

Modelling daily value-at-risk using realized volatility and arch type models

Citation for published version (APA):

Giot, P., & Laurent, S. F. J. A. (2001). *Modelling daily value-at-risk using realized volatility and arch type models*. METEOR, Maastricht University School of Business and Economics. METEOR Research Memorandum No. 026 <https://doi.org/10.26481/umamet.2001026>

Document status and date:

Published: 01/01/2001

DOI:

[10.26481/umamet.2001026](https://doi.org/10.26481/umamet.2001026)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

MODELLING DAILY VALUE-AT-RISK USING REALIZED VOLATILITY AND ARCH TYPE MODELS

Pierre Giot¹ and Sébastien Laurent^{2,3}

April 2001

Abstract

In this paper we show how to compute a daily VaR measure for two stock indexes (CAC40 and SP500) using the one-day-ahead forecast of the daily realized volatility. The daily realized volatility is equal to the sum of the squared intraday returns over a given day and thus uses intraday information to define an aggregated daily volatility measure. While the VaR specification based on an ARFIMAX(0, d ,1)-skewed Student model for the daily realized volatility provides adequate one-day-ahead VaR forecasts, it does not really improve on the performance of a VaR model based on the skewed Student APARCH model and estimated using daily data. Thus, for the two financial assets considered in an univariate framework, both methods seem to be equivalent. This paper also shows that daily returns standardized by the square root of the one-day-ahead forecast of the daily realized volatility are not normally distributed.

Keywords: Value-at-Risk, Realized volatility, skewed Student distribution, APARCH

JEL classification: C52, C53, G15

¹Department of Quantitative Economics, Maastricht University and Center for Operations Research and Econometrics, UCL; email: giot@core.ucl.ac.be or p.giot@ke.unimaas.nl

²Département des Sciences Economiques, Université de Liège and Department of Quantitative Economics, Maastricht University. email: S.Laurent@ulg.ac.be or S.Laurent@ke.unimaas.nl

³Corresponding author.

1 Introduction

The recent widespread availability of databases recording the intraday price movements of financial assets (stocks, indexes, currencies, derivatives) has led to new developments in applied econometrics and quantitative finance as far as the modelling of daily and intradaily volatility is concerned. Focusing solely on the modelling of daily volatility using intraday data, the recent literature suggests at least three possible methods for characterizing volatility and risk at an aggregated level, which we take to be equal to one day in this paper.

The first possibility is to sample the intraday data on a daily basis so that closing prices are recorded from which daily returns are computed. In this setting, the notion of intraday price movements is not an issue, as the method is tantamount to estimating a volatility model on daily data. One of the most famous example is the ARCH model of Engle (1982) and subsequent ARCH type models such as the GARCH model of Bollerslev (1986) (see Palm, 1996, for a recent survey).

The second method is based on the notion of realized volatility which was recently introduced in the literature by Taylor and Xu (1997) and Andersen and Bollerslev (1998) and which is grounded in the framework of continuous time finance with the notion of quadratic variation of a martingale. In this case, a daily measure of volatility is computed as an aggregated measure of volatility defined on intraday returns. More specifically, the daily realized volatility is computed as the sum of the squared intraday returns for the given trading day. We thus make explicit use of the intraday returns to compute the realized volatility, from which the daily volatility is modelled. A third possibility is to estimate a high frequency duration model on price durations for the given asset, and then use this irregularly time-spaced volatility at the aggregated level. Examples are Engle and Russell (1997) or Giot (2000). In this paper we focus on the first two methods as our aggregation level is equal to one day, and it is not clear how duration models could be of any help in this situation.

The recent literature on realized volatility and the huge literature on daily volatility models seem to indicate that a researcher or market practitioner faces two distinct possibilities when daily volatility is to be modelled. Going one way or the other is however not a trivial question. If one decides to model daily volatility using daily realized volatility, then intraday data are needed so that corresponding intraday returns can be computed. Even today, intraday data remain relatively costly and are not readily available for all assets. Furthermore, a large amount of data handling and computer programming is usually needed to retrieve the needed intraday returns from the raw data files supplied by the exchanges or data vendors. On the contrary, working with daily data is relatively simple and the data are broadly available. However, one has the feeling that all the relevant data are not taken into account, i.e. that by going at the intraday level one could get a much better model.

In this paper we aim to address this issue by comparing the performance of a daily ARCH type model with the performance of a model based on the daily realized volatility when the one-step ahead Value-at-Risk (VaR) measure is to be computed for a stock or market index. This exercise is done for two stock indexes (French CAC40 and US SP500 indexes) for which intraday data are available over a long time period (i.e. at least 5 years). VaR modelling is a natural application of volatility models as in a parametric framework the VaR measure (which by definition is a quantile of the conditional distribution) is a deterministic function of the volatility. See Jorion (2000) for a recent review of VaR models. Because we have intraday data over a long time period, we can retrieve the daily closing prices for the indexes and then compute daily VaR measure using ARCH type models. When we make use of all the available data and compute intraday returns and realized volatility, we then have the competing model which uses the intraday information.

Our main results can be summarized in one sentence: yes, an (adequate) ARCH type model can deliver accurate VaR forecasts and this model performs as well as a competing VaR model based on the realized volatility. The key issue is to use a daily ARCH type model that clearly recognizes the full features of the empirical data such as a high kurtosis and skewness in the observed returns. In this paper we use the asymmetric skewed Student APARCH model (see for instance Lambert and Laurent, 2001 and Giot and Laurent, 2001), which delivers excellent results when applied to daily data. It is also true that the model based on the realized volatility delivers equally adequate VaR forecasts but this comes at the expense of using intraday information. Thus, for the two indexes under review, the results clearly indicate that modelling the realized volatility may be useful, but it is far from being the only game in town.

The rest of the paper is organized in the following way. In Section 2, we describe the available intraday data for the two stock indexes and characterize the stylized facts of the corresponding realized volatility. In Section 3, we introduce the two competing models (i.e. the skewed Student APARCH model for the daily returns and the model based on the realized volatility) for computing the one-step-ahead VaR. These two models are applied to the daily stock index data in Section 4 where we assess their performances. Section 5 concludes.

2 Data and stylized facts

2.1 Data

The data are available for two stock indexes on an intraday basis and for a relatively long period of time which allows VaR modelling and testing. For both assets we consider daily returns (which are used by the skewed Student APARCH model) and intraday returns defined on a 5-minute and 15-minute time grid (these intraday returns are used to compute the daily realized volatility).

Our first asset is the French CAC40 stock index for the 1995-1999 years (1249 daily observations). It is computed by the exchange as a weighted measure of the prices of its components and is available in the database on an intraday basis with the price index being computed every 30 seconds (approximately). For the time period under review, the opening hours of the French stock market were 10h am to 5h pm, thus 7 hours of trading per day. With the 5- (15-) minute time grid, this translates into 84 (28) intraday returns used to compute the daily realized volatility. Intraday prices at the 5- and 15-minute level are the outcomes of a linear interpolation between the closest recorded prices below and above the time set in the grid. Correspondingly, all returns are computed as the first difference in the regularly time-spaced log prices of the index. Because the exchange is closed from 5h pm to 10h am the next day, the first intraday return (computed at 10h05 when working with a 5-minute time grid for example) is the first difference between the log price at 10h05 and the log price at 5h pm the day before. Daily returns in percentage are defined as 100 times the first difference of the log of the closing prices.¹

Our second dataset contains 12 years (from January 1989 to December 2000, 3241 daily observations) of tick-by-tick prices for SP500 futures contracts traded on the Chicago Mercantile Exchange. Such SP500 futures contracts can be traded from 8h30 am to 15h10 pm Chicago time, i.e. from 9h30 am to 16h10 pm New York time. To conveniently define 5- and 15-minute returns, we remove all prices recorded after 16h New York time.² As for the CAC40 dataset, intraday prices at the 5- and 15-minute level are the outcomes of a linear interpolation between the closest recorded prices (for the nearest contract to maturity) below and above the time set in the regularly time-spaced sampling grid.³ Returns are computed as the first difference in the regularly time-spaced log prices of the index, with the overnight return included in the first intraday return. Daily returns in percentage are defined as 100 times the first difference of the log of the closing prices.

2.2 Realized volatility: stylized facts

Estimating and forecasting volatility is a key issue in empirical finance. After the introduction of the ARCH model by Engle (1982) or the Stochastic Volatility (SV) model (see Taylor, 1994) and their various extensions, a new generation of conditional volatility models has been advocated recently by Taylor and Xu (1997) and Andersen and Bollerslev (1998), i.e. models making use of the realized volatility. The origin of this concept is not so recent as it would seem at first sight. Merton (1980) already mentioned that, provided data sampled at a high frequency are available,

¹By definition and using the properties of the log distribution, the sum of the intraday returns is equal to the observed daily return based on the closing prices.

²Thus the last recorded price for the futures at 16h corresponds more or less to the closing price of the ‘cash’ SP500 index computed from its constituents traded on the NYSE or NASDAQ.

³The choice of the nearest contract to maturity means that we always select very liquid futures contracts.

the sum of squared realizations can be used to estimate the variance of an *i.i.d.* random variable. Taylor and Xu (1997) and Andersen and Bollerslev (1998) (among others) show that daily realized volatility may be constructed simply by summing up intraday squared returns. Assuming that a day can be divided in N equidistant periods and if $r_{i,t}$ denotes the intraday return of the i^{th} interval of day t , it follows that the daily volatility for day t can be written as:

$$\left[\sum_{i=1}^N r_{i,t} \right]^2 = \sum_{i=1}^N r_{i,t}^2 + 2 \sum_{i=1}^N \sum_{j=i+1}^N r_{j,t} r_{j-i,t}. \quad (1)$$

If the returns have mean zero and are uncorrelated, $E \left[\sum_{i=1}^N r_{i,t}^2 \right]$ is a consistent (see Andersen, Bollerslev, Diebold, and Labys, 1999) and unbiased estimator of the daily variance σ_t^2 .⁴ Because all squared returns on the right side of this equation are observed when intraday data are available, $\left[\sum_{i=1}^N r_{i,t} \right]^2$ is called the daily *realized* volatility.

By summing sufficiently many high-frequency squared returns we may then obtain an error free measure of the daily volatility. However, choosing a very high sampling frequency (30-seconds, 1-minute, etc.) may introduce a bias in the variance estimate due to market microstructure effects (bid-ask bounces, price discreteness or non-synchronous trading). As a trade off between these two biases, Andersen, Bollerslev, Diebold, and Labys (1999a) propose the use of 5-minute returns to compute daily realized volatility. Using the FTSE-100 stock market index (on the period 1990-2000), Oomen (2001) shows that the realized volatility measure increases when the sampling interval decreases while the summation of the cross terms in (1) decreases. Comparing the average daily realized volatility and the autocovariance bias factor, Oomen (2001) argues that the optimal sampling frequency for his dataset suggests using 25-minute returns. For our two datasets, a sampling frequency of about 15-minute was found to be optimal.⁵ By way of illustration, we also present results for 5-minute returns.

Although the empirical work on realized volatility is still in its infancy, some stylized facts have already been ascertained and we highlight these with our datasets.

- First, the unconditional distribution of the realized volatility is highly skewed and kurtosed.

On the other hand, the unconditional distribution of the logarithmic realized volatility is nearly gaussian, while standard tests reject the normality assumption. Figures 1 and 2

⁴Areal and Taylor (2000) show that even if this estimator is consistent and unbiased, it has not the least variance when N is finite. These authors propose to weight the intraday squared returns by a factor proportional to the intraday activity. This deflator may be obtained easily by applying Taylor and Xu's (1997) variance multiplier or the Flexible Fourier Function (FFF) of Andersen and Bollerslev (1997). Due to the strong similarity of the results with the 'non weighted squared returns', we will not report the results using Areal and Taylor's (2000) approach.

⁵To find the optimal sampling frequency, Oomen (2001) proposes to plot both the sum of squared intra-daily returns and the autocovariance bias factor versus the sampling frequency. The "optimal" sampling frequency is chosen as the highest available frequency for which the autocovariance bias term has disappeared.

display the level and the unconditional distribution of the logarithmic realized volatility of the CAC40 and SP500 stock indexes based on 15-minute returns. From Figure 2, both series appear slightly skewed (the unconditional skewness are respectively 0.62 and 0.38) and kurtosed (the unconditional kurtosis are respectively equal to 4.25 and 3.37).

- Secondly, the (logarithmic) realized volatility appears to be fractionally integrated. Indeed, Figure 3 displays the first 200 autocorrelations of the logarithmic realized volatility of the CAC40 and SP500 stock indexes based on 15-minute returns. This figure shows that a shock on volatility dies out very slowly, which is neither in accordance with an ARMA structure (which implies an exponential decay) nor with a unit root process (ADF tests, not reported to save space, all clearly reject the unit root assumption). This is in line with the previous findings of Ding, Granger, and Engle (1993) and Baillie, Bollerslev, and Mikkelsen (1996) (among others) who suggest the modelling of conditional variance of high frequency financial data by the use of an (Asymmetric) Power GARCH (APARCH) or Fractionally Integrated GARCH (FIGARCH) models.

To gain a first insight in the degree of persistence of a shock on the (logarithmic) realized volatility, we computed the Geweke and Porter-Hudak (1983) (GPH) log-periodogram estimate for the fractional integration parameter d .⁶ If $d \in (0, 1/2)$, the process is stationary, has a long memory and is said to be persistent. If $d \in (-1/2, 0)$, the process has a short memory and is said to be antipersistent.⁷ The estimated d are equal to 0.437 (0.038) and 0.430 (0.026) respectively for the CAC40 and SP500 stock indexes based on 15-minute⁸ returns (standard errors are given in parentheses). Thus d is fairly close to the ‘typical value’ of 0.4 (see Andersen, Bollerslev, Diebold, and Labys, 1999, Ebens, 1999 among others) and just significantly lower than 0.5 at the 5% critical level, suggesting that these series might be covariance-stationary.

- Finally, according to Ebens (1999) who analyzes the Dow Jones Industrial portfolio over the January 1993 to May 1998 period, the (logarithmic) realized volatility of stock indexes are non-linear in returns. To show this, consider the following Least-Squares (LS) regression: $\ln RV_t = c_0 + c_1 r_{t-1} + c_2 r_{t-1}^- + u_t$, where $\ln RV_t$ is the logarithm of the realized volatility, r_t is the daily return on day t , r_t^- is equal to 0 when $r_t > 0$ and is equal to r_t when $r_t < 0$ and u_t is a white noise. Figure 4 displays the fitted values of these LS regressions (solid lines) for the CAC40 (top panel) and SP500 (bottom panel) stock indexes based on 15-minute returns

⁶The number of low frequency periodogram points used in the estimation is set to $T^{\frac{4}{5}}$, see Hurvich, Deo, and Brodsky (1998).

⁷Furthermore, if $d \leq -1/2$, the process is non invertible and if $d \geq 1/2$, the process is not stationary but mean reverting if $d < 1$.

⁸Results for the 5-minute returns are very similar and are thus not reported.

as well as a nonparametric estimation (dashed lines).⁹ These graphs suggest that a negative shock on the returns is more likely to be associated with a high volatility (the next day) than for a positive shock.¹⁰ This feature is also well known for ARCH type models and is known as the *leverage effect*¹¹ (see Black, 1976; French, Schwert, and Stambaugh, 1987; Pagan and Schwert, 1990, Zakoian, 1994).

3 Two competing VaR models

Realized volatility was reviewed in the preceding section and we can now introduce a model for the daily VaR based on this measure. Subsection 3.2 is devoted to this topic. As the goal of the paper is to compare the performance of an ARCH type model directly applied to the daily data with the performance of a model based on the realized volatility, we also need to characterize the skewed Student APARCH model for the daily data. This is done in Subsection 3.1.

In both cases the link between the forecasted one-day-ahead volatility and the one-day-ahead VaR is immediate. Indeed, both models are parametric conditional models for volatility and the corresponding VaR measures are easily computed as the product of the square root of the conditional volatility and the quantile at $\alpha\%$ of the underlying distribution for the standardized error term.¹² Thus, for example, if the forecasted volatility at time $t - 1$ is $\hat{h}_t^2$ and one assumes a normal distribution for the error term, then the forecasted one-day-ahead VaR in $t - 1$ is equal to $z_\alpha \hat{h}_t$, with z_α being the left quantile at $\alpha\%$ for the normal distribution.

3.1 The skewed Student APARCH model

To model daily returns r_t , with $t = 1 \dots T$, we use an AR(3)-APARCH(1,1) model:¹³

$$r_t = \rho_0 + \rho_1 r_{t-1} + \rho_2 r_{t-2} + \rho_3 r_{t-3} + \epsilon_t \quad (2)$$

$$\epsilon_t = h_t z_t \text{ with } z_t \sim D(0, 1, \kappa) \quad (3)$$

⁹Quite similar to Ebens (1999), the nonparametric regression estimates are obtained using the Nadaraya-Watson estimator with the Epanechnikov kernel while the bandwidth parameters are determined using cross-validation scores. The plot regions are restricted to returns in the -5 to 5 interval, even if all the sample size was used when estimating this nonparametric regression.

¹⁰The R^2 of these LS regressions are respectively 11.5 and 17.5%, which is very similar to the ones reported by Ebens (1999).

¹¹Past negative (resp. positive) shocks have a different impact on current realized volatility than past positive shocks.

¹²In this paper we consider a forecast for the demeaned VaR which only depends on the level of the volatility.

¹³Based on information criteria and standard serial correlation tests, the AR(3)-APARCH(1,1) specification was found to be adequate in describing our two series. In order to save space, we only report the results concerning the more parsimonious specification.

$$h_t^\delta = \omega + \alpha_1 (|\epsilon_{t-1}| - \alpha_n \epsilon_{t-1})^\delta + \beta_1 h_{t-1}^\delta, \quad (4)$$

where $\rho_0, \rho_1, \rho_2, \rho_3, \omega, \alpha_1, \alpha_n, \beta_1, \delta$ and κ are the parameters to be estimated. κ is a vector of parameters relevant for specifying the shape of the density $D(\cdot)$. δ ($\delta > 0$) plays the role of a Box-Cox transformation of h_t , while α_n ($-1 < \alpha_n < 1$) reflects the so-called leverage effect. A positive (resp. negative) value of α_n indicates that past negative (resp. positive) shocks have a larger impact on current conditional volatility than past positive shocks (see Black, 1976; French, Schwert, and Stambaugh, 1987; Pagan and Schwert, 1990). The properties of the APARCH model have been studied recently by He and Terasvirta (1999a, 1999b). This specification is also motivated by a stylized fact first presented by Taylor (1986) who observed that absolute returns ($|r_t|$) of financial time series are positively autocorrelated, even at long lags. Ding, Granger, and Engle (1993) found that, the closer δ is to 1, the larger the memory of the process.

In VaR applications, the choice of an appropriate distribution for $D(\cdot)$ is an important issue. As in Giot and Laurent (2001), we use the skewed Student distribution introduced by Fernández and Steel (1998).¹⁴ According to Lambert and Laurent (2001) and provided that $v > 2$, the innovation process z_t is said to be (standardized) skewed Student distributed, i.e. $z_t \sim SKST(0, 1, \xi, v)$, if:

$$f(z_t|\xi, v) = \begin{cases} \frac{2}{\xi + \frac{1}{\xi}} sg[\xi(sz_t + m)|v] & \text{if } z_t < -\frac{m}{s} \\ \frac{2}{\xi + \frac{1}{\xi}} sg[(sz_t + m)/\xi|v] & \text{if } z_t \geq -\frac{m}{s} \end{cases}, \quad (5)$$

where $g(\cdot|v)$ is a symmetric (unit variance) Student density and ξ is the asymmetry coefficient.¹⁵
¹⁶ Parameters m and s^2 are respectively the mean and the variance of the non-standardized skewed Student:

$$m = \frac{\Gamma(\frac{v-1}{2})\sqrt{v-2}}{\sqrt{\pi}\Gamma(\frac{v}{2})} \left(\xi - \frac{1}{\xi} \right) \quad (6)$$

and

$$s^2 = \left(\xi^2 + \frac{1}{\xi^2} - 1 \right) - m^2. \quad (7)$$

In short, ξ models the asymmetry, while v accounts for the tail thickness. See Lambert and Laurent (2001) for a discussion of the link between these two parameters and the skewness and the kurtosis.

¹⁴Giot and Laurent (2001) show that an AR-APARCH model with a skewed Student density succeeds in correctly forecasting (both in- and out-of-sample) the VaR of the CAC40, DAX, NASDAQ, NIKKEI and SMI stock indexes on a daily basis. Models based on the normal or Student distributions clearly underperform when applied to the same datasets.

¹⁵The asymmetry coefficient $\xi > 0$ is defined such that the ratio of probability masses above and below the mean is $\frac{\Pr(\epsilon \geq 0|\xi)}{\Pr(\epsilon < 0|\xi)} = \xi^2$. Note also that the density $f(\epsilon|1/\xi, v)$ is the symmetric of $f(\epsilon|\xi, v)$ with respect to the mean. Therefore, working with $\ln(\xi)$ might be preferable to indicate the sign of the skewness.

¹⁶If $D(0, 1, \kappa)$ in (3) is the (standardized) skewed Student density, κ is then defined as $(\ln(\xi), v)$.

Because of the direct relationship between the VaR and the quantile in parametric VaR models, the one-day-ahead VaRs for long and short positions are given by $F_{\alpha,\xi,v}^{-1}\hat{h}_t$ and $F_{1-\alpha,\xi,v}^{-1}\hat{h}_t$, with $F_{\alpha,\xi,v}^{-1}$ being the left quantile at α for the skewed Student distribution with v degrees of freedom and asymmetry coefficient ξ ; $F_{1-\alpha,\xi,v}^{-1}$ is the corresponding right quantile.¹⁷ As formally defined in Giot and Laurent (2001), the long side of the daily VaR is defined as the VaR level for traders having long positions in the relevant equity index: this is the ‘usual’ VaR where traders incur losses when negative returns are observed. Correspondingly, the short side of the daily VaR is the VaR level for traders having short positions, i.e. traders who incur losses when stock prices increase.

3.2 Forecasting realized volatility

Regarding the realized volatility, the main findings of Section 2 are that the logarithmic realized volatility is approximately normal, appears fractionally integrated and correlated with past negative shocks. To take these properties into account, let us consider the following ARFIMAX(0, d ,1) model (initially developed by Granger, 1980 and Granger and Joyeux, 1980 among others):¹⁸

$$(1-L)^d(\ln RV_t - \mu_0 - \mu_1 r_{t-1} - \mu_2 r_{t-1}^-) = (1 + \theta_1 L)\varepsilon_t \quad (8)$$

$$(1-L)^d = \sum_{k=0}^{\infty} \frac{\Gamma(d+1)}{\Gamma(k+1)\Gamma(d-k+1)} L^k,$$

where L is the lag operator, μ_0 , μ_1 , μ_2 , θ_1 and d are parameters to be estimated, $\ln RV_t$ is the logarithm of the realized volatility computed from the intraday returns observed for day t , r_t is the daily return on day t , r_t^- takes the value 0 when $r_t > 0$ and the value r_t when $r_t < 0$.

Estimation of (8) is carried out by exact maximum likelihood (Sowell, 1992) under the normality assumption using ARFIMA 1.0 (see Ooms and Doornik, 1998 and Doornik and Ooms, 1999) and conditional sum-of-squares maximum likelihood¹⁹ (Hosking, 1981) using G@RCH 2.0 (see Laurent and Peters, 2001). Due to the strong similarity between the outcomes of the two estimation procedures, we only report the results obtained with the first method.

When $\varepsilon_t \sim N(0, \sigma^2)$, we have by definition that $\exp(\varepsilon_t) \sim \log N(0, \sigma^2)$ (where $\log N$ denotes the log-normal distribution). Thus, the conditional realized volatility (or in-sample one-step-ahead forecast of the volatility) is computed according to:

$$\hat{RV}_{t|t-1} = \exp\left(\ln RV_t - \hat{\varepsilon}_{t|t-1} + \frac{1}{2}\hat{\sigma}^2\right), \quad (9)$$

¹⁷The quantile function of the (standardized) skewed Student has been derived in Lambert and Laurent (2001) as a mixture of two Student quantile functions. See also Giot and Laurent (2001).

¹⁸As in the previous section, the choice of this specification is based on information criteria and standard serial correlation tests.

¹⁹The finite sample properties of this estimator have been investigated by Chung and Baillie (1993).

where $\hat{\varepsilon}_{t|t-1}$ denotes the estimated value of ε_t by (8) and $\hat{\sigma}^2$ is the estimated variance of ε_t in the same equation.

To compute a one-day-ahead forecast for the VaR of the daily returns r_t using the conditional realized volatility, we specify the following AR(3) model:

$$r_t^* = r_t / \sqrt{h_t^*} \quad (10)$$

$$r_t^* = \rho_0^* + \rho_1^* r_{t-1}^* + \rho_2^* r_{t-2}^* + \rho_3^* r_{t-3}^* + \epsilon_t^* \quad (11)$$

$$\epsilon_t^* \sim D(0, \sigma^{2,*}, \kappa^*) \quad (12)$$

where now $h_t^* = \hat{R}V_{t|t-1}$ and $\rho_0^*, \rho_1^*, \rho_2^*, \rho_3^*, \sigma^{2,*}$ and κ^* are parameters to be estimated. As in (3), κ^* stands for a vector of parameters determining the shape of the density $D(\cdot)$, while $\sigma^{2,*}$ is the variance of ϵ_t^* . This specification is almost identical to the one introduced in Subsection 3.1, but now the conditional volatility for the daily returns is equal to the conditional realized volatility $\hat{R}V_{t|t-1}$. As in Subsection 3.1, an adequate distribution for $D(\cdot)$ should be selected. The recent empirical literature has stressed that the normal distribution is a good candidate for $D(\cdot)$ when $h_t^* = RV_t$, i.e. when one uses realized volatility computed at the end of day t (or ex-post realized volatility). Because we wish to forecast the one-day-ahead VaR, $h_t^* = \hat{R}V_{t|t-1}$ is substituted to $h_t^* = RV_t$ in our framework. In Section 4, we show that this invalidates the choice of the normal distribution as an adequate distribution for $D(\cdot)$. Therefore, we suggest the use of the skewed Student distribution. For reason of comparison we also present results for the normal distribution.²⁰ In both cases, the one-day-ahead (demeaned) VaR for long and short positions are given as the product of the quantile at $\alpha\%$ for each distribution with $\sqrt{\hat{R}V_{t|t-1}}$.

3.3 Assessing the VaR performance of the models

Using a procedure that is now standard in the VaR literature, we assess the models' performance by first computing their empirical failure rate (both for the left and right tails of the distribution of returns) and then performing a Kupiec LR test. By definition, the failure rate is the number of times returns exceed (in absolute value) the forecasted one-day-ahead VaR. If the VaR model is correctly specified, the failure rate should be equal to the prespecified VaR level $\alpha\%$. Because the computation of the empirical failure rate defines a sequence of yes/no observations, it is possible to test $H_0 : f = \alpha$ against $H_1 : f \neq \alpha$, where f is the failure rate (estimated by $\hat{f}$, the empirical failure rate).²¹ At the 5% level and if T yes/no observations are available, an approximate confidence

²⁰Note that if $D(\cdot)$ is the normal density, then κ^* is a null vector, while the choice of the skewed Student distribution for $D(\cdot)$ implies that $\kappa^* = (\ln(\xi^*), v^*)$.

²¹In the literature on VaR models, this test is called the Kupiec LR test, if the hypothesis is tested using a likelihood ratio test. See Kupiec (1995).

interval for $\hat{f}$ is also given by $\left[\hat{f} - 1.96\sqrt{\hat{f}(1-\hat{f})/T}, \hat{f} + 1.96\sqrt{\hat{f}(1-\hat{f})/T}\right]$.

4 Empirical application

In this section, we report estimation results for the two models presented in Section 3. We first focus on the skewed Student APARCH model which is applied to the daily returns; the second model uses the intraday returns via the computation of the realized volatility. Both models are used to forecast the one-day-ahead VaR for the two stock indexes and their performance is assessed by comparing their empirical failure rate with the theoretical threshold.

4.1 VaR, daily returns and the skewed Student APARCH

Our first setting uses daily data only and computes the one-day-ahead daily VaR using these daily observations. The skewed Student APARCH and corresponding one-day-ahead VaR were defined in Subsection 3.1. Tables 1 (estimated parameters) and 2 (assessment of the one-day-ahead VaR) report estimation results when this model is applied to the CAC40 and SP500 daily returns. According to the estimated coefficients for the skewed Student APARCH,

- β_1 is close to 1 but significantly different from 1 for both indexes, which indicates a high degree of volatility persistence.²² Furthermore both APARCH models are stationary in the sense that $\alpha_1 E(|z| - \gamma z) + \beta_1$ is lower than 1. See Ding, Granger, and Engle (1993) and Lambert and Laurent (2001) for more details on the computation of $\alpha_1 E(|z| - \gamma z) + \beta_1$, which depends on the assumption made on the stochastic innovation.
- δ is close to 2 for the CAC40 and close to 1 for the SP500: the APARCH models the conditional variance for the CAC40 and the conditional standard deviation for the SP500;
- α_n is significantly positive: negative returns lead to higher subsequent volatility than positive returns (asymmetry in the conditional variance);
- v is much larger for the CAC40 than for the SP500: daily returns defined on the U.S. data display a much larger kurtosis and exhibit fatter tails than returns for the French data;
- $\ln(\xi)$ is negative in both specifications, albeit not significant from zero for the SP500 and barely significant for the CAC40.²³

²²Tse (1998) extended the APARCH by including a pure long memory feature (FIAPARCH). Likelihood ratio tests between the APARCH and the FIAPARCH clearly reject the FIAPARCH specification.

²³This indicates that, at least for the U.S. data, there is no real need for a skewed Student APARCH; nevertheless, as this specification encompasses the simpler Student APARCH, we stick with the more general model (owing to the large number of observations, the loss of degrees of freedom is minimal).

- the APARCH dynamical structure succeeds in taking into account all the dynamical structure exhibited by the volatility as the Ljung-Box $Q(20)$ on the squared standardized residuals is not significant at the 5% level for both models.

For the skewed Student APARCH model, the P-values for the null hypothesis $f_l = \alpha$ (VaR for the left tail of the distribution of returns) and $f_s = \alpha$ (VaR for the right tail of the distribution of returns) given in Table 2 confirm that this volatility model succeeds in correctly forecasting the one-day-ahead VaR for most of the probability levels α . Indeed, the P-values are larger than 0.05 for all configurations except the VaR for short positions on the SP500 (with α ranging from 0.25% to 1%). Broadly speaking these results are similar to those of Giot and Laurent (2001) reported for five stock market indexes.

4.2 VaR, intraday returns and daily realized volatility

In our second framework we explicitly use the intradaily (5- and 15-minute) returns to compute the daily realized volatility. We first estimate an ARFIMAX(0, d ,1) model on the logarithmic realized volatility $\ln RV_t$ as in equation (8). In a second step, we standardize the daily returns r_t by the one-day-ahead forecast of the realized volatility $\hat{RV}_{t|t-1}$ as in equation (10) and compute the one-day-ahead VaR using an AR(3) model on the $r_t^* = r_t/\sqrt{\hat{h}_t^*}$. As explained below, the choice of the distribution for $D(\cdot)$ is of paramount importance. Table 3 presents estimation results for the ARFIMA specification:

- First, the ARFIMA specification seems to be adequate in modelling the dynamics of $\ln RV_t$. Indeed, the Ljung-Box statistics indicate that all serial correlation in the error term has been removed (at the conventional levels of significance). Parameter d is well above 0 but is not significantly lower than 0.5, indicating that, in contrast to the GPH test of Subsection 2.2, the logarithm of the realized volatility is not covariance-stationary;²⁴
- μ_1 and μ_2 are respectively non significant and significantly positive: negative returns lead to higher subsequent volatility than positive returns (asymmetry in the conditional variance similar to the APARCH model).

Estimation results for the skewed Student AR(3) model are presented in Table 4. As indicated by the Ljung-Box $Q^2(20)$ on the standardized residuals of this model, the $r_t^* = r_t/\sqrt{\hat{RV}_{t|t-1}}$ do not display time dependence in volatility. This justifies the use of a ‘static’ skewed Student AR(3) model. Of course, this is expected as the time dependence in volatility has been captured

²⁴However as argued by Andersson (2000), one has to be careful with the notion of long memory because “(surprisingly) negative moving average parameters (θ_1 is significantly below 0 for both indexes), which alone make no memory contribution, absorb a substantial amount of memory induced by fractional integration.”

by the previous ARFIMA model on the dynamics of $\ln RV_t$. In the usual ARCH framework, the $r_t^* = r_t / \sqrt{\hat{RV}_{t|t-1}}$ would play the role of standardized residuals. This is somewhat true as we do standardize the returns by the square root of forecasted realized volatility.

While the recent literature has stressed that ex-post standardized returns have an almost normal distribution (see Andersen, Bollerslev, Diebold, and Labys, 1999b), this is certainly not true for ex-ante standardized returns. The estimated parameters $\ln(\xi^*)$ and v^* reported in Table 4 suggest that the ex-post standardized returns of the CAC40 are slightly skewed and kurtosed while the SP500 is kurtosed but symmetric. These results are in line with those reported in Table 1 (skewed-Student APARCH on daily returns).²⁵ Furthermore, assessing the VaR performance of a *normal* model (i.e. choosing the normal distribution for $D(\cdot)$ instead of the skewed Student distribution) for the ex-ante standardized returns gives the results shown in the first line of each cell of Table 5:

- for the left tail of the distribution of returns (long VaR), the P-values for the null hypothesis $f_l = \alpha$ are smaller than 0.05 when α is below 1%: the empirical failure rate is significantly higher than α for low VaR levels;
- for the right tail of the distribution of returns (short VaR), the performance of the model is satisfactory;
- there are no real differences between the results for the 5- and 15-minute returns.

However, using the skewed Student distribution instead gives much better results (second line of each cell of Table 5). For the CAC40 data, all P-values are larger than 0.05, both for the long and short VaR. For the SP500 data, all P-values are larger than 0.05 except for the short VaR at level $\alpha = 1\%$ and $\alpha = 0.25\%$. Thus the switch from the normal distribution to the skewed Student distribution yields a significant improvement in the VaR performance of the model.

Finally we also give density plots (empirical vs the normal distribution) for the ex-ante and ex-post standardized returns in Figure 5. While the tails of the ex-post standardized returns closely

²⁵Note that one has to be careful when computing the empirical skewness and the kurtosis on the raw data. Indeed, Table 4 also reports these statistics (lines 1 and 2 for both series). For instance, the empirical skewness of the 5-minute (ex-post) standardized returns of the CAC40 and SP500 equal respectively -0.198 and -1.093. To test the departure from normality, it is common to use the t-test $\frac{sk}{\sqrt{\frac{6}{T}}}$ where sk is the empirical skewness and T the number of observations. Based on the result of this test one could be tempted to conclude that the SP500 is highly skewed while the CAC40 is hardly skewed (which contradicts the results obtained with the skewed-Student density, see lines 4 and 5 of Table 4). However, as shown by De Ceuster and Trappes (1992) and Peiró (1999), this test is not appropriate when the series is fat-tailed. For a sample size of 2000 observations, De Ceuster and Trappes (1992) tabulate that the 95% confidence intervals of the skewness of Student-t distributed observations with a kurtosis of 3.5 and 18 are respectively $(-0.131; 0.127)$ and $(-0.814; 0.787)$, i.e. the higher the kurtosis, the larger the confidence bands of the skewness.

track those of the normal distribution, ex-ante standardized returns feature fat tails, especially for the U.S. data. Estimation results and descriptive statistics given in Table 4 tell the same story.

4.3 Which model is best?

The evidence presented in the two preceding subsections indicates that using an APARCH model with daily data or a two step approach relying on the new concept of realized volatility leads very similar results in terms of VaR. It should be emphasized that to have accurate VaR forecasts, one needs to specify correctly the full conditional density with *both* methods.

This implies that previous results given in the empirical literature must be qualified. For example, Ebens (1999) concludes his paper by stating that the GARCH model underperforms (when volatility must be forecasted) with respect to the model based on the daily realized volatility. However, the author uses a ‘simple’ GARCH model which neither really accounts for the long memory property observed in the realized volatility nor the fat-tails or asymmetry of the returns (even after standardization). Indeed, when estimating the more simple RiskMetrics VaR model on daily returns (the RiskMetrics model is tantamount to an IGARCH model with pre-specified coefficients, under the additional assumption of normality), we have the VaR results given in Table 2: its one-day-ahead forecasting performance is rather poor, especially when α is small.²⁶ With a more ‘sophisticated’ model on the other hand (the skewed Student APARCH model in this paper), VaR results are much better.

Interestingly and as pointed out in the previous subsection by comparing the results obtained with the normal and skewed Student distributions for the ex-ante standardized returns, the same conclusion is true for the more complex model based on the combination of intraday returns and realized volatility.

5 Conclusion

In this paper we show how to compute a daily VaR measure for two stock indexes (CAC40 and SP500) using the one-day-ahead forecast of the daily realized volatility. The daily realized volatility is equal to the sum of the squared intraday returns over a given day and thus uses intraday information to define an aggregated daily volatility measure. While the VaR forecasts which use this method perform adequately over our sample, we also show that a more simple model based solely on daily returns delivers good results too. Indeed, while the VaR specification based on an ARFIMAX(0, d ,1)-skewed Student model for the daily realized volatility provides adequate one-day-ahead VaR forecasts, it does not really improve on the performance of a VaR

²⁶Although the results are not reported in the paper, we also estimated a normal GARCH(1,1) model and its performance was not much better than the RiskMetrics specification.

model based on the skewed Student APARCH model and estimated using daily data. Thus, for the two financial assets considered in an univariate framework, the two methods seem to be rather equivalent. Another important conclusion of this paper is that daily returns standardized by the square root of the one-day-ahead forecast of the daily realized volatility are not normally distributed.

At this stage, one of the most immediate and promising extension of these techniques is to consider corresponding multivariate volatility models to forecast the VaR of a portfolio of financial assets. Multivariate models of the ARCH type are not easy to implement as they often require the estimation of a large number of parameters. Furthermore, these parameters are present in the latent volatility specification and this is one of the main difficulty of the problem. Therefore, multivariate realized volatility models should provide a much easier way to correctly model variances and correlations across financial assets as they assume that volatility is observed. This paves the way for the use of ‘usual’ multivariate models (VAR, ECM) directly applied to realized volatility and correlations.

References

- ANDERSEN, T., AND T. BOLLERSLEV (1997): “Intraday Periodicity and Volatility Persistence in Financial Markets,” *Journal of Empirical Finance*, 4, 115–158.
- (1998): “Answering the skeptics: Yes, standard volatility models do provide accurate forecasts,” *International Economic Review*, 39, 885–905.
- ANDERSEN, T., T. BOLLERSLEV, F. DIEBOLD, AND P. LABYS (1999a): “The Distribution of Exchange Rate Volatility,” NBER Working Paper No. 6961.
- (1999b): “Exchange Rate Returns Standardized by Realized Volatility are (Nearly) Gaussian,” Manuscript in progress.
- ANDERSSON, M. (2000): “Do Long-Memory Models Have Long Memory ?,” *International Journal of Forecasting*, 16, 121–124.
- AREAL, N., AND S. TAYLOR (2000): “The realized volatility of FTSE-100 futures prices,” Manuscript, Department of Accounting and Finance, Lancaster University.
- BAILLIE, R., T. BOLLERSLEV, AND H. MIKKELSEN (1996): “Fractionally integrated generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, 74, 3–30.
- BLACK, F. (1976): “Studies of Stock Market Volatility Changes,” *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, pp. 177–181.

- BOLLERSLEV, T. (1986): “Generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, 31, 307–327.
- CHUNG, C.-F., AND R. BAILLIE (1993): “Sample Bias in Conditional Sum-of-Squares Estimator of Fractionally Integrated ARMA Models,” *Empirical Economics*, 18, 791–806.
- DE CEUSTER, M., AND D. TRAPPERS (1992): “Diagnostic Checking of Estimation with a Student-t Error Density,” Working paper UFSIA, Centrum voor Bedrijfsconomie en Bedrijfsconometrie, 173.
- DING, Z., C. W. J. GRANGER, AND R. F. ENGLE (1993): “A Long Memory Property of Stock Market Returns and a New Model,” *Journal of Empirical Finance*, 1, 83–106.
- DOORNIK, J. A., AND M. OOMS (1999): “A Package for Estimating, Forecasting and Simulating Arfima Models: Arfima package 1.0 for Ox,” Discussion paper, Econometric Institute, Erasmus University Rotterdam.
- EBENS, H. (1999): “Realized Stock Index Volatility,” Working Paper No. 420, Department of Economics, Johns Hopkins University.
- ENGLE, R. (1982): “Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation,” *Econometrica*, 50, 987–1007.
- FERNÁNDEZ, C., AND M. STEEL (1998): “On Bayesian modelling of fat tails and skewness,” *Journal of the American Statistical Association*, 93, 359–371.
- FRENCH, K., G. SCHWERT, AND R. STAMBAUGH (1987): “Expected Stock Returns and Volatility,” *Journal of Financial Economics*, 19, 3–29.
- GEWEKE, J., AND S. PORTER-HUDAK (1983): “The Estimation and Application of Long Memory Time Series Models,” *Journal of Time Series Analysis*, 4, 221–238.
- GIOT, P. (2000): “Intraday Value-at-Risk,” CORE DP 2045, Maastricht University METEOR RM/00/030.
- GIOT, P., AND S. LAURENT (2001): “VALUE-AT-RISK FOR LONG AND SHORT TRADING POSITIONS,” CORE DP xxxx, Maastricht University METEOR RM/01/005.
- GRANGER, C. (1980): “Long Memory Relationships and the Aggregation of Dynamic Models,” *Journal of Econometrics*, 14, 227–238.
- GRANGER, C., AND R. JOYEUX (1980): “An Introduction to Long-Memory Time Series Models and Fractional Differencing,” *Journal of Time Series Analysis*, 1, 15–29.

- HE, C., AND T. TERASVIRTA (1999a): “Higher-order dependence in the general Power ARCH process and a special case,” Stockholm School of Economics, Working Paper Series in Economics and Finance, No. 315.
- (1999b): *Statistical Properties of the Asymmetric Power ARCH Process* chap. 19, pp. 462–474, Cointegration, causality, and forecasting. Festschrift in honour of Clive W.J. Granger. in Engle, Robert F. and Halbert White, oxford university press edn.
- HOSKING, J. (1981): “Fractional differencing,” *Biometrika*, 68, 165–176.
- HURVICH, C. M., R. DEO, AND J. BRODSKY (1998): “The Mean Squared Error of Geweke and Porter-Hudak’s Estimator of the Long-Memory Parameter of a Long-Memory Time Series,” *Journal of Time Series Analysis*, 19, 19–46.
- JORION, P. (2000): *Value-at-Risk*. McGraw-Hill.
- KUPIEC, P. (1995): “Techniques for verifying the accuracy of risk measurement models,” *Journal of Derivatives*, 2, 173–84.
- LAMBERT, P., AND S. LAURENT (2001): “Modelling Financial Time Series Using GARCH-Type Models and a Skewed Student Density,” Mimeo, Université de Liège.
- LAURENT, S., AND J.-P. PETERS (2001): “G@RCH 2.0 : An Ox Package for Estimating and Forecasting Various ARCH Models,” Mimeo, Université de Liège.
- MERTON, R. (1980): “On Estimating the Expected Return on the Market; An Exploratory Investigation,” *Journal of Financial Economics*, 8, 323–361.
- OOMEN, R. (2001): “Using High Frequency Stock Market Index Data to Calculate, Model and Forecast Realized Volatility,” Manuscript, European University Institute, Department of Economics.
- OOMS, M., AND J. A. DOORNIK (1998): “Estimation, simulation and forecasting for fractional autoregressive integrated moving average models,” Discussion paper, Econometric Institute, Erasmus University Rotterdam, presented at the fourth annual meeting of the Society for Computational Economics, June 30, 1998, Cambridge, UK.
- PAGAN, A., AND G. SCHWERT (1990): “Alternative Models for Conditional Stock Volatility,” *Journal of Econometrics*, 45, 267–290.
- PALM, F. (1996): “GARCH Models of Volatility,” in *Maddala, G.S., Rao, C.R., Handbook of Statistics*, pp. 209–240.
- PEIRÓ, A. (1999): “Skewness in financial returns,” *Journal of Banking and Finance*, 23, 847–862.

- SOWEL, F. (1992): “Maximum likelihood estimation of stationary univariate fractionally integrated time series models,” *Journal of Econometrics*, 53, 165–188.
- TAYLOR, S. (1986): *Modelling financial time series*. Wiley, New York.
- (1994): “Modeling stochastic volatility: A review and comparative study,” *Mathematical Finance*, 4, 183–204.
- TAYLOR, S., AND X. XU (1997): “The Incremental Volatility Information in One Million Foreign Exchange Quotations,” *Journal of Empirical Finance*, 4, 317–340.
- TSE, Y. (1998): “The Conditional Heteroscedasticity of the Yen-Dollar Exchange Rate,” *Journal of Applied Econometrics*, 13, 49–55.
- ZAKOIAN, J.-M. (1994): “Threshold Heteroskedasticity Models,” *Journal of Economic Dynamics and Control*, 15, 931–955.

Table 1: Skewed Student APARCH

	CAC40 (daily returns)	SP500 (daily returns)
ω	0.023 (0.013)	0.006 (0.002)
α_1	0.042 (0.015)	0.053 (0.009)
α_n	0.452 (0.193)	0.539 (0.105)
β_1	0.940 (0.018)	0.954 (0.009)
$\ln(\xi)$	-0.075 (0.042)	-0.029 (0.024)
v	12.849 (4.391)	5.462 (0.504)
δ	1.775 (0.568)	0.955 (0.157)
$Q^2(20)$	14.75	17.36
$\alpha_1 E(z - \gamma z) + \beta_1$	0.975	0.991

Estimation results for the volatility specification of the skewed Student APARCH model. Robust standard errors are reported in parentheses. $Q^2(20)$ is the Ljung-Box Q-statistic of order 20 computed on the squared standardized residuals.

Table 2: VaR results for the CAC40 and SP500
(models using daily data)

α	5%	2.5%	1%	0.5%	0.25%
VaR for long positions (CAC40)					
RiskMetrics	0.065	0.005	0	0	0
Skewed Student APARCH	0.764	0.170	0.483	0.609	0.498
VaR for long positions (SP500)					
RiskMetrics	0.335	0.001	0	0.003	0
Skewed Student APARCH	0.682	0.995	0.543	0.962	0.332
VaR for short positions (CAC40)					
RiskMetrics	0.665	0.336	0.671	0.297	0.021
Skewed Student APARCH	0.928	0.336	0.879	0.762	0.948
VaR for short positions (SP500)					
RiskMetrics	0.297	0.152	0.035	0.008	0.060
Skewed Student APARCH	0.625	0.906	0.010	0.024	0.010

P-values for the null hypothesis $f_l = \alpha$ (i.e. failure rate for the long trading positions is equal to α , top of the table) and $f_s = \alpha$ (i.e. failure rate for the short trading positions is equal to α , bottom of the table). α is equal successively to 5%, 2.5%, 1%, 0.5% and 0.25%. The RiskMetrics and skewed Student APARCH models are estimated on the daily returns (i.e. no use is made of the intraday returns).

Table 3: Asymmetric ARFIMA

	CAC40		SP500	
	5-minute	15-minute	5-minute	15-minute
μ_0	-0.019 (0.913)	-0.016 (0.729)	-0.457 (1.758)	-0.565 (1.120)
μ_1	0.027 (0.023)	0.029 (0.026)	-0.007 (0.017)	-0.016 (0.020)
μ_2	-0.188 (0.040)	-0.187 (0.035)	-0.190 (0.028)	-0.215 (0.034)
θ_1	-0.345 (0.045)	-0.341 (0.053)	-0.237 (0.022)	-0.287 (0.030)
d	0.478 (0.025)	0.463 (0.034)	0.492 (0.010)	0.480 (0.019)
σ^2	0.357	0.444	0.289	0.399
$Q(20)$	20.0	15.6	22.8	15.9

Estimation results for the logarithm of the realized volatility (defined on 5- and 15-minute returns) using an ARFIMAX(0,d,1) specification. Standard errors are reported in parentheses. $Q(20)$ is the Ljung-Box Q-statistic of order 20 computed on the residuals.

Table 4: Ex-ante standardized returns
(w.r.t. forecasted realized volatility)

5-minute returns		
	CAC40	SP500
Skewness	-0.198	-1.093
Kurtosis	3.537	17.861
$\sigma^{2,*}$	1.073	1.067
$\ln(\xi^*)$	-0.078 (0.042)	-0.020 (0.024)
v^*	14.516 (5.384)	6.055 (0.618)
$Q^2(20)$	15.87	3.17
15-minute returns		
	CAC40	SP500
Skewness	-0.167	-1.107
Kurtosis	3.441	18.247
$\sigma^{2,*}$	1.024	1.106
$\ln(\xi^*)$	-0.073 (0.041)	-0.022 (0.024)
v^*	15.708 (6.414)	5.987 (0.606)
$Q^2(20)$	15.86	3.01

Descriptive statistics (skewness and kurtosis) and estimation results ($\sigma^{2,*}$, $\ln(\xi^*)$ and v^*) for the skewed Student AR(3) model on the ex-ante standardized returns with respect to the daily realized volatility computed on 5- and 15- minute intraday returns. $Q^2(20)$ is the Ljung-Box Q-statistic of order 20 computed on the squared standardized residuals.

Table 5: VaR results for the CAC40 and SP500
(models using intraday data and daily realized volatility on 5- and 15-minute returns)

α	5%	2.5%	1%	0.5%	0.25%
VaR for long positions (CAC40)					
Normal	0.323, 0.173	0.087, 0.122	0.026, 0.048	0.018, 0.007	0.021, 0.146
Skewed Student	0.541, 0.265	0.607, 0.607	0.333, 0.482	0.927, 0.927	0.630, 0.499
VaR for long positions (SP500)					
Normal	0.086, 0.032	0.572, 0.816	0.051, 0.035	0.001, 0	0.001, 0
Skewed Student	0.800, 0.800	0.432, 0.917	0.775, 0.948	0.840, 0.657	0.974, 0.694
VaR for short positions (CAC40)					
Normal	0.217, 0.336	0.250, 0.338	0.470, 0.897	0.495, 0.760	0.949, 0.949
Skewed Student	0.461, 0.722	0.444, 0.444	0.470, 0.897	0.760, 0.760	0.499, 0.499
VaR for short positions (SP500)					
Normal	0.059, 0.040	0.575, 0.499	0.948, 0.807	0.963, 0.765	0.332, 0.200
Skewed Student	0.514, 0.739	0.917, 0.740	0.010, 0.052	0.170, 0.051	0.039, 0.010

P-values for the null hypotheses $f_l = \alpha$ (i.e. failure rate for the long trading positions is equal to α , top of the table) and $f_s = \alpha$ (i.e. failure rate for the short trading positions is equal to α , bottom of the table). α is equal successively to 5%, 2.5%, 1%, 0.5% and 0.25%. In each cell, the first (second) number corresponds to the case where the daily realized volatility is estimated on the 5- (15-) minute returns. Normal and skewed Student refer to the statistical distribution used to model the daily realized volatility.




Figure 1: Logarithmic realized volatility of the CAC40 (top panel) and SP500 (bottom panel) stock indexes based on 15-minute returns.




Figure 2: Density estimates (dashed line) and corresponding normal density (solid line) for the logarithmic realized volatility of the CAC40 (top panel) and SP500 (bottom panel) stock indexes based on 15-minute returns.




Figure 3: First 200 autocorrelations for the logarithmic realized volatility of the CAC40 (top panel) and SP500 (bottom panel) stock indexes based on 15-minute returns. The horizontal lines show the upper limit 95% Bartlett confidence bands.




Figure 4: Regression lines for the logarithmic realized volatility (y-axis) of the CAC40 (top panel) and SP500 (bottom panel) stock indexes based on 15-minute returns against the previous (i.e. one day before) returns (x-axis).




Figure 5: The graphs display the density distributions, i.e. empirical (dashed lines) vs normal (solid lines), for the daily returns standardized with respect to the square root of the ex-post (left panel) and the ex-ante (right panel) daily realized volatility computed for the CAC40 (top panel) and SP500 (bottom panel) stock indexes based on 15-minute returns.