

Common intraday periodicity

Citation for published version (APA):

Hecq, A. W., Palm, F. C., & Laurent, S. F. J. A. (2011). *Common intraday periodicity*. METEOR, Maastricht University School of Business and Economics. METEOR Research Memorandum No. 010 <https://doi.org/10.26481/umamet.2011010>

Document status and date:

Published: 01/01/2011

DOI:

[10.26481/umamet.2011010](https://doi.org/10.26481/umamet.2011010)

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.

Alain Hecq, Sébastien Laurent, Franz Palm

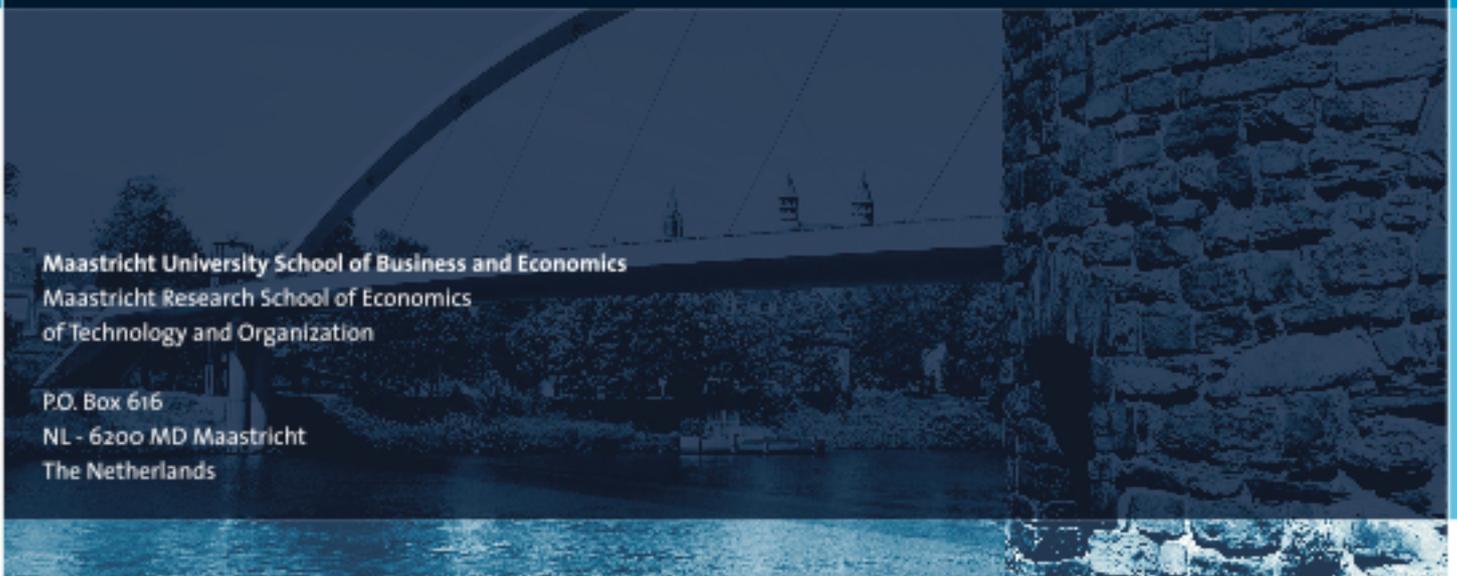
Common intraday periodicity

RM/11/010

METEOR

Maastricht University School of Business and Economics
Maastricht Research School of Economics
of Technology and Organization

P.O. Box 616
NL - 6200 MD Maastricht
The Netherlands



COMMON INTRADAY PERIODICITY*

Alain Hecq

Maastricht University, Department of Quantitative Economics
a.hecq@maastrichtuniversity.nl

Sébastien Laurent[†]

Maastricht University, Department of Quantitative Economics
and CORE, Université Catholique de Louvain
s.laurent@maastrichtuniversity.nl

Franz C. Palm

Maastricht University, Department of Quantitative Economics
f.palm@maastrichtuniversity.nl

December 23, 2010

Abstract

Using a reduced rank regression framework as well as information criteria we investigate the presence of commonalities in the intraday periodicity, a dominant feature in the return volatility of most intraday financial time series. We find that the test has little size distortion and reasonable power even in the presence of jumps. We also find that only three factors are needed to describe the intraday periodicity of thirty US asset returns sampled at the 5-minute frequency. Interestingly, we find that for most series the models imposing these commonalities deliver better forecasts of the conditional intraday variance than those where the intraday periodicity is estimated for each asset separately.

JEL: C10, C32.

Keywords: Intraday periodicity, Realized volatility, Common features.

*We thank the editor Torben Andersen, two anonymous referees, Kris Boudt, Christophe Croux, Jurgen Doornik, Neil Ericsson, David Hendry, Giovanni Urga, Francesco Violante and participants at the ninth OxMetrics conference, Computational and Financial Econometrics (CFE'10) conference and Einaudi Institute for Economics and Finance seminar for providing us useful suggestions on an earlier version of the paper.

[†]Correspondence to: Sébastien Laurent, Department of Quantitative Economics, Maastricht University, School of Business and Economics, P.O. Box 616, 6200 MD Maastricht, The Netherlands. Tel.: +31 43 388 38 43. Fax.: +31 43 388 20 00.

1 Introduction

The returns of most intraday financial times (e.g. at the 5-minutes frequency) are characterised by the presence of periodicity in their volatility. The behaviour of a time series is called intraday periodic if it shows a periodic structure within a day. For instance, the foreign exchange (FX) market exhibits strong periodic effects caused by the varying number of traders present during the day in the three major markets.¹ This translates into a U-shaped pattern in the ACF of absolute and squared intraday returns. Standard volatility models (ARCH or SV models), implying a geometric decay in the squared return autocorrelation structure, cannot accommodate strong regular cyclical patterns of that sort.

These periodic movements can be captured by non-parametric techniques (see Taylor and Xu, 1997) or in a parametric approach by a set of dummy variables (see Baillie and Bollerslev, 1991) or a bunch of trigonometric functions (see Andersen and Bollerslev, 1997). In the latter framework however, the number of parameters to estimate is usually quite large. This number further inflates when considering several assets in a multivariate modelling leading to a potential loss of efficiency.

However, this intraday periodic feature might be common to several series. Testing, discovering and imposing these commonalities can be exploited to improve parameter efficiency and forecasts accuracy. To this goal we first extend to intraday series the testing procedure proposed by Engle and Hylleberg (1996) to extract common deterministic seasonal features in macroeconomic time series.² We propose to use a reduced rank approach to study the presence of commonalities in the intraday periodic movements as well as multivariate information criteria to select the variables explaining the common periodic features. The Monte Carlo simulations indicate that our proposed strategy detects remarkably well both the number of periodic elements to be included and the existence of commonalities. We illustrate our approach using thirty US stock returns observed every five minutes in the period 2000-2008. Our approach shows that three common sources suffice to describe the intraday periodicity of these thirty series and that imposing these commonalities

¹The global FX market consists of three major markets, i.e., Asia, Europe and North America, and the major movements of intradaily return volatility can be attributed to the passage of market activity around the globe.

²Note that this paper does not look at the co-movements in the volatility (Engle and Susmel, 1993; Engle and Marcucci, 2006; Hecq et al., 2010) but at co-movements in the conditional mean equation for the logarithm of the absolute value of standardised asset returns.

improves prediction of not only the intraday periodicity but also the intraday conditional variance.

The approach adopted to extract and forecast the intraday periodicity in return volatility will be useful for modelling intraday Value-at-Risk (IVaR) (see e.g. Dionne et al., 2009 and Giot, 2005) and more generally intraday market risk measurement. The results of this paper are expected to be useful to determine the linkage between markets and for the application of temporal intraday trading rules as discussed by Goodhart and O'Hara (1997).

The remaining of the paper is structured as follows. Section 2 presents the model for univariate high-frequency time series. In Section 3 we propose tools for detecting the existence of common periodicity, the accuracy of which is evaluated in Section 4 with a set of Monte Carlo experiments. Section 5 deals with the empirical analysis and Section 6 concludes.

2 The Model for Univariate High-frequency Time Series

We assume that the sample consists of T days of M equally-spaced and continuously compounded intraday return observations $r_{j,t,i}$ ($t = 1, \dots, T$ and $i = 1, \dots, M$) of a financial asset j , $j = 1, \dots, N$. Hence, $r_{j,t,i}$ equals the i th return on day t of series j . In their seminal papers Andersen and Bollerslev (1997, 1998b), assume that the return $r_{j,t,i}$ is a normally distributed random variable with zero mean and that the standard deviation $\sigma_{j,t,i}$ can be rewritten as the product of a deterministic component $f_{j,t,i}$ representing essentially the calendar features and $s_{j,t,i}$ capturing the remaining volatility components (usually modelled using ARCH or stochastic volatility models), with $f_{j,t,i}$ and $s_{j,t,i} > 0 \forall j, t, i$. This leads to the univariate data generating process (DGP) for the high-frequency return $r_{j,t,i}$ given in Assumption 1.

Assumption 1 (Conditional normality of intraday returns)

$$r_{j,t,i} = \sigma_{j,t,i} u_{j,t,i} \text{ with } u_{j,t,i} \stackrel{i.i.d.}{\sim} N(0, 1) \quad (1)$$

$$\sigma_{j,t,i} = s_{j,t,i} f_{j,t,i}. \quad (2)$$

The periodic factor $f_{j,t,i}$ is assumed to be a deterministic function of periodic variables such as the time of the day and the day of the week. To ensure identifiability of both the periodicity and

the stochastic volatility $s_{j,t,i}$, we impose (see Assumption 2) that $f_{j,t,i}^2$ has mean one over the day.³

Assumption 2 (Normalization of $f_{j,t,i}$)

$$\frac{1}{M} \sum_{i=1}^M f_{j,t,i}^2 = 1 \quad \forall j, t. \quad (3)$$

The returns in (1) can be seen as discrete changes of an underlying continuous-time log-price process. Model (1) is motivated by the idea that this log-price process follows a *Brownian Semi-Martingale* (BSM) diffusion. Under the BSM model the log-price follows a diffusion consisting of the sum of a conditionally normal random process with mean $\mu(s)ds$ and variance $\sigma^2(s)ds$. Let $w(s)$ be a standard Brownian motion, then a BSM log-price diffusion admits the following representation $dp(s) = \mu(s)ds + \sigma(s)dw(s)$.

Throughout, we will study return series observed with a sufficiently high frequency such that the drift can be ignored. Model (1) is thus a discrete time version of the above BSM model where the drift is set to 0.

As mentioned above, Andersen and Bollerslev (1998a) also assume (see Assumption 3) that $s_{j,t,i}$ is constant over the day but can vary from day to day.

Assumption 3 (Constant stochastic volatility over the day)

$$s_{j,t,i} = \frac{s_{j,t}}{\sqrt{M}} \quad \forall i, j. \quad (4)$$

Visser (2010) recently used Assumption 3 in a GARCH context where $s_{j,t}$ is the conditional standard deviation of a GARCH(1,1) on daily returns $r_{j,t} \equiv \sum_{i=1}^M r_{j,t,i}$.

Under Assumptions 1 and 3, a consistent and very efficient estimator of $s_{j,t,i}$ is given by the square root of $\frac{1}{M}$ times the realized volatility of day t , i.e.

$$\hat{s}_{j,t,i} = \sqrt{\frac{1}{M} RV_{j,t}}, \quad (5)$$

$$\text{with } RV_{j,t} = \sum_{i=1}^M r_{j,t,i}^2. \quad (6)$$

³Note that Andersen and Bollerslev (1997) use a slightly different normalization condition, i.e. that $f_{j,t,i}$ has mean one over the day.

To estimate the periodicity factor $f_{j,t,i}$, Andersen and Bollerslev (1997) use the result that, under this model, the standardised returns $\bar{r}_{j,t,i} = r_{j,t,i}/\hat{s}_{j,t,i}$ are normally distributed with mean zero and variance $f_{j,t,i}^2$. Furthermore, they consider the regression equation

$$\log |\bar{r}_{j,t,i}| = \log f_{j,t,i}^* + \varepsilon_{j,t,i}, \quad (7)$$

where $f_{j,t,i}^*$ differs from $f_{j,t,i}$ because it does not necessarily satisfies Assumption 2, the error term $\varepsilon_{j,t,i}$ is i.i.d. distributed with mean zero and having the density function of the centered absolute value of the log of a standard normal random variable, i.e. $g(z) = \sqrt{2/\pi} \exp[z+c-0.5 \exp(2(z+c))]$. The parameter $c = -0.63518$ equals the mean of the log of the absolute value of a standard normal random variable.

Andersen and Bollerslev (1997) model $\log f_{j,t,i}^*$ as a linear function of a $m_j^* \times 1$ vector of variables $\mathbf{x}_{j,t,i}$ (such as sinusoid and polynomial transformations of the time of the day), i.e.

$$\log f_{j,t,i}^* = \omega_j + \boldsymbol{\gamma}_j' \mathbf{x}_{j,t,i}, \quad (8)$$

where $\boldsymbol{\gamma}_j$ is a column vector with m_j parameters.

Combining (7) with (8), we obtain the following regression equation

$$\log |\bar{r}_{j,t,i}| = \omega_j + \boldsymbol{\gamma}_j' \mathbf{x}_{j,t,i} + \varepsilon_{j,t,i}. \quad (9)$$

Despite the fact that $\varepsilon_{j,t,i}$ has a known and non-normal distribution, Andersen and Bollerslev (1997) propose to estimate model (9) by OLS, which corresponds to a Gaussian QML estimator under model (1). Monte Carlo simulation results reported by Boudt et al. (2010) suggest that the loss of efficiency in the estimation of $f_{j,t,i}$ by OLS compared to that of the efficient MLE is not dramatic under this model. Furthermore they also show that the OLS estimator is much less sensitive to jumps in the DGP than the MLE (see also Section 4.3).

Given consistent estimates of $\hat{\omega}_j$ and $\hat{\boldsymbol{\gamma}}_j$, $\log \hat{f}_{j,t,i}^*$ is obtained using Equation (8). Furthermore, following Andersen and Bollerslev (1997), an estimator for $f_{j,t,i}$ that satisfies Assumption 2 is given by

$$\hat{f}_{j,t,i} = \frac{\exp(\log \hat{f}_{j,t,i}^*)}{\sqrt{\frac{1}{M} \sum_{l=1}^M [\exp(\log \hat{f}_{j,t,i}^*)]^2}}, \quad (10)$$

where $\exp(\log \hat{f}_{j,t,i}^*)$ is a consistent estimate of the conditional median of model (7), not of its conditional mean.⁴

3 Testing for Common Intraday Periodic Features

Let us now assume that we observe a $N \times 1$ vector of returns $\mathbf{r}_{t,i}$ whose elements are $r_{j,t,i}$ (for $j = 1, \dots, N, t = 1, \dots, T$ and $i = 1, \dots, M$). Denote by $\bar{\mathbf{r}}_{t,i} = (r_{1,t,i}/\hat{s}_{1,t,i}, \dots, r_{N,t,i}/\hat{s}_{N,t,i})'$ the vector of standardised returns. Under the assumption that the $m_j^* \times 1$ vectors $\mathbf{x}_{j,t,i}$ are identical across j , i.e. $\mathbf{x}_{j,t,i} = \mathbf{x}_{t,i}$ with $m_j^* = m^*$, the multivariate counterpart of model (9) is

$$\mathbf{y}_{t,i} \equiv \log |\bar{\mathbf{r}}_{t,i}| = \boldsymbol{\omega} + \mathbf{\Gamma} \mathbf{x}_{t,i} + \boldsymbol{\varepsilon}_{t,i}, \quad (11)$$

where $\boldsymbol{\omega}$ and $\mathbf{\Gamma}$ are respectively a $N \times 1$ vector and a $N \times m^*$ matrix of parameters. The j -th row of $\mathbf{\Gamma}$ is given by $\boldsymbol{\gamma}'_j$ in (9). For the MT observations, (11) can be rewritten more compactly as follows

$$\mathbf{y} \equiv \log |\bar{\mathbf{r}}| = \boldsymbol{\iota} \otimes \boldsymbol{\omega}' + \mathbf{x} \mathbf{\Gamma}' + \boldsymbol{\varepsilon}, \quad (12)$$

where \mathbf{y} is a $MT \times N$ matrix, $\boldsymbol{\iota}$ is a MT column vector of ones and \otimes denotes the Kronecker product. Notice that the multivariate regression model (12) is in fact a system of seemingly unrelated regressions with identical regressors in each equation. For such a system the generalised least squares estimator is identical to the OLS estimator equation by equation.

In our framework, testing the presence of common periodic features in volatility is equivalent to testing for the rank of the matrix $\mathbf{\Gamma}$, namely investigating $\text{rank}(\mathbf{\Gamma}) = k$, with $0 \leq k \leq \min(N, m^*)$. For instance, when the true number of factors k^* equals 1 there is a unique source of periodicity generating the N returns. There will be commonality whenever $m^* < N$ whether $k^* = m^*$ or $k^* < m^*$. In the case where $N > m^* = k^*$, the m^* variables $\mathbf{x}_{t,i}$ can be interpreted as common factors. The cases where either $N > m^* > k^*$ or $k^* < \min(N, m^*)$ are more interesting as in these cases, $\mathbf{\Gamma}$ has reduced rank. This rank reduction implies that the model (11) can be parameterized by fewer than the Nm^* parameters as $\mathbf{\Gamma} = \boldsymbol{\alpha} \boldsymbol{\beta}'$ where $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are full column rank matrices of dimensions $N \times k^*$ and $m^* \times k^*$ respectively. Let us also denote $\mathbf{x} \boldsymbol{\beta} = \mathbf{F}$ the common periodic series.

⁴A transformation by a continuous function (such as the exponential transformation) of a quantile (median of logged variables) yields a consistent estimate of the quantile of the transformed variable.

One strategy to search for the rank of $\mathbf{\Gamma}$ is to jointly determine the number of periodic elements m to be included in \mathbf{x} and k by minimising the following multivariate information criteria over both the values of m and k :

$$AIC(s, m) = \ln \det \left(\hat{\Omega}_{\varepsilon, s} \right) + \frac{2}{MT} (N \times m - v_{s, m, N}) \quad (13)$$

$$HQ(s, m) = \ln \det \left(\hat{\Omega}_{\varepsilon, s} \right) + \frac{2 \ln \ln MT}{MT} (N \times m - v_{s, m, N}) \quad (14)$$

$$SC(s, m) = \ln \det \left(\hat{\Omega}_{\varepsilon, s} \right) + \frac{\ln MT}{MT} (N \times m - v_{s, m, N}), \quad (15)$$

where $\hat{\Omega}_{\varepsilon, s} = \hat{\Omega}_{\varepsilon} - \sum_{l=l^*}^s \ln(1 - \hat{\lambda}_l)$ for $s = 1, \dots, \min(N, m)$, the estimated covariance matrix of the residuals in the multivariate reduced rank regression, the one with full rank being $\hat{\Omega}_{\varepsilon} = \{\hat{\Sigma}_{\mathbf{y}\mathbf{y}} - \hat{\Sigma}_{\mathbf{y}\mathbf{x}} \hat{\Sigma}_{\mathbf{x}\mathbf{x}}^{-1} \hat{\Sigma}_{\mathbf{x}\mathbf{y}}\}$.⁵ This approach is similar to the one used by Athanasopoulos et al. (2009) for VAR models. This method can be used to detect the true pair (m^*, k^*) , the true number of factors k^* or the true number of periodic elements m^* .

A second strategy for determining k^* , and hence the domain of the spaces generating $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, is to rely on the above information criteria to detect m^* and for a given number of periodic elements, say m , rely on a canonical correlation analysis using a spectral decomposition of

$$\Sigma_{\mathbf{y}\mathbf{y}}^{-1} \Sigma_{\mathbf{y}\mathbf{x}} \Sigma_{\mathbf{x}\mathbf{x}}^{-1} \Sigma_{\mathbf{x}\mathbf{y}}, \quad (16)$$

where $\Sigma_{\mathbf{y}\mathbf{x}}$ are covariance matrices to be estimated by their empirical counterparts

$\hat{\Sigma}_{\mathbf{y}\mathbf{x}} = (\mathbf{y} - \bar{\mathbf{y}})'(\mathbf{x} - \bar{\mathbf{x}})/MT$ where $\bar{\mathbf{y}}$ and $\bar{\mathbf{x}}$ denote the empirical means of \mathbf{y} and \mathbf{x} respectively.

The null hypothesis that there exist at least $s \leq \min(N, m)$ linear combinations that annihilate k common periodic features is tested using

$$\xi_s = -MT \sum_{l=l^*}^s \ln(1 - \hat{\lambda}_l), \quad s = 1, \dots, \min(N, m) \quad (17)$$

with $l^* = \max(1, N - m + 1)$ and where $\hat{\lambda}_l$ is the l -th smallest eigenvalue of the estimated matrix (16). For i.i.d. normally distributed random variables, ξ_s follows asymptotically under the null a

⁵We must be careful however on the bounds for k when the number of periodic elements runs from $m < N$ to $m \geq N$. For instance, consider $N = 5$ returns and no reduced rank in $\mathbf{\Gamma}$, i.e. $k^* = \min(N, m^*)$. We should obtain $k = 2$ with $m^* = 2$, $k = 4$ with $m^* = 4$ and $k = 5$ for $m^* \geq 5$.

χ^2 distribution with $\nu_{s,m,N} = s \times \max(N, m) - s(\min(N, m) - s)$ degrees of freedom. Then, after having determined s , the number of detected factors is $k = \min(N, m) - s$.

Given the non-normal distribution of the disturbance terms here a χ^2 -test will at best be approximately valid. Its accuracy will be evaluated in a Monte Carlo study in the next section.

Finally, once k and m are determined, either by ξ_s or with the help of information criteria, we can form the common periodic components $\mathbf{x}\hat{\boldsymbol{\beta}} = \hat{\mathbf{F}}$. We obtain $\hat{\boldsymbol{\beta}}$ from the k eigenvectors associated with the k largest eigenvalues of $\hat{\boldsymbol{\Sigma}}_{\mathbf{x}\mathbf{x}}^{-1}\hat{\boldsymbol{\Sigma}}_{\mathbf{x}\mathbf{y}}\hat{\boldsymbol{\Sigma}}_{\mathbf{y}\mathbf{y}}^{-1}\hat{\boldsymbol{\Sigma}}_{\mathbf{y}\mathbf{x}}$, the dual problem of (16). Then the loadings coefficients $\hat{\boldsymbol{\alpha}}$ are estimated by regressing each return on an intercept and the k components in $\hat{\mathbf{F}}$. The notion underlying common features is, although similar in spirit, different from the one used in traditional factor models. Indeed, our extracted factors \mathbf{F} are linear combinations of observed variables $\mathbf{x}_{t,j}$. They are such that no significant information is lost when imposing these restrictions contrary to traditional factor models with latent factors where these factors (often estimated by principal components) try to explain a sufficient percentage of the variability of the series with a limited number of combinations of these series.

Finally, note that Model (11) assumes a common left null space of every periodic intraday component. The model can be generalized to include exogenous variables or additional periodic effects $\mathbf{z}_{t,i}$ such that

$$\mathbf{y}_{t,i} \equiv \log |\bar{\mathbf{r}}_{t,i}| = \boldsymbol{\omega} + \boldsymbol{\Gamma}\mathbf{x}_{t,i} + \boldsymbol{\Upsilon}\mathbf{z}_{t,i} + \boldsymbol{\varepsilon}_{t,i}. \quad (18)$$

In this framework we can either test the reduced rank of $[\boldsymbol{\Gamma} : \boldsymbol{\Upsilon}]$ or only of $\boldsymbol{\Gamma}$. In this latter case we can concentrate out the effect of $\mathbf{z}_{t,i}$ from both $\mathbf{y}_{t,i}$ and $\mathbf{x}_{t,i}$ by multivariate least squares and apply the previous approach to the residuals from these regressions. We use this approach in Section 4.2 to account for the presence of serial correlation in $\boldsymbol{\varepsilon}_{t,i}$ induced by a violation of Assumption 3.

In the next section we evaluate the performance of the three information criteria $AIC(s, m)$, $HQ(s, m)$ and $SC(s, m)$ to determine m and/or k as well as ξ_s to determine k .

4 Monte Carlo Simulation

We use simulated data to gauge the quality of the proposed approach in several situations. We generate $T = 100$ or 250 days of $N = 5$ or 15 univariate time series with $M = 288$ intraday observations per day (corresponding to 5-minute data of exchange rate returns). The DGP is a

multiplicative model implying intraday periodicity in volatility as well as GARCH effects.

We carry out three Monte Carlo studies. In the first one Assumptions 1-3 are satisfied while in the second and third ones respectively, Assumption 3 and Assumptions 1 and 3 are violated.

4.1 Case 1: Constant intraday stochastic volatility and conditional normality

The structure of the first DGP is similar to the one employed recently by Visser (2010). The stochastic part of the volatility is constant during the day but varies from day to day in accordance with a GARCH(1,1) structure at the daily level.

More specifically, the DGP is defined as Equations (1)-(2), with

$$s_{j,t,i} = \frac{s_{j,t}}{\sqrt{M}} \quad (19)$$

$$s_{j,t}^2 = \alpha_0 + \alpha_1 r_{j,t-1}^2 + \beta_1 s_{j,t-1}^2, \quad (20)$$

where $j = 1, \dots, N, t = 1, \dots, T, i = 1, \dots, M, r_{j,t} = \sum_{i=1}^M r_{j,t,i}$ and $u_{g,t,i} \perp u_{l,t,i} \forall g \neq l$.

The parameters of the GARCH model, α_0, α_1 and β_1 , have been set to 0.022, 0.068 and 0.898 respectively for all series, which correspond to the estimated parameters of a GARCH(1,1) model reported by Andersen and Bollerslev (1998a) for the daily returns on the Deutschemark-US Dollar exchange rates from 1987 until 1992.

Notice that the impact of the values of the parameters α_0, α_1 and β_1 on the outcome of the test is small as each return series $r_{j,t,i}$ is divided by $\hat{s}_{j,t,i}$.

To simulate a realistic periodic factor we consider four cos and four sin terms depending only on the time of the day, i.e.

$$\log f_{j,t,i}^* = \sum_{l=1}^4 \gamma_{j,l} \cos\left(\frac{i2\pi l}{M}\right) + \sum_{l=1}^4 \gamma_{j,4+l} \sin\left(\frac{i2\pi l}{M}\right) \quad (21)$$

or more compactly in matrix form

$$\log \mathbf{f}^* = \mathbf{x}\Gamma', \quad (22)$$

i.e. there are $m^* = 8$ variables in \mathbf{x} and the constant ω is set to 0. $f_{j,t,i}$ is recovered from $\log f_{j,t,i}^*$ using (10).




Figure 1: Simulated periodicity

With respect to the commonalities in the periodicity, three cases are investigated, i.e. the presence of one, two and three factors. This means that the DGP_i considered in this simulation satisfies the null hypothesis that $rank(\mathbf{\Gamma}) = i, i \in \{1, 2, 3\}$.

The coefficients chosen for the decomposition of $\mathbf{\Gamma} = \boldsymbol{\alpha}\boldsymbol{\beta}'$ are reported here below for DGP_3 (i.e. 3 factors case) for $N = 5$ variables:

$$\begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} -0.24422 & -0.49756 & -0.054171 & 0.073907 & -0.26098 & 0.32408 & -0.11591 & -0.21442 \\ -0.24422 & -0.40000 & -0.054171 & 0.073907 & -0.26098 & 0.32408 & -0.11591 & -0.21442 \\ -0.15000 & 0.40000 & -0.054171 & -0.073907 & -0.56098 & 0.32408 & -0.11591 & -0.21442 \end{pmatrix}$$

Only the first row of $\boldsymbol{\beta}'$ is taken for the one factor case (DGP_1) and the first two rows are considered in the two factor case (DGP_2). The three periodic components, denoted $f(1)$, $f(2)$ and $f(3)$ are plotted in Figure 1.

The parameters of the first factor (i.e. first row of $\boldsymbol{\beta}'$) correspond to the estimated parameters

of (9) estimated by OLS on 3 years of 5-minute returns of the EUR-USD exchange rate and thus implies a realistic intraday periodic pattern in volatility. The second (resp. third) factor implies an arbitrary minor (resp. major) modification of the intraday periodicity.

The loadings on the other hand will depend on the number of series. In DGP_3 the first factor enters with coefficients equal to one for the first $\lfloor (N + 1)/3 \rfloor$ elements only. The second factor enters in the second set of variables of size also $\lfloor (N + 1)/3 \rfloor$. The third factor only influences the $N - 2\lfloor (N + 1)/3 \rfloor$ remaining series. This is what we illustrate above for $N = 5$. For DGP_2 one takes the first $\lfloor (N + 1)/2 \rfloor$ loading coefficients equal to one and the remaining series equal to zero; the second factor enters only in the $N - \lfloor (N + 1)/2 \rfloor + 1$ variables. The $N \times 1$ vector of loadings is equal to one in DGP_1 .

To compute ξ_s , one has first to determine the number m of variables to include in \mathbf{x} , e.g. the number of cos and sin terms. Recall that the true value of m used in the DGP is $m^* = 8$. The same value has been used by Andersen and Bollerslev (1997, 1998b) in their empirical applications. Table 1 reports for the three information criteria the frequencies (over 1000 replications) with which minimization of the criterion over both the values of m and k leads to selecting respectively the true number of periodic elements m^* , the true number of factors k^* and the true pair (m^*, k^*) .⁶ To be clear we choose the pair (m, k) that minimises the information criterion.

It emerges from Table 1 that one cannot rely on information criteria to choose either k or the pair (m, k) because frequencies of determination of the true value(s) are not uniformly satisfactory across the DGPs considered. Indeed, information criteria perform very poorly in this case, except when the number of factors is very small. However, frequencies of determination of the true number of periodic components m^* reach 100% in all cases for the SC information criterion and thus one can safely rely on them to determine m .

Table 2 concerns the finite sample properties of the ξ_s test statistic for the null hypothesis that there exist at least $s \leq \min(N, m)$ linear combinations that annihilate k common periodic features. The value for m used when computing ξ_s is the one obtained in the pair (m, k) that minimises the SC information criterion because this strategy was found to deliver the correct value for m in 100% of the cases. Column $Prob(\xi_{s=s^*+1} > q_{v_{s,m,N}}^{(1-\alpha)})$, also labelled ‘Empirical power’, reports the rejection

⁶All estimations and simulations in this paper have been obtained by the authors using the Ox programming language (Doornik, 2009) and the G@RCH software (Laurent, 2009).

frequencies when the null hypothesis is not satisfied by the DGP, where $q_{df}^{(1-\alpha)}$ is the $(1 - \alpha)\%$ quantile of the χ^2 distribution with df degrees of freedom. We only report results for a 5% nominal size α but results for $\alpha = 1\%$ and 10% were qualitatively similar. For instance, the first element of this column corresponds to the case where $T = 100$, $N = 5$ and there is one factor ($k^* = 1$). Consequently, $s^* = \min(N, m^*) - k^* = 4$ because there are 4 linear combinations annihilating this common factor. In this case, the number reported in this column gives the frequency of rejection of the null assumption of absence of intraday periodicity in volatility (i.e., rejecting $s \leq 4$ in favour of $s = 5$) while there is one common intraday periodic factor. The empirical power of the test in this configuration is thus 100%.⁷

The next column, $Prob(\xi_{s=s^*} > q_{v_{s,m,N}}^{(1-\alpha)})$ corresponds to the empirical size at the 5% nominal level, i.e. the rejection frequency using the test statistic ξ_s under $H_0 : s = s^*(\equiv \min(N, m^*) - k^*)$ for $N = 5$ and $m^* = 8$. The first element of this column equals 4.9 suggesting that there is no evidence of size distortion.

The overall conclusion from this simulation study is that the test has good power properties and does not suffer from any significant size distortion. Hence we recommend to use SC for determining m and then to use ξ_s to determine s (or equivalently k).

4.2 Case 2: Time-varying intraday stochastic volatility and conditional normality

The assumption of constancy of the stochastic volatility during the day (Assumption 3) is questionable and a rejection of this assumption might affect the properties of the test.⁸ Indeed, our test is based on the assumption that ε in (12) and (18) is i.i.d. If $s_{j,t,i}$ is not constant during the day, ε might exhibit serial correlation.

We propose to explicitly take into account this autocorrelation by adding lagged values of $\mathbf{y}_{t,i}$ into $\mathbf{z}_{t,i}$ in (18). Hence, we first concentrate out the effect of lags by multivariate least-squares of \mathbf{y} and \mathbf{x} on a constant and \mathbf{z} , i.e. lagged values of \mathbf{y} . The analysis is then performed on the residuals

⁷Notice that rejection frequencies and simulated power function results are not size-adjusted and that power for other set-ups is not reported because they are almost always equal to 100%.

⁸We thank one referee for bringing this issue to our attention.

Table 1: Frequencies of correct determination of m^* , k^* and the true pair (m^*, k^*) using information criteria

		k^*	AIC			HQ			SC		
			$m = m^*$	$k = k^*$	$m = m^*$ $k = k^*$	$m = m^*$	$k = k^*$	$m = m^*$ $k = k^*$	$m = m^*$	$k = k^*$	$m = m^*$ $k = k^*$
$T = 100$											
$N = 5$	1		80.0	83.6	67.9	98.9	100	98.9	100	100	100
	2		88.0	86.5	77.1	100	99.7	99.7	100	21.5	21.5
	3		90.1	67.6	61.4	99.7	3.90	3.90	100	0.00	0.00
$N = 15$	1		75.9	84.1	67.5	99.4	100	99.4	100	100	100
	2		84.9	85.3	75.5	100	100	100	100	97.4	97.4
	3		86.0	81.1	70.4	99.9	1.10	1.10	100	0.00	0.00
$T = 250$											
$N = 5$	1		78.1	85.2	67.2	99.1	100	99.1	100	100	100
	2		87.7	89.2	78.4	99.9	100	99.9	100	100	100
	3		91.5	89.3	81.5	99.9	46.7	46.7	100	0.10	0.10
$N = 15$	1		78.7	83.9	69.5	99.5	100	99.5	100	100	100
	2		86.7	86.1	76.7	99.9	100	99.9	100	100	100
	3		89.8	88.4	81.0	99.9	67.1	67.1	100	0.00	0.00

Note: the true number of periodic elements $m^* = 8$ and $k^* \in \{1, 2, 3\}$.

Table 2: Empirical power and empirical size of the ξ_s statistic for a 5% nominal size

		Empirical power		Empirical size	
	k^*	s^*	$Prob(\xi_{s=s^*+1} > q_{v_{s,m,N}}^{(1-\alpha)})$	$Prob(\xi_{s=s^*} > q_{v_{s,m,N}}^{(1-\alpha)})$	
$T = 100$					
$N = 5$	1	4	100	4.90	
	2	3	100	5.40	
	3	2	51.5	2.60	
$N = 15$	1	7	100	5.40	
	2	6	100	4.80	
	3	5	67.3	2.40	
$T = 250$					
$N = 5$	1	4	100	5.40	
	2	3	100	3.50	
	3	2	95.3	4.40	
$N = 15$	1	7	100	5.00	
	2	6	100	5.60	
	3	5	99.8	4.50	

Note: the true number of periodic elements $m^* = 8$, $s^* = \min(N, m) - k^*$, $q_{df}^{(1-\alpha)}$ is the $(1 - \alpha)\%$ quantile of the χ^2 distribution with df degrees of freedom while $v_{s,m,N} = s \times \max(N, m) - s(\min(N, m) - s)$. Column ‘Empirical size’ (resp. ‘Empirical power’) correspond to the rejection frequencies when the null hypothesis is (resp. is not) satisfied by the DGP described in Case 1.

of these two multivariate regressions.⁹

To study the performance of this approach and the effects of neglecting the serial correlation in ε , we consider a second simulation design where the stochastic part of the volatility follows a weak GARCH(1, 1) model. To this aim we use an Euler discretization of the continuous time GARCH(1, 1) model proposed by Nelson (1990) with intraday periodicity.

More specifically, the new DGP consists of Equations (1)-(2), where

$$s_{j,t,i}^2 = \theta\sigma^2 \frac{1}{M} + s_{j,t,i-1}^2 \left(1 - \theta \frac{1}{M} + \sqrt{2\lambda\theta \frac{1}{M}} z_{j,t,i} \right), \quad (23)$$

where $z_{j,t,i}$ is i.i.d. $N(0,1)$ and independent of $u_{j,t,i}$ and by convention, $s_{j,t,0}^2 = s_{j,t-1,M}^2$.

This DGP is used to generate 5-minute returns characterised by intraday periodicity and time-varying stochastic volatility. The memory of the volatility process depends on the values of θ and λ while, for given values of the previous parameters, σ^2 controls essentially the level of the unconditional standard variance. As shown by Drost and Werker (1996), there is an exact one to one relationship between these three parameters and the discrete-time weak GARCH(1,1) parameters at the daily frequency, i.e. α_0 , α_1 and β_1 in (20):

$$\theta = -\log(\alpha_1 + \beta_1) \quad (24)$$

$$\sigma^2 = \alpha_0(1 - \alpha_1 - \beta_1)^{-1} \quad (25)$$

$$\lambda = \frac{2\log^2(\alpha_1 + \beta_1)}{\frac{[1 - (\alpha_1 + \beta_1)^2](1 - \beta)^2}{\alpha_1[1 - \beta_1(\alpha_1 + \beta_1)]} + 6\log(\alpha_1 + \beta_1) + 2\log^2(\alpha_1 + \beta_1) + 4(1 - \alpha_1 - \beta_1)} \quad (26)$$

To control for the degree of persistence of the stochastic volatility, we chose several values of θ and λ implying a weak GARCH(1, 1) satisfying the restriction $\alpha_1 + \beta_1 = 0.95$ at the daily frequency with $\alpha_1 = 0.05, 0.1, 0.15, 0.20, 0.25, 0.30, 0.35$ and 0.4 . The higher α_1 is, the less sustainable the assumption of constant stochastic volatility during the day is. Note that the most realistic values for α_1 in this setting are $\alpha_1 = 0.05$ or 0.1 .

For the sake of comparison we also reconsider the constant volatility DGP presented in the previous subsection. Results concerning the frequencies of selection of the right value for m using

⁹An adjustment of the eigenvalues for the presence of a MA component (see Tiao and Tsay, 1989) produces very high size distortions and hence is not recommended (results are not reported to save space).

the SC criterion are not reported here to save space but are in line with those reported in Section 4.1.

Figure 2 plots the size of the test statistic ξ_s , i.e. the rejection frequency using the test statistic ξ_s under $H_0 : s = s^*(\equiv \min(N, m^*) - k^*)$ for $N = 5$ and $m^* = 8$. This figure is divided into six panels corresponding to six different situations where the time dimension varies ($T = 100$ and 250) as well as the number of common factors in periodicity ($k^* = 1, 2$ and 3). We refer to Subsection 4.1 for a description of the common factors in periodicity.

The x-axis corresponds to the number of lags of the endogenous variable that we include in $\mathbf{z}_{t,i}$ to control for the potential presence of autocorrelation in the residuals. The number of lags varies from zero to 10.

It appears that for values of $\alpha_1 \leq 0.2$, the rejection frequencies of the tests are close to the nominal size of 5% even if the number of included lagged endogenous variables and T are small, in particular in presence of 1 or 2 factors. In the presence of 3 common factors the tests appear to over-reject, in particular when a few lagged endogenous variables are included as regressors. For $T = 250$, with high order lags of the endogenous variables, the tests are found to be undersized.

These findings indicate that when $s_{j,t,i}$ is not constant over the day, including several lags (around 5) of the endogenous variable would be sufficient to assure that the test of the number of common factors using the test statistic ξ_s will have the right size even when $T = 100$, but also certainly when T is as large as 250. Including few lagged endogenous variables results in an oversized test whereas going beyond 5 lags of the endogenous variable leads to an undersized test.

To conclude, the size distortions are very small for realistic DGPs. For heavily volatile but less frequently observed series, the correction we propose delivers accurate results. The determination of the optimal number of lags in our correction is however beyond the scope of this paper.

4.3 Case 3: Time-varying intraday stochastic volatility and additive jumps

Prices of financial assets sometimes exhibit large jumps that are not in accordance with the assumption of conditional normality in (1). It is thus more realistic to see intraday returns as realisations of a *Brownian SemiMartingale with Finite Activity Jumps* (BSMFAJ) diffusion process like for




Figure 2: Rejection frequency (i.e. empirical size) using the ξ_s test statistic in the presence of non-constant intraday stochastic volatility

instance in Barndorff-Nielsen and Shephard (2004), and Lee and Mykland (2008).¹⁰

In the last simulation analysis, we study the impact of these jumps on our test by replacing Equation (1) in the system (1)-(2)-(23) by

$$r_{j,t,i} = \sigma_{j,t,i} u_{j,t,i} + a_{j,t,i} \quad (27)$$

$$a_{j,t,i} = q_{j,t,i} \kappa_{j,t,i}, \quad (28)$$

where $u_{g,t,i} \perp u_{l,t,i} \forall g \neq l$ and the parameters in Equation (23) are obtained using formulas (24)-(26)

¹⁰A count process is defined to be of finite activity if the change in the count process over any interval of time is finite with probability one.

and imply a GARCH(1,1) with $\alpha_0 = 0.022$, $\alpha_1 = 0.068$ and $\beta_1 = 0.898$ at the daily frequency.

The additive jumps variable $a_{j,t,i}$ is a random variable that is zero for most of the observations. For the intervals in which jumps occur, $a_{j,t,i}$ is non-zero and can be seen as an additive outlier with respect to $\sigma_{j,t,i}$ $u_{j,t,i}$. More specifically, $q_{j,t,i}$ is a Poisson distributed random variable generating on average \bar{q} jump(s) per day for each series (with $q_{g,t,i} \perp q_{l,t,i} \forall g \neq l$). The jump size $\kappa_{j,t,i}$ is modeled as the product between a uniformly distributed random variable on $\sqrt{h/\bar{q}}([-2, -1] \cup [1, 2])$ and the total instantaneous volatility $\sigma_{j,t,i}$. The parameter h determines the magnitude of the jumps. Note that the lower the intensity of the jump process, the larger the jumps are. In the simulation the average number of jumps per day (\bar{q}) ranges from 1 to 5 while h is set to 0 (no jumps), 0.1, 0.5, 1, 2, 3, 4 and 5, respectively.

In presence of jumps, we follow Lee and Mykland (2008) and Boudt et al. (2010) and estimate $s_{j,t,i}$, when evaluating (12) or (18), as the square root of a normalized version of Barndorff-Nielsen and Shephard (2004)'s realized daily bipower variation, i.e.,

$$\hat{s}_{j,t,i} = \sqrt{\frac{1}{M-1} BV_{j,t}}, \quad (29)$$

$$\text{with } BV_{j,t} = \mu_1^{-2} \sum_{l=2}^M |r_{j,t,l}| |r_{j,t,l-1}|, \quad (30)$$

where $\mu_1 = \sqrt{2/\pi} \approx 0.79788$. Alternatively, one can for instance use the square root of a normalized version of the MinRV and MedRV estimators of Andersen et al. (2009).

Monte Carlo simulation results reported by Boudt et al. (2010) suggest that the log-transformation shrinks the outliers and makes the OLS estimator of model (9) less sensitive to jumps.

In Figure 3, rejection frequencies for testing the presence of one factor ($k = 1$ or $s = \min(N, m) - 1 = 4$) using ξ_s against the alternative that $k > 1$ are plotted against h for different values of the number of jumps per day. The true number of factors k^* equals 1 while $N = 5$, $m^* = 8$ and $T = 100$.¹¹ In absence of jumps, the empirical size equals 4.90% which corresponds to the value reported in column 'Empirical size' in Table 2. It appears from Figure 3 that the presence of jumps, which are not taken into account, leads to a slightly oversized test when h is small. When h is large the tests are slightly undersized. In the presence of fewer jumps, the oversize is larger when jumps

¹¹Results concerning the frequencies of selection of the right value for m using the SC criterion are not reported here to save space but are also in line with those reported in Section 4.1.

occur more frequently. The general conclusion of this simulation study is that the procedure is not heavily affected by the inclusion of jumps in the DGP.

Tests statistics using canonical correlations that are robust to the presence of jumps such as the method of Taskinen et al. (2006) which uses the fast reweighted minimum covariance determinant (MCD) of Rousseeuw and van Driessen (1999) to substitute for the estimated covariances in $\Sigma_{\mathbf{y}\mathbf{y}}^{-1}\Sigma_{\mathbf{y}\mathbf{x}}\Sigma_{\mathbf{x}\mathbf{x}}^{-1}\Sigma_{\mathbf{x}\mathbf{y}}$ in (16) are not appropriate in our setting. Indeed the MCD requires the conditional distribution of the data that are not contaminated by outliers (or jumps) to follow an elliptical distribution while in our case \mathbf{y} is log-normally distributed. This method was found to produce very high size distortions both in absence and presence of jumps. Results are not reported to save space.

5 Application

The data set was obtained from TickData and consists of transaction prices at the 5-minute sampling frequency for $N = 30$ large capitalization stocks from the NYSE, AMEX NASDAQ, covering the period from January 1, 2000 to December 31, 2008 (2239 trading days). A list of ticker symbols and company names is provided in Appendix A. The trading session runs from 9:30 EST until 16:00 EST (390 minutes). Because of the unusual trading activity at the beginning of each day, we start our intraday sampling at 9.35 am, 5 minutes after the market officially opens, such that $M = 77$.

5.1 Testing for common intraday periodicity

For the choice of variables driving the intraday periodicity in volatility, we follow Andersen and Bollerslev (1997) and include both a linear and a quadratic trend in \mathbf{x} as well as $p_j \cos$ and $p_j \sin$ terms such that Equation (9) can be rewritten as

$$\log |\bar{r}_{j,t,i}| = \omega_j + \delta_{j,1} \frac{i}{N_1} + \delta_{j,2} \frac{i^2}{N_2} + \sum_{l=1}^{p_j} \gamma_{j,l} \cos\left(\frac{i2\pi l}{M}\right) + \sum_{l=1}^{p_j} \gamma_{j,p_j+l} \sin\left(\frac{i2\pi l}{M}\right) + \varepsilon_{j,t,i}, \quad (31)$$

where $N_1 = (M + 1)/2$ and $N_2 = (2M^2 + 3M + 1)/6$ are normalizing constants and p_j is the number of cos and sin terms (determined using the SC criterion) for series j . Note that in this case $m_j = p_j \times 2 + 2$ while the multivariate version of (31) imposes $m_j = m \equiv p \times 2 + 2 \forall j = 1, \dots, N$. Because of the presence of jumps in the data, $s_{j,t,i}$ is estimated by (29) for each series.

The outcome of the test is reported in Table 3. The test is applied to three windows of three consecutive years (respectively 742, 750 and 747 days for the periods 2000-2002, 2003-2005 and 2006-2008).

The number m of variables included in \mathbf{x} is reported in column 2 and equals 6 for each sub-period. Recall that \mathbf{x} contains a linear and a quadratic trend as well as a number of cos and sin terms determined by minimising the SC criterion (15).¹² The number of common factors detected at the 5% critical level is reported in the column labelled ' k ' while the p-values of the null hypothesis that there are at least $s = (\max(N, m) - k)$ linear combinations that annihilate k common periodic features are reported in columns ' $\xi_{s=l}$ ' (for $l = 25, 26, \dots, 30$).

It emerges from the reading of this table that, out of the 30 US stocks, only three factors are driving the intraday periodicity in volatility. The common periodicity series $\hat{F} = \mathbf{x}\hat{\beta}$ extracted from the data are plotted in Figure 4, where the factors are ranked in terms of their informativeness (corresponding to the k largest to the smallest eigenvalues, see (17)).

The estimated factors in Figure 4 exhibit similar behaviour across different sampling periods. The factor represented by the solid line corresponds to the typical U-shaped pattern observed in return volatility over the trading day (see e.g. Andersen and Bollerslev, 1997), i.e. volatility is high at the open and close of trading and low in the middle of the day. The factor represented by the broken line is almost constant over a large part of the day and then increases at the end. Finally, the factor given by the dotted line fluctuates most during the trading day.

Factor 1 mimics the behavior of the well-known intraday volatility shape and the factors 2 and 3 capture more erratic fluctuations. Moreover, the shape of the factors changes somewhat over time. The factors should not necessarily have a clear economic interpretation. The factors are linear combinations of the sinusoids and of polynomials of the time of the day. The factors 2 and 3 multiplied by their respective loadings for a given asset could be interpreted as deviations of that asset from the well-known intraday volatility shape given by the first factor, also multiplied by the corresponding loading for the asset considered. Typically, we observe a slowly decreasing intraday trend (factor 2) and a sinusoidal factor (factor 3). Depending on the value and sign of the loadings, factor 2 could either strengthen or weaken the typical intraday volatility shape of an asset whereas

¹²Results reported in Table 3 concern the case where no lagged values of $\mathbf{y}_{t,i}$ are included into $\mathbf{z}_{t,i}$ but similar results have been obtained with 1 or 2 lags.

factor 3 would add a new feature of intraday fluctuations to this typical shape. When the factors are multiplied by their loadings, we get components that will exhibit asset-specific idiosyncracies that could be more easily interpreted.

Table 3 suggests that there is some strong evidence of commonalities in the intraday periodicity in volatility. In the next two subsections, we investigate whether imposing these commonalities can be exploited to better forecast either future values of $\log f_{j,t,i}^*$ or the conditional variance of 5-minute returns.

5.2 Predicting log absolute standardised returns

As explained in the introduction, adequately imposing commonalities in a multivariate model can be exploited to improve parameter efficiency and hopefully for some series also improve forecasts accuracy. The first forecasting exercise considers the problem of predicting the values of $\log |\bar{r}_{j,t,i}|$ for the period 2003-2005 (resp. 2006-2008) on the basis of the values of $\log f_{j,t,i}^*$ obtained for the period 2000-2002 (resp. 2003-2005).

The first model is the unrestricted model where Equation (31) is estimated by OLS series by series. Note that in this case, m_j is chosen by minimising the Schwarz criterion for univariate linear regression models and thus can vary from one series to another (but for each series we include linear and quadratic terms). The second model is the multivariate extension of (31) that imposes the presence of the three detected common factors.

Predicted values $\log \hat{f}_{j,t,i}^*$ are compared to realisations (i.e. $\log |\bar{r}_{j,t,i}|$) for each model and for each series separately by means of the following mean squared (prediction) error (MSE) criterion $\text{MSE}_j = \frac{1}{MT^*} \sum_{t=1}^{T^*} \sum_{i=1}^M l_{j,t,i}$, where $l_{j,t,i} \equiv e_{j,t,i}^2 = (\log |\bar{r}_{j,t,i}| - \log \hat{f}_{j,t,i}^*)^2$ and T^* is the number of days in the forecasting period (about 750 for each period). To test the null hypothesis of equal prediction MSE of the factor and unrestricted models for series j we rely on the Diebold and Mariano (1995) test (denoted as DM hereafter). We also test the null hypothesis of equal MSE across the 30 series. To do so and in order take into account the presence of potential contemporaneous correlation between the prediction errors, we use the following criterion $\text{MSE}_{All} = \frac{1}{MT^*} \sum_{t=1}^{T^*} \sum_{i=1}^M l_{t,i}$, where $l_{t,i} \equiv \sum_{j=1}^N l_{j,t,i}$.

It is well known that the DM test should be applied with care in situations where the competing models are nested, which is the case here. Giacomini and White (2006) have shown that when the

estimation window size is bounded (e.g., for the fixed and rolling schemes) the DM test is still valid. Our setting corresponds to the fixed scheme because the models are estimated on the period 2000-2002 (resp. period 2003-2005) and these values are used to predict $\log |\bar{r}_{j,t,i}|$ on the period 2003-2005 (resp. period 2006-2008).

The outcome of the DM test is summarised in Table 4. Column labelled ‘DM’ contains a plus when the DM statistic is higher than the 5% critical value, suggesting that the non-restricted model significantly under-performs (and possibly a minus when it significantly over-perform, which never happens). This column is left empty when the two models are not statistically different. The first row of this table, labelled ‘All’ corresponds to the null hypothesis of equal MSE across the 30 series while the other 30 rows correspond to the tests for individual assets.

Interestingly, the results suggest that imposing the detected commonalities helps to better predict $\log |\bar{r}_{j,t,i}|$ in most cases and never leads to a deterioration of the forecast accuracy.

We have also implemented the Superior Predictive Ability (SPA) test proposed by Hansen (2005) in addition to the DM test. The advantage of the SPA test over the DM test when applied to pairwise comparisons is that the former approximates the finite sample distribution of the test via block-bootstrap while the latter relies on asymptotic critical values. The SPA naturally accounts for the potential presence of correlation within the blocks. The analysis leads to results very similar to those of the DM test and similar critical values, confirming our findings reported above. For that reason, we have omitted reporting detailed results for the SPA test.

5.3 Forecasting the intraday conditional variance

The finding that imposing common factor restrictions improves forecasts of $\log |\bar{r}_{j,t,i}|$ is encouraging but does not necessarily mean that this strategy will also lead to better forecasts of the intraday conditional variances $\sigma_{j,t,i}^2$. To investigate that issue we consider now four different modelling strategies to obtain one-step-ahead forecasts of the conditional variance of 5-minute returns. For each model we implement two versions, one imposing and one not imposing the three detected common factors in the intraday periodicity in volatility. This leads to a total of 8 competing models. Like in the first forecasting exercise we rely on the 30 US stocks and divide the period into 3 sub-periods of three years.

Table 3: Test of common periodic common features applied to 30 US stocks

Period	m	$\xi_{s=25}$	$\xi_{s=26}$	$\xi_{s=27}$	$\xi_{s=28}$	$\xi_{s=29}$	$\xi_{s=30}$	k
2000-2002	6	0.787	0.370	0.102	0.000	0.000	0.000	3
2003-2005	6	0.795	0.399	0.092	0.000	0.000	0.000	3
2006-2008	6	0.488	0.531	0.071	0.003	0.000	0.000	3

Column m corresponds number of variables included in \mathbf{x} , i.e. linear and quadratic trends as well as $(m-2)/2 \cos$ and $(m-2)/2 \sin$ terms as defined in (21) selected using the SC criterion. Columns $\xi_{s=l}$ (for $l = 25, 26, \dots, 30$) correspond to the p-value of the null hypothesis that there are at least $s = l$ linear combinations that annihilate k common periodic features. The number of periodic factors k obtained by ξ_s is reported in the last column.

The first period (2000-2002) is used to estimate $f_{j,t,i}$ either by estimating (9) equation by equation or with the reduced rank version of the multivariate model (12).¹³ These values are used as forecasts of the intraday periodicity for the second period (2003-2005). Similarly, the intraday periodicity of the third period (2006-2008) is forecasted using the estimates of the second period.

For each model, one-step-ahead forecasts of the 5-minute conditional variance of $r_{j,t,i}$ are obtained as $E(\sigma_{j,t,i+1}^2 | \Omega_{t,i}) = E(s_{j,t,i+1}^2 | \Omega_{t,i}) E(f_{j,t,i+1}^2 | \Omega_{t,i})$, where $\Omega_{t,i}$ is the information set available at the beginning of the i th interval of day t and where by convention, $\sigma_{j,t,M+1}^2 = \sigma_{j,t+1,1}^2$, $s_{j,t,M+1}^2 = s_{j,t+1,1}^2$ and $f_{j,t,M+1}^2 = f_{j,t+1,1}^2$. The models differ in the way they forecasts $s_{j,t,i+1}^2$ and $f_{j,t,i+1}^2$. We present now the four modelling strategies to forecast $s_{j,t,i}^2$.

Model 1 (Daily GARCH) corresponds exactly to Equations (1)-(2)-(19)-(20). The stochastic volatility is assumed to follow a GARCH(1, 1) at the daily frequency and the intraday variations of the conditional variance are entirely due to the deterministic periodic component $f_{j,t,i}$. The first forecast of $s_{j,t,i+1}^2$ is obtained by estimating a GARCH(1, 1) model by QML for the period 2000-2002. The parameters are kept constant during 50 days, whereon the GARCH model is re-estimated on a rolling window (i.e. keeping the number of observations fixed).

Model 2 (Intradaily GARCH) is a GARCH(1, 1) estimated on filtered intraday returns $r_{j,t,i}/f_{j,t,i}$,

¹³Recall that $f_{j,t,i}$ is recovered from $\log f_{j,t,i}^*$ using (10).

Table 4: Out-of-sample forecast analysis

	Period 2003-2005									Period 2006-2008								
	DM	OLS ₁	OLS ₂	OLS ₃	OLS ₄	CF ₁	CF ₂	CF ₃	CF ₄	DM	OLS ₁	OLS ₂	OLS ₃	OLS ₄	CF ₁	CF ₂	CF ₃	CF ₄
All	+							*	*	+	*	*			*	*	*	*
ABT		*		*	*	*		*	*	+	*	*	*	*	*	*	*	*
BAC				*				*	*	+	*	*	*	*	*	*	*	*
BMJ		*	*	*	*	*	*	*	*		*	*	*	*	*	*	*	*
C	+				*						*	*	*		*	*	*	*
CL			*	*	*		*	*	*	+	*	*	*	*	*	*	*	*
CSCO								*	*	+			*	*		*	*	*
DELL				*	*			*	*	+			*	*		*	*	*
DIS	+	*		*		*		*	*	+	*	*	*	*	*	*	*	*
EK	+		*	*	*	*	*	*	*	+			*	*		*	*	*
EXC	+							*	*		*		*	*		*	*	*
F	+				*				*		*	*	*	*	*	*	*	*
GE		*		*	*	*		*	*	+	*		*	*	*	*	*	*
GM		*	*	*	*	*	*	*	*		*	*	*	*	*	*	*	*
HD								*	*	+	*	*	*	*	*	*	*	*
INTC								*	*		*	*	*	*	*	*	*	*
JPM				*	*			*	*	+	*	*	*	*	*	*	*	*
KO				*	*	*	*	*	*	+	*	*	*	*	*	*	*	*
LLY		*		*	*	*		*	*	+	*	*	*	*	*	*	*	*
MCD				*	*			*	*	+	*	*	*	*	*	*	*	*
MRK	+	*	*	*	*	*	*	*	*	+	*	*	*	*	*	*	*	*
PEP	+							*	*	+	*	*	*	*	*	*	*	*
PFE	+	*	*	*	*	*	*	*	*		*	*	*	*	*	*	*	*
PG	+							*	*		*	*	*	*	*	*	*	*
T	+							*	*	+	*	*	*	*	*	*	*	*
TWX		*	*	*	*	*	*	*	*	+	*	*	*	*	*	*	*	*
VZ								*	*		*	*	*	*	*	*	*	*
WFC		*			*	*		*	*	+	*	*	*	*	*	*	*	*
WYE	+	*	*	*	*	*	*	*	*	+	*	*	*	*	*	*	*	*
XOM								*	*		*		*	*	*	*	*	*
XRX	+							*	*		*	*	*	*	*	*	*	*

Note: Column ‘DM’ corresponds to the Diebold and Mariano (1995) test of equal predictive ability of the intraday periodicity as discussed in Subsection 5.2. A + (resp. -) means that the unconstrained univariate model underperforms (resp. underperforms) compared to model imposing the commonalities. Columns ‘OLS_l’ and ‘CF_l’ ($l \in \{1, 2, 3, 4\}$) concerns the outcome of the MCS test of Hansen et al. (2009) for superior predictive ability of the conditional variance. A * means that the model belongs to the set of superior models at the 5% critical level. Columns labelled ‘OLS_l’ refer to the unconstrained univariate model estimated by OLS, equation by equation while columns labelled ‘CF_l’ concern the models imposing the common factors. Models $l \in \{1, 2, 3, 4\}$ correspond respectively to the GARCH(1,1) on daily data, GARCH(1,1) on intradaily data, HAR-RV model and HAR-RV-J model.

i.e.

$$r_{j,t,i}/f_{j,t,i} = s_{j,t,i} u_{j,t,i} \quad (32)$$

$$s_{j,t,i} = \alpha_{j,0} + \alpha_{j,1}r_{j,t,i-1}^2 + \beta_{j,1}\sigma_{j,t,i-1}^2. \quad (33)$$

The model is estimated by QML on a rolling window of 100 days (i.e. 7700 observations). The parameters are also kept constant during 50 days.

Model 3 (Daily HAR-RV) replaces Equation (19) in Model 1 by a forecast of the daily conditional variance $s_{j,t}^2$ based on the Heterogenous Autoregressive Realized Volatility model (HAR-RV) model of Corsi (2009):

$$s_{j,t}^2 = E(RV_{j,t}|\Omega_{j,t-1}) \quad (34)$$

$$RV_{j,t} = \alpha_{j,0} + \alpha_{j,1}RV_{j,t-1} + \alpha_{j,2}RV(5)_{j,t-1} + \alpha_{j,3}RV(22)_{j,t-1} + e_{j,t}, \quad (35)$$

where $RV_{j,t}$ is given in (6) and by convention, $X(m)_{j,t-1} = \frac{1}{m} \sum_{i=1}^m X_{j,t-i}$. The model is an additive cascade model of volatility components defined over different time periods, one day, one week and one month. Corsi (2009) has shown that this model delivers remarkably accurate forecasts on real data. The first forecast of $s_{j,t}^2$ is obtained by estimating Equation (35) by OLS on the period 2000-2002. The model is then re-estimated every 50 days on a rolling window.

Model 4 (Daily HAR-RV-J) is an extension of Model 3 where the HAR-RV specification is extended in order to take into account the effect of past jumps. We adopt the HAR-RV-J specification of Andersen et al. (2007) (where J stands for jumps), i.e.

$$\begin{aligned} RV_{j,t} &= \alpha_{j,0} + \alpha_{j,1}RV_{j,t-1} + \alpha_{j,2}RV(5)_{j,t-1} + \alpha_{j,3}RV(22)_{j,t-1} \\ &+ \gamma_{j,1}J_{j,t-1} + \gamma_{j,2}J(5)_{j,t-1} + \gamma_{j,3}J(22)_{j,t-1} + e_{j,t}, \end{aligned} \quad (36)$$

where $J_{j,t} = I_{j,t}(RV_{j,t} - BV_{j,t})$, $BV_{j,t}$ is given in (30), $I_{j,t} \equiv I[Z_{j,t} > \Phi_{0.999}]$, $Z_{j,t} = \frac{M^2[RV_{j,t}] - BV_{j,t}RV_{j,t}^{-1}}{[(\mu_1^{-4} + 2\mu_1^{-2} - 5)\max\{1, TQ_{j,t}(M)BV_{j,t}^{-2}\}]^{1/2}}$, $TQ_{j,t}$ is the tri-power quarticity,¹⁴ a robust to jumps estimator of the integrated quarticity and $\Phi_{0.999}$ is the 99.9% quantile of the standard normal distribution.

To measure the out-of-sample forecasting performance of the competing models, forecasts have to be compared to ex-post realisations as they become available. This implies choosing both a loss

¹⁴ $TQ_{j,t} \equiv M\mu_{4/3}^{-3} \sum_{l=3}^M |r_{j,t,l}|^{4/3} |r_{j,t,l-1}|^{4/3} |r_{j,t,l-2}|^{4/3}$, where $\mu_{4/3} \equiv 2^{2/3}\Gamma(7/6)\Gamma(1/2)^{-1}$.

function and a proxy for the true conditional variance (which is unobservable even ex-post). The question arises on which volatility proxy and which loss function to use. Hansen and Lunde (2006) provide conditions, for both the loss function and the volatility proxy, under which the ranking of models based on the proxy is consistent for the true ranking (i.e. the one implied by the true but unobserved variance). Starting from this result, Patton (2011) derives necessary and sufficient conditions on the functional form of the loss function for the ranking to be robust to the presence of noise in the proxy, all of which being satisfied by the MSE loss function. This is the reason why we rely on this loss function. About the volatility proxy, we use the 5-minute squared return $r_{j,t,i}^2$ which is known to be an unbiased (but noisy) proxy of $\sigma_{j,t,i}^2$.

Furthermore, instead of just ranking the models in function of their MSE, we use the model confidence set (MCS) approach of Hansen et al. (2009) to compare the forecasts. Given a universe of model based forecasts, the MCS allows us to identify the subset of models that are equivalent in terms of forecasting ability, but outperform all the other competing models. We set the confidence level for the MCS to $\alpha = 5\%$ and used 1000 bootstrap resamples (with block length of 6 observations) to obtain the distribution under the null of equal forecasting performance.¹⁵ The MCS test is summarised in Appendix B.

Table 4 indicates by a * which models belong to the set of superior forecasting models according to the MCS test for the two forecasting periods. Like in the previous section, MSEs are computed for each series separately as the average of the squared forecasting errors $e_{j,t,i}^2$ but also for the 30 series jointly (row labelled ‘All’) as the average of $\sum_{j=1}^N e_{j,t,i}^2$ over the total number of intraday observations in the forecasting period.

Columns labelled ‘OLS_l’ and ‘CF_l’ correspond respectively to the forecasts where $f_{j,t,i}$ is estimated equation by equation (by OLS) or with the multivariate model (12). Sub-strict l ($l \in \{1, 2, 3, 4\}$) refers to the modelling strategies used to forecast $s_{j,t,i}^2$, i.e. respectively the GARCH(1,1) on daily data, GARCH(1,1) on intradaily data, HAR-RV model and HAR-RV-J model.

Results suggest that for the period 2006-2008, models are hardly distinguishable but forecasts based on the reduced rank version of the multivariate model (12) always belong to the set of superior models. This result is in line with the one of Laurent et al. (2010) who also find on similar

¹⁵Implementation of this test has been done using the Ox software package MULCOM of Hansen and Lunde (2007). Note that we got similar results with different block lengths for the block bootstrap and a higher number of resamples.

series that (multivariate) GARCH models (from simple to sophisticated ones) are indistinguishable during extremely volatile periods (e.g. over the 2007-2008 financial crisis). This is essentially due to the fact that large jumps are not forecastable by these models, leading to extremely large forecasting errors (and thus MSEs) for all models. Notice that other criteria that down-weight the effect of these jumps, like the mean absolute deviation (MAD), do not satisfy the conditions stated in Hansen and Lunde (2006) and Patton (2011) to ensure the ranking of models to be robust to the presence of noise in the proxy.

Interestingly, during the more quiet period (2003-2005), forecasts based on the reduced rank version of model (12) clearly dominate the MCS. Indeed, they belong to the MCS in 29 out of 30 cases when considering the individual MSEs. More specifically, the MCS test usually points two models: the HAR-RV of Corsi (2009) and the HAR-RV-J model of Andersen et al. (2007) to forecast $s_{j,t,i}^2$, coupled with the reduced rank version of model (12) to forecast $f_{j,t,i}$. These two models correspond also to the MCS for the join test (row labelled ‘All’). The general message is that for this period models imposing the detected commonalities in the periodicity and using Assumption 3 to forecast $s_{j,t,i}^2$ using a simple linear regression model on the daily realized volatility outperform in most cases those not imposing these commonalities as well as GARCH models fitted on daily and even intradaily data.

6 Conclusion

Using a simple canonical correlation test as well as information criteria we investigate the presence of commonalities in the intraday periodic components. Given the nature of the data and the number of series considered the number of common factors is obtained. A likelihood ratio statistic based on testing that the first set of eigenvalues obtained in a canonical correlation framework works remarkably well. Information criteria determine very accurately the number of periodic elements to be added in the system (by SC) but tend to heavily underestimate the number of factors.

The presence of serial correlation in the disturbances of the model affects the performance of the test based on canonical correlations. However, including lagged values of the endogenous variables can lead to a correctly sized test. The test appears to be fairly robust to the presence of jumps in the DGP, which are not taken into account by the model.

We have illustrated that 30 US asset returns are driven by only three factors in periodicity although in that case only a few periodic elements are needed. Anyway, the reduction in the number of parameters we have when we impose that factor structure can lead to a gain in efficiency and to more accurate forecasts of both the intraday periodicity and the intraday conditional variance of most assets considered in the application. Our framework is flexible enough to include additional exogenous or deterministic variables (e.g. over night returns) sharing or not co-movements with the periodicity.

Appendix A: Stocks used in the empirical application

Symbol	Issue name	Symbol	Issue name
ABT	ABBOTT LABORATORIES	JPM	JP MORGAN CHASE
BAC	BANK OF AMERICA	KO	COCA COLA CO
BMY	BRISTOL MYERS SQ	LLY	ELI LILLY & CO
C	CITIGROUP	MCD	MCDONALDS CORP
CL	COLGATE-PALMOLIVE CO	MRK	MERCK & CO
CSCO	CISCO SYSTEMS	PEP	PEPSICO INC
DELL	DELL INC	PFE	PFIZER INC
DIS	WALT DISNEY CO	PG	PROCTER & GAMBLE
EK	EASTMAN KODAK	T	AT&T CORP
EXC	EXELON CORP	TWX	TIME WARNER
F	FORD MOTOR CO	VZ	VERIZON COMMS
GE	GENERAL ELEC	WFC	WELLS FARGO & CO
GM	GENERAL MOTORS	WYE	WEYERHAEUSER CO
HD	HOME DEPOT INC	XOM	EXXON MOBIL
INTC	INTEL CORP	XRX	XEROX CORP

Appendix B: Model Confidence Set

The MCS approach, introduced by Hansen et al. (2009), is a testing procedure for superior predictive ability based on the reality check for data snooping of White (2000) and the superior predictive ability (SPA) test of Hansen (2005). The test allows to identify a subset of models equivalent in terms of predictive ability, that are superior to the other models. The advantage of the MCS procedure is that it does not require a benchmark model to be specified which is useful for applications without an objective benchmark.

Let us denote \mathfrak{S}^0 the initial set of models for which we compute one-step ahead conditional variance forecasts, denoted by $\hat{\sigma}_{m,T+1}^2, \dots, \hat{\sigma}_{m,T+T^*+1}^2$, $l = 1, \dots, l^*$ where T^* defines the forecasting sample length. For ease of exposition we only use one time index in this section to capture both the daily and intradaily time intervals. The MCS procedure allows to select a subset of models, \bar{l} , which are superior, in terms of predictive ability, with respect to all the other models in \mathfrak{S}^0 . To do this, we need an equivalence test, an elimination rule and an updating algorithm. The starting hypothesis is that all models in \mathfrak{S}^0 have equal forecasting performance as measured by a loss function $L_{l,t} = L(\sigma_t^2, \hat{\sigma}_{l,t}^2)$ that compares the true but unobserved volatility σ_t^2 and the forecasts of model l , i.e. $\hat{\sigma}_{l,t}^2$. If the null of equal predictive ability is rejected, then the elimination rule removes the worst performing model. This process is repeated until the non-rejection of the null occurs (at a given confidence level). The set of surviving models is the MCS. More formally, we start by defining the relative performance at time t as $d_{ij,t} = L_{i,t} - L_{j,t}$ for all $i, j = 1, \dots, l^*$. Under the assumption that $d_{ij,t}$ is stationary, the null hypothesis takes the form $H_{0,\mathfrak{S}^0} : E(d_{ij,t}) = 0$, $\forall i, j \in \mathfrak{S}^0$ and the test statistic

$$T_D = \frac{1}{l^*} \sum_{i \in \mathfrak{S}^0} t_i^2, \quad (37)$$

where $t_i = \frac{\sqrt{T^*} \bar{d}_i}{\omega_i}$ and $\bar{d}_i = \frac{1}{l^*} \sum_{j \in \mathfrak{S}^0} \bar{d}_{ij}$ is the contrast of model i 's sample loss with respect to the average across all models and $\bar{d}_{ij} = \frac{1}{T^*} \sum_{t=1}^{T^*} d_{ij,t}$ is the sample loss difference between model i and j . Hence the name of the statistic T_D where D stands for deviation (from the average loss across models). The variances $\omega_i^2 = \lim_{T^* \rightarrow \infty} \text{Var}(\sqrt{T^*} \bar{d}_i)$ can be estimated by $\hat{\omega}_i^2$ using a bootstrap scheme, e.g., block bootstrap to account for serial dependence in the loss, and the distribution of T_D derived. If the null hypothesis is rejected, then we use as elimination rule $\text{argmax}_i t_i$ to exclude the

weakest model from the set. The elimination rule excludes the model with the largest standardised excess loss relative to the average across models, that is $\bar{d}_i = \bar{L}_i - \bar{L} = \bar{L}_i - \frac{1}{\bar{I}^*} \sum_{j \in \mathfrak{S}^0} \bar{L}_j = \frac{1}{\bar{I}^*} \sum_{j \in \mathfrak{S}^0} (\bar{L}_i - \bar{L}_j)$. The MCS p-value is equal to $p_i = \max_{r \leq i} p(r)$ where $p(r)$ is the p-value of the test under the null H_{0, \mathfrak{S}^r} where r is the number of surviving models at step i of the iteration process. After the necessary iterations, the set of superior models is given by $\{i \in \mathfrak{S}_0 : E(d_{i,j,t}) \leq 0 \forall i \neq j \in \mathfrak{S}^0\}$.

References

- Andersen, T. G. and T. Bollerslev (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance* 4, 115–158.
- Andersen, T. G. and T. Bollerslev (1998a). Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885–905.
- Andersen, T. G. and T. Bollerslev (1998b). Deutsche Mark-Dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance* 53, 219–265.
- Andersen, T. G., T. Bollerslev, and F. Diebold (2007). Roughing it up: including jump components in the measurement, modelling and forecasting of return volatility. *The Review of Economics and Statistics* 89, 701–720.
- Andersen, T. G., D. Dobrev, and E. Schaumburg (2009). Jump-robust volatility estimation using nearest neighbor truncation. *NBER Working Paper No. 15533*.
- Athanasopoulos, G., O. Guillen, J. Issler, and F. Vahid (2009). Model selection, estimation and forecasting in var models with short-run and long-run restrictions. Monash Econometrics and Business Statistics Working Papers 2/09.
- Baillie, R. and T. Bollerslev (1991). Intra-day and inter-market volatility in foreign exchange rates. *Review of Economic Studies* 58, 565–585.

- Barndorff-Nielsen, O. E. and N. Shephard (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2, 1–37.
- Boudt, K., C. Croux, and S. Laurent (2010). Robust estimation of the periodicity in intraday volatility and intraday jump detection. Forthcoming in *Journal of Empirical Finance*.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7, 174–196.
- Diebold, F. and R. Mariano (1995). Comparing predictive accuracy. *Journal of the Business & Economics Statistics* 13, 253–263.
- Dionne, G., P. Duchesne, and M. Pacurar (2009). Intraday value at risk (ivar) using tick-by-tick data with application to the toronto stock exchange. *Journal of Empirical Finance* 16, 777–792.
- Doornik, J. A. (2009). *An Object-Oriented Matrix Programming: Ox 6*. London: Timberlake Consultants Ltd. See <http://www.doornik.com>.
- Drost, C. and B. Werker (1996). Closing the garch gap: Continuous time garch modelling. *Journal of Econometrics* 74, 31–57.
- Engle, R. F. and S. Hylleberg (1996). Common seasonal features : Global unemployment. *Oxford Bulletin of Economics and Statistics* 58, 615–630.
- Engle, R. F. and J. Marcucci (2006). A long-run pure variance common features model for the common volatilities of the dow jones. *Journal of Econometrics* 132, 7–42.
- Engle, R. F. and R. Susmel (1993). Common volatility in international equity markets. *Journal of Business and Economic Statistics* 11, 167–175.
- Giacomini, R. and H. White (2006). Tests of conditional predictive ability. *Econometrica* 74, 1545–1578.
- Giot, P. (2005). Market risk models for intraday data. *European Journal of Finance* 11, 309–324.
- Goodhart, C. A. and M. O’Hara (1997). High frequency data in financial markets: Issues and applications. *Journal of Empirical Finance* 4, 73–114.

- Hansen, P. (2005). A test for superior predictive ability. *Journal of Business and Economic Statistics* 23, 365–380.
- Hansen, P. and A. Lunde (2006). Consistent ranking of volatility models. *Journal of Econometrics* 131, 97–121.
- Hansen, P. and A. Lunde (2007). Mulcom 1.00, econometric toolkit for multiple comparisons. <http://www.hha.dk/~alunde/MULCOM/MULCOM.HTM>.
- Hansen, P., A. Lunde, and J. Nason (2009). Model confidence sets. Forthcoming in *Econometrica*.
- Hecq, A., S. Laurent, and F. Palm (2010). On the univariate representation of multivariate volatility models with common factors. Mimeo, Maastricht University.
- Laurent, S. (2009). *G@RCH 6, Estimating and Forecasting ARCH Models*. London: Timberlake Consultants Ltd. See <http://www.garch.org>.
- Laurent, S., J. Rombouts, and F. Violante (2010). On the forecasting accuracy of multivariate garch models. Mimeo, Maastricht University.
- Lee, S. S. and P. A. Mykland (2008). Jumps in financial markets: a new nonparametric test and jump dynamics. *Review of Financial studies* 21, 2535–2563.
- Nelson, D. (1990). ARCH models as a diffusion approximation. *Journal of Econometrics* 45, 7–38.
- Patton, A. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal Econometrics* 160(1), 246–256.
- Rousseeuw, P. and K. van Driessen (1999). A fast algorithm for the minimum covariance determinant estimator. *Technometrics* 41, 212–223.
- Taskinen, S., C. Croux, A. Kankainen, E. Ollila, and H. Oja (2006). Canonical analysis based on scatter matrices. *Journal of Multivariate Analysis* 92 (2), 359–384.
- Taylor, S. J. and X. Xu (1997). The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance* 4, 317–340.

- Tiao, G. C. and R. S. Tsay (1989). Model specification in multivariate time series (with comments). *Journal of Royal Statistical Society, Series B* 51, 157–213.
- Visser, M. P. (2010). Garch parameter estimation using high-frequency data. *Journal of Financial Econometrics* 8, 1–36.
- White, H. (2000). Reality check for data snooping. *Econometrica* 68, 1097–1126.




Figure 3: Rejection frequency (i.e. empirical size) of the ξ_s test statistic in presence of 1 common factor ($k^* = 1$), non-constant intraday stochastic volatility and jumps. The magnitude of the jumps is controlled by h/\bar{q} where \bar{q} is the expected number of jumps per day (which varies between 0 and 5) and $h = 0.1, 0.5, 1, \dots, 5$.




Figure 4: Estimated intraday periodicity factors $\hat{F} = \mathbf{x}\hat{\beta}$