

The Relationship Between the Volatility of Returns and the Number of Jumps in Financial Markets

Forthcoming: Econometric Reviews

ÁLVARO CARTEA AND DIMITRIOS KARYAMPAS*

This version: May 7, 2014

Abstract

We propose a methodology to employ high frequency financial data to obtain estimates of volatility of log-prices which are not affected by microstructure noise and Lévy jumps. We introduce the ‘number of jumps’ as a variable to explain and predict volatility and show that the number of jumps in SPY prices is an important variable to explain the daily volatility of the SPY log-returns, has more explanatory power than other variables (e.g. high and low, open and close), and has a similar explanatory power to that of the VIX. Finally, number of jumps is very useful to forecast volatility and contains information that is not impounded in the VIX.

JEL CLASSIFICATION CODES: C53, G12, G14, C22.

KEY WORDS: volatility forecasts; high-frequency data; implied volatility; VIX; jumps; microstructure noise.

*Cartea is at University College London and Karyampas at UBS AG, IB Exposure Measurement. a.carte@ucl.ac.uk, dimitrios.karyampas@ubs.com. We are grateful to G. Amromin, D. Amengual, P. E. George, I. Peña, J. Penalva, Z. Psaradakis, A. Rubia, and J. van Bommel for their comments. Karyampas is grateful to Birkbeck College Research Committee and the Economic & Social Research Council for financial support. This paper has benefited from comments of seminar participants at The University of Chicago, Federal Reserve Bank Chicago, University of Oxford, University of Toronto and ESSEC. This work expresses the opinion of its authors and is in no way representing the opinion of the institutions the authors work for.

I. Introduction

Modeling and forecasting volatility of asset prices is a crucial task in finance. The recent financial crisis has highlighted the importance that investors place on the returns and volatilities of assets. During the crisis, the volatility of most financial assets almost doubled, and at the same time changes in volatility (known as the volatility of volatility) also increased, reflecting the “puzzled” expectations and reactions of investors in the risky and uncertain environment.

Many methods have been proposed to estimate daily volatility using data at higher frequencies. One of the best known approaches is known as ‘realized volatility’ where volatility is calculated at a 5-minute sampling frequency, see Andersen and Bollerslev (1998). There are other more recent developments that estimate volatility at even higher frequencies, some of which are also designed to address the problems stemming from the microstructure noise when sampling at high frequencies, see Zhang, Mykland, and Aït-Sahalia (2005), Aït-Sahalia, Mykland, and Zhang (2005), Barndorff-Nielsen, Hansen, and Lunde (2008).

Another generation of papers proposes how to employ high frequency data to measure volatility of returns when prices exhibit jumps, but does not address the problems arising from microstructure noise. See for example Barndorff-Nielsen and Shephard (2004), Barndorff-Nielsen and Shephard (2006), Andersen, Bollerslev, and Diebold (2007), Mancini (2009), and Corsi, Pirino, and Renò (2010). The literature keeps on growing in different directions, for example in Cartea and Karyampas (2011), Corsi and Renò (2012), Bandi and Renò (2012), Cartea and Karyampas (2012), Mykland and Zhang (2012).

The contribution of this paper is two-fold. First we show how to estimate the volatility of log-returns where the estimates are not affected by the problems arising from microstructure noise and the presence of jumps. Here, ‘jumps’ refers to price revisions that are *not* consequence of Brownian motion or Gaussian shocks, but come from large and rare Poisson-type events or from small infinite activity jumps, both of which we consider to be Lévy-type jumps.

The second contribution of our paper focuses on the link between volatility and the jumps in log-returns. We propose, for the first time in the literature, to use the ‘number of jumps’ as a jump activity measure and show that the number of jumps within a trading day helps to explain and forecast future volatility.

We show the empirical performance of our volatility estimator and the link between the number of jumps and volatility by employing minute-by-minute observations of the SPY, the fund tracker of the S&P 500, from January 2000 to December 2006. We highlight two of our empirical findings.

First, in addition to other well-documented variables, such as high, low, open, closing prices, we show that the number of jumps in the SPY is a crucial variable in explaining the SPY volatility. We show that: a) the explanatory power of the jump activity measure ‘number of jumps’ is higher than the explanatory power of other jump measures proposed in the literature when explaining the volatility of log-price innovations, b) the number of jumps in SPY prices has more explanatory power with respect to daily volatility than other variables based on: high and low, and open and close, and c) the number of jumps in SPY prices has a similar explanatory power to that of the VIX when it comes to explaining volatility.

Second, using the number of jumps as an explanatory variable increases the forecasting ability of autoregressive volatility models. Results show that the incorporation of forecasts of the monthly average number of jumps in volatility models leads to better monthly volatility forecasts and contains relevant information which is not impounded in the VIX.

The rest of the paper is organized as follows. Section II reviews the literature on volatility estimation with high frequency data, describes the methodology to detect jumps in prices, and discusses how to produce volatility estimates which are not affected by jumps or microstructure noise. Section III describes the data used in our empirical study. Section IV uses different models to explain and forecast volatility and section V concludes.

II. Volatility and jump detection

We assume that the log-price X_t of a security follows

$$X_t = \sigma W_t, \tag{1}$$

where W_t is a standard Brownian motion, $\sigma > 0$ is a constant, and define the realized variance of (1) as

$$RV_{X,T} = \sum_{i=1}^N (X_{t_i} - X_{t_{i-1}})^2, \tag{2}$$

where N is the number of observations. In equation (1) the drift is not included because at high frequencies it is negligible relative to the diffusion.

One problem arising from high frequency financial data is that instead of observing the true or efficient log-price, denoted by X_t , we observe

$$Y_t = X_t + \varepsilon_t \quad (3)$$

where ε_t is microstructure noise. Thus, at high frequencies, the $RV_{X,T}$ in (2) is dominated by the variance of the noise term when all observations in the sample are used. To overcome this problem, a typical approach is to sparse sample the data at low frequencies or to use the Two-Scale Realized Variance estimator proposed in Zhang, Mykland, and Aït-Sahalia (2005). Another approach is that of Aït-Sahalia, Mykland, and Zhang (2005) who assume that the noisy returns are $r_i = \sigma(W_{t_i} - W_{t_{i-1}}) + \varepsilon_{t_i} - \varepsilon_{t_{i-1}} = \zeta_i + \eta\zeta_{i-1}$ and propose a maximum likelihood estimation method (*MLE*) which produces fully efficient volatility estimates, see Cartea and Karyampas (2011) for an alternative parametric method.

A more realistic model of the log-price is one which includes discontinuities by adding a jump component in (1). One way to do this is to assume that jumps arrive according to a Poisson process and one of the challenges is how to measure the volatility of the continuous part of such a model. A popular estimator that focuses on the continuous part of a log-price process that includes Poisson jumps is the *Bipower Variation*:

$$BPV_t = \mu^{-2} \sum_{i=2}^N |r_{i-1}| |r_i|, \quad (4)$$

where r_i indicates the log-return, N is the number of observations and $\mu \simeq 0.7979$. Moreover, based on this estimator the literature has used

$$J_t = \max(RV_{X,t} - BPV_t, 0), \quad (5)$$

to build jump detection tests and to examine the informational content of jumps in volatility forecasts, see Barndorff-Nielsen and Shephard (2004), Andersen, Bollerslev, and Diebold (2007), Corsi, Pirino, and Renò (2010), Becker, Clements, and McClelland (2009), and the *Threshold Bipower Variation (TBPV)* of Corsi, Pirino, and Renò (2010).

The approaches described above provide a way of estimating the volatility of the diffusion part of the price process when either microstructure noise or Poisson-type jumps is present. However, how can we deal with the biases introduced into the volatility estimates by both microstructure noise and jumps in the log-prices? How can we estimate the volatility of the Brownian component in log-returns when more general processes drive the price dynamics?

Considering only Poisson jumps ignores other Lévy-type jumps that are frequent and small, for which, if conflated with the Gaussian movements of the price, the estimator will produce incorrect volatility estimates. Therefore, our aim is to propose a volatility estimator that is neither affected by Lévy-type jumps (infinite activity and Poisson) nor microstructure noise. We assume that we observe the noisy log-prices

$$Y_t = X_t + \varepsilon_t,$$

where $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ is the microstructure noise and the true log-price is given by

$$X_t = \int_0^t \sigma_s dW_s + \int_0^t dL_s, \quad (6)$$

where σ_t is the volatility, W_t is Brownian motion, and L_t is a pure jump Lévy process.

To produce consistent and unbiased estimates of the volatility parameter σ_t we use high frequency data within every trading day to estimate the intraday volatility and assume that volatility can vary from day-to-day but that it is constant within one trading day, see Oomen (2006). It is also possible to relax this assumption and allow for volatility to change within the day, see Christensen, Oomen, and Podolskij (2010).

We deal with the two problems, jumps in returns and microstructure noise, in sequence. We start with the high frequency observations Y_t and employ the non-parametric tests proposed by Lee and Mykland (2008) and Lee and Hannig (2010), see below, to locate and remove price innovations that come from the jump component $\int_0^t dL_s$ in equation (6). Thus, our new series, which we denote \tilde{Y}_t , is given by

$$\tilde{Y}_t = \int_0^t \sigma_s dW_s + \varepsilon_t \quad (7)$$

which is the Gaussian component of the true log-price plus the microstructure noise. This allows us to employ the *MLE* proposed by Aït-Sahalia, Mykland, and Zhang (2005) on the series \tilde{Y}_t , which produces the most efficient estimate of daily volatility σ_t that we can obtain in the presence of

microstructure noise.

Detecting jumps

Lee and Mykland (2008) propose a non-parametric test based on high frequency data to detect jumps that are generated by a non-homogeneous Poisson-type jump process, and Lee and Hannig (2010) propose a test to detect Lévy-type jumps: jumps that are difficult to locate due to their infinite activity and their small size which makes it difficult to differentiate them from Gaussian price changes. Although both tests have been developed to be applied to high frequency data in the absence of microstructure noise, for practical purposes, when the variance of the microstructure noise σ_ε^2 is of the order of 10^{-7} or smaller, the performance of both tests is not affected by the presence of microstructure noise as shown in Cartea and Karyampas (2012).

Therefore, to ensure that the jump tests can be applied to the data we employ in our study of the SPY (in section III we discuss the data), we use the approach of Gatheral and Oomen (2010) and Aït-Sahalia, Mykland, and Zhang (2005) to estimate the variance of the microstructure noise in the high frequency data of the SPY. The average of the daily variance estimates equals $\hat{\sigma}_\varepsilon^2 = 7.155 \cdot 10^{-8}$ which is small enough to apply the jump detection tests.

The *MLE-F* as an alternative volatility estimator

The first stage of our approach is to remove jumps from the price series to obtain the noisy log-price series \tilde{Y}_t (see equation (7)) which contains a combination of two Gaussian shocks, price innovations in the continuous part of the true log-price process X_t and the microstructure noise ε_t , and then apply the *MLE* approach on \tilde{Y}_t . We arrange the log-price observations in a vector with entries \tilde{Y}_{t_i} , with $t_i = i\Delta$, $t_N = N\Delta = T$ for $i = 0, \dots, N$, where Δ is the time-step between observations and N is the number of (Gaussian) observations in the trading day. Note that when we delete observations because they are considered jumps, the time step between two consecutive \tilde{Y}_t 's is not Δ , it is Δ multiplied by $1 + nj$ where nj is the number of jumps that have been deleted in that time interval. Hence, we denote the time in between observations by Δ_i and we proceed to obtain fully efficient estimates for the volatility of the diffusion part and the variance of microstructure

noise by maximizing the likelihood function

$$l(\sigma^2, \sigma_\varepsilon^2) = -\log \det(\Sigma)/2 - N/2 \log(2\pi) - \frac{1}{2} \tilde{r}' \Sigma^{-1} \tilde{r}, \quad (8)$$

where Σ is the covariance matrix of the returns

$$\Sigma = \begin{pmatrix} \sigma^2 \Delta_1 + 2\sigma_\varepsilon^2 & -\sigma_\varepsilon^2 & 0 & \dots & 0 \\ -\sigma_\varepsilon^2 & \sigma^2 \Delta_2 + 2\sigma_\varepsilon^2 & -\sigma_\varepsilon^2 & \ddots & \vdots \\ 0 & -\sigma_\varepsilon^2 & \sigma^2 \Delta_3 + 2\sigma_\varepsilon^2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & -\sigma_\varepsilon^2 \\ 0 & \dots & 0 & -\sigma_\varepsilon^2 & \sigma^2 \Delta_N + 2\sigma_\varepsilon^2 \end{pmatrix},$$

and $\tilde{r} = (\tilde{r}_1, \dots, \tilde{r}_N)'$ is the returns vector that does not contain jumps. Properties of this estimator have been extensively tested in Cartea and Karyampas (2012) under several assumptions and different models.

III. Data

The data used in this paper are from the Trade and Quote (TAQ) and CRSP databases and our analysis is based on the S&P 500 index. Instead of using S&P 500 data, we proceed as in Patton and Verardo (2012) and employ the S&P 500 SPDR traded on Amex with ticker SPY, available on the TAQ database, which is the exchange traded fund tracking the S&P 500. The SPY transactions included here are trades from 9.30am to 4.00pm on a sample period that runs from January 2000 to December 2006. Below we also use the VIX index over the same period. The total period consists of 1,759 trading days. The first 12 days are used as an initial window required by the test to detect jumps in the SPY log-returns. Therefore the sample used to compute realized variances and jump activity consists of 1,747 days, i.e. 681,330 high frequency (minute-by-minute) observations.

Because high frequency data are irregularly spaced, and it is not possible to find a sequence of trades with exactly $\Delta = 60$ seconds between them, we prepare the data by designating the last trade within the preceding 60 seconds as the observation for the minute in question and then apply the jump detection tests to obtain the position and total number of jumps within each trading day. The total number of jumps for the entire period is 1,899, thus the daily average is 1.087. The

maximum number of jumps in one day is 32, which represents 8.2% of the minute-by-minute trades taking place on 28 August 2001. Figure 1 shows the SPY return where red circles and black dots indicate big and small jumps respectively.

Figure 1 and Figure 2 About Here

Finally, in Figure 2 we show the QQ Plots of the SPY returns before and after removing the jumps. The blue circles depict the QQ Plot of the SPY returns and it is evident that they are not Normal. The light blue dots represent the QQ Plot of the filtered SPY returns where it is clear that these filtered returns are very ‘close’ to being Normally distributed.

IV. Explaining and forecasting volatility

Our objective is to show the relation between the volatility of log-returns and the number of jumps. As discussed above, the literature has proposed a number of approaches to estimate volatility and has also proposed different ways to measure jump activity. Therefore, to analyze the explanatory power of the total number of jumps during day t , and denoted by NJ_t , we measure volatility using the estimator proposed here which deals with jumps and microstructure noise, as well as three other alternative estimators. Moreover, we also compare the explanatory power of NJ_t against that of the widely used measure J_t given in equation (5).

The volatility estimators used here are: (i) *MLE*, (ii) *BPV*, (iii) *TBPV*, and (iv) *MLE-F*. The *MLE* approach assumes that there is i.i.d. microstructure noise in the log-prices and, consequently, the returns follow an *MA*(1) process. In Figure 4 we show that this is the case where it is clear that there is significant autocorrelation of the SPY returns at lag 1.

Figure 4 About Here

As a preliminary result in Table I we present the correlations between the daily volatility estimates that are obtained from the four estimators. Moreover, although below we only focus on these four estimators, Table I also shows the correlation of the volatility estimates when $|r_t|$ and *GARCH* are used as alternative volatility estimators which are estimated with daily data. The table also shows the correlation of the VIX with the other volatility estimates.

From Table I we observe that the two estimators which show the lowest correlations with the other measurements of volatility are the ones based on daily data, while the correlation between them ($|r_t|$, *GARCH*) is equal to 96%. We also see that the highest correlation is between *MLE-F* and *TBPV* which equals 0.97.

Figure 3 shows the time series of the daily volatility estimates using the *MLE*, *BPV*, *TBPV* and *MLE-F*. It is interesting to observe that the *MLE-F* does not exhibit too many ‘spikes’ whereas the *BPV* volatility estimates exhibit large fluctuations which could be due to the fact that in theory the *BPV* estimates are not affected by big and rare jumps, but are affected by small jumps and by microstructure noise.

Table I and Figure 3 About Here

A. The informational content of the realized number of jumps

In this section we investigate whether the jump activity of the SPY can explain the volatility of its log-returns. To do so, we apply several models that have jump activity measures as explanatory variables. We proceed in the following way: (i) we use as benchmark an autoregressive model with no jump activity measure as an independent variable, (ii) we extend the autoregressive model by including the jump activity measure J_t as an additional explanatory variable, (iii) instead of including J_t as explanatory variable, we include the number of jumps NJ_t as an alternative jump activity measure, and finally (iv) we extend the autoregressive model to include both J_t and NJ_t .

Consistent with the literature, all models in this section examine the relation between the logarithmic value of volatility and the independent variables. The autoregressive model *AR*(1) is

$$\log \sigma_t = c + \beta_1 \log \sigma_{t-1} + u_t, \quad (9)$$

where σ_t are the daily estimates provided by the four alternative volatility estimators (*MLE*, *BPV*, *TBPV* and *MLE-F*) and u_t is the noise term of the regression.

Table II shows that the coefficients are significant and that volatility is a highly mean-reverting process. The adjusted- R^2 for all regressions is very high, with values close to 75%. For all regressions we use the Heteroscedasticity and Autocorrelation Consistent covariance estimates proposed by

Newey and West (1987).

Table II About Here

At this point we incorporate into our analysis the effects of the jump activity as another explanatory variable in the regressions of the SPY return volatility. We do this in the following set of regressions where we extend the $AR(1)$ process by incorporating a jump activity measure and its past value. The first jump activity measure we include is J_t given above in equation (5). The other jump activity measure, which is proposed for the first time in this paper, is given by the actual number of jumps occurring within each trading day.

To test the explanatory power of J_t we employ

$$\log \sigma_t = c + \beta_1 \log \sigma_{t-1} + \beta_2 J_t + \beta_3 J_{t-1} + u_t, \quad (10)$$

where σ_t are the daily estimates provided by the four alternative volatility estimators and u_t is the noise term of the regression. Similarly, by replacing J_t with NJ_t we run the regression

$$\log \sigma_t = c + \beta_1 \log \sigma_{t-1} + \beta_2 NJ_t + \beta_3 NJ_{t-1} + u_t, \quad (11)$$

and finally, we run the regression that extends the $AR(1)$ to include both jump activity measures

$$\log \sigma_t = c + \beta_1 \log \sigma_{t-1} + \beta_2 J_t + \beta_3 J_{t-1} + \beta_4 NJ_t + \beta_5 NJ_{t-1} + u_t. \quad (12)$$

The estimates of the coefficients are presented in Table III.

Table III About Here

The J_t variable and its lagged value J_{t-1} are significant for the MLE volatility estimator, all at a 5% significance level. J_t does not provide any informational content in the new $MLE-F$ volatility estimator introduced in this paper. By looking at the adjusted- R^2 s of the $AR(1)$ in Table II and the adjusted- R^2 s in Table III, we see that the incorporation of J_t and its lagged value increase slightly the models' explanatory power for the MLE , BPV , $TBPV$ and $MLE-F$. Moreover, by looking at the information criteria we see that they all have smaller values than in the previous pure $AR(1)$ regressions, while the log-likelihood maximum values ($\log L$) are greater in this case.

The coefficients for the number of jumps NJ_t and NJ_{t-1} are positive and negative respectively and highly significant in all cases, even for the *MLE-F* volatility estimate. For instance, we see that for *BPV* the adjusted- R^2 increases from 69.99%, in the *AR*(1) model (Table II), to 80.9% with the inclusion of the NJ_t and NJ_{t-1} regressors. This adjusted- R^2 is even higher than in the case where we only had J_t , where the adjusted- R^2 is 70.4%, as shown in the first column of *BPV* results in Table III.

In the pure *AR*(1) the $\log L$ value for the *BPV* is -298.44 , but when the *NJ*'s are included, see equation (11), the $\log L$ increases to 98.39 which implies that the extended *AR*(1) is a much better model to explain the market's variation. Moreover, in this extended model, the information criteria values decrease when NJ_t and NJ_{t-1} are included. This indicates that NJ_t and its lagged value increase the explanatory power of the volatility models more than J_t and J_{t-1} .

Similarly, for the other models, we observe that the adjusted- R^2 's and information criteria show that the extended model (11) performs better. For example, in the case of the *MLE*, the adjusted- R^2 increases from 72.4% in the *AR*(1) model to 77.7% in the model with NJ_t and NJ_{t-1} , while the $\log L$ value for the pure *AR*(1) is -42.09 , see Table II, and for the extended *AR*(1) is 145.13 as shown in Table III. Furthermore, the adjusted- R^2 when daily volatility is calculated using the *TBPV*, is 86.8% in the extended model whereas it is 84.7% in the pure *AR*(1) model. Finally, when we use the *MLE-F* to calculate daily volatility, the $\log L$ and R^2 in the pure *AR*(1) case are 235.80 and 79.2% respectively, and in the extended case the $\log L$ and R^2 are 304.38 and 80.8% respectively.

In general, including the number of jumps as an explanatory variable increases the adjusted- R^2 of the four estimators. This increase is higher for the *MLE* and the *BPV*, while not as large as that for the *TBPV* and the *MLE-F*. But for all estimators the increase of the explanatory power due to the number of jumps NJ_t and its lagged value is much higher than that of J_t and J_{t-1} as jump activity measures.

Finally, by looking at the last column of each of the four estimators in Table III we observe that in most cases NJ_t and NJ_{t-1} overlap the effect that J_t and J_{t-1} have on volatility. J_t and J_{t-1} remain statistically insignificant for *BPV*, *TBPV* and *MLE-F*, while NJ_t and NJ_{t-1} still have significant coefficients for all estimators. The explanatory power of the model where both J_t and NJ_t , and their lagged values are taken into account, is close to the model for which only number

of jumps is used to explain volatility.

B. Alternative variables: VIX, high and low, open and close

So far we have focused on two jump activity measures when extending the pure $AR(1)$ log-volatility process to identify which of these two competing measures has more power to explain the daily volatility of the SPY returns. There are, however, many other well studied factors that have been successfully employed to explain the behavior of the volatility of assets' log-returns. Therefore, here we analyze the explanatory power of NJ_t in the presence of other variables such as VIX, close and open, and high and low, see Alizadeh, Brandt, and Diebold (2002), Garman and Klass (1980), Rogers and Satchell (1991), Yang and Zhang (2000). The regression we use is the following

$$\log \sigma_t = c + \beta_1 \log \sigma_{t-1} + \beta_2 NJ_t + \beta_3 NJ_{t-1} + \beta_4 \log HL_t + \beta_5 OC_t + \beta_6 \log VIX_t + u_t, \quad (13)$$

where $HL_t = \log(\text{high}_t/\text{low}_t)$, $OC_t = |\text{close}_t/\text{open}_t|$, and results are shown in Table IV.

Table IV About Here

When comparing the results of running regression (13) with those of regression (11) we find that all independent variables are statistically significant for the MLE estimator and the adjusted- R^2 for (13) is 87.4%, which is higher than the 77.7% obtained from running regression (11).

Regarding the BPV estimator, the adjusted- R^2 also increases from 80.9%, in the case where we had only the NJ s as explanatory variables, to 86.3% with the incorporation of the extra variables in (13). In addition, the variable OC_t is insignificant at a 5% significance level.

The explanatory power of the model for the $TBPV$ case becomes also higher with the extra explanatory variables. The new adjusted- R^2 is 90.8%, which is higher than 86.8% from model (11). The new $\log L$ value has nearly doubled, increasing from 487.47 to 802.09.

For the $MLE-F$ volatility estimator the incorporation of the extra variables increases the adjusted- R^2 from 80.8% to 88.8%. Moreover, the coefficient of NJ_{t-1} is insignificant.

Finally, information criteria for all four estimators become lower when the extra explanatory variables are incorporated in the model described by (11), implying that this model better explains

the SPY volatility.

Therefore, it is interesting to see that our proposed jump activity measure, given by the number of jumps within each day, is still statistically significant after the incorporation of the extra explanatory variables in the model.

C. HAR model and jumps

In this section we use the HAR model proposed by Corsi (2009) to analyze the explanatory power of J_t and NJ_t when the long memory effects of volatility are taken into account. We estimate

$$\log \sigma_{t+1} = c + \beta^d \log \sigma_t + \beta^w \log \sigma_t^{(5)} + \beta^m \log \sigma_t^{(22)} + \beta_1 J_{t+1} + \beta_2 J_t + u_t, \quad (14)$$

where σ_t are the daily estimates provided by the four alternative volatility estimators,

$$\log \sigma_t^{(h)} = \frac{1}{h} \sum_{j=1}^h \log \sigma_{t-j+1},$$

and $\beta^d, \beta^w, \beta^m$ are the coefficients for daily, weekly, and monthly volatility respectively. Moreover, proceeding as above, we replace J_t with NJ_t to run the regression

$$\log \sigma_{t+1} = c + \beta^d \log \sigma_t + \beta^w \log \sigma_t^{(5)} + \beta^m \log \sigma_t^{(22)} + \beta_1 NJ_{t+1} + \beta_2 NJ_t + u_t, \quad (15)$$

and finally include both jump activity measures:

$$\log \sigma_{t+1} = c + \beta^d \log \sigma_t + \beta^w \log \sigma_t^{(5)} + \beta^m \log \sigma_t^{(22)} + \beta_1 NJ_{t+1} + \beta_2 NJ_t + \beta_3 J_{t+1} + \beta_4 J_t + u_t. \quad (16)$$

We estimate these three models using OLS with Newey-West covariance correction for serial correlation. The estimates of the coefficients are presented in Table V where we observe that for all volatility estimators NJ_{t+1} is significant. Hence, the number of jumps is still relevant in the presence of $\log \sigma_t^{(5)}$ and $\log \sigma_t^{(22)}$ variables of the HAR model.

Table V About Here

D. Forecasting volatility

At this point we examine whether the number of jumps in the SPY improves the ability to forecast SPY volatility. We follow a similar approach to that of Corsi, Pirino, and Renò (2010) and Andersen, Bollerslev, and Diebold (2003), but propose the use of the number of jumps as a key variable to explain and forecast volatility. From this point onwards, we only use the *MLE-F* to produce volatility figures because it is the only estimator that deals with both jumps and microstructure noise.

The goal is to obtain, at time t , a model-based forecast (MBF) of the volatility of log-returns of the S&P 500 for the next month. Our notation for the monthly volatility forecast at time t is $\hat{\sigma}_{t \rightarrow t+22}$. This means that at time t we produce a forecast which is the square root of the average variance for the next 22 trading days. In our study, the first day for which we produce a forecast is October 5 2006 and the last day is 9 November 2006.

In the literature we can find several MBFs. For example, Becker, Clements, and McClelland (2009) use MBFs that incorporate the jump activity measure J_t as one of the variables in the models. The jump activity measure we use here to forecast volatility is based on $\overline{NJ}_{t \rightarrow t+22}$, which denotes the average of the number of jumps for a period of 22 trading days starting at time t .

When forecasting volatility, VIX is an important variable to consider because it is a measure of the implied volatility of options on the S&P 500 index and therefore it conveys (risk-neutral) forward looking information about what market participants expect the volatility of the S&P 500 to be over the coming month. Therefore, we run

$$\log \sigma_{t \rightarrow t+22} = c + \beta_1 \log VIX_t + u_t, \quad (17)$$

where $\sigma_{t \rightarrow t+22}$ are obtained using the *MLE-F* over the period from t to $t+22$. The estimate of β_1 in equation (17) is close to 1.07 and the adjusted- R^2 is 75.26%.

Alternatively we may also use

$$\log \sigma_{t \rightarrow t+22} = c + \beta_1 \log \sigma_{t-22 \rightarrow t} + \beta_2 \log VIX_{t-1} + u_t, \quad (18)$$

where $\sigma_{t-22 \rightarrow t}$ denotes the square root of the average variance between time $t-22$ and t . Our

main interest is to extend model (18) to include the average number of jumps between time t and $t + 22$ as an explanatory variable. However, at time t we do not know the number of jumps that will occur between t and $t + 22$. Instead, we need to use a forecast of the average number of jumps. Before doing this, we first show that the realized average number of jumps between t and $t + 22$, $\overline{NJ}_{t \rightarrow t+22}$, explains monthly volatility. Thus, we run the following regression

$$\log \sigma_{t \rightarrow t+22} = c + \beta_1 \log \sigma_{t-22 \rightarrow t} + \beta_2 \log VIX_{t-1} + \beta_3 \overline{NJ}_{t \rightarrow t+22} + u_t. \quad (19)$$

Table VII About Here

Table VII presents the results from estimating models (18) and (19). Once again, the average number of jumps plays an important role in explaining monthly volatility. The explanatory power of the model that includes VIX and the average number of jumps is 83.50%, whilst in the model described by (18), where only VIX is employed, the adjusted- R^2 is 78.30%. The high adjusted- R^2 values indicate that model (18) can be used to forecast monthly SPY volatility. In the case of model (19) it shows that the average number of jumps is an important variable to explain monthly SPY volatility.

In the second column of the table we present the model where instead of VIX we use the average number of jumps to explain volatility. The interesting point in this regression is that the adjusted- R^2 is 80.60%, higher than the 78.30% for the model where only VIX is employed.

We emphasize that model (19) employs the number of jumps to explain monthly volatility at time t . However, to predict volatility we must forecast $\overline{NJ}_{t \rightarrow t+22}$ so we employ a moving average process of order 5 ($MA(5)$) and denote these forecasts $\widehat{\overline{NJ}}_{t \rightarrow t+22}$, see Table VI. The adjusted- R^2 of 90% indicates that the moving average specification can explain most of the variability of the average jumps activity measure.

Table VI About Here

Finally, we forecast volatility based on

$$\log \sigma_{t \rightarrow t+22} = c + \beta_1 \log \sigma_{t-22 \rightarrow t} + \beta_2 \log VIX_{t-1} + \beta_3 \widehat{\overline{NJ}}_{t \rightarrow t+22} + u_t. \quad (20)$$

The monthly volatility forecasts are for the period 5 October 2006 to 9 November 2006, i.e. 26

monthly volatility forecasts. Note that this model only differs from (19) in that we use the $MA(5)$ forecasts $\widehat{NJ}_{t \rightarrow t+22}$ as the regressor.

We present four indicators of how good our forecasts are relative to the true values: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and Theil Inequality index which is defined as

$$TI = \frac{\sqrt{\sum_{T+1}^{T+h} (\hat{y}_t - y_t)^2 / h}}{\sqrt{\sum_{T+1}^{T+h} \hat{y}_t^2 / h} + \sqrt{\sum_{T+1}^{T+h} y_t^2 / h}},$$

where \hat{y}_t is the forecast, y_t the observed value, and h the total number of forecasts. The TI index lies between zero and one, with zero indicating a perfect fit.

The last column of Table VII presents the estimates of model (20). We see that when using the forecast of the average number of jumps, the adjusted- R^2 is 83.5% which incidentally is the same adjusted- R^2 that we obtained when the actual average number of jumps was used as regressor in model (19).

The last four rows of Table VII present the four indicators for models (18) and (20) for the last 22 days of our sample. We see that the values of the RMSE, MAE and MAPE are smaller in the case where the forecast of the average number of jumps is incorporated in the model to forecast monthly volatility. Model (20) has an adjusted- R^2 equal to 83.5% which is higher than 78.30% of model (18). Furthermore, the TI coefficient of model (20) is 0.016 which indicates a very good fit and it is also lower than the TI of model (18).

Therefore, our results show that forecasts of the average number of jumps provides crucial information to forecast the SPY one-month volatility. And we remark that our results in Table VII show that the explanatory power of the realized average number of jumps is on a par with that of the VIX and, moreover, the expected number of jumps contains relevant information when forecasting monthly volatility which is not impounded in the VIX.

V. Conclusions

The first contribution of the paper is to propose the $MLE-F$ volatility estimator, an estimator that is neither affected by jumps (Lévy-type that include infinite activity and Poisson) nor microstructure noise when the variance of the microstructure noise is sufficiently small. We propose a two-step procedure that takes into account recent developments in jump detection tests and high frequency volatility estimation to obtain a fully efficient volatility estimator in the presence of noise and price discontinuities. Properties of this estimator have been extensively tested in Cartea and Karyampas (2012) under several assumptions and different models.

The second contribution is to examine the relationship between the jump activity of the SPY, the exchange traded fund tracking the S&P 500 index, and its volatility. We employ high frequency data and decompose the SPY high frequency log-returns into its Lévy-type jumps and Gaussian components. This decomposition of the dynamics of log-prices allows us to propose for the first time the realized number of jumps NJ_t as a new jump activity measure which we use to explain and to forecast the volatility of the SPY.

We find that the number of jumps in the SPY has more explanatory power with respect to daily volatility than other widely used variables such as the difference between log-high and log-low, and the ratio between open and close. Furthermore, we show that our jump activity measure NJ_t has more explanatory power than the well studied jump activity measure J_t , which is based on the difference between the quadratic variation of the log-prices and the Bipower Variation estimator. Our results also show that the number of jumps NJ_t has a similar explanatory power than that of the VIX when explaining daily volatility.

We emphasize that in our analysis the number of jumps NJ_t is statistically significant in all models, even after the inclusion of combinations of all, or some, of the explanatory variables that we have mentioned above. We also show that the number of jumps in the SPY log-returns help us to better forecast SPY monthly volatility. Finally, we show that the explanatory power of the realized average number of jumps is on a par with that of the VIX and we show that the expected number of jumps contains relevant information when forecasting monthly volatility which is not impounded in the VIX.

References

- Aït-Sahalia, Y., P.A. Mykland, and L. Zhang, 2005, How often to sample a continuous-time process in the presence of market microstructure noise, *Review of Financial Studies* 18, 351–416.
- Alizadeh, S., M.W. Brandt, and F.X. Diebold, 2002, Range-based estimation of stochastic volatility models, *The Journal of Finance* 57, 1047–1091.
- Andersen, T.G., and T. Bollerslev, 1998, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* 39, 885–905.
- Andersen, T.G., T. Bollerslev, and F.X. Diebold, 2003, Some like it smooth, and some like it rough: Untangling continuous and jump components in measuring, modeling, and forecasting asset return volatility, *working paper*.
- Andersen, T.G., T. Bollerslev, and F.X. Diebold, 2007, Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility, *The Review of Economics and Statistics* 89, 701–720.
- Bandi, F.M., and R. Renò, 2012, Time-varying leverage effects, *Journal of Econometrics* 169, 94 – 113 Recent Advances in Panel Data, Nonlinear and Nonparametric Models: A Festschrift in Honor of Peter C.B. Phillips.
- Barndorff-Nielsen, O.E., P.R. Hansen, and A. Lunde, 2008, Designing realized kernels to measure the ex post variation of equity prices in the presence of noise, *Econometrica* 76, 1481–1536.
- Barndorff-Nielsen, O.E., and N. Shephard, 2004, Power and bipower variation with stochastic volatility and jumps, *Journal of Financial Econometrics* 2, 1–37.
- Barndorff-Nielsen, O.E., and N. Shephard, 2006, Econometrics of testing for jumps in financial economics using bipower variation, *Journal of Financial Econometrics* 4, 1–30.
- Becker, R., A.E. Clements, and A. McClelland, 2009, The jump component of S&P 500 volatility and the VIX index, *Journal of Banking and Finance* 33, 1033–1038.
- Cartea, Á., and D. Karyampas, 2011, Volatility and covariation of financial assets: A high-frequency analysis, *Journal of Banking & Finance* 35, 3319–3334.
- Cartea, Á., and D. Karyampas, 2012, Assessing the performance of different volatility estimators: A Monte Carlo analysis, *Applied Mathematical Finance* 19, 535–552.

- Christensen, K., R. Oomen, and M. Podolskij, 2010, Realised quantile-based estimation of the integrated variance, *Journal of Econometrics* 159, 74–98.
- Corsi, F., 2009, A simple approximate long-memory model of realized volatility, *Journal of Financial Econometrics* 7, 174–196.
- Corsi, F., D. Pirino, and R. Renò, 2010, Threshold Bipower variation and the impact of jumps on volatility forecasting, *Journal of Econometrics* 159, 276–288.
- Corsi, F., and R. Renò, 2012, Discrete-Time Volatility Forecasting With Persistent Leverage Effect and the Link With Continuous-Time Volatility Modeling, *Journal of Business & Economic Statistics* 30, 368–380.
- Garman, M.B., and M.J. Klass, 1980, On the estimation of security price volatilities from historical data, *Journal of Business* pp. 67–78.
- Gatheral, J., and R. Oomen, 2010, Zero-intelligence realized variance estimation, *Finance and Stochastics* 14, 249–283.
- Lee, S., and J. Hannig, 2010, Detecting jumps from Lévy jump diffusion processes, *Journal of Financial Economics* 96, 271–290.
- Lee, S., and P.A. Mykland, 2008, Jumps in financial markets: A new nonparametric test and jump dynamics, *Review of Financial Studies* 21, 2535.
- Mancini, C., 2009, Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps, *Scandinavian Journal of Statistics* 36, 270–296.
- Mykland, P.A., and L. Zhang, 2012, The econometrics of high frequency data, *Statistical Methods for Stochastic Differential Equations* 124, 109.
- Newey, W.K., and K.D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* pp. 703–708.
- Oomen, R.C.A., 2006, Properties of realized variance under alternative sampling schemes, *Journal of Business & Economic Statistics* 24, 219–237.
- Patton, A.J., and M. Verardo, 2012, Does beta move with news? Firm-specific information flows and learning about profitability, *Review of Financial Studies* 25, 2789–2839.
- Rogers, L.C.G., and S.E. Satchell, 1991, Estimating variance from high, low and closing prices, *The Annals of Applied Probability* 1, 504–512.

Yang, D., and Q. Zhang, 2000, Drift-independent volatility estimation based on high, low, open, and Close Prices, *The Journal of Business* 73, 477–492.

Zhang, L., P.A. Mykland, 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.

Tables and Figures

Table I
Correlation matrix of volatility estimates from high and low frequency estimators and VIX

	<i>MLE-F</i>	<i>MLE</i>	<i>BPV</i>	<i>TBPV</i>	<i>RV_{5min}</i>	<i>GARCH</i>	$ r_t $	<i>VIX</i>
<i>MLE-F</i>	1	0.96	0.86	0.97	0.80	0.83	0.77	0.83
<i>MLE</i>	0.96	1	0.85	0.94	0.85	0.79	0.73	0.80
<i>BPV</i>	0.86	0.85	1	0.89	0.85	0.73	0.68	0.73
<i>TBPV</i>	0.97	0.94	0.89	1	0.84	0.84	0.77	0.83
<i>RV_{5min}</i>	0.80	0.85	0.85	0.84	1	0.69	0.64	0.69
<i>GARCH</i>	0.83	0.79	0.73	0.84	0.69	1	0.96	0.89
$ r_t $	0.77	0.73	0.68	0.78	0.64	0.96	1	0.86
<i>VIX</i>	0.83	0.80	0.73	0.83	0.69	0.89	0.86	1

Table II
***AR*(1) model for volatility estimators**

Volatility estimator		Coefficient	Std. Error	Prob	adjusted- R^2	$\log L$
$\log MLE$	c	0.376	0.033	0.000	0.724	-42.09
	$\log MLE(-1)$	0.851	0.013	0.000		
$\log BPV$	c	0.430	0.041	0.000	0.699	-298.44
	$\log BPV(-1)$	0.837	0.015	0.000		
$\log TBPV$	c	0.203	0.019	0.000	0.847	359.59
	$\log TBPV(-1)$	0.921	0.007	0.000		
$\log MLE-F$	c	0.275	0.024	0.000	0.792	235.80
	$\log MLE - F(-1)$	0.890	0.009	0.000		

Table III
Incorporating the jump activity measures. For each volatility estimator, each column represents the results from running regressions (10), (11), (12), p-values shown in brackets.

	log MLE			log BPV			log $TBPV$			log $MLE-F$		
c	0.327 (0.026)	0.309 (0.027)	0.280 (0.023)	0.424 (0.040)	0.281 (0.035)	0.281 (0.035)	0.200 (0.019)	0.173 (0.017)	0.173 (0.017)	0.272 (0.023)	0.254 (0.022)	0.254 (0.022)
$\log \sigma_{t-1}$	0.869 (0.010)	0.871 (0.011)	0.883 (0.010)	0.838 (0.015)	0.885 (0.013)	0.886 (0.014)	0.920 (0.008)	0.927 (0.007)	0.927 (0.007)	0.890 (0.009)	0.893 (0.009)	0.893 (0.008)
J_t	0.052 (0.010)	- (0.006)	0.039 (0.006)	0.019 (0.010)	- (0.004)	-0.004 (0.004)	0.011 (0.013)	- (0.009)	0.000 (0.009)	0.013 (0.011)	- (0.008)	0.004 (0.008)
J_{t-1}	-0.040 (0.000)	- (0.005)	-0.032 (0.005)	-0.011 (0.008)	- (0.004)	0.006 (0.004)	-0.000 (0.009)	- (0.006)	-0.000 (0.006)	-0.006 (0.006)	- (0.005)	-0.001 (0.005)
NJ_t	- (0.004)	0.047 (0.004)	0.040 (0.004)	- (0.007)	0.072 (0.007)	0.073 (0.007)	- (0.003)	0.032 (0.003)	0.032 (0.003)	- (0.004)	0.026 (0.004)	0.026 (0.004)
NJ_{t-1}	- (0.004)	-0.030 (0.004)	-0.026 (0.004)	- (0.007)	-0.054 (0.007)	-0.055 (0.007)	- (0.003)	-0.019 (0.003)	-0.020 (0.003)	- (0.003)	-0.014 (0.003)	-0.014 (0.003)
adjusted- R^2	0.763	0.777	0.798	0.704	0.809	0.809	0.849	0.868	0.868	0.794	0.808	0.808
Akaike info	-0.099	-0.162	-0.262	0.334	-0.108	-0.108	-0.414	-0.554	-0.554	-0.274	-0.343	-0.343
Schwarz	-0.086	-0.149	-0.243	0.346	-0.096	-0.089	-0.401	-0.541	-0.535	-0.262	-0.332	-0.324
Hannan-Quinn	-0.094	-0.157	-0.255	0.338	-0.103	-0.100	-0.409	-0.547	-0.547	-0.269	-0.339	-0.336
$\log L$	90.05	145.13	234.77	-287.38	98.39	100.10	365.21	487.47	489.39	243.56	304.38	305.25

Table IV
The incremental information of VIX, Open-Close and High-Low, p-values in brackets

Independent Variables	$\log MLE$	$\log BPV$	$\log TBPV$	$\log MLE-F$
c	2.707 (0.496)	1.781 (0.580)	2.619 (0.381)	3.350 (0.448)
$\log \sigma_{t-1}$	0.384 (0.024)	0.497 (0.037)	0.609 (0.023)	0.458 (0.022)
NJ_t	0.031 (0.003)	0.061 (0.007)	0.022 (0.002)	0.012 (0.004)
NJ_{t-1}	-0.012 (0.004)	-0.029 (0.006)	-0.011 (0.003)	-0.005 (0.003)
$\log HL_t$	0.537 (0.056)	0.354 (0.049)	0.357 (0.042)	0.489 (0.053)
OC_t	-1.199 (0.439)	-0.973 (0.546)	-1.721 (0.354)	-1.968 (0.394)
$\log VIX_t$	0.413 (0.043)	0.431 (0.051)	0.299 (0.035)	0.537 (0.039)
adjusted- R^2	0.874	0.863	0.908	0.888
Akaike info	-0.732	-0.435	-0.911	-0.883
Schwarz	-0.710	-0.413	-0.888	-0.861
Hannan-Quinn	-0.724	-0.426	-0.903	-0.875
$\log L$	646.11	386.50	802.09	785.16

Table V
Incorporating the jump activity measures, p-values in brackets

	$\log MLE$		$\log BPV$		$\log TBPV$		$\log MLE-F$	
c	-0.385 (0.065)	-0.239 (0.173)	-0.249 (0.175)	-1.215 (0.000)	-1.059 (0.000)	-1.104 (0.000)	-0.457 (0.003)	-0.331 (0.044)
$\log \sigma_t$	0.355 (0.000)	0.346 (0.000)	0.427 (0.000)	0.068 (0.124)	0.103 (0.085)	0.099 (0.085)	0.418 (0.000)	0.469 (0.000)
$\log \sigma_t^{(5)}$	0.316 (0.000)	0.331 (0.000)	0.271 (0.000)	0.402 (0.000)	0.393 (0.000)	0.391 (0.000)	0.325 (0.003)	0.295 (0.001)
$\log \sigma_t^{(22)}$	0.279 (0.000)	0.273 (0.001)	0.254 (0.000)	0.470 (0.000)	0.447 (0.000)	0.452 (0.000)	0.223 (0.001)	0.205 (0.001)
J_{t+1}	1.007 (0.000)	- (0.000)	1.099 (0.000)	1.900 (0.000)	- (0.000)	0.156 (0.398)	0.811 (0.000)	0.189 (0.306)
J_t	-0.384 (0.001)	- (0.000)	-0.453 (0.000)	0.109 (0.262)	- (0.194)	0.128 (0.194)	0.065 (0.309)	0.088 (0.189)
NJ_{t+1}	- (0.000)	1.083 (0.000)	0.877 (0.000)	- (0.000)	1.895 (0.000)	1.868 (0.000)	- (0.000)	0.771 (0.000)
NJ_t	- (0.002)	-0.244 (0.002)	-0.453 (0.000)	- (0.177)	-0.135 (0.177)	-0.153 (0.152)	- (0.012)	-0.171 (0.007)
adjusted- R^2	0.785	0.784	0.831	0.753	0.754	0.755	0.877	0.880
							0.824	0.825
								0.826

Table VI
Forecasting average monthly number of jumps, p-values in brackets

	$NJ_{t \rightarrow t+22}$		
	Coefficient	Std. Error	Prob
c	1.091	0.032	
$MA(1)$	1.056	0.032	0.000
$MA(2)$	0.963	0.040	0.000
$MA(3)$	0.919	0.041	0.000
$MA(4)$	0.832	0.038	0.000
$MA(5)$	0.466	0.030	0.000
adjusted- R^2	0.900		

Table VII
Forecasting monthly volatility, p-values in brackets

Independent Variables	$\log MLE-F_{t \rightarrow t+22}$	$\log MLE-F_{t \rightarrow t+22}$	$\log MLE-F_{t \rightarrow t+22}$	$\log MLE-F_{t \rightarrow t+22}$
c	-0.192 (0.104)	0.142 (0.068)	-0.355 (0.097)	-0.357 (0.096)
$\log MLE - F_{t-22 \rightarrow t}$	0.483 (0.073)	0.855 (0.026)	0.487 (0.052)	0.407 (0.054)
$\log VIX_{t-1}$	0.517 (0.086)	-	0.497 (0.067)	0.576 (0.067)
$\overline{NJ}_{t \rightarrow t+22}$	-	0.200 (0.020)	0.196 (0.020)	-
$\widehat{\overline{NJ}}_{t \rightarrow t+22}$	-		-	0.172
adjusted- R^2	0.783	0.806	0.835	0.821
Akaike info	-0.423	-0.532	-0.692	-0.702
Schwarz	-0.413	-0.522	-0.680	-0.686
Hannan-Quinn	-0.419	-0.528	-0.688	-0.697
$\log L$	363.29	455.58	594.06	595.22
Forecasting statistics				
Root Mean Squared Error	0.249			0.237
Mean Absolute Error	0.210			0.194
Mean Absolute Percent Error	2.793			2.591
Theil Inequality Index	0.017			0.016




Figure 1. S&P 500 Lévy and Poisson-type jumps.




Figure 2. QQ plot for raw (blue circles) and filtered (light blue dots) SPY returns




Figure 3. SPY daily volatilities using the MLE , BPV , $TBPV$ and $MLE-F$.




Figure 4. SPY returns autocorrelogram