SCHOOL OF ECONOMICS AND MANAGEMENT
FACULTY OF SOCIAL SCIENCES
AARHUS UNIVERSITY

CREATES
Center for Research in Econometric Analysis of Time Series

CREATES Research Paper 2010-45

The Role of Realized Ex-post Covariance Measures and Dynamic Model Choice on the Quality of Covariance Forecasts

Rasmus Tangsgaard Varneskov and Valeri Voev

School of Economics and Management
Aarhus University
Bartholins Allé 10, Building 1322, DK-8000 Aarhus C
Denmark

The Role of Realized Ex-post Covariance Measures and Dynamic Model Choice on the Quality of Covariance Forecasts

Rasmus Varneskov*

Aarhus University
CREATES

Valeri Voev†

Aarhus University
CREATES

August 26, 2010

Abstract

Recently, consistent measures of the ex-post covariation of financial assets based on noisy high-frequency data have been proposed. A related strand of literature focuses on dynamic models and covariance forecasting for high-frequency data based covariance measures. The aim of this paper is to investigate whether more sophisticated estimation approaches lead to more precise covariance forecasts, both in a statistical precision sense and in terms of economic value. A further issue we address, is the relative importance of the quality of the realized measure as an input in a given forecasting model vs. the model's dynamic specification. The main finding is that the largest gains result from switching from daily to high-frequency data. Further gains are achieved if a simple sparse-sampling covariance measure is replaced with a more efficient and noise-robust estimator.

JEL classification: C32, C53, G11

Keywords: Forecast evaluation, Volatility forecasting, Portfolio optimization, Mean-variance analysis

---

*School of Economics and Management, Aarhus University, 8000 Aarhus C, Denmark. Phone +45-2367-2009, email: rasmus.tangsgaard.varneskov@post.au.dk. Financial support from the Faculty of Social Sciences, Aarhus University, and from CREATES funded by the Danish National Research Foundation is gratefully acknowledged.

†School of Economics and Management, University of Aarhus, 8000 Aarhus C, Denmark. Phone +45-8942-1539, email: vvoev@creates.au.dk. Financial support by the Center for Research in Econometric Analysis of Time Series, CREATES, funded by the Danish National Research Foundation, is gratefully acknowledged.

1 Introduction

The availability of high-frequency (HF) data in finance has fueled a large body of research on volatility measurement.[1] Many of the recently proposed techniques possess impressive statistical properties (e.g., close to maximum likelihood efficiency despite of their non-parametric nature) even in the presence of a so-called market microstructure (MMS) noise incorporating features such as bid-ask bounce, price discreteness, execution of block trades, etc. In multivariate applications, where an estimate of a covariance matrix is needed, non-synchronicity of observation times poses an added challenge, which has also been thoroughly addressed.[2]

From a more practical point of view, a growing number of recent studies assess the gains of improved volatility estimation in terms of better volatility prediction, compared to standard GARCH-type models using only daily data.[3] Generally, and not surprisingly, it is found that incorporating the information contained in HF data does indeed facilitate better forecasting. Arguably, the second component, i.e., the time series model, plays a key role in obtaining good forecasts and is the focus in many of the studies cited above. In virtually all of the studies, however, the results are based on a single estimation methodology which produces the volatility series of interest. Given that we observe a lot of research effort being invested in the development of more advanced estimation techniques, we believe it is of interest to assess, whether more sophisticated and efficient methods give an edge over simpler approaches with potentially worse statistical properties. After all, in the move from daily to HF data, the informational gain is intuitively enormous. However, given that researchers and businesses nowadays routinely possess HF datasets containing a given fixed amount of information, it then becomes important to assess how we can use this information effectively and extract gains from it.

The aim of this paper is to address the issue of whether a practitioner should necessarily apply the latest and most efficient estimator, which usually comes with a fair amount of computational complication, or computing a simple "sum of squares" type

[1] see, e.g., Zhang, Mykland & Aït-Sahalia (2005), Zhang (2006), Barndorff-Nielsen, Hansen, Lunde & Shephard (2008), Barndorff-Nielsen, Hansen, Lunde & Shephard (2010), Jacod, Li, Mykland, Podolskij & Vetter (2009), among others.

[2] see, e.g., Hayashi & Yoshida (2005), Zhang (2010), Griffin & Oomen (2010), Barndorff-Nielsen, Hansen, Lunde & Shephard (2009), Christensen, Kinnebrock & Podolskij (2010), Voev & Lunde (2007), Nolte & Voev (2008), among others.

[3] see, e.g., Fleming, Kirby & Ostdiek (2001, 2003), Hansen & Lunde (2006), Şerban, Brockwell, Lehoczky & Srivastava (2007), Patton (2009), Liu (2009), Clements, Doolan, Hurn & Becker (2009), Ghysels, Rubia & Valkanov (2009), Chiriac & Voev (2009), Voev (2009), etc.

of statistic is sufficient to extract the information to a desirable extent. Obviously, the advances toward more efficient and noise-robust estimators is of fundamental importance and the search in this direction is only natural. Nevertheless, given the importance of volatility in terms of its applied value in many finance-related problems, we feel that it is interesting to have an assessment of the potential practical benefits resulting from these advances.

2 Volatility Estimation

Modelling time varying volatility can be traced back at least to Engle (1982) and his autoregressive conditional heteroscedasticity (ARCH) model. Multivariate extensions to the ARCH idea originated in the works of Bollerslev, Engle & Wooldridge (1988), Bollerslev (1990), Engle & Kroner (1995), etc. Later, models for time-varying correlations were proposed by Engle (2002) and Tse & Tsui (2002). With the advances of computational technology and the advent of HF datasets, focus shifted towards using intermittent intraday data for volatility measurement and the development of multivariate models for "vast" dimensions. In this paper we consider some of the most recently proposed models for the estimation of the covariance matrix of financial assets by means of HF data, but we restrict our attention to approaches suitable for moderate dimensions. In particular, in our empirical application we consider 10 assets. The models we estimate can be applied to higher dimensions as well, but we do not necessarily advocate that they are still applicable in cases of hundreds or thousands of assets. One of the reasons to consider a relatively small-scaled problem is that the literature on vast-dimensional models is still in its infancy.[4] A second reason is that the HF-based models we consider, have been designed with a view on extracting statistical precision from HF data, rather than being able to handle very large dimensions. The so-called multivariate kernel approach (Barndorff-Nielsen et al. (2010)) can in theory be applied to an arbitrary number of assets, but is universe-dependent and involves an increasing loss of information as more assets are added. Alternative approaches based on non-synchronized data involving minimal loss of information (e.g., Hayashi & Yoshida (2005), Christensen et al. (2010), Voev & Lunde (2007), Nolte & Voev (2008), etc.) require the estimation of $n(n+1)/2$ separate elements of the covariance matrix, where n is the number of assets. The computational aspects aside, these methods do

[4]Two references include Engle, Shephard & Sheppard (2008) and Bannouh, Martens, Oomen & van Dijk (2009)

not, per se, guarantee a positive definite estimate, in particular for very large n.
Let r_t be a vector of daily log returns of dimension $n \times 1$. The process r_t can be written as:

$$r_t = E[r_t|\mathcal{F}_{t-1}] + \varepsilon_t, \tag{1}$$

where $\mathcal{F}_{t-1}$ is the information set consisting of all relevant information up to and including $t-1$. We assume that the innovation term can be expressed as $\varepsilon_t = H_t^{1/2} z_t$, where H_t is a symmetric positive definite matrix of dimension $n \times n$, $H_t^{1/2}$ is its Cholesky decomposition and z_t is an $n \times 1$ vector assumed to be iid with $E[z_t] = 0$ and $V[z_t] = I_n$. In the GARCH framework, the latent conditional covariance matrix H_t is specified parametrically, and is typically estimated based on daily data. Recently, GARCH models have been enhanced to include a realized volatility measure as an exogenous variable (see the review of Hansen & Lunde (2010)).
High-frequency volatility estimation is based on continuous-time models for the log price process, which is typically modeled as a (Brownian) semimartingale with stochastic volatility. The observable process is given by $p_\tau = p_\tau^* + u_\tau$, where the efficient $n \times 1$ vector price process is defined as[5]

$$p_\tau^* = \int_0^\tau a_s ds + \int_0^\tau \sigma_s dW_s \tag{2}$$

where a_s is a vector of predictable drifts, σ_s is a cadlag volatility matrix process and W is a vector of independent Brownian motions. The term u_τ represents noise due to MMS effects. The central object of interest in these class of models is the integrated (co)variation, $\int_{t-1}^t \sigma_s \sigma_s' ds$, which is a measure of the asset returns' (co)variance over a particular period (say, a day). For a detailed definition of these concepts, we refer the reader to Barndorff-Nielsen & Shephard (2004). Denote the j-th observation time for asset i on day t at the tick frequency as $\tau_{i,j,t}$, $j = 1, \ldots, m_{i,t}$, where $m_{i,t}$ is the number of observations for asset i on day t. For simplicity, we will use the triple i, j, t to denote a process observed at time $\tau_{i,j,t}$. At the tick frequency, we observe $p_{i,j,t} = p_{i,j,t}^* + u_{i,j,t}$ for $i = 1, \ldots, n$ and $j = 1, \ldots, m_{i,t}$, where $u_{i,j,t}$ is the i-th element of u_τ at time $\tau_{i,j,t}$. The intra-daily tick return series for asset i on day t is defined as $r_{i,j,t} = p_{i,j,t} - p_{i,j-1,t}$, $j = 1, \ldots, m_{i,t}$ which can be decomposed as the sum of the efficient return $r_{i,j,t}^* = p_{i,j,t}^* - p_{i,j-1,t}^*$ and the noise return $e_{i,j,t} = u_{i,j,t} - u_{i,j-1,t}$. Since assets trade asynchronously, it is generally not possible to define an n-vector of high

[5] We use $t = 1, \ldots, T$ as a discrete-time index and τ to index continuous time.

frequency returns unless some sort of synchronization scheme is used. Furthermore, at the highest frequencies the impact of market microstructure noise becomes dominant, so that either sparse sampling, or some noise-robust estimation technique needs to be employed.

Three different HF estimators of integrated (co)variation that exploit two types of synchronization schemes are applied in this paper. The first two are the realized covariance estimators based on using either sparse sampling, or a combination of sparse- and subsampling. Subsampling techniques are developed and discussed in the univariate case by Zhang et al. (2005) and Zhang (2006) and applied in a multivariate setting by Voev & Lunde (2007). The third estimator is the multivariate realized kernel presented in Barndorff-Nielsen et al. (2010). The former two estimators employ a calendar time sampling scheme for synchronization while the latter estimator requires the use of a so-called refresh time sampling scheme, detailed below.

2.1 Calendar Time Sampling

Calendar time sampling described in synchronizes tick-by-tick observations of multivariate time series by determining a sampling frequency and employing previous tick interpolation (see Dacorogna, Gençay, Müller, Olsen & Pictet (2001)) to construct a multivariate time series of prices on the sparse grids. The sampling frequency must be determined to mitigate the tradeoff between a strong negative bias due to non-synchronicity of the individual stocks' trading times (the so-called Epps (1979) effect), and explosive behavior in both the mean and the variances/covariances as the sampling frequency increases due to MMS noise as shown by Hansen & Lunde (2006). The empirical analysis in this paper is based on equidistant sampling at five-minute intervals similar to the comparison of estimators in Barndorff-Nielsen et al. (2010).

The calendar time sampling setup is used to define the intra-daily return vector $r_{j_s,t} \equiv p_{j_s,t} - p_{j_s-1,t}$ for $j_s = 2, \ldots, m_s$ where s indicates sparse sampling at a given sampling frequency, and m_s is independent of i and t. Using sampling at five-minute intervals with data from the NYSE that has a trading period of 6.5 hours each trading day results in $m_s - 1 = 78$ high-frequency returns.

This framework can be applied similarly to construct the intra-daily return vector for the realized covariance estimator that employs both sparse- and subsampling. The full sampling grid is partitioned into K non-overlapping subgrids, where the realized covariance is computed over each subgrid and then averaged across grids. In this case, we explicitly define the sparse grids through j_s^k, for $k = 1, \ldots, K$ and the corresponding returns by $r_{j_s^k,t} \equiv p_{j_s^k,t} - p_{j_s^k-1,t}$ for $j_s^k = 1, \ldots, m_s$, $k = 1, \ldots, K$. For the empirical implementation of the estimator, we choose $K = 300$, which is the maximal degree of subsampling with 5-minute returns and data stamped to the nearest second.

2.2 Refresh Time Sampling

Refresh time sampling contrary to calendar time sampling adjusts to some extent to the trading frequency of the assets by sampling at times at which at least one of the assets has traded. Thus, refresh time scale ensures that the prices at each sampling point are not too stale (at least not more than by a tick). The first point on the grid is obtained when all assets have traded at least once and the last known price for each asset is recorded. When all assets have traded at least once more, the latest price for each asset is recorded again, and so on. In this manner, the number of observations on the refresh time grid is at most equal (and often less than) to the number of observations of the least traded asset.

The sequence of refresh time observations is formally defined in Barndorff-Nielsen et al. (2010). The first time on the refresh time scale is given by $\tau_{1,t} = \max(\tau_{1,1,t}, \ldots, \tau_{n,1,t})$. The subsequent times are given by $\tau_{j_r,t} = \max\left(\tau^{(1)}_{N^{(1)}(\tau_{j_r-1,t})+1}, \ldots, \tau^{(n)}_{N^{(n)}(\tau_{j_r-1,t})+1}\right)$, $j_r = 2, \ldots, m_{r,t}$, where $N^{(i)}(\tau)$ is a process counting the number of observations of asset i up to (and including) time τ. $m_{r,t}$ is the number of synchronized observations on day t, which is not necessarily constant across days as opposed to the number of sparse grid observations m_s. The refresh time returns are then defined as $r_{j_r,t} \equiv p_{j_r,t} - p_{j_r-1,t}$ for $j_r = 2, \ldots, m_{r,t}$.

The degree of non-synchronicity and the frequency of trading for each asset determine the refresh time sample size, $m_{r,t}$. More information is lost for faster trading assets compared to their slower trading counterparts. A measure of data loss on trading day t can be defined as $1 - q_t$ with $q_t = \frac{n m_{r,t}}{\sum_{i=1}^{n} m_{i,t}}$.

2.3 High-Frequency Covariance Estimation

The sparse-drid (subsampled) realized covariance estimators can be defined straightforwardly after having defined the $n \times 1$ vector return series $r_{j_s,t}$ $(r_{j_s^k,t})$. The realized covariance estimator is consistent for the underlying integrated (co)variance if prices are semimartingales. However, the estimator will be severely biased if non-synchronicity and MMS noise are present. Sparse sampling is used to mitigate the

trade-off between bias and variance of the estimators.
The realized covariance estimator with sparse sampling is defined as

$$RC_t = \sum_{j_s=1}^{m_s} r_{j_s,t} r_{j_s,t}{}'$$

where the precision of the estimator depends on how well sparse sampling controls the tradeoff between non-synchronicity and MMS noise. Given that the estimator is based on sparse sampling, subsampling and averaging is a natural extension. The subsampled realized covariance estimator is given by

$$RC_t^{(s)} = \frac{1}{K} \sum_{k=1}^{K} \sum_{j_s^k=1}^{m_s} r_{j_s^k,t} r'_{j_s^k,t}.$$

Zhang et al. (2005) show that using subsampling and averaging reduces considerably the variance of the estimator and can eventually result in consistency (in a combination with bias correction). The multivariate realized kernel is an alternative to the realized covariance estimators that is consistent under the assumption of a covariance stationary noise process, and guarantees a positive semi-definite estimate of the integrated (co)variance. Having synchronized the return sequence using refresh time sampling, we can define the class of positive semi-definite multivariate realized kernels as[6]

$$RK_t = \sum_{h=-H_t}^{H_t} k\left(\frac{|h|}{H_t+1}\right) \Gamma_{h,t}$$

where $k(x)$ for $x \in R$ is the weight function. H_t is the length of the kernel, i.e. the bandwidth parameter, and $\Gamma_{h,t}$ is the h-order realized autocovariance on day t, that takes the following form

$$\Gamma_{h,t} = \begin{cases} \sum_{j_r=|h|+1}^{H_t} r_{j_r,t} r'_{j_r-h,t} & \text{if } h \geq 0 \\ \sum_{j_r=|h|+1}^{H_t} r_{j_r-h,t} r'_{j_r,t} & \text{if } h < 0 \end{cases}$$

Two critical issues for the empirical implementation of the multivariate realized kernel are the choice of kernel function, and the determination of the bandwidth parameter. The class of kernel functions we consider is characterized by the following four

[6]For an extensive treatment of the multivariate realized kernel, the reader is referred to Barndorff-Nielsen et al. (2010).

properties

1. $k(0) = 1$, $k'(0) = 0$;
2. k is twice differentiable with continuous derivatives;
3. $k_{\bullet}^{0,0} \equiv \int_0^1 k(x)^2 dx$, $k_{\bullet}^{1,1} \equiv \int_0^1 k'(x)^2 dx$, $k_{\bullet}^{2,2} \equiv \int_0^1 k''(x)^2 dx < \infty$;
4. $\int_{-\infty}^{\infty} k(x) \exp(tx\lambda) dx \geq 0$, for all $\lambda \in R$.

The first property implies that $\Gamma_{0,t}$ has a unit weight while $\Gamma_{h,t}$ gets a weight close to unity for small values of $|h|$. The second and the third property ensure that it is possible to derive the asymptotic distribution of the estimator. Furthermore, they are used to determine the optimal length of the bandwidth parameter. The fourth property guarantees RK_t to be a positive semi-definite covariance matrix.
The preferred kernel by Barndorff-Nielsen et al. (2010) that satisfies these four properties is the Parzen kernel, which is defined as

$$k_P(x) = \begin{cases} 1 - 6x^2 + 6x^3 & \text{if } \ 0 \leq x \leq 1/2 \\ 2(1-x)^3 & \text{if } \ 1/2 < x \leq 1 \\ 0 & \text{if } \ x > 1 \end{cases}$$

The Parzen kernel has the advantage that only a finite number, H_t, autocovariances need to be computed compared to theoretically infinite number required for the more efficient kernels such as the quadratic spectral and the Fejér kernel (where the gain in efficiency is negligible in practice).
Essential to the empirical implementation of the multivariate realized kernel is the bandwidth parameter, H_t. Under an iid noise assumption Barndorff-Nielsen et al. (2010) derive the optimal (for asset i) kernel length as $H_{i,t} = c^* \xi_{i,t}^{4/5} m_{i,t}^{3/5}$, where c^* is a constant[7], and $\xi_{i,t}^2 = \frac{\Omega_{ii,t}}{\sqrt{IQ_{ii,t}}}$ where $\Omega_{ii,t}$ is the i-th diagonal element of the covariance matrix of the noise process u_τ, $\Omega_t = \mathrm{V}[u_\tau]$ and $IQ_{ii,t}$ is the i-th diagonal element of the integrated quarticity matrix, given in Barndorff-Nielsen & Shephard (2004). In practice, $\sqrt{IQ_{ii,t}}$ is approximated by $\int_{t-1}^{t} \sigma_{ii,u}\sigma'_{ii,u} du$, the integrated variance for asset i. Finally, the global H_t is determined by averaging over the n optimal asset-specific $H_{i,t}$'s, i.e., $H_t = n^{-1} \sum_{i=1}^{n} H_{i,t}$.
In practice, the two quantities $\Omega_{ii,t}$ and $\int_{t-1}^{t} \sigma_{ii,u}\sigma'_{ii,u} du$ must be estimated. Barndorff-Nielsen et al. (2009) propose estimating the integrated variance for stock i using a subsampled realized variance estimator with 20-minute sampling. Furthermore, to

[7]For the Parzen kernel, $c^* = 3.5134$.

estimate the asset-specific noise component the authors propose a realized variance estimator based on a denser, one-minute sampling.

3 Volatility Forecasting and Evaluation

The question we have set out to answer is whether more sophisticated approaches to volatility measurement using high-frequency data, taking account of MMS noise explicitly deliver substantial gains compared to a more rudimentary sparse sampling approach. In terms of gains over daily data, these have been already documented in Fleming, Kirby & Ostdiek (2003), Liu (2009), Chiriac & Voev (2009), etc. Since we take a volatility forecasting perspective, a crucial ingredient in our study is how to specify the process for the volatility dynamics, from which we will eventually obtain the forecasts. A natural starting point would be a multivariate GARCH specification. To connect to the above cited literature, we consider an exponentially weighted GARCH model used in Fleming et al. (2003), which can easily incorporate the information contained in high-frequency data. The forecasts of the covariance matrix on day $t+1$ given some information at time t will be termed $H_{t+1|t}^{d,GARCH}$, and $H_{t+1|t}^{hf,GARCH}$, for the model using daily and high-frequency data, respectively. Formally, these are given by the following two equations:

$$
\begin{aligned}
H_{t+1|t}^{d,GARCH} &= \exp(-\alpha_d)H_t^{d,GARCH} + \alpha_d \exp(-\alpha_d)e_t e_t' && (3) \\
H_{t+1|t}^{hf,GARCH} &= \exp(-\alpha_{hf})H_t^{hf,GARCH} + \alpha_{hf} \exp(-\alpha_{hf})Y_t, && (4)
\end{aligned}
$$

where $e_t = r_t - \mu_t$ and Y_t is a realized covariance measure (the sparse RC_t, the subsampled $RC_t^{(s)}$ or the realized kernel RK_t). The conditional mean return process μ_t can be specified as generally as needed, but for obvious reasons we assume a constant mean, i.e., $\mu_t = \mu$. In practice, this will be estimated in a first step as the sample average of r_t. The specification in Equation (4) is motivated by the fact that Y_t is a much more precise measure of the volatility in period $t-1$ than $e_t e_t'$. Recursive substitution in Equation (4) shows that the forecast is implicitly given by a weighted sum of (lags of) Y_t with exponentially declining weights. The α parameters control the decay of the weights and given the discussion above, we expect $\alpha_{hf} > \alpha_d$, that is if Y_t indeed contains more information than $e_t e_t'$, it should get a higher weight in the forecast.

More recently, explicit reduced-form models for the dynamics of the matrix process

Y_t have been proposed. Chiriac & Voev (2009) advocate a long memory ARFIMA specification and compare its performance to the heterogeneous autoregressive (HAR) model of Corsi (2009), the Wishart autoregressive (WAR) model of Gourieroux, Jasiak & Sufana (2009) and the HAR extension of the WAR model by Bonato, Caporin & Ranaldo (2009). They find that using $RC_t^{(s)}$ as the realized measure, the ARFIMA and HAR specifications provide better forecasts than the remaining alternatives, but do not explore the issue of the sensitivity of the model performance to the choice of realized measure. Defining an unconstrained ARFIMA (HAR) model for the conditional expectation of a positive-definite valued matrix process is not unproblematic, as it can lead to negative forecasts out-of-sample. As a way around this problem, Chiriac & Voev (2009) propose to model the $n(n+1)/2 \times 1$ vector of Cholesky factors $X_t = vech(Chol(Y_t))$ and subsequently "square" the forecast, to obtain a positive definite covariance prediction.[8] The ARFIMA(1,d,1) model for the Cholesky factors X_t we employ is given by:

$$(1 - \phi L)D(L)[X_t - c] = (1 - \theta L)\varepsilon_t, \qquad \varepsilon_t \sim N(0, \Sigma), \tag{5}$$

where c is a vector of constants, ϕ and θ are the AR- and MA- coefficients, and $D(L) = (1 - L)^d \otimes I_{n(n+1)/2}$, where d is the (scalar) common degree of fractional integration for all of the $n(n+1)/2$ elements of the vector X_t and $I_{n(n+1)/2}$ is the identity matrix of the indicated dimension. The ARFIMA model exhibits long memory, but the process is stationary as long as $d < 0.5$. The HAR model is a mixed-frequency AR model which represents a simple alternative to fractionally integrated models for modeling persistent processes. In this paper, we consider the following specification

$$X_{t+1,d} = c^{(d)} + \beta^{(d)} X_t^{(d)} + \beta^{(w)} X_t^{(w)} + \beta^{(bw)} X_t^{(bw)} + \beta^{(m)} X_t^{(m)} + \omega_{t+1,d}, \tag{6}$$

where d stands for the daily, w – for the weekly (5 days), bw – for the bi-weekly (10 days) and m – for the monthly (20 days) frequency, $c^{(d)}$ is an $n(n+1)/2 \times 1$ parameter vector and the β's are scalar parameters, which can easily be estimated by OLS. The regressors $X_t^{(\cdot)}$ are averages of past values of X_t scaled to match the frequency of the left-hand-side (LHS) variable, e.g., in the equation above $X_t^{(bw)} = \frac{1}{10}\sum_{i=0}^{9} X_{t-i}$. By including AR components at different frequencies, the autocorrelation function

[8]The $vech$ operator stacks the upper triangular portion of a matrix into a vector. $Chol(Y_t)$ is the Cholesky decomposition of Y_t, an upper triangular matrix with the property that $(Chol(Y_t))'Chol(Y_t) = Y_t$.

of the process X_t mimics the hyperbolic decay of a long-memory process. For a detailed description of the models, the estimation and forecasting procedures, we refer the reader to Chiriac & Voev (2009). The ARFIMA and HAR models enable forecasting of X_t, e.g., a one-step ahead prediction $X_{t+1|t}$. The associated forecast of the conditional covariance matrix H_t, denoted for the ARFIMA model by $H_{t+1|t}^{hf,ARFIMA}$, can be obtained from the forecast $X_{t+1|t}$ by applying the inverse of the Cholesky decomposition, i.e, by "squaring". More formally, the ij-element of $H_{t+1|t}^{hf,ARFIMA}$ can be obtained from the elements of $X_{t+1|t}$ using the following equation:

$$H_{ij,t+1|t}^{hf,ARFIMA} = \sum_{l=1+\frac{i(i-1)}{2}}^{\frac{i(i+1)}{2}} X_{l,t+1|t} X_{l+\frac{j(j-1)}{2}-\frac{i(i-1)}{2},t+1|t}, \quad i,j=1,\ldots,n,\ j \geq i, \tag{7}$$

where $X_{l,t+1|t}$ is the l-th element of $X_{t+1|t}$. The HAR forecast, $H_{t+1|t}^{hf,HAR}$, is obtained from the HAR prediction of X_t in exactly the same manner.

The forecast evaluation is split into two parts. The first part consists of a statistical evaluation of the precision of one-step-ahead forecasts using a multivariate version of the RMSE criterion, which has been shown by Laurent, Rombouts & Violante (2009) to be robust to noise in the volatility proxy. The version we adopt is based on the Frobenius norm[9] of the matrix error term defined as the difference between the forecast and the ex-post covariance proxy:

$$e_{t+1}^{descr} \equiv Y_{t+1} - H_{t+1|t}^{descr}, \tag{8}$$

where $descr$ is the superscript describing the particular forecasting model, e.g., $hf, ARFIMA$. The main emphasis is put on the second part of the empirical evaluation, which is an analysis of the potential economic benefits associated with more sophisticated estimators of integrated covariation. The evaluation is based on a mean-variance portfolio optimization strategy similar to Chiriac & Voev (2009) and Voev (2009). We assume that at each period, the volatility forecast is used to determine the portfolio weights as

$$\begin{aligned} w_{t+1|t}^{*} &= \underset{w_{t+1|t}}{\operatorname{argmin}}\, w_{t+s|t}' H_{t+1|t}^{descr} w_{t+1|t} \\ \text{s.t.} \quad w_{t+1|t}' E_t[r_{t+1}] &= \frac{\mu_p}{250} \quad \text{and} \quad w_{t+1|t}' \iota = 1, \end{aligned} \tag{9}$$

[9] The Frobenius norm of a real $m \times n$ matrix A is defined as $||A|| = \sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij}^2$.

where $w_{t+1|t}$ is the $n \times 1$ vector of portfolio weights chosen at t, ι is an $n \times 1$ vector of ones, and μ_p is an annualized target return (scaled down to the daily frequency in the optimization problem above). In the empirical study, we vary μ_p over a range of values to derive the mean-variance frontier.
To test whether HF-data based models are significantly superior in terms forecasting performance, and whether there is significant practical benefits of applying more sophisticated estimators in terms of mean-variance efficiency, the Model Confidence Set (MCS) methodology of Hansen, Lunde & Nason (2009) is applied.

4 Empirical Study

In this section, we present estimation and forecasting results for three HF estimators of integrated (co)variation: realized covariance with sparse sampling, realized covariance with subsampling, and the multivariate realized kernel. The empirical evaluation is based on one-step-ahead forecasts of ex-post covariance. We have selected three different forecasting models to ensure consistency of our results: a exponentially weighted GARCH model, a HAR model, and an ARFIMA model, where the latter two are based on a Cholesky factorization of the covariance estimates. A common feature of the three dynamic specifications is that they are all guaranteed to provide positive semi-definite forecasts out-of-sample given that the in-sample covariance matrices on which the models are estimated are positive semi-definite.

4.1 The Data

The data consists of tick-by-tick trades from the New York Stock Exchange Trade and Quotation (TAQ) database for the period 01.01.2001 to 30.07.2008 during trading hours from 9.30 a.m. to 4 p.m, for a total of $T = 1904$ days.[10] We have selected ten representative stocks from the Dow Jones Industrial index: American International Group, Inc. (AIG), Boing Co. (BA), Bank of America Corporation (BAC), General Motors Corporation (GM), International Business Machines Coporation (IBM), The Coca-Cola Company (KO), McDonald's Corporation (MCD), Merck and Co., Inc. (MRK), Verizon Communications Inc. (VZ), and Wal-Mart Stores Inc. (WMT). Summary statistics of the daily returns for the ten stocks are presented in Table A.1 of Appendix A.1. Similarly, summary statistics of the multivariate realized kernel

[10]We are grateful to Asger Lunde for providing us with cleaned data.

estimates of integrated (co)variation are presented in Appendix A.1, Table A.2.[11] In the table we see that the estimated variances and covariances are highly right skewed and extremely leptokurtic. These results are similar to findings of Andersen, Bollerslev, Diebold & Ebens (2001) and Chiriac & Voev (2009) for realized variances and covariances.

4.2 Forecast Evaluation

It is instructive to first consider the differences in the informational content embedded in the volatility estimators we consider. The subsampled realized covariance estimator uses $m_s = 78$ observations with $K = 300$ subsamples (implying that we exhaust all available datapoints), while the realized kernel uses the Refresh Time Sampling scheme where the average degree of data loss in the sample is $1 - \bar{q}_t = 0.2916$, and the average number of observations used on a given trading day is $\bar{m}_{\tau,t} = 2851$.[12] Comparing this to the number of observations for the sparse realized covariance ($m_s = 78$) and the daily returns estimators (one daily observation), respectively, we expect $RC_t^{(s)}$ and RK_t to produce superior performance both statistically and economically.

For the purposes of out-of-sample evaluation, the observation period is split into two, one for estimation of the parameters in the dynamic models ($t_0 = 1256$) and one for validation ($T - t_0 = 648$). The forecasts are carried out in a recursive manner, i.e. the parameters of the dynamic models are re-estimated at each step with all the available data. We consider one-step-ahead forecasts since the investor is assumed to rebalance his portfolio on each trading day. Due to the leptokurtic distribution of returns, and the skewed and leptokurtic estimates of integrated variances and covariances, the GARCH and HAR model parameters' standard errors are estimated using QMLE for robustness. The ARFIMA model is estimated using the conditional maximum likelihood methodology developed in Beran (1995), which is also applicable for non-stationary processes with $d > 0.5$. The estimated rates of decay α_d and α_{hf} in the daily and the high-frequency GARCH model, respectively, substantiate the claim that the more informative volatility measure Y_t should get a higher weight in the forecast compared to $e_t e_t'$. The estimated values of α associated with the HF volatility proxies

[11] Similar tables of the other two HF estimators are available from the authors upon request.

[12] Figure A.1 in Appendix A.2 displays the time series of the number of observations and the degree of data loss for the Refresh Time Sampling scheme. While the degree of data loss fluctuates around its mean, the number of observations are upward trending throughout the sample, suggesting that an increasing amount of information is used to estimate integrated (co)variation in the multivariate realized kernel framework throughout the sample.

are given by $\alpha_{hf} = (0.1453, 0.1775, 0.1735)'$ for RC_t, $RC_t^{(s)}$, and RK_t, respectively, compared to $\alpha_d = 0.0137$. Interestingly, the values of α_{hf} could indicate that $RC_t^{(s)}$ and RK_t are more informative than RC_t.

4.2.1 Statistical Evaluation

The statistical evaluation is carried out as described in Section 3, based on the RMSE of the forecast error in Equation (8). We have selected RK_t as an ex-post proxy for H_t, due to its consistency.[13] The results for the different dynamic model/realized measure combinations are presented in Table 1. The statistical significance of the differences in the reported performance is evaluated across volatility estimators (column-wise comparison), and across dynamic models (row-wise comparison). The numbers in parenthesis are the MCS p-values of Hansen et al. (2009), where the first and second entry relate to the comparison across volatility estimators and dynamic models, respectively. We use (^{a}) and (^{b}) to denote the estimators or dynamic models that belong to the corresponding 5% model confidence sets. To assist interpretation, consider the combination of $RC_t^{(s)}$ and GARCH dynamics, which has a RMSE of 12.186. The MCS p-value of 0.7710^a suggests that holding the dynamic model fixed, $RC_t^{(s)}$ belongs to the 5% MCS of volatility estimators. Similarly, the p-value of 0.0128 shows that holding the volatility estimate (input) fixed, and varying across dynamic models, the GARCH model does not belong to the 5% MCS.

	Daily	RC_t	$RC_t^{(s)}$	RK_t
GARCH	14.802	12.433	12.186	12.173
	$(0.0000, -)$	$(0.0132, 0.0636^b)$	$(0.7710^a, 0.0128)$	$(1.0000^a, 0.0676^b)$
HAR	–	12.344	12.048	12.085
		$(0.0028, 0.1149^b)$	$(1.0000^a, 0.1507^b)$	$(0.4456^a, 0.1927^b)$
ARFIMA	–	12.265	11.980	12.028
		$(0.0040, 1.0000^b)$	$(1.0000^a, 1.0000^b)$	$(0.3492^a, 1.0000^b)$

Table 1: RMSE's based on Frobenius norm of the forecasting error for one-step-ahead forecasts, where RK_t is selected as an ex-post proxy for the true H_t. The MCS p-values of Hansen et al. (2009) are given in parenthesis, with the first referring to the column-wise (across realized measures) comparison and the second to the row-wise (across forecasting models) comparison. We use (^{a}) to denote the volatility measures that belong to the 5% MCS. Similarly, (^{b}) is used to denote the dynamic models that belong to the 5% MCS.

[13] The ranking of RMSE's are qualitatively the same using RC_t and $RC_t^{(s)}$ as proxy variables for H_t, adding robustness to our results.

From the first row of Table 1 we see that for the GARCH model there are substantial and significant gains from switching from daily to HF data in terms of statistical precision of out-of-sample forecasts. This result can be combined with the estimated values of the rate of decay α_d and α_{hf} to conclude that HF estimates of integrated (co)variation contain more information than daily estimates, and consequently get a higher weight in forecasting, which leads to better out-of-sample performance similar to results in Fleming et al. (2003), Liu (2009), Chiriac & Voev (2009). If we consider the comparison of HF estimators for each of the dynamic models, we see that only $RC_t^{(s)}$ and RK_t belong to the 5% MCS. Hence, there are significant gains in statistical precision from switching from a HF estimator based on a "sum of squares" type of computation to more sophisticated estimators that take MMS noise into account and exploit the available information more fully. This gain in statistical precision was similarly indicated by the estimated rate of decay for RC_t, which was slightly lower compared to the other two HF estimators. Holding the input variable fixed, there is a tendency for the ARFIMA model to deliver superior forecasts, followed by the HAR model, and lastly the GARCH model, which is to be excluded from the MCS at a 10% significance level. These results are similar to the findings of Chiriac & Voev (2009), but not as clear-cut. We can conclude that the largest forecasting gains result from improving the quality of the estimators of integrated (co)variance, i.e. the input in the dynamic models. Furthermore forecast precision is enhanced both by switching from daily to HF data, and by switching from simple to more sophisticated HF estimators.

4.2.2 Economic Evaluation

We base our economic evaluation on an investor that is assumed to minimize portfolio volatility subject to a target expected return μ_p as described in equation (9) in Section 3. The target portfolio return μ_p is varied, and the optimization problem is solved for each value to trace the efficiency frontier, which characterizes the best mean-variance tradeoff. The efficient frontiers for all dynamic model/realized measure combinations are depicted in Figure 1. The global minimum variance portfolio (GMVP) is emphasized and used to test for significance in the differences of economic performance.

The upper panel in Figure 1 shows the efficient frontiers for the daily squared return, RC_t, $RC_t^{(s)}$, and RK_t as inputs to the GARCH model.

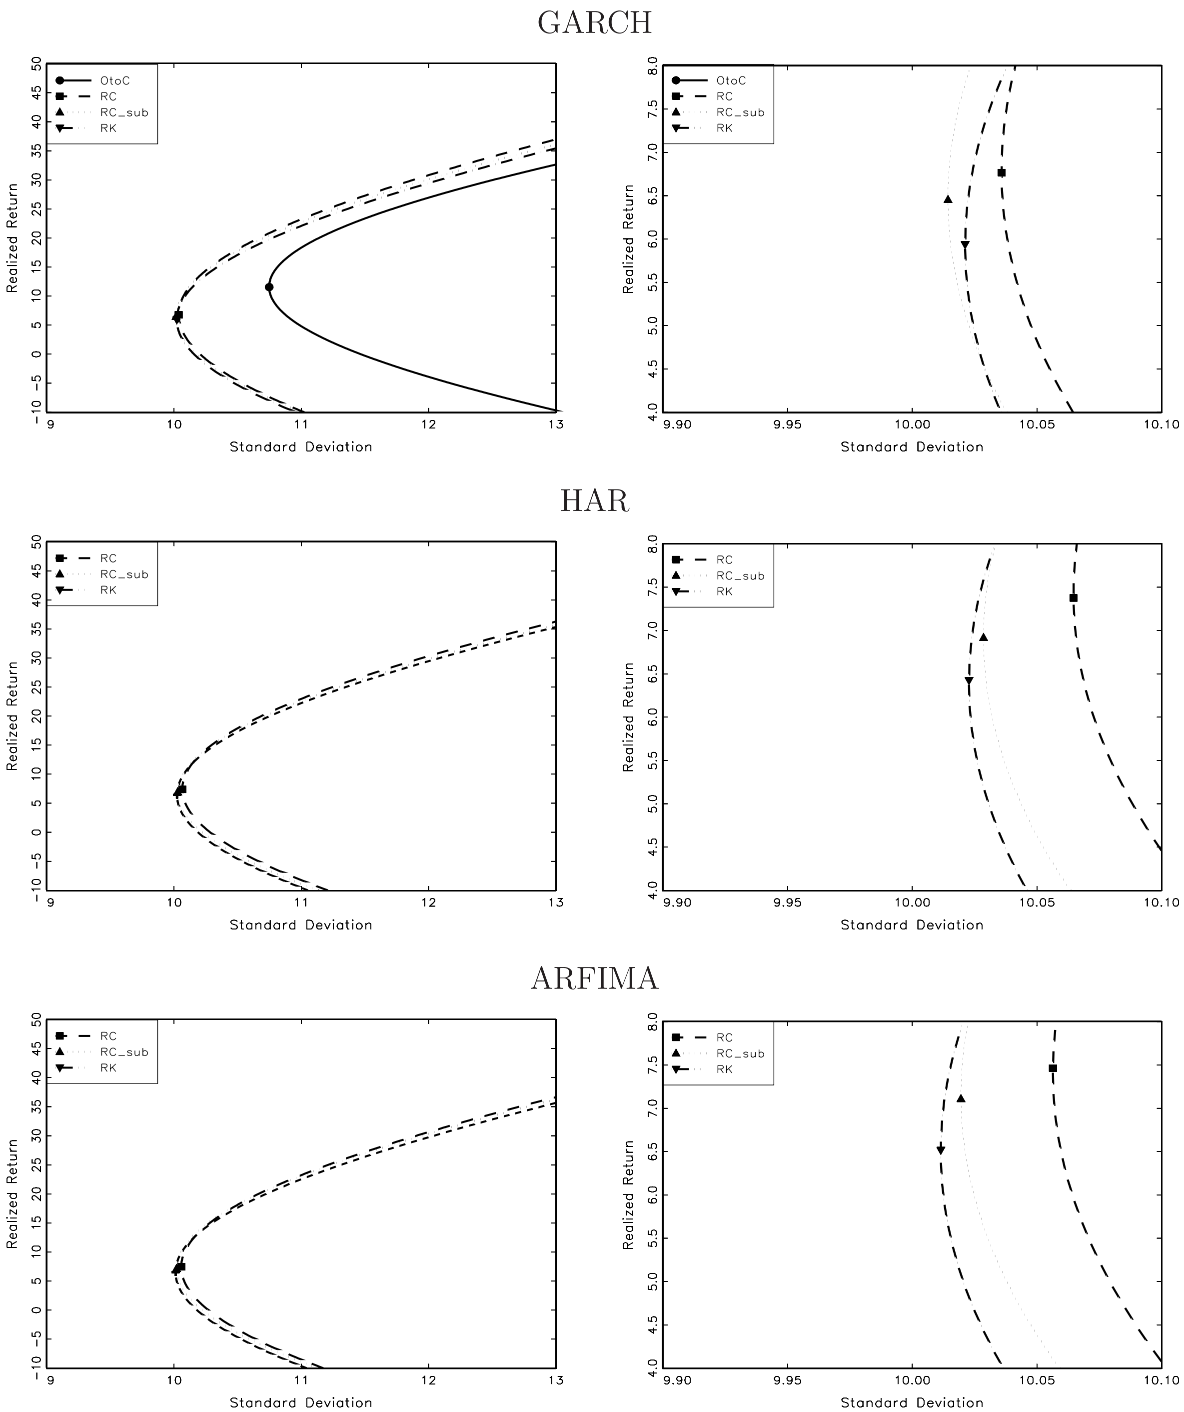

Figure 1: Mean-variance frontiers for the ex-post realized conditional mean (on the y-axis in %, annualized) against conditional standard deviation (on the y-axis in %, annualized). The global minimum variance portfolio is symbolized as follows: circle (daily open-to-close returns ($OtoC$)), square (RC), triangle ($RC^{(s)}$), and inverted triangle (RK). Panels 1, 2, and 3 correspond to GARCH, HAR and AFRIMA dynamics, respectively. The plots in the second column are zoomed into the area around the GMVP's. All plots are averaged over 648 out-of-sample periods.

Global Minimum Variance				
	Daily	RC_t	$RC_t^{(s)}$	RK_t
GARCH	10.747	10.036	10.014	10.021
	$(0.0000, -)$	$(0.0008, 1.0000^b)$	$(1.0000^a, 1.0000^b)$	$(0.3255^a, 0.1634^b)$
HAR	–	10.065	10.029	10.023
		$(0.0003, 0.1065^b)$	$(0.3786^a, 0.4413^b)$	$(1.0000^a, 0.0039)$
ARFIMA	–	10.056	10.020	10.011
		$(0.0004, 0.1065^b)$	$(0.2328^a, 0.4709^b)$	$(1.0000^a, 1.0000^b)$

Table 2: Standard deviation in annualized % for the global minimum variance portfolios formed on the basis of one-step-ahead forecasts. The MCS p-values of Hansen et al. (2009) are given in parenthesis, with the first referring to the column-wise (across realized measures) comparison and the second to the row-wise (across forecasting models) comparison. We use (^{a}) to denote the volatility measures that belong to the 5% MCS. Similarly, (^{b}) is used to denote the dynamic models that belong to the 5% MCS.

We see that there is a clear difference between the efficient frontier found on the basis of HF estimators and the one found using daily returns. Furthermore, when we zoom around the area of the GMVP (second column), we can more easily distinguish between the performance of the HF estimators. $RC_t^{(s)}$ and RK_t are seen to provide a lower global minimum portfolio variance compared to the RC_t estimator. The significance of these differences is tested by the MCS similar to the testing procedure for the RMSE's. The results are presented in Table 2, which is to be interpreted similarly to Table 1. There are significant economic gains from utilizing HF data instead of daily returns. Furthermore, by moving away from the simple RC_t estimator to more sophisticated and data-efficient HF estimators, the global minimum variance portfolio exhibits significantly lower volatility. This result is substantiated by considering both the middle and the lower panel of Figure 1, which show the efficient frontiers for the three HF estimators for the HAR and ARFIMA dynamics, respectively. In both cases, $RC_t^{(s)}$ and RK_t deliver a significantly superior economic tradeoff compared to the RC_t estimator. Hence, the rankings of the integrated covariance estimators are in line with the RMSE criterion.

When comparing across dynamic models, the results are not as clear-cut. The GARCH model is seen to perform surprisingly well considering its simple parameterization and rather weak statistical performance in terms of out-of-sample forecasting. It has the highest p-value when both RC_t and $RC_t^{(s)}$ are used as input. However, the statistical significance of these results is rather vague. The HAR model, when holding RK_t fixed as input, is the only example of a dynamic model that does not belong to the 5% MCS. Generally, there are some indications that the ranking of dynamic models does

no longer coincide with the RMSE criterion.
Since the RMSE criterion and mean-variance portfolio optimization strategy are two completely different approaches of model evaluation we did not expect to see exactly the same rankings for all combinations. The GMVP solution depends on the inverse of the covariance matrix. The inverse can become rather unstable if the original matrix has an eigenvalue close to zero which can be caused by an estimation error in a single element of the matrix. In such situations the GMVP loss can be substantially affected, while the Frobenius norm RMSE will increase only marginally. Furthermore, Chiriac & Voev (2009) find that the rankings of the two evaluation criteria are affected oppositely if the portfolio rebalancing horizon is increased to five or ten days.
Despite the differences in evaluation criteria, there is clear evidence that HF estimators outperform estimators based on daily returns both statistically and economically. Similarly, it is evident that more sophisticated estimators of integrated (co)variation are superior to the simple RC_t estimator. Hence, the quality of the realized measure is shown to be relatively more important than the dynamic model specification.

5 Conclusion

In this paper, we investigated whether more sophisticated estimators of integrated covariation lead to more precise covariance forecasts, both in terms of statistical precision and in terms of economic value, represented by the global minimum variance portfolio. Three HF estimators were selected for this purpose: the realized covariance with sparse sampling, the realized covariance with subsampling, and the multivariate realized kernel, where the latter two estimators are based on sampling schemes that utilize the large number of observations to a much greater extent.
Furthermore, we have addressed the relative importance of the quality of the realized measure as an input in a given forecasting model compared to the model's dynamic specification. Hence, three different dynamic models were used, an exponentially weighted GARCH model, the HAR model, and the ARFIMA model, where the latter two are based on a Cholesky factorization to ensure positive definite forecasts.
The forecast evaluation is split into two parts. The first part consists of a statistical evaluation of the precision of one-step-ahead forecasts using a multivariate version of the RMSE criterion. The economic evaluation is based on a mean-variance portfolio optimization strategy using one-step-ahead forecasts to minimize portfolio volatility subject to a target expected return.

The main finding in this paper is that the largest gains result from in the switch from daily to high-frequency data, which is similar to the findings of Fleming et al. (2003), Liu (2009) and Chiriac & Voev (2009). In addition, we show that further gains can be achieved if a simple sparse sampling covariance measure is replaced with a more efficient and noise-robust estimator. We show that for a given dynamic model, only the realized covariance with subsampling and the multivariate realized kernel belong to the 5% model confidence set of Hansen et al. (2009). Furthermore, we show that the choice of realized measure as an input to a dynamic model is relatively more important than the dynamic specification itself, encouraging practitioners to keep pace with the developments in the academic literature.

References

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Ebens, H. (2001), 'The distribution of realized stock return volatility', *Journal of Financial Economics* **61**, 43–76.

Bannouh, K., Martens, M., Oomen, R. & van Dijk, D. (2009), Realized mixed-frequency factor models for vast dimensional covariance estimation. Working Paper, Erasmus University Rotterdam.

Barndorff-Nielsen, O. E., Hansen, P., Lunde, A. & Shephard, N. (2008), 'Designing realized kernels to measure the ex-post variation of equity prices in the presence of noise', *Econometrica* **76**, 1481–1536.

Barndorff-Nielsen, O. E., Hansen, P., Lunde, A. & Shephard, N. (2009), 'Realized kernels in practice: Trades and quotes', *Econometrics Journal* **12**(3), C1–C32.

Barndorff-Nielsen, O. E., Hansen, P., Lunde, A. & Shephard, N. (2010), 'Multivariate realised kernels: consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading', *Journal of Econometrics* . forthcoming.

Barndorff-Nielsen, O. E. & Shephard, N. (2004), 'Econometric analysis of realised covariation: High frequency based covariance, regression and correlation in financial economics', *Econometrica* **72**, 885–925.

Beran, J. (1995), 'Maximum likelihood estimation of the differencing parameter for invertible short and long memory autoregressive integrated moving average models', *Journal of the Royal Statistical Society* **57**, 659–672.

Bollerslev, T. (1990), 'Modeling the coherence in short-run nominal exchange rates: A multivariate generalized arch model', *Review of Economics and Statistics* **72**, 498–505.

Bollerslev, T., Engle, R. F. & Wooldridge, J. M. (1988), 'A capital asset pricing model with time varying covariances', *Journal of Political Economy* **96**, 116–131.

Bonato, M., Caporin, M. & Ranaldo, A. (2009), Forecasting realized (co)variances with a block wishart autoregressive model. Working paper.

Chiriac, R. & Voev, V. (2009), 'Modelling and forecasting multivariate realized volatility', *Journal of Applied Econometrics* . forthcoming.

Christensen, K., Kinnebrock, S. & Podolskij, M. (2010), 'Pre-averaging estimators of the ex-post covariance matrix in noisy diffusion models with non-synchronous data', *Journal of Econometrics* . forthcoming.

Clements, A., Doolan, M., Hurn, S. & Becker, R. (2009), On the efficacy of techniques for evaluating multivariate volatility forecasts, NCER Working Paper Series 41, National Centre for Econometric Research.

Corsi, F. (2009), 'A simple approximate long-memory model of realized volatility', *Journal of Financial Econometrics* **7**, 174–196.

Şerban, M., Brockwell, A., Lehoczky, J. & Srivastava, S. (2007), 'Modelling the dynamic dependence structure in multivariate financial time series', *Journal of Time Series Analysis* **28**, 763–782.

Dacorogna, M. M., Gençay, R., Müller, U. A., Olsen, R. B. & Pictet, O. V. (2001), *An Introduction to High-Frequency Finance*, San Diego Academic Press.

Engle, R. (1982), 'Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation', *Econometrica* **50**, 987–1007.

Engle, R. (2002), 'Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroscedasticity models', *Journal of Business and Economic Statistics* **20**, 339–350.

Engle, R. F. & Kroner, K. F. (1995), 'Multivariate simultaneous generalized arch', *Econometric Theory* **11**(1), 122–50.

Engle, R. F., Shephard, N. & Sheppard, K. (2008), Fitting vast dimensional time-varying covariance models. Oxford-Man Institute, University of Oxford.

Epps, T. (1979), 'Comovements in stock prices in the very short run', *Journal of the American Statistical Association* **74**, 291–298.

Fleming, J., Kirby, C. & Ostdiek, B. (2001), 'The economic value of volatility timing', *Journal of Finance* **56**(1), 329–352.

Fleming, J., Kirby, C. & Ostdiek, B. (2003), 'The economic value of volatility timing using "realized" volatility', *Journal of Financial Economics* **67**, 473–509.

Ghysels, E., Rubia, A. & Valkanov, R. (2009), Multi-period forecasts of volatility: Direct, iterated and mixed-data approaches. Working Paper.

Gourieroux, C., Jasiak, J. & Sufana, R. (2009), 'The wishart autoregressive process of multivariate stochastic volatility', *Journal of Econometrics* **150**, 167–181.

Griffin, J. E. & Oomen, R. C. A. (2010), 'Covariance measurement in the presence of non-synchronous trading and market microstructure noise', *Journal of Econometrics* . forthcoming.

Hansen, P. & Lunde, A. (2010), Forecasting volatility using high frequency data, *in* M. P. Clements & D. F. Hendry, eds, 'Oxford Handbook of Economic Forecasting', Oxford University Press. forthcoming.

Hansen, P. R. & Lunde, A. (2006), 'Consistent ranking of volatility models', *Journal of Econometrics* **131**, 97–121.

Hansen, P. R., Lunde, A. & Nason, J. M. (2009), The model confidence set. Working Paper, Aarhus University.

Hayashi, T. & Yoshida, N. (2005), 'On covariance estimation of non-synchronously observed diffusion processes', *Bernoulli* **11**, 359–379.

Jacod, J., Li, Y., Mykland, P. A., Podolskij, M. & Vetter, M. (2009), 'Microstructure noise in the continuous case: The pre-averaging approach', *Stochastic Processes and their Applications* **119**(7), 2249 – 2276.

Laurent, S., Rombouts, J. & Violante, F. (2009), Consistent ranking of multivariate volatility models. ECORE Discussion Paper 2009/15.

Liu, Q. (2009), 'On portfolio optimization: How and when do we benefit from high-frequency data?', *Journal of Applied Econometrics* **24**, 560–582.

Nolte, I. & Voev, V. (2008), Estimating high-frequency based (co-) variances: A unified approach. CREATES Working paper 2008-31, Aarhus University.

Patton, A. (2009), 'Volatility forecast comparison using imperfect volatility proxies', *Journal of Econometrics* . forthcoming.

Tse, Y. & Tsui, A. (2002), 'A multivariate generalized auto-regressive conditional heteroscedasticity model with time-varying correlations', *Journal of Business and Economic Statistics* **20**, 351–362.

Voev, V. (2009), On the economic evaluation of volatility forecasts. CREATES Working Paper 2009-56.

Voev, V. & Lunde, A. (2007), 'Integrated covariance estimation using high-frequency data in the presence of noise', *Journal of Financial Econometrics* **5**, 68–104.

Zhang, L. (2006), 'Efficient estimation of stochastic volatility using noisy observations: A multi-scale approach', *Bernoulli* **12**, 1019–1043.

Zhang, L. (2010), 'Estimating covariation: Epps effect and microstructure noise', *Journal of Econometrics* . forthcoming.

Zhang, L., Mykland, P. A. & Aït-Sahalia, Y. (2005), 'A tale of two time scales: Determining integrated volatility with noisy high frequency data', *Journal of the American Statistical Association* **100**, 1394–1411.

A Appendix

A.1 Supplementary Tables

Daily Open-to-Close Returns						
	Mean	Max	Min	SD	Skewness	Kurtosis
AIG	-0,09	11,16	-12,06	1,72	-0,19	9,83
BA	-0,03	9,39	-8,02	1,57	-0,20	5,44
BAC	0,04	15,87	-12,50	1,58	1,11	21,82
GM	-0,21	13,64	-12,76	2,17	0,05	6,91
IBM	0,07	12,21	-6,39	1,36	0,36	8,40
KO	0,05	7,51	-4,39	1,05	0,12	5,75
MCD	0,09	6,01	-11,33	1,40	-0,21	7,09
MRK	0,02	9,75	-11,13	1,42	-0,36	9,68
VZ	-0,01	7,12	-7,63	1,46	0,26	5,70
WMT	-0,01	8,61	-5,74	1,26	0,41	6,23

Table A.1: Summary statistics of daily returns for the 10 stocks scaled to the power of 10^4 for the full sample period 01.01.2001 to 30.07.2008.

Daily Estimates of Integrated (Co)variation						
	Mean	Max	Min	SD	Skewness	Kurtosis
AIG	3.07	47.56	0.13	4.47	4.64	33.24
BA	3.30	38.20	0.17	3.51	3.51	21.97
BAC	2.55	61.65	0.11	4.60	6.52	62.58
GM	4.28	117.38	0.21	6.47	7.04	81.21
IBM	2.22	47.10	0.11	2.65	5.00	54.99
KO	1.70	24.23	0.06	1.99	4.64	36.94
MCD	2.74	45.34	0.23	3.43	5.40	46.97
MRK	2.59	109.26	0.15	4.09	12.12	262.25
VZ	3.05	32.62	0.19	3.58	3.42	19.79
WMT	2.36	27.63	0.15	2.69	4.02	27.19
AIG-BA	0.91	13.51	-1.06	1.36	4.08	27.94
AIG-BAC	1.22	37.60	-5.69	2.58	6.49	63.61
AIG-GM	1.07	23.37	-0.77	1.97	5.37	43.18
AIG-IBM	0.90	15.59	-1.94	1.33	4.39	32.22
AIG-KO	0.70	14.52	-0.87	1.18	5.65	50.48
AIG-MCD	0.75	18.96	-1.26	1.27	6.64	72.86
AIG-MRK	0.78	16.66	-2.18	1.38	5.16	43.31
AIG-VZ	0.99	16.97	-0.89	1.53	4.59	36.24
AIG-WMT	0.99	17.65	-1.25	1.55	5.11	43.71
BA-BAC	0.85	15.37	-1.61	1.30	4.68	37.53
BA-GM	0.90	15.24	-1.24	1.32	3.93	26.12
BA-IBM	0.88	14.84	-1.75	1.32	4.88	38.02
BA-KO	0.70	16.13	-1.80	1.22	5.86	57.04
BA-MCD	0.77	15.83	-2.29	1.22	5.72	53.07
BA-MRK	0.78	18.33	-1.50	1.40	5.69	52.03
BA-VZ	0.91	14.06	-0.81	1.32	3.78	24.91
BA-WMT	0.94	18.48	-1.03	1.51	5.55	49.75

Table A.2: Summary statistics of the variances and covariances using the multivariate realized kernel for the sample period 01.01.2001 to 30.07.2008. (continued on the next page)

Daily Estimates of Integrated (Co)variation continued						
	Mean	Max	Min	SD	Skewness	Kurtosis
BAC-GM	1.04	28.29	-1.46	2.06	6.03	53.77
BAC-IBM	0.87	17.78	-0.77	1.28	4.53	36.19
BAC-KO	0.63	14.07	-2.02	1.04	5.75	56.43
BAC-MCD	0.71	10.82	-2.10	1.06	4.26	29.71
BAC-MRK	0.69	13.28	-7.62	1.15	4.40	39.07
BAC-VZ	0.89	11.85	-1.79	1.26	3.40	20.21
BAC-WMT	0.94	16.86	-1.32	1.42	4.47	33.93
GM-IBM	0.87	28.67	-0.99	1.35	7.07	107.16
GM-KO	0.62	12.37	-1.46	0.97	4.45	32.81
GM-MCD	0.74	10.31	-2.00	1.07	3.73	23.17
GM-MRK	0.69	11.19	-2.61	1.11	4.05	26.93
GM-VZ	0.90	11.37	-1.78	1.27	3.27	18.03
GM-WMT	0.92	12.08	-1.73	1.29	3.55	21.11
IBM-KO	0.66	10.98	-1.24	0.98	4.36	31.40
IBM-MCD	0.72	10.39	-1.39	1.03	4.72	35.78
IBM-MRK	0.74	11.91	-3.62	1.19	4.42	30.48
IBM-VZ	0.93	16.95	-1.34	1.28	3.79	26.94
IBM-WMT	0.93	17.70	-1.71	1.34	4.59	35.84
KO-MCD	0.59	10.99	-0.60	0.97	5.73	51.28
KO-MRK	0.65	15.21	-1.02	1.20	6.18	58.93
KO-VZ	0.74	11.46	-0.89	1.15	4.25	28.67
MCD-MRK	0.62	15.51	-1.94	1.12	6.22	61.93
MCD-VZ	0.76	11.97	-1.45	1.13	4.29	32.57
MCD-WMT	0.79	15.09	-2.14	1.20	5.91	55.60
KO-WMT	0.74	17.28	-1.26	1.24	5.91	58.34
MRK-VZ	0.83	16.22	-2.50	1.36	4.47	34.50
MRK-WMT	0.79	20.74	-1.50	1.41	6.67	73.63
VZ-WMT	0.96	13.56	-0.79	1.38	3.96	26.42

Table A.2 ***(cont'd)*****:** Summary statistics of the variances and covariances using the multivariate realized kernel for the sample period 01.01.2001 to 30.07.2008.

A.2 Supplementary Figures

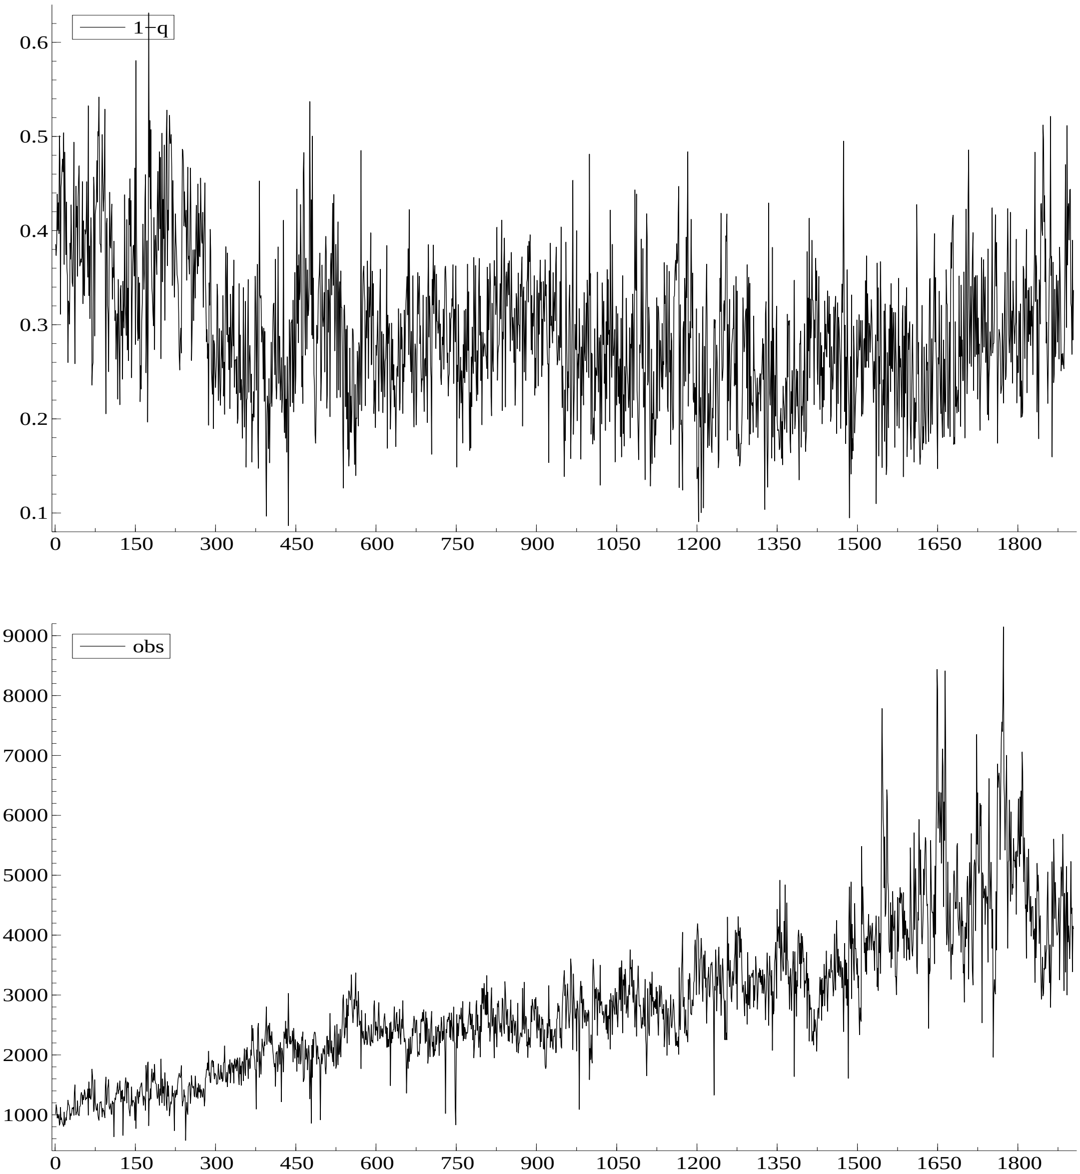

Figure A.1: Time Series of the degree of data loss when applying Refresh Time Sampling (upper panel), and the number of observations used in the multivariate kernel on each trading day (lower panel) in the sample period 01.01.2001 to 30.07.2008 .

Research Papers 2010

CREATES
Center for Research in Econometric Analysis of Time Series

2010-31: Morten Ørregaard Nielsen and Per Frederiksen: Fully Modified Narrow-Band Least Squares Estimation of Weak Fractional Cointegration

2010-32: Mogens Bladt and Michael Sørensen: Simple simulation of diffusion bridges with application to likelihood inference for diffusions

2010-33: Fernando Baltazar-Larios and Michael Sørensen: Maximum likelihood estimation for integrated diffusion processes

2010-34: Leonidas Tsiaras: The Forecast Performance of Competing Implied Volatility Measures: The Case of Individual Stocks

2010-35: Leonidas Tsiaras: Dynamic Models of Exchange Rate Dependence Using Option Prices and Historical Returns

2010-36: Robinson Kruse and Rickard Sandberg: Linearity Testing in Time-Varying Smooth Transition Autoregressive Models under Unknown Degree of Persistency

2010-37: Tom Engsted and Thomas Q. Pedersen: The log-linear return approximation, bubbles, and predictability

2010-38: Thomas Q. Pedersen: Predictable return distributions

2010-39: Rasmus Tangsgaard Varneskov: The Role of Dynamic Specification in Forecasting Volatility in the Presence of Jumps and Noisy High-Frequency Data

2010-40: Antonis Papapantoleon and David Skovmand: Picard Approximation of Stochastic Differential Equations and Application to Libor Models

2010-41: Ole E. Barndorff–Nielsen, Fred Espen Benth and Almut E. D. Veraart: Modelling electricity forward markets by ambit fields

2010-42: Robinson Kruse and Philipp Sibbertsen: Long memory and changing persistence

2010-43: Dennis Kristensen: Semi-Nonparametric Estimation and Misspecification Testing of Diffusion Models

2010-44: Jeroen V.K. Rombouts and Lars Stentoft: Option Pricing with Asymmetric Heteroskedastic Normal Mixture Models

2010-45: Rasmus Tangsgaard Varneskov and Valeri Voev: The Role of Realized Ex-post Covariance Measures and Dynamic Model Choice on the Quality of Covariance Forecasts