

Forecasting Daily Volatility with Intraday Data

Bart Frijns

Nijmegen School of Management

Radboud University Nijmegen, The Netherlands

Dimitri Margaritis

Faculty of Business

Auckland University of Technology, New Zealand

SEPTEMBER 2006

Corresponding Author: Professor Dimitri Margaritis, Department of Finance, AUT, 42 Wakefield St, Auckland, New Zealand; e-mail: dmargaritis@aut.ac.nz

Forecasting Daily Volatility with Intraday Data

ABSTRACT

The aim of this paper is to assess to what extent intraday data can explain and predict end of day volatility. Using a realized volatility measure as proposed by Andersen et al. (2001), we hypothesize that volatility generated at the start of the day is an important predictor of daily volatility either on its own accord or in conjunction with information about the seasonal pattern characterizing intraday volatility. We address the question of how much information needs to arrive to the market before a good predictor can be formed. Using data from a specialist market (NYSE), a dealer market (Nasdaq) and a continuous auction market (Paris Bourse) we investigate how different trading structures may affect intraday volatility formation. As a preview to our results, we find that the explanatory power of first hour volatility for daily volatility is as high as 68%, while the average volatility generated during this first hour is less than 30%.

Keywords: intraday return volatility, volatility forecasting, realized volatility, quadratic variation.

JEL classification: G10

1. Introduction

Volatility plays a pivotal role in the pricing of many financial instruments and portfolio management decisions. Being able to forecast volatility and understanding the nature and determinants of volatility is therefore essential. Since the introduction of ARCH-type models (Engle, 1982 and Bollerslev, 1986), a plethora of specifications have been proposed to describe the time series properties of and to forecast volatility. However, these ARCH-type models have several drawbacks. First, the volatility remains latent and second, estimates depend on model specification. More recently the concept of *realized volatility* has been introduced, which is an observed and non-parametric measure of volatility, and is based on the theory of Quadratic Variation. Andersen et al. (2003) show that under suitable conditions, which are applicable to a wide set of models, the Quadratic Variation *is* the key ingredient for volatility measurement and forecasting. The Quadratic Variation is not a directly observed quantity but an estimate of conditional return volatility can be constructed by cumulating sums of squared returns of high frequency data (Andersen et al., 2001). The theory of Quadratic Variation suggests that these (ex-post) realized Quadratic Variation measures are an unbiased estimator of return volatility, and in spite of its simplicity, realized volatility is a fairly powerful tool for the measurement, modelling, and forecasting of high frequency data (see Andersen et al., 2002 and Andersen et al., 2003).

The usage of intraday high frequency data in the estimation of the realized volatility acknowledges that information arriving during the trading day contributes to the daily realized volatility. Indeed, Schwert (1990) and Ané and Geman (2000) demonstrate the importance of intraday data in explaining the level of daily volatility. More recently, Avramov et al. (2006) show the importance of intraday events in explaining the asymmetry in daily volatility, while Martens et al. (2002) investigate the extent to which explicit modelling of intraday patterns improves out-of-sample foreign exchange market volatility forecasting.

In this paper we focus on a different aspect of the importance of intraday data for daily volatility. Instead of focussing on all events within a trading day we address the issue of how important the opening hours in the market are in explaining daily volatility. We build on two empirical regularities - viz. that return volatility is (1) high during the start and end of the trading day and (2) it gradually declines reaching a low towards the middle of the day - to

develop a model for forecasting end of day (daily) volatility. We argue that start of the day volatility, measured by the realized variance is a sufficient, but not complete statistic for predicting end of day volatility.

We exploit the intraday seasonal volatility pattern to propagate observed or realized start of day volatility into a forecast of daily volatility. Recognizing the inadequacies of standard time series approaches to capturing the pervasive periodicity patterns of high frequency intraday return volatility, Andersen and Bollerslev (1997) were the first to use a Flexible Fourier Form to model the dynamics of intraday volatility. We follow their approach by adopting two different propagation devices: (1) a simple model in which the intraday periodic components are averages of past realized volatilities measured at a particular interval; (2) a Flexible Fourier Form (FFF) model for intraday seasonality.

To our knowledge this is the first study in which start of the day realized volatility, a measure of return variability observed at the beginning of the trading day, is used to explain/predict end of day volatility. We find that start of trading day volatility encompasses sufficient statistical content as a predictor of end of day volatility. The explanatory power of first hour volatility for the daily volatility is as high as 68%, while the average volatility generated during this first hour is less than 30%. By scaling these forecasts by seasonal factors we find further improvement in the model's forecasting performance. We assess out-of-sample forecasting performance using a variety of forecasting assessment criteria, including regression of observed realized ex-post volatility on volatility forecasts. Second, we illustrate the robustness of the proposed procedure by carrying out an analysis on stocks traded under different market structures.

2. Data

Data for this study is provided by SIRCA (Securities Industry Research Centre of Asia-Pacific). In total we collect data for stocks listed on three different exchanges, NYSE, Nasdaq, and Paris Bourse (Euronext France). For each exchange we randomly select 5 actively traded stocks that are in the main indices of that market. For the NYSE we select stocks from the DJIA - Alcoa (AA), Boeing (BA), General Electric (GE), International Business Machines (IBM) and AT&T (T). The Nasdaq stocks are all in the Nasdaq-100 and include Cisco

(CSCO), Dell (DELL), Intel (INTC), Microsoft (MSFT) and Sun Microsystems (SUNW). The stocks selected from the Paris Bourse are all in the CAC-40 index - Credit Agricole (CAGR), Alcatel (CGEP), France Telecom (FTE), Total (TOTF) and Vivendo (VIE). In total we collect data for a 10 year period, from January 1996 to December 2005. However, not all stocks have data available over this sample period and for those companies that list on later dates, we include the data from that date.

We collect data sampled at a five minute frequency, and the data we obtain from SIRCA includes the prices at the end of each five minute interval, the volume traded within each 5 minute interval, the total number of trades within each interval, the difference between the highest and lowest price within the interval, and the average bid-ask spread. For the NYSE and Nasdaq, where stocks trade from 9.30 – 16.00, we have 78 observations within the day. For the Paris Bourse trading hours are different. Before September 20, 1999, stocks on the Paris Bourse traded from 10.00-17.00, giving a total of 84 observations per trading day. Consecutively, two changes were made to the trading hours at the Paris Bourse. On September 20, 1999, the opening of the market was moved from 10.00 am to 9.00 am. Second, on April 3, 2000, the exchange extended its trading hours by another half hour. Hence from 3 April, 2000 onwards trading occurred from 9.00-17.30, which gives a total of 102 observations within the trading day. Since the period between these two changes is too short to make forecasts, we ignore this period in our analysis and focus on both sub-periods, before and after the changes in trading hours.

INSERT TABLE 1 ABOUT HERE

In table 1 we present some summary statistics for the stocks in our sample. Given the length of the sample period we have almost 200,000 observations per security. The Nasdaq securities in the sample are more risky than their NYSE counterparts, as is expected for this exchange. We further observe that unconditional return distributions are all clearly non-symmetric with more probability mass in the centre and in the tails of the distribution in comparison to the normal distribution.

3. Methodology

In this section we discuss the approach followed and the forecasting models used. Before discussing these forecasting models, we first discuss the benchmark used to evaluate the performance of the forecasts.

To determine how well our forecasts of daily volatility are, we need to set a benchmark for performance evaluation. This benchmark is the realized variance as proposed by Andersen et al. (2001). This realized variance is computed as the sum of squared intraday returns, and has been used as a benchmark to evaluate forecasting models (see e.g. Martens, 2001).

Let t refer to time measured in days and p_t refer to the log *closing* price of day t . Within this day, log-prices can be sampled at equidistant intervals and we refer to these intraday prices as $p_{t,n}$, where $n = 1, \dots, N$, with N the total number of intervals. Intraday returns are computed as the difference of these log prices, $r_{t,n} = p_{t,n} - p_{t,n-1}$. From these intraday returns we can compute the daily realized variance as

$$RV_t = \sum_{n=1}^N r_{t,n}^2, \quad (1)$$

where RV_t is the daily realized variance at the end of day t . This realized variance is shown to be a consistent estimator of the true, integrated variance (see Andersen et al., 2001). Further properties, for example the rate of convergence of the estimator, can be found in e.g. Barndorff-Nielsen and Shephard (2002).

The main research question in this paper is how information from the start of the day can be used to forecast the daily volatility. A first question that needs to be answered in that respect is how informative this information at the start of the day is about the volatility at the end of the day. That is, how informative is the volatility generated in the first N^* intervals of the trading day about the daily volatility. Similar to (1), the realized variance generated during this period is equal to

$$RV_t^{N^*} = \sum_{n=1}^{N^*} r_{t,n}^2, \quad (2)$$

where $RV_t^{N^*}$ represents the realized variance generated up to the N^* th interval. To investigate how informative the first N^* intervals are about the daily volatility, we run the following simple regression

$$RV_t = \alpha + \beta RV_t^{N^*} + \varepsilon_{t+1}, \quad (3)$$

where the coefficient of determination (R^2) of this regression will indicate how much of the daily realized variance can be explained by the realized variance generated over the first N^* intervals. Of course we cannot expect $\beta = 1$, which is preferable for a well functioning forecasting model, as $RV_t^{N^*}$ is not an unbiased estimate of RV_t .

Another feature we need to acknowledge is the part of the RV_t that has already been observed at the start of the day ($RV_t^{N^*}$). Therefore to fully assess whether the volatility generated at the start of the day is informative about the daily volatility, we need to determine what fraction of daily volatility has already been generated during the first N^* intervals. This fraction can be measured by the variance ratio

$$VR^{N^*} = \frac{1}{T} \sum_{t=1}^T \frac{RV_t^{N^*}}{RV_t}. \quad (4)$$

An unbiased forecast of daily volatility

The next step in forming a forecast of the daily volatility is to use typical patterns observed during the trading day to form an unbiased forecast. These typical patterns often referred to as U-shapes or J-shapes, have been observed in several intraday variables including the return volatility during the trading day (see Wood et al., 1985; Harris, 1986 and Andersen and Bollerslev, 1994). In the following subsections we discuss two approaches we follow to forecast end-of-day volatility. Firstly, we follow a very simple approach that determines the intraday seasonality by averaging over past observations. Second, we fit Flexible Fourier

Forms to past data, to parameterize and smoothen the intraday seasonality. A great advantage of this second approach is that we can make the seasonality a function of other variables, such as past volatility, forecasted daily volatility, or the volatility generated at the start of the day.

Using the Average Shape over past trading days

Our first approach to forecast end of day volatility is to consider the past K trading days, and infer the intraday seasonal from that data. Let $s_{t,n}$ denote the expected seasonal on day t for the intraday interval n . This seasonal can be determined as

$$s_{t,n} = \frac{1}{K} \sum_{k=1}^K r_{t-k,n}^2. \quad (5)$$

Thus, the seasonal is obtained from the average squared return for the interval n over the past K trading days. Once the seasonal for every interval during the day has been determined, we can use its information to form an unbiased forecast of the daily volatility.

According to (5) the variance observed over the first N^* intervals is equal to $\sum_{n=1}^{N^*} s_{t,n}$. Similarly,

the total variance for the day is equal to $\sum_{n=1}^N s_{t,n}$, the sum of the seasonal over all N intervals.

Given these numbers we can compute the fraction of volatility we have observed already assuming the seasonal is a good indicator for this. Combining this information with the information from the realized variance already observed during the first N^* intervals, we can form an unbiased prediction of the daily realized variance as

$$R\tilde{V}_t = RV_t^{N^*} \times \frac{\sum_{n=1}^N s_{t,n}}{\sum_{n=1}^{N^*} s_{t,n}} \quad (6)$$

and measure the performance of this prediction using the simple regression

$$RV_t = \alpha + \beta R\tilde{V}_t + \varepsilon_{t+1}. \quad (7)$$

The performance of this prediction can again be measured by the R^2 of the regression. However, given that $R\tilde{V}_t$ is merely a rescaled version of $RV_t^{N^*}$, the R^2 will not differ substantially from the R^2 obtained from (3), unless the intraday seasonal over the past K trading days contains additional information. As an alternative performance measure we also consider the heteroskedasticity consistent mean squared error (HMSE)

$$HMSE = \frac{1}{T} \sum_{t=0}^{T-1} \left(1 - \frac{R\tilde{V}_{t+1}}{RV_{t+1}} \right)^2. \quad (8)$$

Flexible Fourier Form

In addition to the seasonals we determined using the average of the past K trading days, we can also fit a functional form to describe the return volatility shape during the trading day. A previously followed approach (e.g. Andersen and Bollerslev, 1997 and Martens, 2001) is to fit a Flexible Fourier Form (FFF) to the intraday return volatility. Some advantages of fitting an FFF are that (1) all data is being used to determine the shape of the seasonal (instead of using only those observation in the same interval in the day); (2) the shape will be smoother, thus mitigating the effect of outliers; (3) the shape of the seasonal can be a function of other variables e.g. previous days volatility, or the volatility at the beginning of the trading day. Following Andersen and Bollerslev we consider the following decomposition for intraday returns

$$r_{t,n} = \bar{r}_{t,n} + \frac{\sigma_t s_{t,n} Z_{t,n}}{N^{1/2}} \quad (9)$$

where $Z_{t,n}$ is a standardised i.i.d. random variate. By squaring both sides of (9) and taking a logarithmic transformation, we can deduce the following relationship,

$$x_{t,n} \equiv 2 \ln[|r_{t,n} - \bar{r}_{t,n}|] - \ln \sigma_t^2 + \ln N = \ln s_{t,n}^2 + \ln Z_{t,n}^2. \quad (10)$$

The functional form for the seasonal, $f(\theta; \sigma_t, n)$ can be obtained by running a nonlinear regression of $x_{t,n}$ on parameterization of this form. We specify this functional form similar to Andersen and Bollerslev as

$$\begin{aligned}
f(\theta; \sigma_t, n) = & \sum_{j=0}^J \sigma_t^j \left[\mu_{0j} + \mu_{1j} \frac{n}{N_1} + \mu_{2j} \frac{n^2}{N_2} + \sum_{d=1}^D \lambda_{dj} I_{n=D_d} \right] \\
& + \sum_{p=1}^P \left(\gamma_{pj} \cos \frac{pn2\pi}{N} + \delta_{pj} \sin \frac{pn2\pi}{N} \right).
\end{aligned} \tag{11}$$

The first part of this functional form relates to linear and quadratic terms that may affect the shape of the intraday seasonal (relating to the μ -coefficients), where N_1 and N_2 are normalizing constants.¹ We also include dummy variables to capture irregularities in the seasonal shape located near the opening and close of the trading day. The last part of the function approximates the shape of the seasonal by fitting a numbers of sinusoids (captured by the γ - and δ -coefficients).

In the empirical application we set J equal to one, thus allowing the shape to be a function of the standard deviation of the trading day. Specifically, we use the amount of volatility generated during the first N^* intervals as the volatility variable, σ_t . Because the shape of the seasonal for stocks typically has a very pronounced U- or J- shape, the number of sinusoids needed is small and hence we set $P = 2$. As mentioned, any irregularities occurring near the start and close of the trading day can be captured by the inclusion of a few dummy variables.

4. Results

In this section we present the results for the forecasts of end of day volatility. We first address the issue of how much information the start of day volatility carries about the end-of-day volatility. Second, we consider how much information other intraday variables contribute to explaining end-of-day volatility. Finally, we present results for the forecasts using either simple averages to determine the intraday shape or using the FFF.

To assess how informative the start of the day volatility is about the end-of day volatility, we compute daily realized variances, as well as the realized variance at the start of the day. These start of day volatilities are computed over the first 30, 60, 90 and 120 minutes of the trading day, respectively. We present Variance Ratios and coefficients of determination (R^2 (adj)) for

¹ The two normalizing constants are defined as $N_1 = (N + 1)/2$ and $N_2 = (N + 1)(N + 2) / 6$.

the regression in (3) in table 2. The results show a very interesting and consistent picture. When using only the first 30 minutes of the trading day, variance ratios are all below 0.2, indicating that the first 30 minutes of the trading day generate less than 20% of the daily volatility. However, the R^2 values of these regressions are up to 0.55, suggesting that the volatility generated during this first half hour is highly informative about the end-of-day volatility. Similar pictures are observed for the other intervals. When using the first 60 minutes of the trading day we observe up to 30.6% of the daily volatility. Again the information we can obtain about the daily volatility is much higher, up to 68% as measured by R^2 . For the first 90 and 120 minutes we observe approximately up to 40% and 48% respectively, while the coefficients of determination are as high as 73.5% and 78.4% for the 90 and 120 minute intervals, respectively.

INSERT TABLE 2 ABOUT HERE

Two other interesting observations can be made from table 2. First, the first two intervals (30 and 60 minutes) are marginally much more informative than the latter two intervals providing a very clear indication of the importance of the first hour of trading. Second, when considering results for the Paris Bourse before and after the change in trading hours, we see that after the change in trading hours, both variance ratios and coefficients of determination increase. This indicates that in the later period (1) more volatility is being generated at the start of the trading day, and (2) the activity at the start of the trading day contains more information about the end-of-day volatility.

Table 2 shows that the sum of squared returns during the initial hours of the trading day contains important information about the realized volatility at the end of the day. However, the realized volatility at the start of the day as such is a poor predictor of end of day volatility and needs to be rescaled to form a good predictor of end-of day volatility. The rescaling of the start of day volatility can be done using the intraday shape typically observed in squared returns and other microstructure variables. These intraday shapes, commonly referred to as U-shapes or J-shapes, are typically very pronounced.

INSERT FIGURE 1 ABOUT HERE

In figure 1 we show these intraday shapes in squared returns for representative stocks in the sample (solid line)², together with a fitted Flexible Fourier Form (dashed line). Both NYSE and Nasdaq stocks reveal this pronounced U-shape and open with about the same amount of volatility. The stocks listed on the Paris Bourse have a slightly different seasonal and we can observe an increase in the volatility during the later part of the trading day due to the opening of the US markets.

When forming these seasonals, as defined in (5), to determine the scale factor, we use the past 200 trading days and calculate the average squared return within each five minute interval. This scale factor is then multiplied by the variance that was generated during the first N^* intervals. Again we use start of the day volatilities computed at the initial 30, 60, 90 and 120 minutes of the trading day.

The forecast evaluations are presented in table 3. We show the evaluations using regressions of the observed realized variance on the forecasted realized variance as well as the HMSE. For the regression evaluation we expect the coefficient $b_0 = 0$ and $b_1 = 1$.

INSERT TABLE 3 ABOUT HERE

Panel A shows the results for the stocks listed on the NYSE. Overall, we find that b_1 is in most cases smaller than one, indicating that our forecast model generally overestimates the actual volatility. Interesting though is the fact that R^2 (adj) increases slightly, indicating that the addition of the shape, which changes over time, still adds information about the end of day volatility. The HMSE values confirm the results of the R^2 (adj). Note these values decrease considerably when the initial day interval increases. In panel B results are reported for the Nasdaq stocks. Overall, the findings are in line with the findings for the NYSE stocks, i.e. the forecast based on the intraday seasonal estimated using past averages overestimates the actual realized volatility. However, in contrast to the NYSE we observe that this overestimation is in general less than for the NYSE stocks. Panels C and D show the results for the Paris Bourse stocks before and after the change in trading hours. The results before the change show that the overestimation of volatility is considerable. However, after the change in trading hours the results are in line with those from the other two markets.

² All shapes are normalized to 1.

The usage of the intraday shape estimated on the basis of an average of past squared returns has led to reasonable forecasts of the actual daily realized volatility. However, up to now we have assumed that high and low volatility days evolve in the same way. An interesting question is to see whether the extent of volatility generated at the start of the day has an impact on the intraday seasonal for the rest of the day. In figure 2 we show the intraday seasonals for the same stocks presented in figure 1 for high and low first hour realized volatility days. For all stocks and all market we see very distinct differences between these shapes. For the US market we find that when the first hour volatility is high, the intraday seasonal takes the form of a J-shape, whereas when the first hour volatility is low the graphs show very distinct U-shapes. For the stocks on the Paris Bourse we observe a similar effect, although the distinction between J-shapes and U-shapes is less pronounced. We do observe that on low first hour volatility days, the volatility generated during the latter part of the day increases.

INSERT FIGURE 2 ABOUT HERE

These graphs clearly demonstrate the importance of the volatility that is generated at the start of the day and also show that if the volatility near the opening is high the rest of the day tends to be calmer. For this reason, when fitting the FFF, we allow the intraday seasonal to be a function of the volatility generated in the first hour. In setting up the forecasting model we estimate this shape on data over the past 200 trading days and then use the volatility generated during the first hour of the new trading day as the input to forecast the seasonal for that day. Once the seasonal has been forecasted we proceed as before by computing the scale factor and the forecast of the realized volatility.

INSERT TABLE 4 ABOUT HERE

In table 4 we present the forecast evaluation for the forecasts based on the Flexible Fourier Forms in conjunction with the forecasts made using the first hour of trading information. In comparison to the forecasts based on simple averages, the performance of the FFF adjusted forecasts is substantially better. In general, coefficients for b_1 increase and move closer to one, indicating that the magnitude of the overestimation is decreasing. Hence, incorporating the start of day volatility as an input to determine the intraday seasonal is an important factor

to consider. Also in most cases we observe increases in $R^2(\text{adj})$, which indicates that the prediction of the shape also increases the prediction of the end of day volatility.

4. Conclusion

In this study we use start of the day realized volatility to predict end of day volatility. By scaling these forecasts by seasonal factors we have been able to further improve the model's forecasting performance. We assess out-of-sample forecasting performance using a variety of forecasting assessment criteria, including regression of observed realized ex-post volatility on volatility forecasts. Second, we illustrate the robustness of the proposed procedure by carrying out an analysis on stocks traded under different market structures. Our findings show that the start of the trading day is highly informative about the volatility observed at the end of the trading day and that end-of-day forecasts can be made fairly accurately. Our research further contributes to the extensive literature that considers the influence of intraday variables on daily volatility. While these studies consider the impact of these variables over the whole trading day, we consider the role of the first (few) trading hours. The results show that this opening period contains important information about the formation of daily volatility.

References

- Andersen, T.G. and T. Bollerslev (1994), "Intraday Seasonality and Volatility Persistence in Foreign Exchange and Equity Markets," Working Paper, Finance Department, Kellogg School of Management, Northwestern University.
- Andersen, T. and T. Bollerslev (1997), "Intraday Periodicity and Volatility Persistence in Financial Markets", *Journal of Empirical Finance* 4, 115-158.
- Andersen, T., T. Bollerslev, F.X. Diebold and P. Labys (2001), "The Distribution of Realized Exchange Rate Volatility," *Journal of the American Statistical Association* 96, 42-55.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2003), "Modeling and Forecasting Realized Volatility," *Econometrica*, 71 (2), 579-625.
- Andersen, T.G., T. Bollerslev and N. Meddahi (2002), "Analytic Evaluation of Volatility Forecasts," Working Paper, Finance Department, Kellogg School of Management, Northwestern University.
- Andersen, T.G., L. Benzoni and J. Lund (2002), "An Empirical Investigation of Continuous-Time Models for Equity Returns," *Journal of Finance*, 57, 1239-1284.
- Ané, T. and H. Geman (2000), "Order Flow, Transaction Clock, and Normality of Asset Returns", *Journal of Finance* 55, 2259-2284.
- Avramov, D., T. Chordia and A. Goyal (2006), "The Impact of Trades on Daily Volatility", *Review of Financial Studies*, forthcoming.
- Barndorff-Nielsen, O., N. Shephard (2002), "Estimating Quadratic Variation using Realized Variance", *Journal of Applied Econometrics* 17, 457-477.
- Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics* 31, 307-327.

Engle, R.F. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation", *Econometrica* 50, 987-1008.

Harris, L. (1986), "A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns", *Journal of Financial Economics* 16, 99-117.

Martens, M. (2001), "Forecasting Daily Exchange Rate Volatility using Intraday Returns", *Journal of International Money and Finance* 20, 1-23.

Martens, M., Y.-C. Chang and S. Taylor (2002), "A Comparison of Seasonal Adjustment Methods when Forecasting Intraday Volatility", *Journal of Financial Research* 25, 283-299.

Schwert, G.W. (1990), "Stock Volatility and the Crash of '87", *Review of Financial Studies* 3, 77-102.

Wood, R., T. McInish and J. Ord (1985), "An Investigation of Transaction Data for NYSE Stocks", *Journal of Finance* 25, 723-739.

Table 1: Summary Statistics

	Obs	Mean	St. dev.	Kurtosis	Skewness
		(x10,000)	(x10,000)		
<i>NYSE</i>					
Alcoa (AA)	193,363	-0.026	21.303	14.15	-0.015
Boeing (BA)	193,207	0.028	20.830	14.42	-0.034
GE	193,285	0.006	19.277	19.97	0.357
IBM	193,285	0.009	19.345	22.85	0.525
AT&T (T)	192,739	-0.050	21.762	11.78	0.179
<i>Nasdaq</i>					
CSCO	195,469	-0.110	30.088	20.42	0.076
DELL	195,313	0.097	30.918	13.94	0.099
INTC	195,391	-0.061	27.135	12.84	0.099
MSFT	195,469	0.013	21.825	22.85	0.446
SUNW	195,313	-0.268	37.299	11.16	0.012
<i>Paris Bourse</i>					
CAGR	105,987	0.074	24.029	45.30	-0.303
Alcatel (CEGP)	227,470	-0.139	32.415	31.22	-0.075
Telecom (FTE)	190,307	-0.018	31.404	44.35	0.093
TOTF	227,470	0.071	21.215	17.43	0.043
Vivendo (VIE)	142,243	0.077	28.340	43.64	-0.126

Table 2: Variance Ratios and Forecast Evaluation Regressions for Realized Variances Generated during Initial Trading Hours

	30 min VR ^{30 min}	R ² (adj)	60 min VR ^{60 min}	R ² (adj)	90 min VR ^{90 min}	R ² (adj)	120 min VR ^{120 min}	R ² (adj)
<i>NYSE</i>								
AA	0.171	0.394	0.299	0.596	0.391	0.694	0.463	0.754
BA	0.189	0.347	0.314	0.618	0.408	0.735	0.479	0.784
GE	0.150	0.399	0.264	0.624	0.354	0.723	0.419	0.772
IBM	0.175	0.476	0.292	0.635	0.373	0.670	0.435	0.705
T	0.179	0.450	0.304	0.623	0.385	0.713	0.453	0.768
<i>Nasdaq</i>								
CSCO	0.187	0.461	0.299	0.550	0.375	0.603	0.438	0.640
DELL	0.200	0.468	0.317	0.581	0.396	0.623	0.457	0.657
INTC	0.186	0.554	0.304	0.680	0.384	0.733	0.446	0.764
MSFT	0.175	0.430	0.295	0.556	0.372	0.633	0.437	0.676
SUNW	0.194	0.435	0.306	0.569	0.388	0.633	0.450	0.669
<i>Paris Bourse – before</i>								
CAGR	-	-	-	-	-	-	-	-
CEGP	0.068	0.451	0.151	0.602	0.221	0.602	0.288	0.620
FTE	0.064	0.213	0.121	0.254	0.184	0.291	0.246	0.347
TOTF	0.066	0.163	0.128	0.222	0.205	0.425	0.273	0.463
VIE	-	-	-	-	-	-	-	-
<i>Paris Bourse – after</i>								
CAGR	0.162	0.476	0.239	0.592	0.302	0.694	0.352	0.722
CEGP	0.131	0.552	0.209	0.606	0.269	0.680	0.326	0.739
FTE	0.134	0.431	0.212	0.539	0.290	0.735	0.348	0.773
TOTF	0.116	0.512	0.191	0.603	0.253	0.641	0.307	0.677
VIE	0.150	0.296	0.232	0.456	0.316	0.693	0.374	0.706

Notes: (1) Variance Ratios refer to the portion of daily volatility generated at the start of the day; (2) R² is the coefficient of determination for the regression of daily realized variances on the realized variance generated during the first part of the trading day. Results are reported for the first 30, 60, 90 and 120 minutes of the trading day. Results for stocks listed on the Paris Bourse are reported before and after the change in trading hours.

Table 3: Forecast Evaluation using Simple Averages for Seasonals

This table presents ex-post evaluation results, where the daily forecasts of volatility are adjusted using simple averages of intraday seasonals over the past 200 trading days. Forecasts are made based on the first 30, 60 90, and 120 minutes of the trading day. The table reports results for both OLS regressions of realised volatility on a constant and the forecast and the associated heteroskedasticity consistent mean squared error (HMSE) statistic.

Panel A: NYSE

	b0	S.E.	b1	S.E.	R ² (adj)	HMSE
AA						
30 min	0.00021	(0.00001)	0.416	(0.027)	0.405	0.632
60 min	0.00014	(0.00001)	0.601	(0.024)	0.604	0.285
90 min	0.00011	(0.00001)	0.691	(0.027)	0.702	0.180
120 min	0.00008	(0.00001)	0.760	(0.025)	0.760	0.128
BB						
30 min	0.00021	(0.00003)	0.392	(0.082)	0.347	0.493
60 min	0.00014	(0.00002)	0.588	(0.066)	0.625	0.236
90 min	0.00016	(0.00003)	0.526	(0.077)	0.705	0.147
120 min	0.00013	(0.00003)	0.612	(0.076)	0.762	0.111
GE						
30 min	0.00014	(0.00001)	0.515	(0.040)	0.405	0.560
60 min	0.00008	(0.00001)	0.705	(0.044)	0.641	0.260
90 min	0.00015	(0.00003)	0.479	(0.099)	0.666	0.170
120 min	0.00013	(0.00003)	0.565	(0.099)	0.727	0.127
IBM						
30 min	0.00015	(0.00001)	0.486	(0.057)	0.444	0.453
60 min	0.00009	(0.00001)	0.682	(0.058)	0.615	0.232
90 min	0.00007	(0.00001)	0.739	(0.057)	0.656	0.159
120 min	0.00006	(0.00001)	0.802	(0.055)	0.696	0.120
T						
30 min	0.00020	(0.00001)	0.461	(0.039)	0.434	0.518
60 min	0.00016	(0.00002)	0.574	(0.052)	0.599	0.250
90 min	0.00012	(0.00002)	0.675	(0.048)	0.696	0.170
120 min	0.00010	(0.00001)	0.736	(0.045)	0.758	0.125

<i>Panel B: Nasdaq</i>						
	b0	S.E.	b1	S.E.	R ² (adj)	HMSE
CSCO						
30 min	0.00025	(0.00003)	0.624	(0.042)	0.462	0.414
60 min	0.00016	(0.00002)	0.756	(0.039)	0.547	0.218
90 min	0.00012	(0.00002)	0.817	(0.036)	0.601	0.152
120 min	0.00009	(0.00002)	0.861	(0.033)	0.639	0.111
DELL						
30 min	0.00027	(0.00003)	0.616	(0.052)	0.464	0.416
60 min	0.00018	(0.00002)	0.742	(0.033)	0.575	0.218
90 min	0.00014	(0.00002)	0.794	(0.030)	0.614	0.151
120 min	0.00011	(0.00002)	0.841	(0.025)	0.650	0.114
INTC						
30 min	0.00024	(0.00002)	0.598	(0.039)	0.542	0.391
60 min	0.00016	(0.00002)	0.725	(0.036)	0.674	0.189
90 min	0.00012	(0.00001)	0.791	(0.028)	0.730	0.131
120 min	0.00010	(0.00001)	0.823	(0.025)	0.763	0.109
MSFT						
30 min	0.00014	(0.00002)	0.629	(0.048)	0.423	0.476
60 min	0.00012	(0.00004)	0.677	(0.099)	0.553	0.233
90 min	0.00009	(0.00003)	0.767	(0.080)	0.630	0.165
120 min	0.00008	(0.00002)	0.777	(0.062)	0.672	0.124
SUNW						
30 min	0.00048	(0.00004)	0.551	(0.036)	0.433	0.423
60 min	0.00029	(0.00003)	0.726	(0.032)	0.566	0.217
90 min	0.00024	(0.00003)	0.779	(0.034)	0.631	0.147
120 min	0.00019	(0.00003)	0.824	(0.034)	0.667	0.112

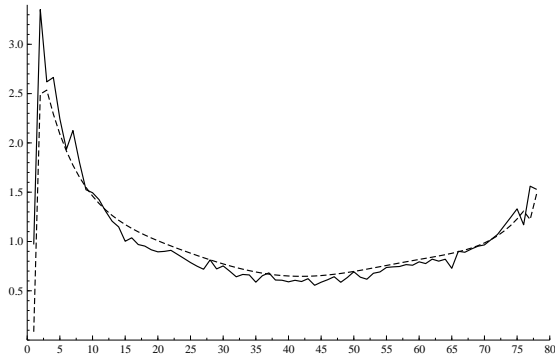
<i>Panel C: Paris Bourse – Before</i>						
	b0	S.E.	b1	S.E.	R ² (adj)	HMSE
CGEP						
30 min	0.00033	(0.00003)	0.402	(0.073)	0.492	0.864
60 min	0.00041	(0.00003)	0.276	(0.021)	0.574	0.517
90 min	0.00041	(0.00004)	0.277	(0.023)	0.571	0.354
120 min	0.00036	(0.00003)	0.364	(0.021)	0.608	0.279
FTE						
30 min	0.00058	(0.00007)	0.356	(0.090)	0.230	1.365
60 min	0.00050	(0.00008)	0.478	(0.124)	0.265	0.697
90 min	0.00046	(0.00009)	0.525	(0.138)	0.308	0.513
120 min	0.00040	(0.00011)	0.612	(0.171)	0.315	0.384
TOTF						
30 min	0.00047	(0.00003)	0.259	(0.039)	0.149	1.136
60 min	0.00041	(0.00002)	0.365	(0.034)	0.231	0.628
90 min	0.00037	(0.00003)	0.418	(0.035)	0.426	0.443
120 min	0.00039	(0.00003)	0.397	(0.034)	0.460	0.330
<i>Panel D: Paris Bourse – After</i>						
CAGR						
30 min	0.00023	(0.00003)	0.440	(0.050)	0.480	0.551
60 min	0.00015	(0.00002)	0.619	(0.060)	0.595	0.321
90 min	0.00010	(0.00002)	0.750	(0.065)	0.694	0.230
120 min	0.00009	(0.00002)	0.779	(0.072)	0.722	0.182
CGEP						
30 min	0.00077	(0.00008)	0.380	(0.026)	0.545	0.600
60 min	0.00072	(0.00010)	0.421	(0.067)	0.585	0.328
90 min	0.00060	(0.00010)	0.514	(0.069)	0.662	0.231
120 min	0.00052	(0.00008)	0.581	(0.053)	0.727	0.69
FTE						
30 min	0.00048	(0.00008)	0.482	(0.058)	0.421	0.564
60 min	0.00035	(0.00007)	0.624	(0.055)	0.532	0.335
90 min	0.00040	(0.00007)	0.568	(0.055)	0.720	0.229
120 min	0.00036	(0.00007)	0.613	(0.057)	0.760	0.175
TOTF						
30 min	0.00013	(0.00002)	0.553	(0.062)	0.504	0.461
60 min	0.00010	(0.00002)	0.674	(0.072)	0.602	0.256
90 min	0.00009	(0.00002)	0.703	(0.085)	0.639	0.174
120 min	0.00007	(0.00002)	0.761	(0.079)	0.677	0.135
VIE						
30 min	0.00036	(0.00004)	0.447	(0.035)	0.304	0.747
60 min	0.00024	(0.00004)	0.620	(0.072)	0.453	0.411
90 min	0.00022	(0.00005)	0.657	(0.070)	0.691	0.268
120 min	0.00026	(0.00005)	0.599	(0.064)	0.701	0.203

Table 4: Forecasts using Flexible Fourier Forms

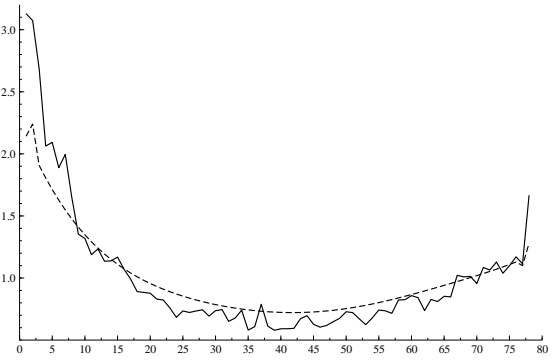
This table presents ex-post evaluation results, where the daily forecasts of volatility are adjusted for seasonality using a Flexible Fourier Form estimated over the past 200 trading days. Forecasts are made using the first 60 minutes of the trading day. The table reports results for both the OLS forecast evaluation regression and the associated heteroskedasticity consistent mean squared error (HMSE) statistic.

	b0	S.E.	b1	S.E.	R ² (adj)	HMSE
<i>NYSE</i>						
AA	0.00006	(0.00001)	0.765	(0.037)	0.605	0.371
BA	0.00006	(0.00001)	0.724	(0.039)	0.598	0.374
GE	0.00004	(0.00001)	0.848	(0.047)	0.585	0.282
IBM	0.00004	(0.00002)	0.773	(0.057)	0.653	0.353
T	0.00005	(0.00002)	0.839	(0.050)	0.636	0.286
<i>Nasdaq</i>						
CSCO	0.00006	(0.00002)	0.860	(0.037)	0.544	0.274
DELL	0.00009	(0.00006)	1.114	(0.137)	0.558	0.394
INTC	0.00007	(0.00001)	0.796	(0.025)	0.694	0.269
MSFT	0.00003	(0.00001)	0.869	(0.044)	0.611	0.273
SUNW	0.00013	(0.00004)	0.820	(0.038)	0.572	0.254
<i>Paris Bourse – Before</i>						
CEGP	0.00025	(0.00006)	0.657	(0.061)	0.092	0.262
FTE	0.00010	(0.00011)	0.841	(0.159)	0.497	0.419
TOTF	0.00020	(0.00005)	0.767	(0.108)	0.309	0.231
<i>Paris Bourse – after</i>						
CAGR	-0.00001	(0.00003)	0.751	(0.074)	0.693	0.749
CGEP	0.00041	(0.00006)	0.564	(0.044)	0.671	0.582
FTE	0.00032	(0.00007)	0.474	(0.058)	0.542	0.842
TOTF	0.00004	(0.00001)	0.718	(0.057)	0.546	0.491
VIE	0.00000	(0.00003)	0.713	(0.053)	0.530	1.149

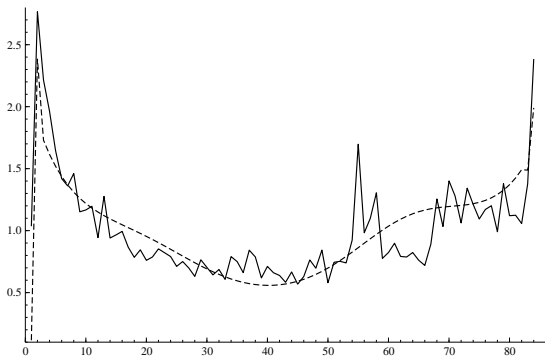
Figure 1: Intraday Seasonals for Selected Stocks




Panel A: NYSE: T



Panel B: Nasdaq: SUNW



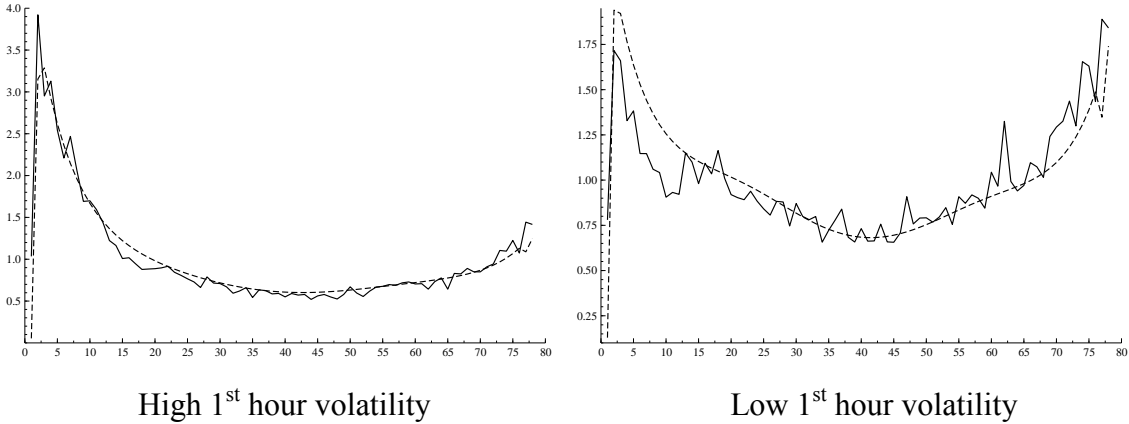
Panel C: Paris Bourse – before: TOTF



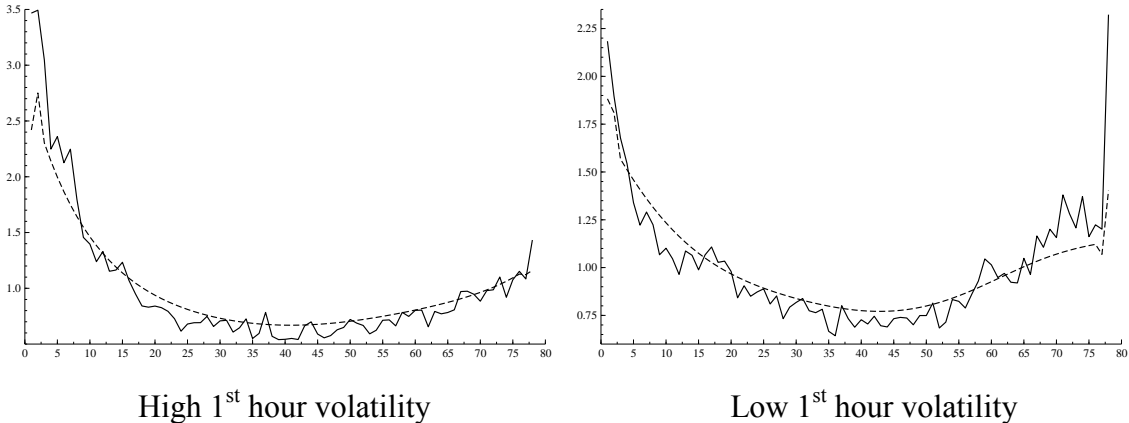
Panel D: Paris Bourse – after: TOTF

Figure 2: Intraday Seasonals for Selected Stocks - High and Low First Hour Volatility


Panel A: NYSE: T




Panel B: Nasdaq: SUNW



Panel C: Paris Bourse – before: TOTF

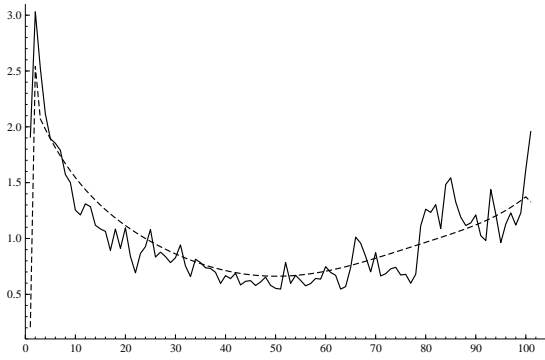


High 1st hour volatility

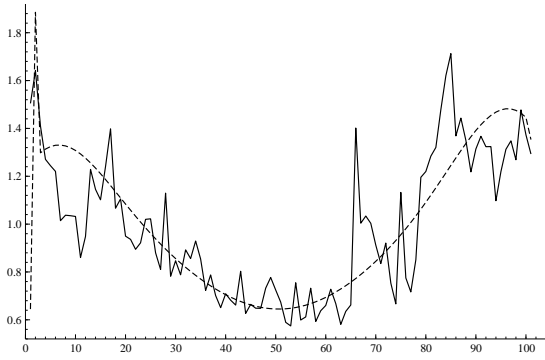


Low 1st hour volatility

Panel D: Paris Bourse – after: TOTF



High 1st hour volatility



Low 1st hour volatility