Modeling CAC40 Volatility Using Ultra-high Frequency Data

Degiannakis, Stavros and Floros, Christos

Department of Economics, Portsmouth Business School, University of Portsmouth, Department of Finance and Insurance, TEI of Crete

2013

Modeling CAC40 Volatility Using Ultra-high Frequency Data

Stavros Degiannakis[a], Christos Floros[a,b]
[a] Department of Economics, Portsmouth Business School, University of Portsmouth, PO1 3DE, 2 UK
[b] Department of Finance and Insurance, TEI of Crete, 72100, 3 Greece

Abstract

Fractionally integrated autoregressive moving average (ARFIMA) and Heterogeneou Autoregressive (HAR) models are estimated and their ability to predict the one-trading-day-ahead CAC40 realized volatility is investigated. In particular, this paper follows three steps: (i) The optimal sampling frequency for constructing the CAC40 realized volatility is examined based on the volatility signature plot. Moreover, the realized volatility is adjusted to the information that flows into the market when it is closed. (ii) We forecast the one-day-ahead realized volatility using the ARFIMA and the HAR models. (iii) The accuracy of the realized volatility forecasts is investigated under the superior predictive ability framework. According to the predicted mean squared error, a simple ARFIMA model provides accurate one-trading day-ahead forecasts of CAC40 realized volatility. The evaluation of model's predictability illustrates that the ARFIMA$(1,d,0)$ forecasts of realized volatility (i) are statistically superior compared to its competing models, and (ii) provide adequate one-trading-day-ahead Value-at-Risk forecasts.

Keywords: intra-day data, long memory, predictability, realized volatility, ultra-high frequency modeling, Value-at-Risk.

JEL Classifications: G17; G15; C15; C32; C53.

1. Introduction

Several studies have used realized volatilities (computed by summing squared returns from ultra-high frequency data over short time intervals during the trading day) due to the consideration of more accurate observations of the actual volatility compared to the traditional variances based on daily frequency data (see Andersen et al., 2001a; Bollerslev et al., 2009; Bollerslev et al., 2011a). Most studies have focused on the theoretical and empirical properties of the realized volatility and conclude that accurate forecasts can be obtained by using ultra-high frequency time-series models (see Andersen and Bollerslev, 1998; Andersen et al., 2001b; Andersen et al., 2005; Ghysels and Sinko, 2006; Koopman et al., 2005; among others).

Autoregressive Fractionally Integrated Moving Average, or ARFIMA, model has been considered to capture the long memory property of the realized volatility; see, for example, Andersen and Bollerslev (1997, 1998), Andersen et al. (2003), Andersen et al. (2005), Thomakos and Wang (2003), Angelidis and Degiannakis (2008), Giot and Laurent (2004), Koopman et al. (2005).

Corsi (2009) suggested the Heterogeneous Autoregressive, or HAR, model which is an autoregressive structure of the volatilities realized over different interval sizes. Its economic interpretation stems from the Heterogenous Market Hypothesis presented by Müller et al. (1993). The basic idea is that market participants have a different perspective of their investment horizon. The heterogeneity, which originates from the difference in the time horizon, creates volatility. The model is represented either in terms of the sum of the realized volatility or in terms of the squared root of the sum of squared realized volatility (Corsi, 2009).

Few papers have studied the realized volatility of CAC40 index. Giot and Laurent (2004) provide evidence of equivalent performance for the daily ARCH type model and the realized volatility ARFIMA model when the 1-day-ahead Value-at-Risk (VaR) for CAC40 is to be computed. Angelidis and Degiannakis (2008) find that a realized volatility ARFIMA model produces more accurate one-day-ahead CAC40 variance forecasts compared to a daily ARCH type model. However, when the performance of the models is investigated in VaR forecasting and simulated option pricing, then a realized volatility ARFIMA model does not provide statistically enhanced predictive power compared to a daily ARCH type model.

Due to data availability and liquidity considerations, we restrict our attention to CAC40 index from Paris stock exchange. The sample considers data from a high volatile period (2000-2009), which is often called as 'lost decade' (see Bollerslev et al., 2011b; Jones, 2011). In addition, the CAC40 index was the poorest performer compared to other major indices from Europe and US[1]; therefore, the current study examines CAC40 as it is the highly volatile major index with a poor performance over the last decade.

Angelidis and Degiannakis (2008, p.464) note that "*The effects of overnight returns and intra-day noise in the high frequency datasets are still an open area of study. An interesting issue for future research is whether different empirical measures of realized volatility affect the evaluation of volatility specifications' predictability*".

In this study, we consider straight forward implemented approaches to (i) construct, (ii) estimate, and (iii) forecast the one-trading-day CAC40 realized volatility using data over the period 2000-2009. To the best of our knowledge, this is the first investigation of comparing intra-day volatility models for the CAC40 index.

The aims of the study are as follows: (i) We construct CAC40 volatility based on one-minute frequency data. For the 13th of June, 2000, to 13th of October, 2009 time period (2392 trading days), the optimal sampling frequency is defined to 7 minutes. Giot and Laurent (2004) show that a sampling frequency of about 15 minutes is optimal for the CAC40; their study considers data for the time period from 3rd of January, 1995 to 31st of December, 1999 (1249 trading days). Hence, our analysis shows that, for different periods in time, there are differences in the optimal sampling frequency.

Further, we compute CAC40 realized volatility accounting for changes in CAC40 during the hours that the stock market is closed, by taking into consideration the inter-day adjustment of realized volatility, proposed by Hansen and Lunde (2005).

(ii) We forecast the one-day-ahead logarithmic realized volatility using the ARFIMA(k,d',l) for $k=0,1,2$ and $l=0,1$, as well as two versions of the HAR model, i.e. HAR with the sum of the realized volatility, and HAR with the squared root of the sum of squared realized volatility.

[1] According to Silbun (2009), the CAC40 (market value-weighted Continuous Assisted Quotation index) is the poorest performer with a 34pc decline, with US Dow Jones shows 8pc decline and German DAX declines 14pc over the decade.

(iii) We conduct a comparison of intra-day volatility models such as the ARFIMA model and the HAR-RV and HAR-sqRV models. We consider these two popular model frameworks, because they are straightforward and easy to implement. At each point in time, the models are re-estimated in order to compute one-trading-day-ahead forecasts of the realized volatility. Then, the forecasts are evaluated measuring the squared distance between actual realized volatility and predicted realized volatility. The loss function computed as the average squared distance between actual and predicted realized volatility provides evidence in favor to the simple ARFIMA model. We conclude that the ARFIMA$(1,d',0)$ model provides the most accurate one-trading day-ahead forecasts of the logarithmic realized volatility, as well as accurate VaR forecasts.

The structure of the paper is as follows: Section 2 provides information about the construction of the intra-day based realized volatility, while Section 3 describes the ARFIMA model as well as the HAR model. Section 4 compares the predictive ability of the realized volatility models based on the predicted mean squared error and investigates the ability of the superior model in forecasting the one-trading-day ahead VaR accurately. Finally, Section 5 summarizes and concludes the paper.

2. Measuring Realized Volatility

Hansen and Lunde (2006) and Patton (2011) showed that the use of a volatility proxy can lead to an evaluation appreciably differing from what would be obtained if the true volatility were used. According to Andersen et al. (2002), realized volatility is widely used in empirical finance as it is straightforward computed and it is a consistent estimator of integrated volatility under general nonparametric conditions. They argue that an important advantage of realized volatility is that it provides asymptotically unbiased measures and therefore approximately serially uncorrelated measurement errors. According to Andersen et al. (2002), "*the realized volatility approach exploiting intraday return observations allow for directly observable return volatility measures that are consistent*".

Under the assumption that the logarithmic price of a financial asset, i.e. $\log(P(t))$, conforms with a diffusion process

$$d\log(P(t)) = \sigma(t)dW(t), \tag{1}$$

with $\sigma(t)$ denoting the volatility of the instantaneous returns process and $W(t)$ being the Wiener process, the *integrated volatility*, $\sigma_t^{2(IV)}$, aggregated over the time interval $(t-1,t)$ is $\sigma_t^{2(IV)} = \int_{t-1}^{t} \sigma^2(x)dx$. The realized volatility, RV_t , based on the theory of quadratic variation[2] of semi-martingales (see Barndorff-Nielsen and Shephard, 2001), is considered a consistent estimator of the integrated volatility. RV_t is defined as the sum of squared returns observed over very small time intervals.

In this paper, we consider ultra-high frequency data[3] to predict the one-trading-day-ahead CAC40 realized volatility. According to Ait-Sahalia et al. (2011, p.161), by considering ultra-high frequency data, we should not compute realized volatility at too high a frequency.

Table 1 presents information about the one-minute intra-day data for the CAC40 index. The dataset is available for 2392 trading days, from 13th of June, 2000, to 13th of October, 2009. The one-trading-day realized variance, at trading day t, based on τ equidistance points in time, is defined as:

$$RV_t^{(\tau)} = \sum_{j=1}^{\tau} \left(\log P_{t_j} - \log P_{t_{j-1}}\right)^2 . \tag{2}$$

Consider the case of a trading day which starts at 09:00 and ends at 15:00, then we get the following: (i) For a sampling frequency of 1 minute, there are $\tau = 361$ equidistance points in time, (ii) for a sampling frequency of 30 minutes, there are $\tau = 13$ equidistance points in time, etc. The sampling frequency should be as high as the market microstructure features do not induce bias to the realized volatility estimator (for details about market microstructure, see Alexander, 2008). The sampling frequency is selected according to the volatility signature plot proposed by Andersen et al. (2006). The volatility signature plot provides a graphical representation of the average realized volatility against the sampling frequency. The accuracy improves as the sampling frequency increases but on the other hand, at a high sampling frequency the market frictions is a source of additional noise in the estimate of volatility. The inter-day variance, $\left(\log P_{t_\tau} - \log P_{t_0}\right)^2$, is decomposed into

---

[2] The theory of quadratic variation as discussed by Andersen et al. (2001b) suggests that realized volatility is an unbiased and highly efficient estimator of return volatility.

[3] According to Engle (2000), "*ultra-high frequency data is defined to be a full record of transactions and their associated characteristics*".

the realized volatility and the intra-day auto-covariances, or $\left(\log P_{t_\tau} - \log P_{t_0}\right)^2 = RV_t^{(\tau)} + 2\sum_{j=1}^{\tau-1}\sum_{i=j+1}^{\tau-1}\left(\log P_{t_{j+1}} - \log P_{t_j}\right)\left(\log P_{t_{j-i+1}} - \log P_{t_{j-i}}\right).$

<< TABLE 1 – HERE >>

<< FIGURE 1 – HERE >>

Further, the optimal sampling frequency is chosen as the highest frequency for which the auto-covariance bias term minimises. In the case of CAC40 index, the optimal sampling frequency is defined to 7 minutes. Figure 1 depicts the volatility signature plot, i.e. the average intra-day auto-covariances against the sampling frequency f, for $f = 1,2,...,40$. We consider the construction of the volatility signature plot which is a useful tool for the financial analysts[4].

Based on Hansen and Lunde (2005), we compute the inter-day adjusted realized volatility as:

$$RV_t^{*(\tau)} = \omega_1\left(\log P_{t_1} - \log P_{t-1_\tau}\right)^2 + \omega_2 RV_t^{(\tau)}. \qquad (3)$$

The $\left(\log P_{t_1} - \log P_{t-1_\tau}\right)^2$ term *measures* the closed-to-open inter-day volatility, whereas the $RV_t^{(\tau)}$ term *measures* the open-to-closed intraday volatility. The parameters ω_1 and ω_2 must be estimated from $\min_{(\omega_1,\omega_2)} E\left(RV_t^{*(\tau)} - \sigma_t^{2(IV)}\right)^2$. As the $\sigma_t^{2(IV)}$ is unobservable, Hansen and Lunde (2005) suggested solving $\min_{(\omega_1,\omega_2)} V\left(RV_t^{*(\tau)}\right)$, as $\arg\min E\left(RV_t^{*(\tau)} - \sigma_t^{2(IV)}\right)^2 = \arg\min V\left(RV_t^{*(\tau)}\right)$. Figure 2 plots the CAC40 one-trading-day realized standard deviation, $\sqrt{RV_t^{*(\tau)}}$, from 13th June 2000 to 13th October 2009, for the optimal sampling frequency of 7 minutes.

Table 2 presents the descriptive statistics of (i) annualized one-trading-day inter-day adjusted realized daily variances, $252RV_t^{*(\tau)}$, (ii) annualized one-trading-day inter-day adjusted realized daily standard deviations, $\sqrt{252RV_t^{*(\tau)}}$, (iii) annualized inter-day adjusted realized daily logarithmic standard deviations, $\log\sqrt{252RV_t^{*(\tau)}}$, and (iv) standardized log-returns, standardized with the annualized one-trading-day inter-day adjusted realized standard deviation, $y_t\big/\sqrt{252RV_t^{*(\tau)}}$, whereas Figure 3

[4] The program that constructs the volatility signature plot from one-minute frequency intra-day dataset is available upon request.

plots the relative estimated densities based on Kernel bandwidths method. We observe that the $\sqrt{252RV_t^{*(\tau)}}$ is approximately log-normally distributed, as well as that the $y_t \big/ \sqrt{252RV_t^{*(\tau)}}$ is approximately unconditionally normally distributed. The findings are in line with Andersen et al. (2000) and 2003, Giot and Laurent (2004).

<< FIGURE 2 – HERE >>

<< FIGURE 3 – HERE >>

<< TABLE 2 – HERE >>

3. Intra-Day Volatility Models

There is significant evidence of long memory in time series, hence most recent studies consider the ARFIMA(k, d', l) model to capture the long memory in realized volatility. The economic motivation behind the ARFIMA model determines from the long memory property of the realized volatility; see for example Andersen et al. (2001a, 2001b, 2003, 2006). Corsi (2009) suggested an alternative autoregressive framework of the volatilities realized over different interval sizes. The economic motivation behind the HAR model stems from the Heterogenous Market Hypothesis, i.e. market participants have a different perspective of their investment horizon. It is widely accepted that the commonly used reduced-form realized volatility models such as the ARFIMA and HAR are superior representations in modeling time-varying volatility (see Corsi et al., 2008). Corsi et al. (2008) report several extensions of the ARFIMA and HAR models and argue that both models have similar performance, but the HAR model might be preferable in practice. Based on the arguments made by Corsi et al. (2008) and the critical review of the major theoretical and empirical developments on realized volatility modeling made by McAleer and Medeiros[5] (2008), we consider these two popular models (ARFIMA and HAR) to test their ability to predict the one-trading-day-ahead CAC40 realized volatility. Both models are straightforward and easy to implement as they often require the estimation of a small number of parameters; therefore, they're are expected to give robust results (see Giot and Laurent, 2004).

[5] McAleer and Medeiros (2008a) extended the HAR-RV model by proposing a flexible multiple regime smooth transition model to capture nonlinearities and long-range dependence in the time series dynamics, but in their recent paper they report several problems in modelling realized volatility (see McAleer and Medeiros, 2008b).

The ARFIMA(k, d', l) model, initially developed by Granger (1980) and Granger and Joyeux (1980), for the logarithmic of the realized volatility, $\log\left(RV_t^{(\tau)}\right)$ is defined as:

$$\left(1 - C(L)\right)\left(1 - L\right)^{d'}\left(\log\left(RV_t^{*(\tau)}\right) - \beta_0\right) = \left(1 + D(L)\right)\varepsilon_t , \tag{4}$$

where $\varepsilon_t \sim N\left(0, \sigma_\varepsilon^2\right)$, $C(L) = \sum_{i=1}^{k} c_i L^i$, $D(L) = \sum_{i=1}^{l} d_i L^i$.

Corsi (2009) suggested the Heterogeneous Autoregressive for the realized volatility (HAR-RV) model,

$$\log\left(RV_t^{*(\tau)}\right) = w_0 + w_1 \log\left(RV_{t-1}^{*(\tau)}\right) + w_2\left(5^{-1}\sum_{j=1}^{5}\log\left(RV_{t-j}^{*(\tau)}\right)\right) + w_3\left(22^{-1}\sum_{j=1}^{22}\log\left(RV_{t-j}^{*(\tau)}\right)\right) + \varepsilon_t , \quad \varepsilon_t \sim N\left(0, \sigma_\varepsilon^2\right), \tag{5}$$

with current trading day's realized volatility explained by the daily, weekly and monthly realized volatilities. The heterogeneity, which originates from the difference in the time horizon, creates volatility.

The HAR-sqRV model can alternatively be represented in terms of the square root of the sum of the realized variances:

$$\log\left(RV_t^{*(\tau)}\right) = w_0 + w_1 \log\left(RV_{t-1}^{*(\tau)}\right) + w_2\sqrt{5^{-1}\sum_{j=1}^{5}\log\left(RV_{t-j}^{*(\tau)}\right)^2} + w_3\sqrt{22^{-1}\sum_{j=1}^{22}\log\left(RV_{t-j}^{*(\tau)}\right)^2} + \varepsilon_t , \quad \varepsilon_t \sim N\left(0, \sigma_\varepsilon^2\right). \tag{6}$$

4. Forecasting CAC40 Realized Volatility

Loss Function - Evaluation Criterion

The ability of the models to forecast the one-trading day-ahead realized volatility is evaluated according to the predicted mean squared error loss function. The predicted mean squared error is the most widely applied loss function in comparing the ability of a model to predict the realized volatility:

$$\overline{\Psi}_{(SE)} = \tilde{T}^{-1}\sum_{t=1}^{\tilde{T}}\left(\log\left(RV_{t+1|t}^{(\tau)}\right) - \log\left(RV_{t+1}^{*(\tau)}\right)\right)^2 , \tag{7}$$

where $RV_{t+1|t}^{(\tau)}$ denotes the trading day's $t+1$ forecasting realized volatility computed at trading day t, and $RV_{t+1}^{*(\tau)}$ denotes the measure of the realized volatility at trading day $t+1$. Because of high non-linearity in volatility models, there is a variety of statistical functions, such as the median absolute error, the heteroskedasticity adjusted mean squared error, the Gaussian likelihood loss function, etc., that measure the distance between actual and predicted volatility (see for example Walsh and Tsou, 1998, Andersen et al., 1999, Saez, 1997, Bollerslev et al., 1994).

The realized volatility, $\log\left(RV_{t+1}^{*(\tau)}\right)$, is a proxy measure for the integrated variance, $\log\left(\sigma_{[t,t+1]}^{2(IV)}\right)$, over the one-day time interval $[t,t+1]$. Hansen and Lunde (2006) derived conditions which ensure that the ranking of any two variance forecasts by a loss function is the same (i.e. *consistent ranking*) whether the ranking is done via the true and unobserved variance, $\log\left(\sigma_{[t,t+1]}^{2(IV)}\right)$, or via a conditionally unbiased volatility proxy such as the realized volatility, $\log\left(RV_{t+1}^{*(\tau)}\right)$. A sufficient condition for a consistent ranking is that, for the loss function $\overline{\Psi}\left(\log\left(\sigma_{[t,t+1]}^{2(IV)}\right),\log\left(RV_{t+1|t}^{(\tau)}\right)\right)$, the quantity $\dfrac{\partial^2\overline{\Psi}\left(\log\left(\sigma_{[t,t+1]}^{2(IV)}\right),\log\left(RV_{t+1|t}^{(\tau)}\right)\right)}{\partial\left(\log\left|\sigma_{[t,t+1]}^{2(IV)}\right|\right)^2}$ does not depend on $\log\left(RV_{t+1|t}^{(\tau)}\right)$. The predicted mean squared error loss function ensures the equivalence of the ranking of volatility models.

Numerous forecast evaluation criteria exist in the literature but none is generally acceptable. For a detailed investigation about evaluation of volatility models, see Xekalaki and Degiannakis (2010, p.357). Of course another approach could be the evaluation of the forecast density instead of the point forecasts. More information about density forecasts is available in Berkowitz (2001) and Diebold et al. (1998).

One-Trading-Day Ahead Volatility Forecasts

The ARFIMA(k,d',l) model is estimated, for $k=0,1,2$, and $l=0,1$. In total, 8 models are considered. Each model is re-estimated every trading day, for $\tilde{T}=1392$ days, based on a rolling sample of $\breve{T}=1000$ days. The one-day-ahead logarithmic realized volatility forecasts, $\log\left(RV_{t+1|t}^{(\tau)}\right)$, are computed as:

ARFIMA(k,d',l)

$$\log\left(RV_{t+1|t}^{(\tau)}\right)=\beta_0^{(t)}\left(1-\sum_{i=1}^{k}c_i^{(t)}\right)+\sum_{i=1}^{k}c_i^{(t)}L^i\log\left(RV_{t-i+1|t}^{(\tau)}\right)+(1-L)^{-d'^{(t)}}\sum_{i=1}^{l}d_i^{(t)}L^i\varepsilon_{t-i+1|t}, \tag{8}$$

where[6] $(1-L)^{-d'^{(t)}}=1+\frac{1}{1!}d'^{(t)}L+\frac{1}{2!}d'^{(t)}\left(1+d'^{(t)}\right)L^2-\ldots$.

HAR-RV

$$\log\left(RV_{t+1|t}^{(\tau)}\right)=w_0^{(t)}+w_1^{(t)}\log\left(RV_t^{*(\tau)}\right)+w_2^{(t)}5^{-1}\sum_{j=1}^{5}\log\left(RV_{t-j+1}^{*(\tau)}\right)+w_3^{(t)}22^{-1}\sum_{j=1}^{22}\log\left(RV_{t-j+1}^{*(\tau)}\right), \tag{9}$$

HAR-sqRV

$$\log\left(RV_{t+1|t}^{(\tau)}\right)=w_0^{(t)}+w_1^{(t)}\log\left(RV_{t-1}^{*(\tau)}\right)+w_2^{(t)}\sqrt{5^{-1}\sum_{j=1}^{5}\log\left(RV_{t-j+1}^{*(\tau)}\right)^2}+w_3^{(t)}\sqrt{22^{-1}\sum_{j=1}^{22}\log\left(RV_{t-j+1}^{*(\tau)}\right)^2}, \tag{10}$$

Consider for example the ARFIMA$(1,d,1)$ model; the parameter vector to be estimated at each trading day t is $\theta^{(t)}=\left(d'^{(t)},c_1^{(t)},d_1^{(t)},\beta_0^{(t)}\right)'$. For each model the vector $\theta^{(t)}$ is re-estimated every trading day, for $t=\breve{T},\breve{T}+1,\ldots,\breve{T}+\tilde{T}-1$ days.

Table 3 presents, for the 8 models, the average of the squared one-step-ahead prediction errors, $\tilde{T}^{-1}\sum_{t=1}^{\tilde{T}}\left(\log\left(RV_{t+1|t}^{(\tau)}\right)-\log\left(RV_{t+1}^{(\tau)}\right)\right)^2$, whereas Figure 4 depicts the one-trading-day-ahead logarithmic realized volatility forecast from the ARFIMA$(1,d',0)$ model, for the total of the $\tilde{T}=1392$ trading days. The ARFIMA$(1,d',0)$ has the lowest value of the loss function, 0.272630.

<< TABLE 3 – HERE >>

<< FIGURE 4 – HERE >>

[6] The ARFIMA(k,d',l) model can be written as:

$$\left(1-\sum_{i=1}^{k}c_iL^i\right)(1-L)^{d'}\log\left(RV_t^{*(\tau)}\right)=\left(1-\sum_{i=1}^{k}c_i\right)\beta_0+\left(1+\sum_{i=1}^{l}d_iL^i\right)\varepsilon_t\rightarrow$$

$$\left(1-\sum_{i=1}^{k}c_iL^i\right)\log\left(RV_t^{*(\tau)}\right)=\beta_0\left(1-\sum_{i=1}^{k}c_i\right)+(1-L)^{-d'}\varepsilon_t+(1-L)^{-d'}\sum_{i=1}^{l}d_iL^i\varepsilon_t\rightarrow$$

$\log\left(RV_t^{*(\tau)}\right)=\sum_{i=1}^{k}c_iL^i\log\left(RV_t^{*(\tau)}\right)+\beta_0\left(1-\sum_{i=1}^{k}c_i\right)+(1-L)^{-d'}\varepsilon_t+(1-L)^{-d'}\sum_{i=1}^{l}d_iL^i\varepsilon_t$. If we denote the estimated parameters at time t as $\left(d'^{(t)},c_1^{(t)},\ldots,c_k^{(t)},d_1^{(t)},\ldots d_l^{(t)},\beta_0^{(t)}\right)'$, then the $\log\left(RV_{t+1|t}^{(\tau)}\right)$ is computed as

$$\log\left(RV_{t+1|t}^{(\tau)}\right)=\sum_{i=1}^{k}c_i^{(t)}L^i\log\left(RV_{t-i+1|t}^{(\tau)}\right)+\beta_0^{(t)}\left(1-\sum_{i=1}^{k}c_i^{(t)}\right)+(1-L)^{-d'^{(t)}}\sum_{i=1}^{l}d_i^{(t)}L^i\varepsilon_{t-i+1|t}$$

Volatility Forecasting Ability Evaluation

Hansen (2005) extended the work of White (2000) and provided a framework for comparing statistically the predictive ability of a seemingly best performing forecasting model, from a large set of potential models. Hansen's test, named Superior Predictive Ability, or SPA, test, investigates the null hypothesis that the best performing forecasting model is not outperformed by the competing models against the alternative hypothesis that best performing forecasting model is inferior to one or more of its competing models.

Let us denote as $\Psi_{t(SE)}^{(i)}$ the value of the predicted squared error of model i at time t, or $\Psi_{t(SE)}^{(i)} = \left(\log\left(RV_{(i)t+1|t}^{(\tau)}\right) - \log\left(RV_{t+1}^{*(\tau)}\right)\right)^2$. The best performing forecasting model i^* is tested against the i, for $i = 1,...,M = 7$, competing models. The SPA statistic equals to:

$$SPA = \max_{i=1,\ldots,M}\left(\sqrt{M}\left(\tilde{T}^{-1}\sum_{t=1}^{\tilde{T}}\Psi_{t(SE)}^{(i^*)} - \Psi_{t(SE)}^{(i)}\right)\left(V\left(\sqrt{M}\left(\tilde{T}^{-1}\sum_{t=1}^{\tilde{T}}\Psi_{t(SE)}^{(i^*)} - \Psi_{t(SE)}^{(i)}\right)\right)\right)^{-1/2}\right).$$ [7]

A high p-value indicates evidence in support of the hypothesis that the benchmark model is superior to one or more of the opponent models. As the p-value of the test is 0.8577, it appears that there is evidence supporting the hypothesis that the forecasting ability of the ARIMA$(1,d,0)$ model is superior to its competitors[8].

Under the assumption of $\varepsilon_t \sim N\left(0,\sigma_\varepsilon^2\right)$, the quantity e^{ε_t} is log-normally distributed. Hence the unbiased estimator of the realized variance is computed as $RV_{(un),t|t}^{(\tau)} = \exp\left(\log\left(RV_{t|t}^{(\tau)}\right) + 0.5\sigma_\varepsilon^2\right)$. Moreover, the one-day-ahead realized variance forecasts are computed as:

$$RV_{(un),t+1|t}^{(\tau)} = \exp\left(\log\left(RV_{t+1|t}^{(\tau)}\right) + \frac{1}{2}\sigma_\varepsilon^{(t)2}\right). \quad (11)$$

Figure 5 plots the CAC40 one-trading-day-ahead realized standard deviation forecast, $\sqrt{RV_{(un),t+1|t}^{(\tau)}}$, from the ARFIMA$(1,d',0)$, for the $\tilde{T} = 1392$ trading days.

<< FIGURE 5 – HERE >>

[7] The value of $V\left(\sqrt{M}\left(\tilde{T}^{-1}\sum_{t=1}^{\tilde{T}}\Psi_{t(SE)}^{(i^*)} - \Psi_{t(SE)}^{(i)}\right)\right)$ and the p-values of the SPA statistic are estimated by Politis and Romano’s (1994) stationary bootstrap method.

[8] The results obtained from the Superior Predictive Ability test are available upon request.

VaR Forecasting - A Financial Application

An important application of volatility forecasts is the prediction of the Value-at-Risk (VaR) measure. VaR quantifies the maximum loss for a portfolio of assets under normal market conditions over a given period of time and at a certain confidence level (95% or 99%). Having estimated the one-trading-day realized volatility, the one-trading-day ahead VaR can be calculated as:

$$VaR_{t+1|t}^{(1-\rho)} = N(\rho)RV_{(un),t+1|t}^{(\tau)}, \tag{8}$$

where $N(\rho)$ is the ρ^{th} quantile of the standard normal distribution. In the continuous time case of equation (1) the instantaneous returns and the integrated volatility process are related via the Wiener process[9]. Therefore, in the discrete time case, the daily log-returns and the realized volatility would be related via the normal distribution. The observed 95% VaR failure rate[10] is 3.95%, whereas the 99% VaR failure rate is 1.15%. Figures 6 and 7 plot the one-trading-day ahead 95% and 99% VaR forecasts, respectively. Based on Kupiec (1995) test we examine the null hypothesis that the observed violation rate, $N/\tilde{T}$, for N denoting the number of days on which a violation occurred, is statistically equal to the expected violation rate, ρ. The likelihood ratio statistic equals to $LR_{UC} = 2\log\left(\left(1-\frac{N}{\tilde{T}}\right)^{\tilde{T}-N}\left(\frac{N}{\tilde{T}}\right)^{N}\right) - 2\log\left((1-\rho)^{\tilde{T}-N}\rho^{N}\right)$ and is chi-squared distributed with one degree of freedom. The p-values of Kupiec test for 95% and 99% VaR are 0.0627 and 0.5841, respectively. We also conduct Christoffersen's (1998) test which examines the null hypothesis that the VaR failures are independently distributed over time against the alternative hypothesis that the failures tend to be clustered. The likelihood ratio statistic is computed as

[9] I.e. $W(t)-W(t') \sim \sqrt{t-t'}N(0,1)$.

[10] Percentage of trading days that the log-returns are lower than the VaR measure.

$$LR_{IN} = 2\log\left((1-\pi_{01})^{n_{00}}\pi_{01}^{\ n_{01}}(1-\pi_{11})^{n_{10}}\pi_{11}^{\ n_{11}}\right) - 2\log\left(\left(1-\frac{N}{\tilde{T}}\right)^{n_{00}+n_{10}}\left(\frac{N}{\tilde{T}}\right)^{n_{01}+n_{11}}\right),$$ and is chi-squared distributed with one degree of freedom[11]. The p-values of Christoffersen test for 95% and 99% VaR are 0.3569 and 0.5417, respectively[12]. Both tests provide p-values which do not reject the hypotheses i) that the observed violation rate is statistically equal to the expected violation rate as well as ii) that the VaR failures are independently distributed over time.

The period analyzed spans the start of the financial crisis, since January, 2008. According to Figure 2, apparently, the realized volatility is higher in the period of the crisis than during the period of 2000-2007. The descriptive statistics inform us that, for the period 2000-2007, the mean and median of the annualized realized volatility, $\sqrt{252RV_t^{*(\tau)}}$, are 12.08% and 10.87%, respectively. On the other hand, during the financial crisis, the mean and median of $\sqrt{252RV_t^{*(\tau)}}$, are 29.97% and 25.72%, respectively.

However, focusing on the period from January, 2008 up to October, 2009, the p-values of Kupiec test for 95% and 99% VaR are 0.6328 and 0.2829, respectively. Therefore, the ARFIMA($1,d',0$) model does not fail to provide accurate one-trading-day ahead VaR forecasts, during the financial crisis, despite the increase of the magnitude of volatility. For more information about the performance of widely-accepted approaches to estimate VaR before and after the financial crisis of 2008 you are referred to Degiannakis et al. (2012).

<< FIGURE 6 – HERE >>

<< FIGURE 7 – HERE >>

5. Conclusion

Many high-frequency forecasting approaches have been proposed (Bollerslev et al., 1994; Andersen et al., 2003) but those approaches suffer from a lack of inter-day adjusted realized volatility measures. Hence it is interesting to see straight

---

[11] n_{ij} is the number of observations with value i followed by j, for $i,j=0,1$, and $\pi_{ij} = n_{ij}\left(\sum_j n_{ij}\right)^{-1}$ are the corresponding probabilities. For $i,j=1$ a violation occurred, π_{ij} indicates the probability that $j(=0,1)$ occurs at time t, given that $i(=0,1)$ occurred at time $t-1$.

[12] Pearson's goodness-of-fit test and Engle and Manganelli's (2004) dynamic quantile test provide qualitatively similar results.

forward implemented frameworks for volatility forecasting using ultra-high frequency datasets.

In this paper we (i) construct, (ii) estimate, and (iii) forecast the one-trading-day CAC40 realized volatility, using data over the period 2000-2009, by employing ARFIMA and HAR models.

According to the robust loss function that measures the squared distance between forecast and actual realized volatility, the ARFIMA$(1,d',0)$ model provides, for the CAC40 index, the most accurate one-trading day-ahead forecasts of the logarithmic realized volatility. The ARFIMA$(1,d',0)$ model not only minimizes the predictive mean squared error loss function but it is superior to its competitors in terms of forecasting accuracy. The present paper provides evidence in favor to a simple ARFIMA model for predicting the one-trading day-ahead CAC40 logarithmic realized volatility.

Our findings should be of direct interest to market participants, analysts and policymakers who deal with ultra-high frequency intraday datasets. For further research, we suggest the comparison of a wider set of volatility models. The volatility of realized volatility, which can be consider as an estimate of the integrated quarticity, may also exhibit time-variation (for details about integrated quarticity, see Barndorff-Nielsen and Shephard, 2006). Our approach can easily be extended to allow for the adequacy of the first order autoregressive long memory model to provide adequate forecasts of realized volatility for other financial assets (indices, stocks, and exchange rates).

References

Ait-Sahalia, Y., Mykland, P.A. and Zhang, L., 2011. Ultra high frequency volatility estimation with dependent microstructure noise. *Journal of Econometrics*. 160, 160-175.

Alexander, C.O., 2008. Market Risk Analysis: Quantitative Methods in Finance. Volume 1. John Wiley and Sons, New York.

Andersen, T. and Bollerslev, T., 1997. Intraday Periodicity and Volatility Persistence in Financial Markets. *Journal of Empirical Finance*. 4, 115-158.

Andersen, T. and Bollerslev, T., 1998. DM-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements and Longer-Run Dependencies. *Journal of*

Finance. 53, 219-265.

Andersen, T., Bollerslev, T. and Lange, S., 1999. Forecasting Financial Market Volatility: Sample Frequency vis-à-vis Forecast Horizon. *Journal of Empirical Finance*. 6, 457-477.

Andersen, T., Bollerslev, T., Diebold, F.X. and Ebens, H., 2001a. The distribution of realized stock return volatility. *Journal of Financial Economics*. 61, 43-76.

Andersen, T., Bollerslev, T., Diebold, F.X. and Labys, P., 2001b. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*. 96, 42-55.

Andersen, T., Bollerslev, T. and Diebold, F.X., 2002. Parametric and Nonparametric Volatility Measurement. In (eds.) Lars Peter Hansen and Yacine Aït-Sahalia, *Handbook of Financial Econometrics*. Amsterdam, North Holland.

Andersen, T., Bollerslev, T., Diebold, F.X. and Labys, P., 2003. Modeling and Forecasting Realized Volatility. *Econometrica*. 71, 529-626.

Andersen, T., Bollerslev, T. and Meddahi, N., 2005. Correcting the errors: volatility forecast evaluation using high-frequency data and realized volatilities. *Econometrica*. 73, 279-296.

Andersen, T., Bollerslev, T., Christoffersen, P. and Diebold, F.X., 2006. Volatility and Correlation Forecasting. In (eds.) Elliott, G. Granger, C.W.J. and Timmermann, A. *Handbook of Economic Forecasting*. North Holland Press, Amsterdam.

Angelidis, T. and Degiannakis, S., 2008. Volatility Forecasting: Intra-day vs. Inter-day Models. *Journal of International Financial Markets, Institutions and Money*. 18, 449-465.

Barndorff-Nielsen, O.E. and Shephard, N., 2001. Non-Gaussian Ornstein-Uhlenbeck based Models and Some of their Uses in Financial Economics. *Journal of the Royal Statistical Society, Series B*. 63, 197-241.

Barndorff-Nielsen, O.E. and Shephard, N., 2006. Econometrics of Testing for Jumps in Financial Economics using Bipower Variation. *Journal of Financial Econometrics*. 4(1), 1-30.

Berkowitz, J., 2001. Testing Density Forecasts, with Applications to Risk Management. *Journal of Business and Economic Statistics*. 19, 465-474.

Bollerslev, T., Engle, R.F. and Nelson, D., 1994. ARCH Models, in (eds.) Engle, R.F. and McFadden, D. *Handbook of Econometrics*. Volume 4, Elsevier Science, Amsterdam, 2959-3038.

Bollerslev, T., Tauchen, G. and Zhou, H., 2009. Expected stock returns and variance risk premia. *Review of Financial Studies*. 22, 4463-4492.

Bollerslev, T., Gibson, M. and Zhou, H., 2011a. Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. *Journal of Econometrics*. 160, 235-245.

Bollerslev, T., Marrone, J., Xu, L. and Zhou, H. 2011b. Stock Return Predictability and Variance Risk Premia: Statistical Inference and International Evidence, Federal Reserve Board, Finance and Economics Discussion Series, 2011-52, Washington D.C.

Corsi, F., 2009. A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics*. 7(2), 174-196.

Corsi, F., Mittnik, S. Pigorsch, C. and Pigorsch, U., 2008. The Volatility of Realised Volatility. *Econometric Reviews*. 27(1-3), 46-78.

Christoffersen, P., 1998. Evaluating Interval Forecasts. *International Economic Review*. 39, 841-862.

Degiannakis, S., Floros, C. and Livada, A., 2012. Evaluating Value-at-Risk Models before and after the Financial Crisis of 2008: International Evidence. *Managerial Finance*, 38(4), 436-452.

Diebold, F.X., Gunther, T.A. and Tay, A.S., 1998. Evaluating Density Forecasts with Applications to Financial Risk Management. *International Economic Review*. 39(4), 863-883.

Engle, R.F., 2000. The Econometrics of ultra-high-frequency data. *Econometrica*. 68, 1-22.

Engle, R.F. and Manganelli, S., 2004. CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles. *Journal of Business and Economic Statistics*. 22(4), 367-381.

Ghysels, E. and Sinko, A., 2006. Comment. *Journal of Business and Economic Statistics*. 24, 192-194.

Giot, P. and Laurent, S., 2004. Modelling Daily Value-at-Risk Using Realized Volatility and ARCH Type Models. *Journal of Empirical Finance*. 11, 379 – 398.

Granger, C.W.J., 1980. Long Memory Relationships and the Aggregation of Dynamic Models. *Journal of Econometrics*. 14, 227-238.

Granger, C.W.J. and Joyeux, R., 1980. An Introduction to Long Memory Time Series Models and Fractional Differencing. *Journal of Time Series Analysis*. 1, 15-39.

Hansen, P.R., 2005. A Test for Superior Predictive Ability. *Journal of Business and Economic Statistics*. 23, 365-380.

Hansen, P.R. and Lunde, A., 2005. A Realized Variance for the Whole Day Based on Intermittent High-Frequency Data. *Journal of Financial Econometrics*. 3(4), 525-554.

Hansen, P.R. and Lunde, A., 2006. Consistent Ranking of Volatility Models. *Journal of Econometrics*. 131, 97-121.

Jones, C.P., 2011. Can Recent Long-Term Investors Recover from Their 2000–2009 Stock Losses? *Journal of Investing*, 20(2), 9-14.

Koopman, S., Jungbacker, B., and Hol, E., 2005. Forecasting Daily Variability of the S&P100 Stock Index Using Historical, Realised and Implied Volatility Measurements. *Journal of Empirical Finance*. 12, 445–475.

Kupiec, P.H., 1995. Techniques for Verifying the Accuracy of Risk Measurement Models. *Journal of Derivatives*. 3, 73-84.

McAleer, M., Medeiros, M. C., 2008a. A Multiple Regime Smooth Transition Heterogeneous Autoregressive Model for Long Memory and Asymmetries. *Journal of Econometrics*. 147, 104-109.

McAleer, M. and Medeiros, M.C., 2008b. Realized Volatility: A Review, *Econometric Reviews*. 27(1), 10-45.

Müller, U.A., Dacorogna, M.M., Davé, R.D., Pictet, O.V., Olsen, R.B., Ward, J.R., 1993. Fractals and Intrinsic Time - A Challenge to Econometricians. International AEA Conference on Real Time Econometrics, 14-15 October 1993, Luxembourg.

Patton, A.J., 2011. Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*. 160, 246-256.

Politis, D.N. and Romano, J.P., 1994. The Stationary Bootstrap. *Journal of the American Statistical Association*. 89, 1303-1313.

Saez, M., 1997. Option Pricing Under Stochastic Volatility and Stochastic Interest Rate in the Spanish Case. *Applied Financial Economics*. 7, 379-394.

Silbun, J., 2009. FTSE 100's recent rally fails to make up for a lost decade, *The Telegraph*, 31 December 2009, U.K.

Thomakos, D.D. and Wang, T., 2003. Realized Volatility in the Futures Markets. *Journal of Empirical Finance*. 10, 321-353.

Walsh, D.M. and Tsou, G.Y-G., 1998. Forecasting Index Volatility: Sampling Interval and Non-trading Effects. *Applied Financial Economics*. 8, 477, 485.

White, H., 2000. A Reality Check for Data Snooping. *Econometrica*. 68, 1097–1126.

Xekalaki, E. and Degiannakis, S., 2010. ARCH Models for Financial Applications. John Wiley & Sons Ltd., New York.

Figures & Tables

Figure 1. CAC40 volatility signature plot. Average daily squared log-returns and average intra-day auto-covariance against sampling frequencies of $f = 1,2,...,40$ minutes.

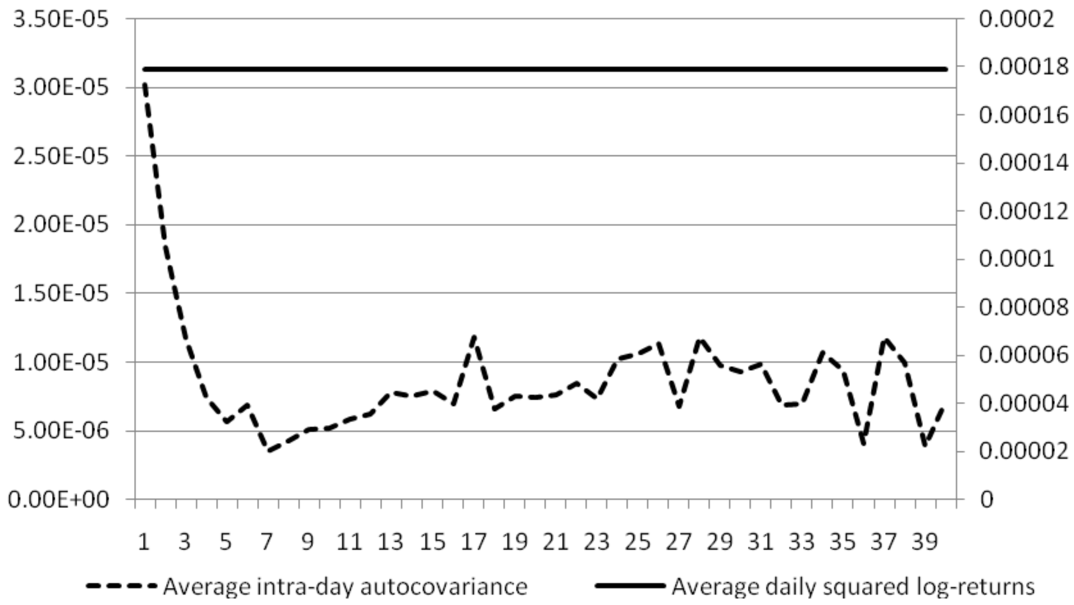

Figure 2. The CAC40 one-trading-day realized standard deviation, $\sqrt{RV_t^{*(\tau)}}$, from 13th June 2000 to 13th October 2009. Optimal sampling frequency of 7 minutes.

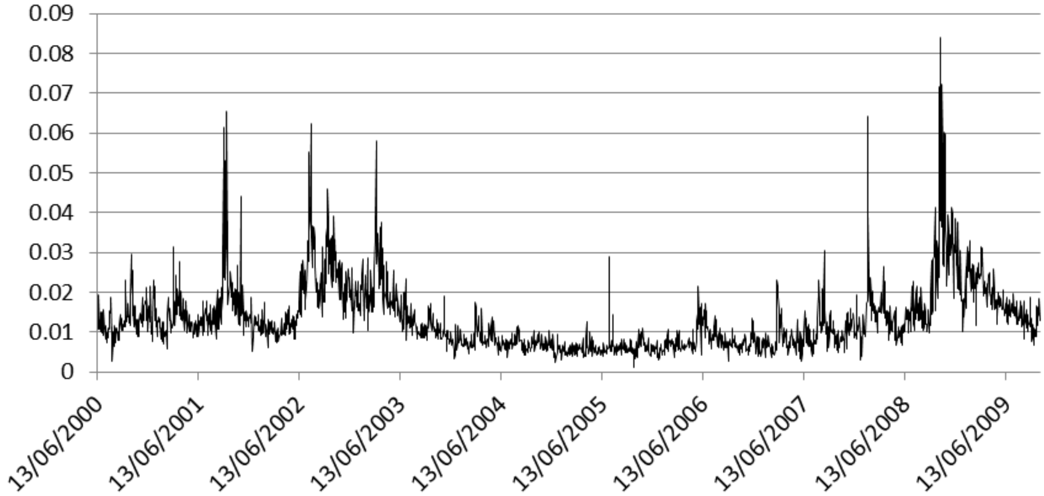

Figure 3. The estimated density of (i) annualized one-trading-day inter-day adjusted realized daily variances, $252RV_t^{*(\tau)}$, (ii) annualized one-trading-day inter-day adjusted realized daily standard deviations, $\sqrt{252RV_t^{*(\tau)}}$, (iii) annualized inter-day adjusted realized daily logarithmic standard deviations, $\log\sqrt{252RV_t^{*(\tau)}}$, and (iv) standardized log-returns, standardized with the annualized one-trading-day inter-day adjusted realized standard deviation, $y_t\big/\sqrt{252RV_t^{*(\tau)}}$.

(i)

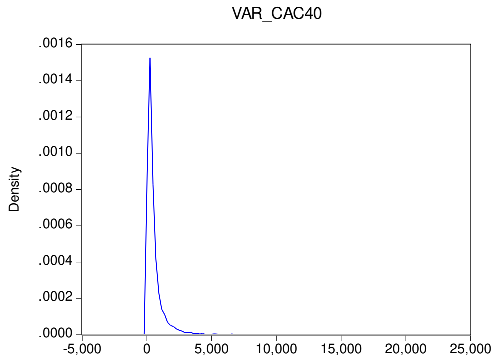

(ii)

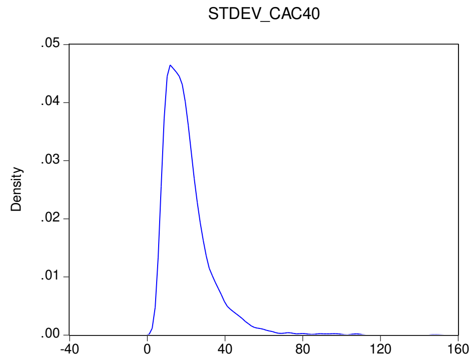

(iii)

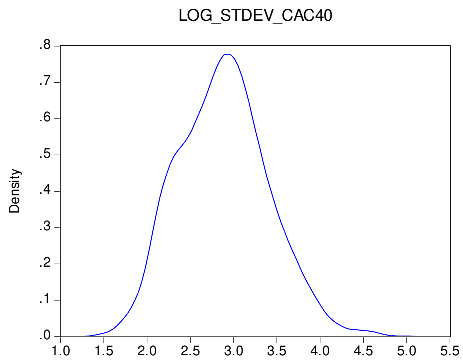

(iv)

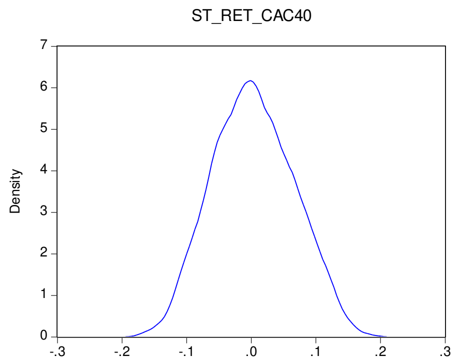

Figure 4. The CAC40 one-trading-day-ahead logarithmic realized volatility, $\log\left(RV_{t+1|t}^{(\tau)}\right)$, from the ARFIMA$(1,d',0)$ model, for $\tilde{T}=1392$ trading days.

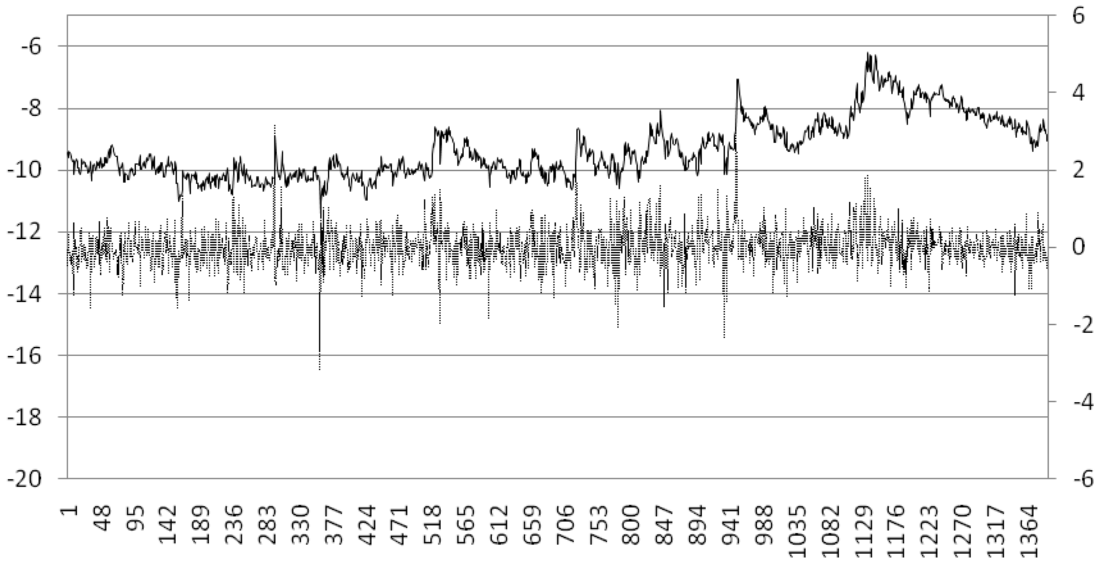

The dotted line presents the difference $\log\left(RV_{t+1}^{*(\tau)}\right)-\log\left(RV_{t+1|t}^{(\tau)}\right)$.

Figure 5. The CAC40 one-trading-day-ahead realized standard deviation forecast, $\sqrt{RV_{(un),t+1|t}^{(\tau)}}$, from the ARFIMA$(1,d',0)$ model, for $\tilde{T}=1392$ trading days.

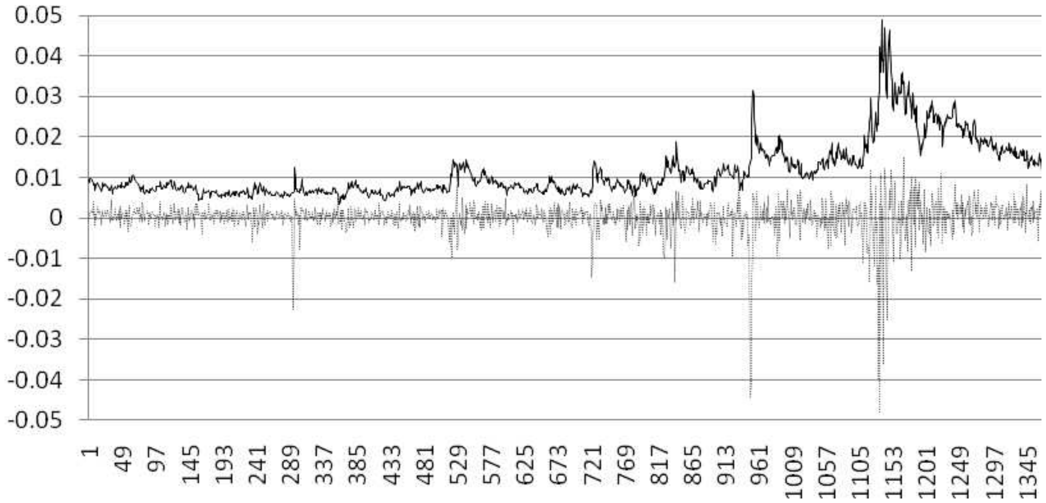

The dotted line presents the difference $\sqrt{RV_{(un),t+1|t}^{(\tau)}}-\sqrt{RV_{t+1}^{*(\tau)}}$.

Figure 6. The CAC40 daily log-returns against the one-trading-day ahead 95% VaR forecasts, from the ARFIMA$(1, d', 0)$ model, for $\tilde{T} = 1392$ trading days.

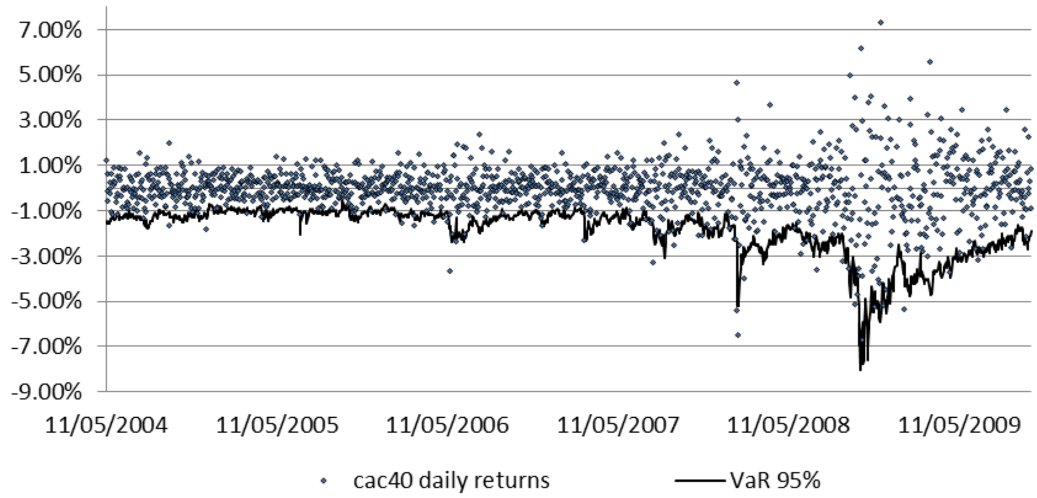

The line presents the 95% VaR, and the dots present the daily log-returns.

Figure 7. The CAC40 daily log-returns against the one-trading-day ahead 99% VaR forecasts, from the ARFIMA$(1,d',0)$ model, for $\tilde{T}=1392$ trading days.

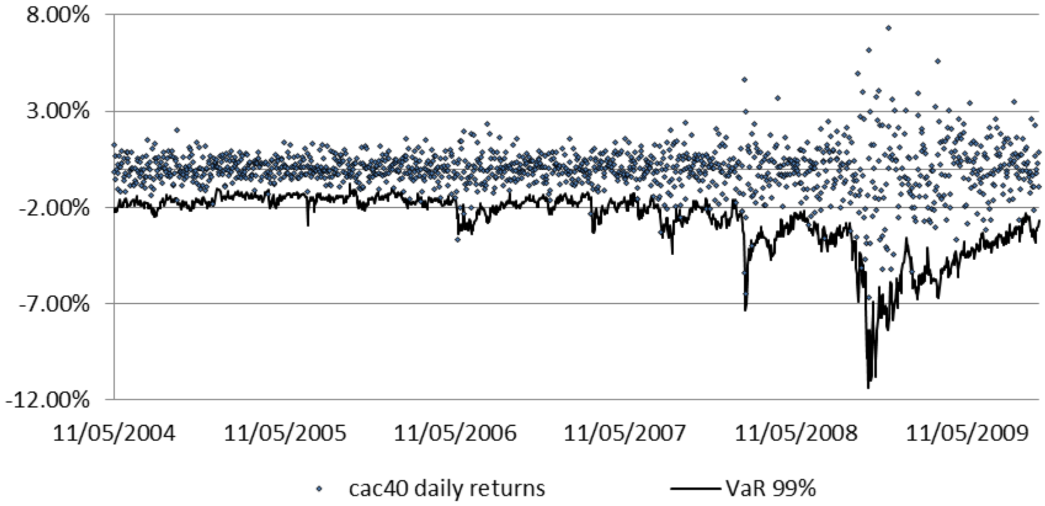

The line presents the 99% VaR, and the dots present the daily log-returns.

Table 1. Information about the intra-day data.

Index	Number of intra-day (1 minute) observations	Number of days	First day	Last day	*Optimal* sampling frequency
CAC40	1.285.783	2392	13th June 2000	13th October 2009	7 minutes

Table 2. Descriptive statistics of (i) annualized one-trading-day inter-day adjusted realized daily variances, $252RV_t^{*(\tau)}$, (ii) annualized one-trading-day inter-day adjusted realized daily standard deviations, $\sqrt{252RV_t^{*(\tau)}}$, (iii) annualized inter-day adjusted realized daily logarithmic standard deviations, $\log\sqrt{252RV_t^{*(\tau)}}$, and (iv) standardized log-returns, standardized with the annualized one-trading-day inter-day adjusted realized standard deviation, $y_t \big/ \sqrt{252RV_t^{*(\tau)}}$.

	Mean	Median	Maximum	Minimum	Std.Dev	Skewness	Kurtosis
$252RV_t^{*(\tau)}$	578.8	317.4	21924.7	16.6	993.7	8.1	115.2
$\sqrt{252RV_t^{*(\tau)}}$	20.6	17.8	148.1	4.1	12.5	2.5	14.6
$\log\sqrt{252RV_t^{*(\tau)}}$	2.88	2.88	5.00	1.40	0.52	0.27	3.01
$y_t \big/ \sqrt{252RV_t^{*(\tau)}}$	0.003	0.001	0.189	-0.177	0.06	0.08	2.56

Table 3. The average of the squared one-step-ahead prediction errors, $\tilde{T}^{-1}\sum_{t=1}^{\tilde{T}}\left(\log\left(RV_{t+1|t}^{(\tau)}\right)-\log\left(RV_{t+1}^{*(\tau)}\right)\right)^2$, for $\tilde{T}=1392$.

ARFIMA $(0,d',0)$	0.274550	ARFIMA $(2,d',0)$	0.272767
ARFIMA $(0,d',1)$	0.272637	ARFIMA $(2,d',1)$	0.273602
ARFIMA $(1,d',0)$	0.272630	HAR-RV	0.273167
ARFIMA $(1,d',1)$	0.272686	HAR-sqRV	0.273706