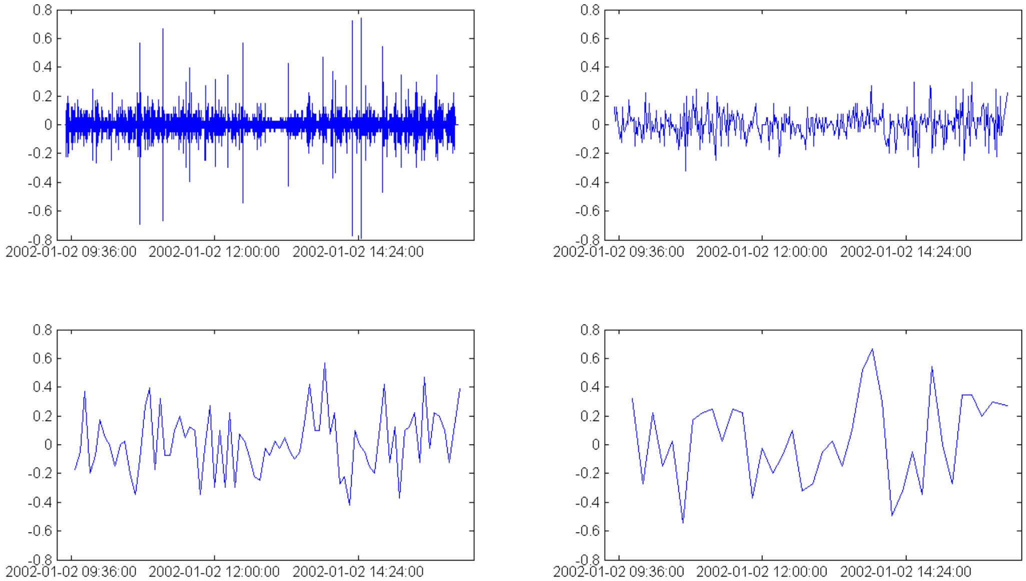

Figure 2: Time series profile of returns during one randomly chosen day for GE - different sampling frequency. Tick–by–tick (top left); 1 minute (top right); 5 minutes (bottom left); 10 minutes (bottom right).

where $q(t)$ is a process keeping track of how many jumps have occurred at time t, and $\kappa(t)$ indicates the size of the jump at time t when one occurs. Correspondingly, over a day, we would then have

$$r(t) = p(t) - p(t-1) = \int_{t-1}^{t} \mu(s)ds + \int_{t-1}^{t} \sigma(s)dW(s) + \sum_{t-1 \leq s \leq t} \kappa(s),$$

and the quantity that the realized variance rv_t converges to is not the integrated variance but

$$qv_t \equiv \int_{t-1}^{t} \sigma^2(s)ds + \sum_{t-1 \leq s \leq t} \kappa^2(s).$$

In a series of papers, Barndorff-Nielsen and Shephard (e.g. 2006), suggest a different measure of volatility which is robust under the presence of jumps.

The Bi-power realized variance is defined as

$$bv_t(\tau) = \frac{\pi}{2} \sum_{i=1}^{1/\tau} |r_{t-1+i\tau}||r_{t-1+(i-1)\tau}| \tag{5}$$

which is also function of the intradaily time interval τ. As τ converges to zero, $bv_t(\tau)$ converges uniformly to the integrated variance iv_t.

For illustrative purposes, it is instructive to look at the elements that are used to compute the two measures. In Figure 3 we report the pattern of squared returns and of cross absolute returns for a frequency of five and ten minutes.

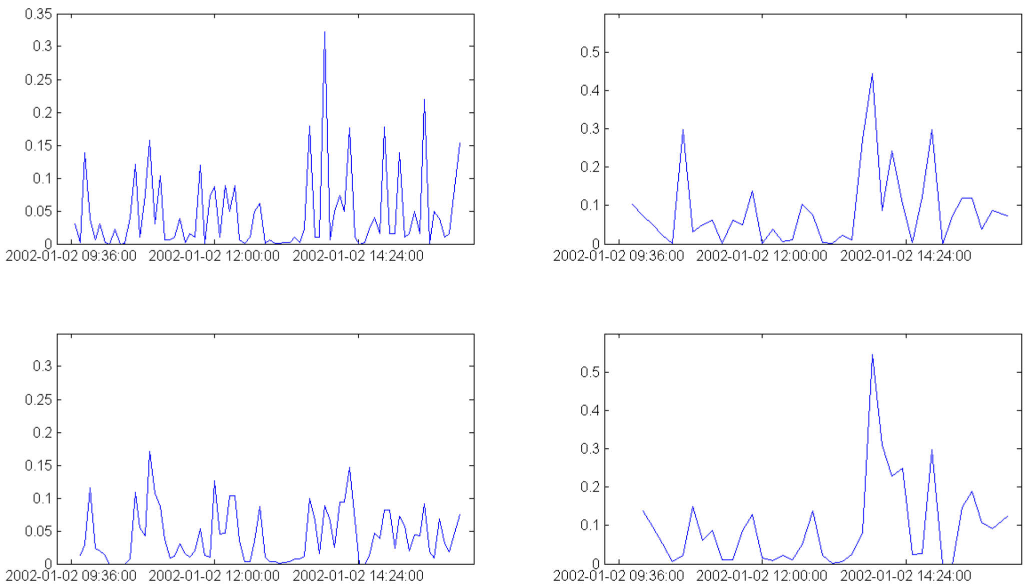

Figure 3: Time series profile of squared returns (top) and cross absolute returns (bottom) during one randomly chosen day for GE. Different sampling frequency: 5 minutes (left) and 10 minutes (right)

The time series of the resulting daily measures are depicted in Figure 4 for the days in Sep. 1999; the bivariate scatterplots (reported in Figure 5) show that the two ultra–high frequency measures of variability are strongly correlated with one another, but much less so with the squared returns.

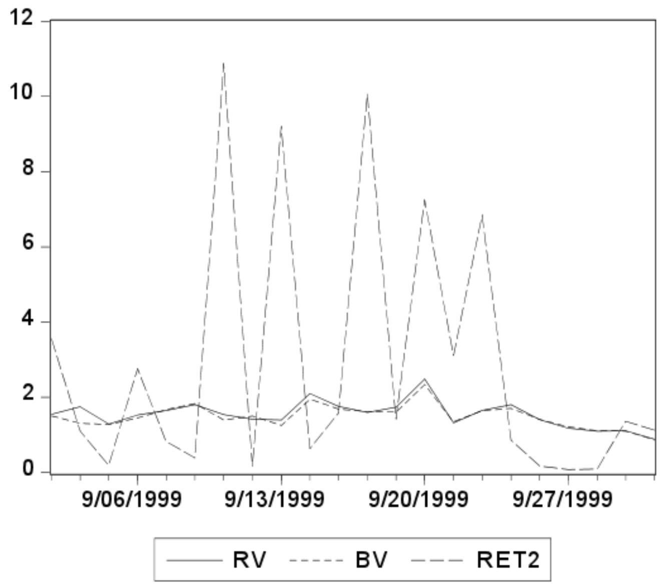

Figure 4: Time series profile of Realized Variance, Bi–power Variation and Squared Returns. Sep. 1999.

What we are interested in is the fact that an estimate of the jumps at time t (Barndorff–Nielsen and Shephard, 2004) can be taken to be the difference between the two measures when it is positive, namely

$$j_t(\tau) \equiv max(0, rv_t(\tau) - bv_t(\tau)). \tag{6}$$

An example of the time profile of jumps is given for the GE stock in Figure 6 (left panel). The series exhibits significant autocorrelation (first–order coefficient around 0.10, Ljung–Box Q(12) test statistic equal to 201.11) and a zeromodal, highly skewed distribution (right panel).

A further dimension of interest is the sampling frequency, as too high of a frequency will produce some biases which are connected with market microstructure (cf. the recent paper by Hansen and Lunde, 2006, with the discussion by several scholars). We will not pursue the issue of optimal sampling here, nor will we reproduce our results for different time intervals. In what follows, we will adopt a five-minute interval as it seems to be a good compromise between the need to have enough sample points within a day and the danger of inserting further biases in the analysis.

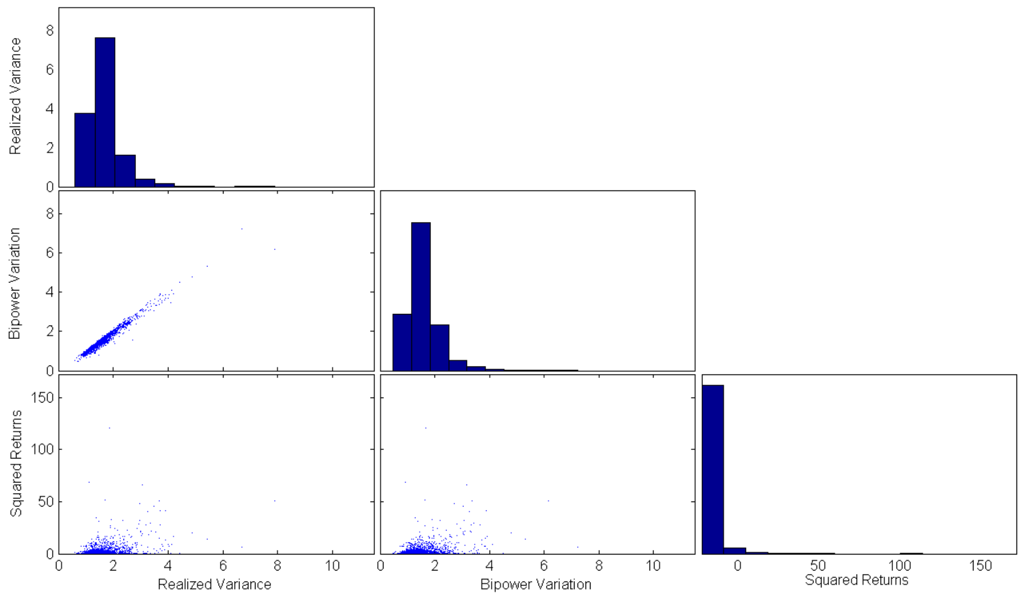

Figure 5: Realized Variance, Bi–power Variation and Squared Returns: Unconditional Distributions and Cross-plots. Jan. 4, 1995 – Dec. 31, 2000.

In concluding this discussion of the theoretical results we have all the ingredients to proceed in our analysis, except for a further look at the daily squared returns. So far, we have assumed that the ultra–high frequency measures of volatility relate to a fine division of the day into intervals of the same size. However, stock markets are closed for a long period during which no (significant) trading occurs. In practice, therefore, researchers prefer to compute both $rv_t(\tau)$ and $bv_t(\tau)$ on the basis of the prices recorded during market activity time. Recorded daily prices, on the contrary, are measured from closing time to closing time, so that, if we define the (log) closing price at time t as c_t, and the (log) opening price[3] at time t as o_t, we have

$$r_t = c_t - c_{t-1} = (c_t - o_t) + (o_t - c_{t-1})$$

that is, the sum of intradaily return and overnight return. Squared daily

[3] In the calculation of $rv_t(\tau)$ and $bv_t(\tau)$ the first few observations of the trading activity of the day may even be excluded. The argument does not change if one moves the first return considered a few minutes into the trading day.

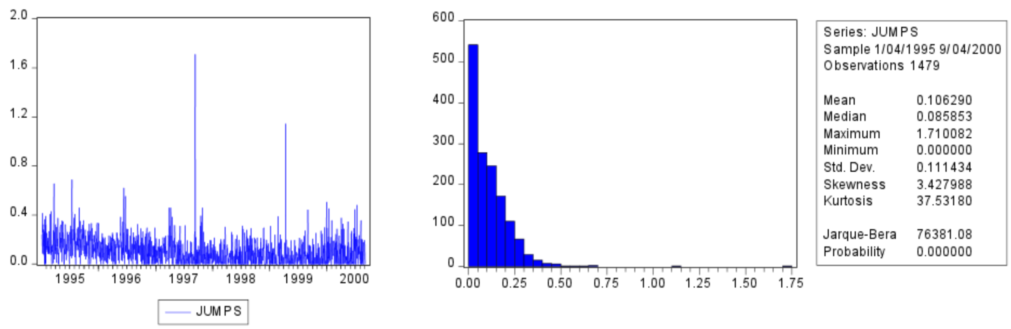

Figure 6: Estimated Jumps. Jan. 4, 1995 – Dec. 31, 2000: Time series and distribution.

returns, therefore, are certainly a noisy measure of integrated variance, as seen above, but they contain some extra information that is not considered in the ultra–high frequency measures of volatility. In fact,

$$\begin{aligned} r_t^2 &= (c_t - o_t)^2 + (o_t - c_{t-1})^2 + 2(c_t - o_t)(o_t - c_{t-1}) \\ &= rv_t(1) + (o_t - c_{t-1})^2 + 2(c_t - o_t)(o_t - c_{t-1}) \end{aligned} \qquad (7)$$

Gallo (2002) shows that in modeling intradaily variance with a GARCH model there may be a significant contribution from overnight innovations which are available in the information set at the beginning of the day. We will return on this aspect in the discussion of the empirical results: for the time being, let us examine the correlations for the GE stock, reported in Table 1. Intradaily returns are clearly more correlated with ultra–high frequency realized measures, especially jumps.

3 The Multiplicative Error Model

We address the issue of analyzing the two ultra–high frequency based measures of volatility in the framework of the Multiplicative Error Model used by Engle and Gallo (2006). Such a model extends the GARCH approach to studying positive valued phenomena (exchanged volumes, durations between trades, daily ranges, etc.): the measures described above fall in this category.

Let us consider a generic indicator as a random variable with positive support $\sigma_{i,t}^2$. The MEM logic is that this variable is the result of the product

Table 1: Correlations between the Variance Indicators: GE. Jan. 4, 1995 – Dec. 31, 2000.

Indicator	rv_t	bv_t	r_t^2	$(c_t - o_t)^2$	$(o_t - c_{t-1})^2$
Bi–power	0.977				
Sq. Returns	0.348	0.336			
Sq. Intraday	0.466	0.418	0.576		
Sq. Overnight	0.142	0.134	0.615	0.086	
Jumps	0.132	−0.073	0.061	0.231	0.040

between a scale factor evolving in a conditionally autoregressive fashion (the expected value) and a innovation term with positive support and unit expectation. In detail, $\sigma^2_{i,t} = \mu_{i,t}\ \varepsilon_{i,t}$, where $\varepsilon_t \sim\ D(1,\varphi)$, i.e. has unit mean and variance φ. As a consequence, the conditional expected value of $\sigma^2_{i,t}$, $\mu_{i,t}$, is specified as

$$\mu_{i,t} = \omega^i + \sum_{j=1}^{p} \alpha^i \sigma^2_{i,t-j} + \sum_{k=1}^{q} \beta^i \mu_{i,t-k} + \mathbf{c}'\ \mathbf{z}_{t-1}.$$

What is of interest in general, and for the case at hand in particular, is the presence in the model of a vector of weakly exogenous variables $\mathbf{z}_{t-1}$ in the information set at time $t-1$ (stationarity conditions are discussed by Engle, 2002).

Since the process $\sigma^2_{i,t}$ represents the behavior of a specific variance indicator at time t it is interesting to see what explanation is added by inserting the values of functions of past indicators among the predetermined variables. By simplifying notation to the case of just one lag on the right–hand side, and splitting lagged observed values into two components corresponding to whether lagged returns were positive or negative, we obtain

$$\begin{aligned} \mu_{i,t} &= \omega^i + \beta^i_1 \mu_{i,t-1} + \delta^{i,i}_1 \sigma^2_{i,t-1} d^+_{t-1} + \psi^{i,i}_1 \sigma^2_{i,t-1} d^-_{t-1} + \\ &\quad + \sum_{s \neq i} (\delta^{i,s}_1 \sigma^2_{s,t-1} d^+_{t-1} + \psi^{i,s}_1 \sigma^2_{s,t-1} d^-_{t-1}). \end{aligned}$$

In the expression, d^+_{t-1} is a dummy variable which assumes value 1 when the stock return at time $t-1$ is positive and $d^-_{t-1} = 1 - d^+_{t-1}$. Such a specification lends itself to estimation and inference using test statistics built with the Bollerslev and Wooldridge (1992) robust coefficient covariance matrices.

Using the short–hand notation $(\sigma^2_{i,t})^+ = \sigma^2_{i,t} d_t^+$, and, analogously, $(\sigma^2_{i,t})^- = \sigma^2_{i,t} d_t^-$, a possible specification is

$$
\begin{aligned}
& \text{Model 1 - Squared returns } r_t^2 = \mu_t^r \epsilon_t^r \\
\mu_t^r \;=\;& \omega^r + \beta_1^r \mu_{t-1}^r + \delta_1^{r,r}(r_{t-1}^2)^+ + \psi_1^{r,r}(r_{t-1}^2)^- + \\
& +\delta_1^{r,v} rv_{t-1}^+ + \psi_1^{r,v} rv_{t-1}^- + \delta_1^{r,b} bv_{t-1}^+ + \psi_1^{r,b} bv_{t-1}^- \\
& \text{Model 2 - Realized variance } rv_t = \mu_t^v \epsilon_t^v \\
\mu_t^v \;=\;& \omega^v + \beta_1^v \mu_{t-1}^v + \delta_1^{v,v} rv_{t-1}^+ + \psi_1^{v,v} rv_{t-1}^- + \\
& +\delta_1^{v,r}(r_{t-1}^2)^+ + \psi_1^{v,r}(r_{t-1}^2)^- + \delta_1^{v,b} bv_{t-1}^+ + \psi_1^{v,b} bv_{t-1}^- \\
& \text{Model 3 - Bi–power variance } bv_t = \mu_t^b \epsilon_t^b \\
\mu_t^b \;=\;& \omega^b + \beta_1^b \mu_{t-1}^b + \delta_1^{b,b} bv_{t-1}^+ + \psi_1^{b,b} bv_{t-1}^- + \\
& +\delta_1^{b,r}(r_{t-1}^2)^+ + \psi_1^{b,r}(r_{t-1}^2)^- + \delta_1^{b,v} rv_{t-1}^+ + \psi_1^{b,v} rv_{t-1}^-
\end{aligned}
$$

In view of the fact that $rv_t \approx bv_t + j_t$ alternative models to be estimated can be seen as reparameterization of the previous ones, highlighting the separate role of jumps:

$$
\begin{aligned}
& \text{Model 1 - Squared returns } r_t^2 = \mu_t^r \epsilon_t^r \qquad (8) \\
\mu_t^r \;=\;& \omega^r + \beta_1^r \mu_{t-1}^r + \alpha_1^r(r_{t-1}^2)^+ + \gamma_1^r(r_{t-1}^2)^- + \\
& +(\delta_1^{r,v} + \delta_1^{r,b}) bv_{t-1}^+ + (\psi_1^{r,v} + \psi_1^{r,b}) bv_{t-1}^- + \delta_1^{r,v} j_{t-1}^+ + \psi_1^{r,v} j_{t-1}^- \\
& \text{Model 2 - Realized variance } rv_t = \mu_t^v \epsilon_t^v \qquad (9) \\
\mu_t^v \;=\;& \omega^v + \beta_1^v \mu_{t-1}^v + \delta_1^{v,r}(r_{t-1}^2)^+ + \psi_1^{v,r}(r_{t-1}^2)^- + \\
& +(\delta_1^{v,v} + \delta_1^{v,b}) bv_{t-1}^+ + (\psi_1^{v,v} + \psi_1^{v,b}) bv_{t-1}^- + \delta_1^{v,v} j_{t-1}^+ + \psi_1^{v,v} j_{t-1}^- + \\
& \text{Model 3 - Bi–power variance } bv_t = \mu_t^b \epsilon_t^b \qquad (10) \\
\mu_t^b \;=\;& \omega^b + \beta_1^b \mu_{t-1}^b + (\delta_1^{b,b} + \delta_1^{b,v}) bv_{t-1}^+ + (\psi_1^{b,b} + \psi_1^{b,v}) bv_{t-1}^- + \\
& +\delta_1^{b,r}(r_{t-1}^2)^+ + \psi_1^{b,r}(r_{t-1}^2)^- + \delta_1^{b,v} j_{t-1}^+ + \psi_1^{b,v} j_{t-1}^-
\end{aligned}
$$

In each of the three expressions (9), (10), and (11) we can analyze the individual significance of the estimated coefficients, but also the following restrictions (for each indicator $i = r, v, b$):

1. if both $\delta^{i,v}$ and $\delta^{i,b}$ are equal to zero, there is no lagged influence from either bi–power variation, jumps or realized variance.

2. If $(\delta_1^{i,v} + \delta_1^{i,b}) = 0$ but $\delta^{i,v} \neq 0$, lagged bv_{t-1} does not affect μ_t^i but the lagged jump j_{t-1} does.

3. If $\delta^{i,v} = 0$ but $\delta^{i,b} \neq 0$, only lagged bv_{t-1} affects μ_t^i but lagged jump j_{t-1} does not.

4. If $\delta^{i,v} \neq 0$ but $\delta^{i,b} = 0$ then the autonomous effect of bi–power variation is not present, since the coefficient on this variable is the same as the one associated with jumps. In other words, this is the case when realized variance enters as a significant variable *in lieu* of bi–power variation.

4 The Estimation Results

The three models were estimated and the coefficients are reported in Tables 2 to 4. In each table there are three columns: for comparison purposes, the first reports the results for the so–called *base* specification where only the own terms as in $\mu_{i,t} = \omega^i + \delta_1^{i,i}(\sigma^2_{i,t-1})^+ + \psi_1^{i,i}(\sigma^2_{i,t-1})^- + \beta_1^i \mu_{i,t-1}$ are estimated. The second one is the full specification where all variables are included, and the third reports the results of a model selected according to a General–to–Specific (G–to–S) criterion.

A few comments are in order: there is a striking regularity of the significance of the lagged squared returns when returns are negative (coefficient values are around 0.10); the counterpart for positive returns is consistently dropped from all specification as insignificant. This is consistent with the findings by Forsberg and Ghysels (2007) who find that lagged absolute returns have explanatory power for realized volatility and expands those findings to bi–power variance. The coefficients on lagged μ are all high and statistically different from zero (more so for the squared returns, around 0.75, than for the other indicators, around 0.6). Lagged bi–power variance is present in all retained specifications with coefficients which are around 0.27 in most specifications (a little lower in the squared returns specification – around 0.16). The significant presence of jumps is detected only in the model for realized variance. The asymmetric effects for these two variables are not present, in that tests for pairwise coefficient equality fail to reject the null hypothesis.

There is no evidence about realized variance influencing itself with a lag on the basis of the tests described before: coefficient restrictions on the G–to–S specification show that the equality of the coefficients on bi–power variation and jumps is always rejected (only marginally so in the case of negative effects). There is a stronger evidence favoring absence of asymmetric effects (the test statistic for the equality of coefficients on bi–power and jumps has a p-value of 0.25[4]), and when coefficient equality across asymmetric effects is imposed, the two coefficients on bi-power and jumps are significantly different from one another.

All evidence, therefore, points to the important role played, even dynamically, by the jumps component of the realized variance which have significant influence on future realized variance separately from the bi-power variation.

The results on the bi–power variance parallel, in a way, the ones on squared returns, with the same significant variables being retained and values of the coefficients which are pairwise very similar to one another.

The usefulness of the model with predetermined indicators in the specification lies in the evaluation of the informational content of the estimated expected value according to various models. The forecasts produced by the models are generally highly correlated with one another, as shown in Table 5: the only exception is the column related to the base version of expected squared returns. The additional information included in the lagged values of the indicators seems to make a difference in the values of the forecasts.

Although limited to in–sample at this stage, the comparison across models that one can make can be performed via two sets of Mincer–Zarnowitz regressions for each indicator in which the left hand–side is the indicator of interest and on the right–hand side an estimated μ_t. The first set of forecasts uses the estimated μ_t according to the base specification, the second one the G-to-S specification detailed in the tables above.

Table 6 reports in the first row the p-values of a Wald test statistics for unbiased forecasts of the form $\sigma^2_{i,t} = a + b\hat{\mu}_{j,t} + u_t$ where the $\sigma^2_{i,t}$ is the indicator reported in the column of the table and $\hat{\mu}_{j,t}$ is the fitted value of the model estimated in Tables 2, 3 and 4 above, and in the second row the R^2 of the corresponding regression. In a lot of cases the test statistics falls in the rejection region signifying biasedness of the corresponding forecast. When we fail to reject, the corresponding p-value is highlighted in italics and among those results we report in boldface the highest R^2. It is striking to see that the best results (unbiasedness *cum* highest R^2) are obtained for the G–to–S specification of the model for its own indicator, a sign that the persistence

[4]Separately, the test statistics for the bi-power coefficients has a p-value of 0.10 and the one for jumps has a p-value of 0.59.

measured by the own lagged μ term and by the extra information available at $t-1$ has valuable content.

The R^2 for the expected squared returns are (unsurprisingly) the lowest in the lot: in a way the findings by Andersen and Bollerslev about this indicator being a noisy measure of volatility get confirmed even in this context where all forecasts deliver a low R^2 even when the bi-power variation enters the G–to–S specification. Interestingly, the R^2 are fairly similar to one another across forecasts but the bi–power specifications do not pass the unbiasedness test.

The forecasts for the realized variance have a strong improvement (as measured by the R^2) when past bi-power and jumps are considered. The role of asymmetric squared returns is also to be pointed out since the R^2 of the G–to–S squared returns forecasts is also fairly high.

Finally, bi–power variance is often a biased forecast of other indicators and is also biasedly predicted by other specifications. The highest fit is obtained by its own specification when lagged squared returns are (asymmetrically) considered.

5 Conclusions

In this paper, we presented some evidence about the dynamic interrelationships among different measures of integrated variance, as estimated within a Multiplicative Error Model framework. The behavior of several variance measures is modeled as the product of a scale factor and a unit mean innovation allowing for the former to have all lagged indicators on the right–hand side. We focus especially on the decomposition of realized variance into a jump component and a bi–power variation component, and on the role of lagged squared returns as an explanatory variable. The results show that lagged jumps are not relevant for either squared returns or bi–power variance (as expected), that lagged squared returns are relevant when they are associated with asymmetric effects (lagged negative returns), and that bi–power variations are entering significantly in all specifications (although not differentiated by the sign of returns). Realized variance gets separate impacts by lagged jumps and bi–power variations.

Mincer–Zarnowitz regressions show that the augmented MEM specifications (selected by a general–to-specific strategy) provide estimates which are unbiased in–sample predictions of the corresponding indicator.

It goes beyond the scope of this paper to investigate the possible reasons of why a noisy measure of variance such as squared returns still provides valuable information to predict high–frequency based measures of volatility. The conjecture is that a decomposition of returns into overnight and intradaily returns (measured on the basis of opening and closing prices) may reveal the impact of market information accumulation during closing times on the size of volatility the following day.

References

[1] Andersen T.G. and Bollerslev T. (1998), Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts, *International Economic Review*, 39, 885-905.

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

[3] Andersen, T.G., Bollerslev, T., Christoffersen, P.F., and Diebold, F.X., (2005), Volatility and Correlation Forecasting, in G. Elliott, C.W.J. Granger, and A. Timmermann (eds.), *Handbook of Economic Forecasting*. Amsterdam: North-Holland, forthcoming.

[4] Andersen T.G., Bollerslev, T., and Diebold, F.X., (2007) Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility, NBER WP 11775, forthcoming *Review of Economics and Statistics*.

[5] Bandorff-Nielsen, O., and Shephard, N., (2002), Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society*, Series B 64, 253280.

[6] Bandorff-Nielsen, O., and Shephard, N. (2004), Power and bipower variation with stochastic volatility and jumps (with discussion) *Journal of Financial Econometrics*, 2, 1-48; ; doi:10.1093/jjfinec/nbh001.

[7] Barndorff-Nielsen, O., and Shephard, N. (2006), Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation, *Journal of Financial Econometrics*, 4, 1-30; doi:10.1093/jjfinec/nbi022.

[8] Bollerslev T. and Wooldridge J., (1992), Quasi-Maximum Likelihood Estimation of Dynamic Models with Time-Varying Covariances, *Econometric Reviews*, 11, 143-172.

[9] Brownlees C.T. and Gallo, G.M. (2005), Financial Econometric Analysis at Ultra–High Frequency: Data Handling Concerns, *Computational Statistics and Data Analysis*, 51, 2232-2245 (Hdoi:10.1016/j.csda.2006.09.030H).

[10] Engle, R.F. (2002), New Frontiers for ARCH Models. *Journal of Applied Econometrics*, 17, 425446.

[11] Engle, R.F. and Gallo G.M. (2006), A Multiple Indicators Model for Volatility using Intra-daily Data, *Journal of Econometrics*, 131, 3-27.

[12] Forsberg L. and Ghysels, E. (2007), Why Do Absolute Returns Predict Volatility So Well? *Journal of Financial Econometrics*, 5, 31-67.

[13] Gallo, G.M. (2001), Modelling the Impact of Overnight Surprises on Intra-daily Volatility, *Australian Economic Papers*, 40, 567-580.

[14] Hansen, P.R. and Lunde, A. (2006), Realized Variance and Market Microstructure Noise (with discussion), *Journal of Business and Economic Statistics*, 24, 127–161.

Table 2: MEM Models: Estimation Results

Squared Returns			
Variables	Base	Complete	G-to-S
Constant	0.068	0.199	0.123
	(0.022)	(0.092)	(0.062)
μ^r_{t-1}	0.940	0.784	0.759
	(0.014)	(0.052)	(0.053)
$(r^2_{t-1})^+$	-0.013	-0.012	
	(0.008)	(0.036)	
$(r^2_{t-1})^-$	0.108	0.085	0.091
	(0.023)	(0.011)	(0.036)
bv^+_{t-1}		0.161	0.166
		(0.072)	(0.083)
bv^-_{t-1}		0.167	0.203
		(0.069)	(0.069)
J^+_{t-1}		-0.746	
		(0.617)	
J^-_{t-1}		-0.519	
		(0.981)	
rv^+_{t-1}			
rv^-_{t-1}			
AIC	3.883	3.871	3.874
ARCH(4)	2.745	2.357	2.376
	(0.601)	(0.670)	(0.667)
Q-stat(12)	12.074	11.830	11.398
	(0.440)	(0.459)	(0.495)
LogL	-2865.212	-2855.756	-2857.956

Note: Robust standard errors in parentheses below the coefficients. ARCH(4) is the LM test for ARCH with 4 lags (p-value in parenthesis); Q-stat(12) is the Ljung Box statistic for residual autocorrelation up to lag 12 (p-value in parenthesis).

Table 3: MEM Models: Estimation Results

Realized Variance			
	Base	Complete	G-to-S
Constant	0.238	0.198	0.199
	(0.036)	(0.038)	(0.039)
μ^v_{t-1}	0.561	0.600	0.596
	(0.032)	(0.030)	(0.029)
$(r^2_{t-1})^+$		-0.003	
		(0.005)	
$(r^2_{t-1})^-$		0.107	0.108
		(0.020)	(0.020)
bv^+_{t-1}		0.263	0.263
		(0.033)	(0.033)
bv^-_{t-1}		0.313	0.315
		(0.034)	(0.034)
J^+_{t-1}		1.020	1.022
		(0.285)	(0.286)
J^-_{t-1}		0.796	0.814
		(0.264)	(0.266)
rv^+_{t-1}	0.297		
	(0.033)		
rv^-_{t-1}	0.425		
	(0.039)		
Tests			
a: $\delta^{v,v} \neq 0$, $\delta^{v,b} = 0$			0.008
b: $\psi^{v,v} \neq 0$, $\psi^{v,b} = 0$			0.068
c: a. and b.			0.003
AIC	3.871	3.872	3.870
ARCH(4)	2.636	3.320	3.250
	(0.620)	(0.506)	(0.517)
Q-stat(12)	30.633	28.423	28.242
	(0.002)	(0.005)	(0.005)
LogL	-2856.836	-2853.276	-2853.288

Note: Robust standard errors in parentheses. The tests section reports the p-values of the test statistics used to ascertain whether lagged realized variance rv_{t-1} in this model should replace the two variables bv_{t-1} and j_{t-1} considered separately, first by the sign of the returns (a., positive, and b., negative) and then jointly.

Table 4: MEM Models for GE: Estimation Results

Bi–power Variance			
Variables	Base	Complete	G-to-S
Constant	0.186	0.153	0.191
	(0.030)	(0.034)	(0.027)
μ^b_{t-1}	0.607	0.628	0.630
	(0.029)	(0.028)	(0.027)
$(r^2_{t-1})^+$		-0.001	
		(0.017)	
$(r^2_{t-1})^-$		0.096	0.094
		(0.004)	(0.017)
bv^+_{t-1}	0.264	0.236	0.236
	(0.028)	(0.029)	(0.027)
bv^-_{t-1}	0.385	0.283	0.277
	(0.034)	(0.031)	(0.029)
J^+_{t-1}		0.407	
		(0.230)	
J^-_{t-1}		0.211	
		(0.246)	
rv^+_{t-1}			
rv^-_{t-1}			
AIC	3.754	3.755	3.752
ARCH(4)	6.861	6.191	6.257
	(0.143)	(0.185)	(0.181)
Q-stat(12)	38.356	35.696	35.808
	(0.000)	(0.000)	(0.000)
LogL	-2770.521	-2767.132	-2767.373

Note: Robust standard errors in parentheses below the coefficients. ARCH(4) is the LM test for ARCH with 4 lags (p-value in parenthesis); Q-stat(12) is the Ljung Box statistic for residual autocorrelation up to lag 12 (p-value in parenthesis).

Table 5: Correlations between the Variance Forecasts: GE. Jan. 4, 1995–Dec. 31, 2000

Indicator	$\hat{r}_t^2$ Base	$\hat{r}_t^2$ G–to–S	$\widehat{rv}_t$ Base	$\widehat{rv}_t$ G–to–S	$\widehat{bv}_t$ Base
$\hat{r}_t^2$ G–to–S	0.873				
$\widehat{rv}_t$ Base	0.724	0.941			
$\widehat{rv}_t$ G–to–S	0.783	0.969	0.975		
$\widehat{bv}_t$ Base	0.767	0.966	0.986	0.977	
$\widehat{bv}_t$ G–to–S	0.812	0.983	0.968	0.995	0.983

Table 6: MEM Models: Biasedness Wald Test and R^2 for Mincer–Zarnowitz Regressions. In–sample GE: Jan. 4, 1995–Dec. 31, 2000

Regression: $\sigma^2_{i,t} = a + b\hat{\mu}_{j,t} + u_t$			
	Indicator $\boldsymbol{\sigma}^2_{i,t}$		
Forecasts $\hat{\boldsymbol{\mu}}_{j,t}$	Sq. Returns	Realized	Bi–power
Sq. Returns (base)	*0.548*	*0.338*	0.000
	0.086	0.251	0.290
Sq. Returns (G-to-S)	*0.719*	*0.555*	0.000
	0.100	0.322	0.360
Realized (base)	*0.536*	*0.808*	*0.227*
	0.100	0.306	0.340
Realized (G-to-S)	*0.594*	*0.719*	0.000
	0.100	**0.342**	0.375
Bi–power (base)	0.015	0.000	*0.861*
	0.101	0.304	0.339
Bi–power (G-to-S)	0.011	0.000	*0.812*
	0.100	0.348	**0.384**

For each estimated model as labeled, the table reports in the first row the p-value of the Wald test statistic for a null joint hypothesis $H_0 : a = 0, b = 1$ in a Mincer–Zarnowitz regression with Newey–West robust variance covariance matrix (italics, shows failure to reject). The second row reports the R^2 associated with the same M-Z regression. Conditional on the p-value being greater than any customary significance level, the highest R^2 by indicator (column) is in boldface font.