VU Research Portal

Fractional integration and fat tails for realized covariance kernels

Opschoor, Anne; Lucas, Andre

published in

Journal of Financial Econometrics
2019

DOI (link to publisher)

[10.1093/jjfinec/nby029](https://doi.org/10.1093/jjfinec/nby029)

document version

Publisher's PDF, also known as Version of record

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Opschoor, A., & Lucas, A. (2019). Fractional integration and fat tails for realized covariance kernels. *Journal of Financial Econometrics*, 17(1), 66-90. [17]. <https://doi.org/10.1093/jjfinec/nby029>

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 ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Fractional Integration and Fat Tails for Realized Covariance Kernels*

Anne Opschoor and André Lucas

Vrije Universiteit Amsterdam and Tinbergen Institute

Address correspondence to Anne Opschoor, Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV, Amsterdam, The Netherlands, or e-mail: a.opschoor@vu.nl

Received June 29, 2017; revised October 5, 2018; editorial decision October 8, 2018; accepted October 11, 2018

Abstract

We introduce a new fractionally integrated model for covariance matrix dynamics based on the long-memory behavior of daily realized covariance matrix kernels. We account for fat tails in the data by an appropriate distributional assumption. The covariance matrix dynamics are formulated as a numerically efficient matrix recursion that ensures positive definiteness under simple parameter constraints. Using intra-day stock data over the period 2001–2012, we construct realized covariance kernels and show that the new fractionally integrated model statistically and economically outperforms recent alternatives such as the multivariate HEAVY model and the multivariate HAR model. In addition, the long-memory behavior is more important during non-crisis periods.

Key words: fractional integration, heavy tails, matrix- F distribution, multivariate volatility, realized covariance matrices, score dynamics

JEL classification: C32, C58

1 Introduction

Various fields in financial econometrics, such as risk- and portfolio management, require the use of an adequate multivariate volatility model to estimate or forecast the covariance matrix of financial asset returns. We can distinguish two main lines of literature proposing

* We thank the associate editors, two anonymous referees, participants at the 69th European meeting of the Econometric Society (Geneva, August 2016), the Financial Econometrics and Empirical Asset Pricing Conference (Lancaster, July 2016), the 3rd annual meeting of the IAAE (Milan, June 2016), the 2015 NBER-NSF Time Series Conference (Vienna, September 2015), the 2nd International Workshop on Financial Markets and Nonlinear Dynamics (Paris, June 2015), and seminar participants at the Econometrics Brown Bag Seminar Series at Vrije Universiteit Amsterdam for useful comments. Both authors thank the Dutch National Science Foundation [NWO, grant VICI453-09-005] for financial support.

different types of models, namely multivariate generalized autoregressive conditional heteroskedasticity (GARCH)-type models (for an overview, see [Silvennoien and Teräsvirta, 2009](#)) and stochastic volatility (SV)-type models (for an overview, see [Asai, McAleer, and Yu, 2006](#)). More recently, the availability of intraday high-frequency data has led to a new class of volatility models including realized (co)variance measures. Such realized measures help to describe and forecast volatility more precisely than traditional measures such as squares and cross-products of daily returns; see for instance [Andersen et al. \(2001\)](#). Typically, either realized variance measures ([Barndorff-Nielsen and Shephard, 2002](#)) or realized kernel measures ([Barndorff-Nielsen et al., 2008](#)) are included. Examples of the former include the Wishart autoregressive (WAR) model of [Gourieroux, Jasiak, and Sufana \(2009\)](#) and the Conditional Autoregressive Wishart (CAW) model of [Golosnoy, Gribisch, and Liesenfeld \(2012\)](#), while examples of the latter include the high-frequency-based volatility (HEAVY) model of [Shephard and Sheppard \(2010\)](#) and its multivariate extension by [Noureddin, Shephard, and Sheppard \(2012\)](#), and the multiplicative error model (MEM) of [Englo and Gallo \(2006\)](#).

Volatilities are typically strongly persistent, which has led to the introduction of volatility models with long-memory features. A seminal reference is the fractionally integrated GARCH (FI-GARCH) model of [Baillie, Bollerslev, and Mikkelsen \(1996\)](#), which is based on squared daily returns. Realized variance measures exhibit even stronger long memory features than squared daily returns. [Andersen et al. \(2001\)](#) for instance find that realized measures are highly persistent and behave as fractionally integrated processes that can be modeled by autoregressive fractionally integrated moving average (ARFIMA) models; see also [Koopman, Jungbacker, and Hol \(2005\)](#). [Corsi \(2009\)](#) also captures long-memory volatility dynamics, but does so using a heterogeneous autoregressive (HAR) model, which relates realized volatility to a linear combination of lagged daily, weekly, and monthly realized volatilities. [Proietti \(2016\)](#) introduces an alternative integrated moving average model for realized variances and assesses its predictive power against other univariate models.

The univariate volatility models with long-memory features available in the literature face two main challenges that complicate their application to a multivariate context. First, these models do not account for fat-tailed returns and outliers in either the realized measures, the returns, or both. Although fat-tailed distributions are often used to describe returns, thin-tailed distributions such as the Wishart are typically used for the realized measures despite the fact that also data for the realized measures can be subject to outliers and influential observations. For example, the Flash Crash in 2010 led to a spike in the realized (co)variance of a large number of assets. Ignoring the possible occurrence of such influential events in the specification of the volatility propagation mechanism and the likelihood function can have a huge impact on the estimated volatility dynamics for each of the volatility models discussed above. Second, multivariate models that incorporate the long-memory features of (realized) (co)variances face the challenge to simultaneously avoid the curse of dimensionality and solve the requirement of ensuring positive definite covariance matrices. [Chiriac and Voev \(2011\)](#) deal with these issues by proposing vector ARFIMA (VARFIMA) models for the Cholesky decomposition of the realized covariance matrix. The same study also extends the HAR model of [Corsi \(2009\)](#) to the multivariate setting. [Bauer and Vorkink \(2011\)](#) solve the issue differently by modeling the matrix-logarithm of the realized covariance matrix. Both studies, however, model the vectorized (vech) matrix of interest.

This may become computationally intensive in higher dimensions. Moreover, neither model accounts for the possible fat-tailedness of realized measures and its impact on the volatility dynamics as discussed earlier.

In this paper, we address both problems at the same time by introducing a new multivariate volatility model for realized (kernel) covariance matrices. Our model allows for both the long-memory behavior and the fat-tailedness of (realized) covariances by combining fractionally integrated dynamics with the generalized autoregressive score (GAS) dynamics of the theoretical results of [Conrad and Haag \(2006\)](#). The only paper to our knowledge that combines long-memory and GAS is [Janus, Koopman, and Lucas \(2014\)](#), but this paper covers only a bivariate setting with univariate long-memory models for the time-varying variances and correlations. Our model, by contrast, is set in the general multivariate matrix setting. Moreover, unlike our paper, [Janus, Koopman, and Lucas \(2014\)](#) do not incorporate realized measures in their analysis and do not provide parameter constraints to ensure positive definiteness of the covariance matrix.

To account for fat tails of the realized covariance matrices, we use score-driven dynamics based on a matrix- F distributional assumption as in [Opschoor et al. \(2018\)](#). The resulting score-driven steps for the time-varying true covariance matrix automatically reduce the impact of outlying realized covariance matrices in an intuitive and empirically relevant way. Due to the matrix formulation of the covariance matrix dynamics, the introduction of long-memory features can be done in a parsimonious, yet flexible way. We do so following a very similar line of argument as proposed in [Baillie, Bollerslev, and Mikkelsen \(1996\)](#) to set up the original FIGARCH model, but then applied to our specific matrix-variate, score-driven context. The parsimony of our approach is a major asset in the multivariate context, where the curse of dimensionality looms large.

A well-known additional issue in modeling dynamic covariance matrices is to ensure positive definiteness of the covariance matrix at all times. Interestingly, our new model directly allows us to apply the theoretical results of [Conrad and Haag \(2006\)](#). Using these results, we can formulate easy parameter restrictions that ensure positive definiteness over the entire sample period.

We provide an empirical application of our multivariate fractionally integrated model based on the matrix- F distribution (FIGAS model from now on) on daily realized kernels for 15 equities from the S&P 500 index. Our sample spans the period January 2001 to December 2012. Using a forecasting horizon of 1, 5, 10, and 22 days ahead, we compare both statistically and economically the performance of our new dynamic covariance matrix model to several strong benchmarks, such as the HEAVY model ([Noureldin, Shephard, and Sheppard, 2012](#)), the GAS tF model ([Opschoor et al., 2018](#)), and the multivariate extension of the HAR model of [Corsi \(2009\)](#).

Using a quasi-likelihood loss function, the FIGAS model outperforms the competing HEAVY and HAR models, both inside and outside crisis periods. Interestingly, we find that the fractionally integrated part of our new model only outperforms the short-memory GAS model during non-crisis periods. Hence, the long-memory property seems particularly relevant during calm periods. We assess the economic significance of our results by considering mean-variance efficient portfolios based on the forecasts. Again we find that the FIGAS model outperforms its competitors by producing statistically significantly lower *ex post* conditional portfolio standard deviations, combined with lower portfolio concentration, lower total number of short positions, and equal or less turnover.

The rest of this paper is set up as follows. In [Section 2](#), we introduce the new FIGAS model for realized covariance matrices under fat-tails. In [Section 3](#), we provide a simulation experiment to show the performance of the model and estimation procedure. In [Section 4](#), we apply the model to a panel of daily realized kernels. We conclude in [Section 5](#).

2 Modeling Framework

2.1 The FI-GARCH Model

Before we introduce our new score-driven fractionally integrated volatility model, we first briefly review the main steps in the development of the univariate FI-GARCH(1, 1) or FIGARCH(1, d , 1) model of [Baillie, Bollerslev, and Mikkelsen \(1996\)](#). This paves the way to the new fractionally integrated dynamics in the score-driven framework. The FIGARCH(1, d , 1) model is obtained by rewriting the standard GARCH(1, 1) model of [Bollerslev \(1986\)](#) as

$$\sigma_{t+1}^2 = \omega + \alpha \epsilon_t^2 + \beta \sigma_t^2 \Leftrightarrow (1 - \alpha L - \beta L) \epsilon_{t+1}^2 = \omega + (1 - \beta L) v_{t+1}, \quad (1)$$

with L the lag operator $L\sigma_{t+1}^2 = \sigma_t^2$, σ_t^2 the conditional variance of ϵ_t , and $v_t = \epsilon_t^2 - \sigma_t^2$ a martingale difference. [Baillie, Bollerslev, and Mikkelsen \(1996\)](#) introduce the FIGARCH(1, d , 1) model by replacing the left-hand side lag polynomial $(1 - \alpha L - \beta L)$ by $(1 - L)^d(1 - \phi L)$, with $|\phi| < 1$ and $(1 - L)^d$ the fractional difference operator defined by the binomial expansion

$$(1 - L)^d = 1 - dL + \frac{d(d-1)}{2!} L^2 - \frac{d(d-1)(d-2)}{3!} L^3 + \dots, \quad (2)$$

for any real order of fractional integration $d > -1$. Using $v_t = \epsilon_t^2 - \sigma_t^2$, the FI-GARCH(1, d , 1) model can also be rewritten in its ARCH(∞) representation

$$(1 - L)^d(1 - \phi L) \epsilon_{t+1}^2 = \omega + (1 - \beta L) v_{t+1} \Leftrightarrow \sigma_{t+1}^2 = \tilde{\omega} + \Psi(L) \epsilon_{t+1}^2, \quad (3)$$

with $\tilde{\omega} = \omega/(1 - \beta)$ and

$$\Psi(L) = 1 - \frac{(1 - L)^d(1 - \phi L)}{(1 - \beta L)} = \sum_{i=1}^{\infty} \psi_i L^i. \quad (4)$$

Thus, the conditional variance σ_{t+1}^2 depends on lags of ϵ_{t+1}^2 , where the weight assigned to each lag declines hyperbolically according to $\Psi(L)$.

An important issue related to FI-GARCH models is whether these models are strictly stationary or not. According to [Conrad and Haag \(2006\)](#) and the references therein, this is still an open question. The same authors show that the original [Baillie, Bollerslev, and Mikkelsen \(1996\)](#) specification is not covariance stationary. To circumvent this problem, [Karanasos, Psaradakis, and Sola \(2004\)](#) introduced the closely related LMGARCH model, changing [Equation \(3\)](#) into

$$(1 - L)^d(1 - \phi L) (\epsilon_{t+1}^2 - \omega) = (1 - \beta L) v_{t+1} \Leftrightarrow \sigma_{t+1}^2 = \omega + \Psi(L) (\epsilon_{t+1}^2 - \omega), \quad (5)$$

where it can be shown that $\mathbb{E}[\sigma_{t+1}^2] = \omega$. The beauty of this model is that it combines the covariance stationary property of ϵ_t (for any $0 < d < 1$) with long-memory in ϵ_t^2 .

The LMGARCH model is one of the basic blocks for our model presented in the following subsection.

2.2 Score-Driven Fractionally Integrated Volatility Dynamics

Analogously to the FI-GARCH/LMGARCH case, we can now introduce the fractionally integrated score-driven multivariate volatility model. Consider a $(k \times k)$ -matrix process RK_t , $t = 1, \dots, T$, generated by

$$RK_t = V_t^{1/2} Z_t \left(V_t^{1/2} \right)', \quad Z_t | \mathcal{F}_{t-1} \sim D_Z(I_k), \quad (6)$$

where $\mathcal{F}_{t-1}$ is the information set containing all information up to time $t - 1$, V_t denotes the conditional covariance matrix, RK_t denotes the realized kernel covariance matrix measure, and Z_t denotes a $(k \times k)$ -matrix-valued innovation.¹ The matrix root $V_t^{1/2}$ is defined such that $V_t^{1/2} (V_t^{1/2})' = V_t$. The realized kernel process RK_t is a consistent and robust estimator of V_t correcting for market-microstructure noise; for more details, see [Barndorff-Nielsen et al. \(2011\)](#).

We assume V_t follows the score-driven dynamics as introduced by [Creal, Koopman, and Lucas \(2011, 2013\)](#) and [Harvey \(2013\)](#). Score dynamics adjust the time varying parameter V_t in the direction of steepest ascent of the local log-likelihood function. The approach is computationally easy to implement given its explicit form for the likelihood function. Score-driven dynamics also possess information theoretic optimality properties; see [Blasques, Koopman, and Lucas \(2015\)](#). Let $p(RK_t | V_t)$ denote the predictive conditional density for RK_t . Then the score-driven dynamics for V_t are driven by the scaled score

$$s_t = S_t \cdot (\partial \log p(RK_t | V_t) / \partial V_t) \cdot S_t', \quad (7)$$

where S_t is a scaling matrix to correct for the curvature of the log predictive density at time t . We come back to the precise form of the conditional observation density $p(RK_t | V_t)$, the choice of scaling S_t , and the expression for the scaled score s_t in [Section 2.3](#). For now, it suffices to note that our distributional and scaling choices allow us to write s_t as $s_t = s_t^* - V_t$, where s_t^* is positive definite for all t . For other distributional or scaling choices, positive definiteness of s_t^* may no longer be ensured.

To introduce fractionally integrated dynamics for the score-driven model, we first note that due to the standard properties of a predictive density score, s_t is a martingale difference by construction. It thus automatically takes the role of the martingale difference ν_t in [Equation \(1\)](#). Similarly, s_t^* is always positive definite and takes the role of ϵ_t^2 in [Equation \(1\)](#). Consider the standard GAS(1, 1) dynamics $V_{t+1} = \Omega + \alpha s_t + \beta V_t$ of [Creal, Koopman,](#)

1 One could add a measurement equation for the daily returns r_b for instance a Student's t -distribution $r_t | \mathcal{F}_{t-1} \sim t(\mu_t, V_t, \nu)$ with conditional mean μ_b , conditional covariance matrix V_t , and degrees of freedom parameter ν . We also estimated such an extension of the model. The returns typically contain much less information on V_t than do the RK_t observations. Omitting the returns therefore typically does not reduce the model fit substantially, which is why we only report the results for the model without the daily return equation.

and Lucas (2013), where α and β are scalar parameters. Using $V_t = s_t^* - s_t$, we first rewrite the GAS(1, 1) model into a GARCH type model

$$V_{t+1} = \Omega + \alpha s_t + \beta V_t \Leftrightarrow V_{t+1} = \Omega + A s_t^* + B V_t \quad (8)$$

with $B = \beta - \alpha$ and $A = \alpha$.

The fractionally integrated score-driven dynamics can now be derived analogously to the FI-GARCH setting, namely

$$\begin{aligned} V_{t+1} = \Omega + A s_t^* + B V_t &\Leftrightarrow (1 - AL - BL)s_{t+1}^* = \Omega + (1 - BL)s_{t+1} \\ &\Rightarrow (1 - L)^d(1 - \phi L)s_{t+1}^* = \Omega + (1 - BL)s_{t+1}, \end{aligned} \quad (9)$$

where we again replaced the standard GAS polynomial $(1 - (A + B)L)$ by the fractionally integrated polynomial $(1 - L)^d(1 - \phi L)$, where ϕ is a scalar parameter and Ω is a fixed positive definite parameter matrix. Finally, we follow Karanasos, Psaradakis, and Sola (2004) by slightly redefining Equation (9) as

$$(1 - L)^d(1 - \phi L)(s_{t+1}^* - \Omega) = (1 - BL)s_{t+1}, \quad (10)$$

such that for appropriate d we obtain that V_t is covariance stationary with $\mathbb{E}[V_t] = \Omega$. We label the model the fractionally integrated GAS model of order $(1, d, 1)$, or in short FIGAS $(1, d, 1)$.

Using the definition $s_t = s_t^* - V_t$, we can rewrite Equation (10) as

$$V_{t+1} = \Omega + \left(1 - \frac{(1 - L)^d(1 - \phi L)}{1 - BL}\right)(s_{t+1}^* - \Omega) = \Omega + \Psi(L)(s_{t+1}^* - \Omega), \quad (11)$$

with $\Psi(L)$ as defined in Equation (4). Thus, also for the FIGAS $(1, d, 1)$ model, the conditional covariance matrix V_{t+1} is an infinite weighted sum of current and past s_t^* , where the weight assigned to each lag declines hyperbolically according to $\Psi(L)$.

The fractionally integrated score dynamics introduced in Equation (10) are substantially different from those introduced in Janus, Koopman, and Lucas (2014). Whereas Janus, Koopman, and Lucas (2014) impose a fractional polynomial directly on the (in their case univariate) volatility parameter V_t , we follow the original approach of Baillie, Bollerslev, and Mikkelsen (1996) much more closely and impose the fractional polynomial on s_t^* . An important advantage of our current FIGAS specification compared with that of Janus, Koopman, and Lucas (2014) is not only that we allow for a multivariate setting, but also that we can immediately establish the positive definiteness of the sequence of covariance matrices V_t for all times t using simple parameter restrictions. This is stated in the following proposition for the FIGAS $(0, d, 1)$ model, which we use in the empirical application later on.

Proposition 1. Assume that Ω and s_t^* in Equation (11) are positive definite for all t . Then the conditional covariance matrices V_t from the FIGAS $(0, d, 1)$ model are positive definite if

- Case 1. $0 < B < 1, d - B \geq 0$;
Case 2. $-1 < B < 0, \left(d - \sqrt{2(2 - d)}\right)/2 \leq B$.

The proof follows directly from Corollary 3 of Conrad and Haag (2006). The proposition is stated for the FIGAS $(0, d, 1)$ model, which is the model we use in the empirical application later on. It is straightforward, however, to apply the results of Conrad and Haag

(2006) also for more general forms of the FIGAS model, such as the FIGAS(1, d , 1). This feature is quite convenient and follows from the way we have set up the fractionally integrated dynamics in contrast to earlier papers on fractionally integrated score-driven dynamics.

The conditions of Proposition 1 are easily checked for our model. The assumption on Ω is easily enforced through the model's parameterization. We show in the next subsection that also the second assumption on the positive definiteness of s_t^* is automatically satisfied for the fat-tailed distributional choice made in this paper. The restrictions in Case 1 or Case 2 of the proposition can then again be easily imposed by the model's parameterization.

2.3 Score for Matrix- F Distribution

We now turn to our choice for the conditional observation density $D_Z(\cdot)$ in Equation (6) to complete the FIGAS specification under fat tails. To account for possible fat tails of the realized kernel distribution, we assume that RK_t has a matrix- F distribution. The use of a matrix- F distribution for realized measures was first proposed in Opschoor et al. (2018) and is given by

$$p_{\text{RK}}(\text{RK}_t | V_t, \mathcal{F}_{t-1}; \nu_1, \nu_2) = K(\nu_1, \nu_2) \times \frac{|\frac{\nu_1}{\nu_2 - k - 1} V_t^{-1}|^{\frac{\nu_1}{2}} |\text{RK}_t|^{(\nu_1 - k - 1)/2}}{|\mathbb{I}_k + \frac{\nu_1}{\nu_2 - k - 1} V_t^{-1} \text{RK}_t|^{(\nu_1 + \nu_2)/2}}, \quad (12)$$

with positive definite expectation $\mathbb{E}_t[\text{RK}_t | \mathcal{F}_{t-1}] = V_t$, and degrees of freedom parameters $\nu_1, \nu_2 > k + 1$, where

$$K(\nu_1, \nu_2) = \frac{\Gamma_k((\nu_1 + \nu_2)/2)}{\Gamma_k(\nu_1/2) \Gamma_k(\nu_2/2)}, \quad (13)$$

and $\Gamma_k(x)$ is the multivariate Gamma function

$$\Gamma_k(x) = \pi^{k(k-1)/4} \cdot \prod_{i=1}^k \Gamma(x + (1-i)/2); \quad (14)$$

see, for example, Konno (1991).

Given the density in Equation (12), Opschoor et al. (2018) show that the scaled score² of the matrix- F distribution with respect to V_t is given by

$$s_t = \frac{\nu_1}{\nu_1 + 1} [W_t \cdot \text{RK}_t - V_t], \quad W_t = \frac{\nu_1 + \nu_2}{\nu_2 - k - 1} \cdot \left(\mathbb{I}_k + \frac{\nu_1}{\nu_2 - k - 1} \text{RK}_t \cdot V_t^{-1} \right)^{-1}, \quad (15)$$

using scaling matrix $S_t = \sqrt{2/(\nu_1 + 1)} V_t$. Given $\nu_1, \nu_2 > k + 1$ and given that V_t and RK_t are positive definite, we obtain that $s_t^* = s_t + V_t$ is positive semi-definite. The choice of scaling matrix S_t to obtain the scaled score in Equation (15) corrects the raw score for the most important curvature aspects. In particular, for the limiting case of the Wishart

- 2 Another parameterization could be to model the (time-varying) Cholesky matrix of V_t . This would result in more complex expressions for the derivatives, however, that would be numerically less efficient than the matrix recursion derived here. In addition, the parameterization of V_t is not sensitive to the order of the variables, whereas the Cholesky decomposition is. Finally, note that Proposition 3 in Opschoor et al. (2018) shows that the scaled score in Equation (15) with respect to a general matrix V_t coincides with the scaled score expression for a symmetric matrix V_t .

distribution S_t collapses up to a constant of proportionality to the inverse conditional Fisher information matrix scaling as introduced (as one of the scaling possibilities) in Creal, Koopman, and Lucas (2013). For finite ν_2 , scaling by S_t no longer coincides with precise inverse information matrix scaling, but still corrects for the most important curvature aspects of the score, while retaining high numerical efficiency; see Opschoor et al. (2018) for further comments and details.

The score s_t is highly intuitive. If the (weighted) realized kernel matrix RK_t is above the current estimate V_t , the estimate is updated upward. Because of the fat-tailedness of the matrix- F distribution, the observations RK_t are weighted in a matrix sense by the weight matrix W_t . This reduces the impact of outlying RK_t observations and generalizes the well-known scalar weights one obtains for Student's t -distributed returns with time-varying volatility as in Creal, Koopman, and Lucas (2011). For example, if $RK_t V_t^{-1}$ grows large, the weight matrix W_t tends to zero. The presence of W_t thus gives the model a robust feature for the measurement for V_t . For the Wishart distribution, we have $\nu_2 \rightarrow \infty$ and thus $W_t = I_k$, such that V_t reacts linearly to the realized kernel value.

Using the expression for the scaled score s_t , we obtain automatically that $s_t^* = s_t + V_t$ is positive semi-definite for all t . To see this, note that $W_t \cdot RK_t$ can be written as $\frac{\nu_1 + \nu_2}{\nu_2 - k - 1} (RK_t^{-1} + \frac{\nu_1}{\nu_2 - k - 1} V_t^{-1})^{-1}$. As the sum of two positive definite matrices is again positive definite, this expression is positive definite if ν_1 is positive, $\nu_2 > k + 1$, and RK_t and V_t are both positive definite. As a result, the parameter restrictions formulated in Proposition 1 can be used to enforce positive definiteness of V_t for all t in the FIGAS model.

2.4 Estimation

We estimate the parameters of the FIGAS model by maximum likelihood. To estimate the entries of Ω , we use a targeting approach and estimate it as $\hat{\Omega} = (1/T) \sum_{t=1}^T RK_t$, as under covariance stationarity our current parameterization implies $\mathbb{E}[RK_t] = \mathbb{E}[V_t] = \Omega$. We estimate the remaining static parameter vector $\theta = \{B, \nu_1, \nu_2, d\}$ of the FIGAS model by maximum likelihood. To do so, we maximize the log-likelihood $\mathcal{L}(\theta) = \sum_{t=1}^T \mathcal{L}_t$, where $\mathcal{L}_t$ is defined as the log-likelihood of the matrix- F distributions of Equation (12). This standard prediction error decomposition of the likelihood function is made possible due to the observation-driven nature of the FIGAS model using the classification of Cox (1981). The starting value V_1 can be either estimated or set equal to RK_1 .

The maximum-likelihood estimation for the fractionally integrated model requires truncation of the infinite distributed lags in Equation (2). We choose the maximum number of lags, which equals $T - 1$. Finally, we put the pre-sample innovations equal to zero, guided by the finding of Bollerslev and Mikkelsen (1996) that the effect of pre-sample values has a negligible effect on the parameter estimates, provided that the sample size is sufficiently large. Pre-sample values of V_t are put to the sample analog of the unconditional covariance matrix.

3 Simulation Experiment

Before presenting the empirical results, we perform a Monte Carlo study to investigate the statistical properties of the maximum-likelihood estimator for θ . To that end, we simulate T time series observations of $k \times k$ daily realized covariance matrices using the FIGAS model as the true data generating process (DGP). We set T equal to 500 and 1500,

Table 1 Parameter estimations of FIGAS DGP

Coef.	True	$T = 500$		$T = 1500$	
Panel A: $k = 5$					
d	0.60	0.574	(0.023)	0.590	(0.013)
B	−0.10	−0.121	(0.048)	−0.109	(0.030)
ν_1	50.00	49.647	(3.230)	49.871	(1.948)
ν_2	35.00	35.431	(1.651)	35.160	(0.967)
Ω_{11}	4.00	4.015	(0.840)	3.982	(0.905)
Ω_{12}	2.80	2.823	(0.729)	2.790	(0.802)
Panel B: $k = 15$					
d	0.60	0.566	(0.009)	0.586	(0.005)
B	0.15	0.118	(0.014)	0.137	(0.008)
ν_1	70.00	69.489	(1.114)	69.699	(0.671)
ν_2	60.00	60.550	(0.821)	60.243	(0.492)
Ω_{11}	4.00	3.995	(0.567)	3.985	(0.618)
Ω_{12}	2.80	2.811	(0.485)	2.790	(0.525)

Notes: This table shows Monte-Carlo averages and standard deviations (in parentheses) of parameter estimates from simulated FIGAS processes. The table reports the mean and the standard deviation in parentheses based on 1000 replications.

respectively, and choose k equal to 5 and 15. The chosen parameters are based on the FIGAS $(0, d, 1)$ model estimated in the application in Section 4. For $k = 5$, we choose $\beta = -0.10$, $\nu_1 = 50$, $\nu_2 = 35$, $d = 0.60$, and Ω the unconditional covariace matrix with $\Omega_{ii} = 4$ ($i = 1, \dots, k$) and $\Omega_{ij} = 2.8$ ($i \neq j$). For $k = 15$, we set $\beta = 0.15$, $\nu_1 = 70$, $\nu_2 = 60$, while the other parameters remain the same. For each simulated series, we estimate θ by numerically maximizing the likelihood function.

Table 1 presents the results based on 1000 replications. Clearly, all parameters are estimated near their true values. The standard deviations decrease if the sample size T increases or if the cross-dimensional dimension k becomes larger. There is a slight downward bias in d and B , which is partly resolved by increasing the sample size. This could also be due to a long lasting effect of the initialization of the initial covariance matrices to their unconditional expectation. This bias tapers off (not shown) if the sample size grows substantially larger. We also observe that the targeting approach works well as both Ω_{11} and Ω_{12} are estimated accurately. Similar results hold for the other entries of Ω .

4 Empirical Application

In this section, we apply the FIGAS model to an empirical data set of 15 U.S. equities. Our aim is to describe the covariance dynamics both in-sample and out-of-sample. All equities are part of the S&P 500 index. We first provide some of the stylized facts of the data. Next, we introduce our competing benchmark models. Finally, we test the in-sample and out-of-sample performance of the different models.

Table 2 S&P 500 constituents

No.	Ticker	Permno	Name	Subsector
1.	AA	24643	Alcoa Inc.	Materials
2.	AXP	59176	American Express Company	Financials
3.	BA	19561	The Boeing Company	Industrials
4.	CAT	18542	Caterpillar Inc.	Industrials
5.	GE	12060	General Electric Company	Industrials
6.	HD	66181	The Home Depot	Consumer discretionary
7.	HON	10145	Honeywell International	Industrials
8.	IBM	12490	International Business Machines	IT
9.	JPM	47896	JP Morgan	Financials
10.	KO	11308	Coca-Cola	Consumer staples
11.	MCD	43449	McDonald's	Consumer discretionary
12.	PFE	21936	Pfizer	Health care
13.	PG	18163	Procter & Gamble	Consumer staples
14.	WMT	55976	Wal-Mart Stores Inc.	Consumer staples
15.	XOM	11850	Exxon Mobil	Energy

Notes: This table lists 15 companies listed at the S&P 500 index during the period January 2, 2001–December 31, 2012. Ticker symbol and Permno (CRSP identifier) are also provided.

4.1 Data

The data consist of daily returns and realized covariances measures for 15 randomly chosen U.S. equities, where we have ensured that they come from various industries, such as Materials, Financials, Energy, IT, etc. Table 2 provides an overview of the companies considered in our data set. The data span the period January 2, 2001 until December 31, 2012 and contain $T=3017$ trading days. We observe consolidated trades (transaction prices) extracted from the Trade and Quote (TAQ) database from 9:30 A.M. until 4:00 P.M. with a time-stamp precision of 1 s. We first clean the high-frequency data following the guidelines of Barndorff-Nielsen et al. (2009) and Brownlees and Gallo (2006).³ Next, we construct realized kernels using the refresh-time-sampling methods of Barndorff-Nielsen et al. (2011).

Figure 1 shows a snapshot of the data by plotting the realized variances (based on the kernel approach) of Alcoa Inc. (AA) and Caterpillar Inc. (CAT) in the diagonal panels, and the realized correlation and covariance in the off-diagonal panels. The figure shows that both the realized (co)variance(s) and the realized correlation contain a substantial number of spikes. The spikes do not only occur during the global financial crisis, but also during other periods such as the early 2000s. This motivates the use of our FIGAS framework based on the fat-tailed matrix- F distribution, which automatically downweights the impact of such incidental observations on the volatility and covariance dynamics.

The autocorrelation functions in Figure 2 strongly suggest that the realized covariance matrix displays long-memory behavior. After lag 50, the autocorrelation is around 0.4 for the realized kernel volatilities of AA and CAT. Likewise, the autocorrelation of the realized

3 See the Online Appendix A for more details.

Figure 1 Realized kernel estimates of AA/CAT. This figure shows daily realized kernel volatilities (square root of the variance) of Alcoa Inc. (AA) and Caterpillar Inc. (CAT) returns on the diagonal panels. The off-diagonal panels contain the realized kernel covariance (upper-right) and correlation (lower-left) between the two asset returns. The sample spans the period from January 2, 2001 to December 31, 2012 ($T = 3017$ days).

covariance and correlation is equal to 0.25 and 0.3 at this long lag length. This provides an empirical motivation to incorporate long-memory features into the model.

4.2 Alternative Forecasting Models

To benchmark the performance of our FIGAS model, we use three relevant alternative models: the multivariate extension of the HAR model (Corsi, 2009), put forward by Chiriac and Voev (2011), the multivariate HEAVY model of Noureldin, Shephard, and Sheppard (2012), and the short-memory GAS tF model of Opschoor et al. (2018). As a fourth benchmark, we also considered the long-memory extension of the RiskMetrics model (Zumbach, 2006). This model, however, turned out to be substantially inferior to all other models considered, which is why we omit this model from our remaining analysis.

Our first benchmark does not directly model RK_t , but first computes the matrix' Choleski decomposition $RK_t = P_t P_t'$, where P_t is lower triangular. The multivariate HAR model then considers $X_t = \text{vech } P_t$ as a function of lagged daily, weekly, and monthly (transformed) volatilities, where $\text{vech } P_t$ stacks the lower triangular elements of P_t into a vector:

$$X_{t+1} = \alpha + \beta_1 X_t + \beta_2 X_t^w + \beta_3 X_t^m + u_{t+1}, \quad (16)$$

Figure 2 Empirical autocorrelation functions of realized kernels. This figure shows the autocorrelation function (ACF) for lag 1 until 50 of daily realized kernel volatilities (square root of variance) of Alcoa Inc. (AA) and Caterpillar Inc. (CAT) in the diagonal panels. The off-diagonal panels contain the ACF of the realized kernel covariance (upper-right) and correlation (lower-left) between the two asset returns. The sample is January 2, 2001 until December 31, 2012 ($T = 3017$ days).

where X_t^w and X_t^m are defined as $N^{-1} \sum_{i=0}^{N-1} X_{t-i}$ with $N = 5$ (weekly) and 22 (monthly), respectively. Finally, α represents a $k(k+1)/2$ vector of coefficients and β_j ($j = 1, 2, 3$) are scalar parameters. All parameters are estimated by OLS.

The multivariate HEAVY model incorporates realized measures into the volatility specification by proposing a system of two multivariate GARCH equations for the quantities $V_t = \mathbb{E}_t[y_t y_t' | \mathcal{F}_{t-1}]$ and $M_t = \mathbb{E}_t[\text{RK}_t | \mathcal{F}_{t-1}]$, where y_t denotes a $k \times 1$ vector of daily returns. The innovations in both of these equations are the realized (co)variance measures as gathered in the matrix RK_t . The dynamics are given by

$$V_{t+1} = C_V C_V' + \alpha_V \text{RK}_t + \beta_V V_t, \quad (17)$$

$$M_{t+1} = C_M C_M' + \alpha_M \text{RK}_t + \beta_M M_t, \quad (18)$$

where $\alpha_V, \alpha_M, \beta_V$, and β_M are scalar parameters, and C_V and C_M are lower triangular matrices. The scalar parameters of both equations are estimated separately by maximum likelihood, assuming a singular Wishart distribution for $y_t y_t'$ and a standardized Wishart distribution with k degrees of freedom for RK_t . The matrices C_V and C_M are typically

estimated by covariance targeting, as discussed by [Noureldin, Shephard, and Sheppard \(2012\)](#). We follow this approach when implementing the model in the remaining analysis.

Our final benchmark is the GAS tF model, which is similar to the short memory equivalent of our FIGAS model:

$$V_{t+1} = \Omega + As_t + BV_t, \quad (19)$$

with the difference that the scaled score s_t depends not only on the matrix- F distribution, but also on Student's t -distribution. This is due to the fact that in [Opschoor et al. \(2018\)](#), there are two observation densities, both for the returns y_t and for the realized kernel covariance matrix. Both observation densities depend on the latent covariance matrix V_t . The score s_t is now defined as the sum of the score from the matrix- F and Student's t -distributions.

All benchmark models allow for easy h -step ahead prediction of V_t . In case of the HEAVY model, the second transition equation delivers forecasts of RK_{t+b} for $b = 1, 2, \dots$, which can subsequently be inserted into the first equation to obtain V_{t+b} . The h -step ahead forecast of V_t of the FIGAS model follows directly from [Equation \(11\)](#): V_{t+b} depends on $s_{t+b-1}^*, s_{t+b-2}^*, \dots, s_t^*, s_{t-1}^*, \dots$, with $s_t^* = s_t + V_t$ by definition. Given the property that $\mathbb{E}_t[s_{t+b}|\mathcal{F}_t] = 0_k$ for any value of $b \geq 1$, V_{t+b} is obtained recursively by setting the values of future score matrices s_t to zero. Similar results hold for the GAS tF model.

We follow [Bollerslev, Patton, and Quaadvlieg \(2018\)](#) by considering direct forecasts in case of the multivariate HAR model. These forecasts are obtained by running the following regression:

$$X_{t+b} = \alpha + \beta_1 X_t + \beta_2 X_t^w + \beta_3 X_t^m + u_{t+b}, \quad (20)$$

where b stands for the forecast horizon. As indicated by [Bollerslev, Patton, and Quaadvlieg \(2018\)](#), direct forecasts might be more adequate than iterative forecasts due to the possibility of model misspecification.

4.3 Model Evaluation Procedure

We follow [Noureldin, Shephard, and Sheppard \(2012\)](#) and compare the in-sample and out-of-sample statistical fit of the models by computing the quasi-likelihood loss function:

$$\text{QLIK}_{t,b} \left(RK_{t+b}, V_{t+b|t}^a \right) = \log |V_{t+b|t}^a| + \text{tr} \left(\left(V_{t+b|t}^a \right)^{-1} RK_{t+b} \right), \quad (21)$$

with $V_{t+b|t}^a$ the covariance matrix forecast for time $t+b$ given all information up to time t based on model a . Note that we use RK_{t+b} as a proxy of the true covariance matrix. In-sample, b is set to 1. Since V_{t+1} is known at time t , the criteria can also be interpreted as one-step ahead forecasting criteria. As indicated by [Laurent, Rombouts, and Violante \(2013\)](#), the QLIK loss-function implies a consistent ranking of volatility models since it is robust to noise in the proxy RK_t .

We additionally test the predictive performance of the models using the framework of [Giacomini and White \(2006\)](#). We start by computing the difference in loss functions between two competing models a and b ,

$$d_{t,b}(a, b) = \text{QLIK}_{t,b} \left(RK_{t+b}, V_{t+b|t}^a \right) - \text{QLIK}_{t,b} \left(RK_{t+b}, V_{t+b|t}^b \right), \quad (22)$$

for $t = R + 1, \dots, T - h$, where the parameters are estimated based on a rolling window of $T_w = 1500$ observations. The difference d_t can be interpreted as a difference between two Kullback–Leibler (KL) divergences. Even if the underlying two models are both misspecified, the difference in their KL divergences still provides a valid assessment criterion. The corresponding null-hypothesis of equal predictive ability is given by $H_0 : \mathbb{E}[d_{t,b}(a, b)] = 0$ for all $T - h - R$ out-of-sample forecasts, which can be tested using the Diebold and Mariano (1995) (DM) test-statistic

$$DM_b(a, b) = \frac{\bar{d}_b(a, b)}{\sqrt{\hat{s}_b^2(a, b)/(T - h - R)}}, \quad (23)$$

with $\bar{d}_b(a, b)$ the out-of-sample average of the loss differences, and $\hat{s}_b^2(a, b)$ a HAC-consistent variance estimator of $d_{t,b}(a, b)$. A significantly negative value of $DM_b(a, b)$ means that model a has a superior forecast performance over model b . The QLIK test can be used in-sample (interpreted as a “one-step-ahead prediction”) and out-of-sample. In the out-of-sample test, we choose $h = 1, 5, 10$, and 22 . In addition, we consider the cumulative forecasts $V_{t:t+N|t} = \sum_{i=1}^N V_{t+i|t}$, where N equals $5, 10$, and 22 , respectively.

As the above evaluation criteria are statistical in nature, we finally also assess the forecasting performance from an economic point of view. Motivated by the mean–variance optimization setting of Markowitz (1952), we do so by considering global minimum variance portfolios (GMVP); see, for example, Chiriac and Voev (2011); Engle and Kelly (2012), among others, who perform a similar analysis. The best forecasting model should provide portfolios with the lowest *ex post* variance. Assuming that the investor’s aim is to minimize the h -step portfolio volatility at time t subject to a fully invested portfolio, the resulting GMVP weights $w_{t+h|t}$ are obtained by the solution of the quadratic programming problem

$$\min w'_{t+h|t} V_{t+h|t} w_{t+h|t}, \quad \text{s.t. } w'_{t+h|t} \mathbf{1} = 1. \quad (24)$$

with $\mathbf{1}$ a $k \times 1$ vector of ones. Similar as Chiriac and Voev (2011), we assess the predictive ability of the different models by comparing the results to the *ex post* realizations or “oracle forecasts” of the conditional standard deviation, which are given by $\sigma_{p,t} = \sqrt{w'_{t+h|t} \text{RK}_{t+h} w_{t+h|t}}$. We again test for significantly different portfolio standard deviations by means of the DM-test statistic.

Finally, we follow Bollerslev, Patton, and Quaedvlieg (2018) by reporting practical portfolio quantities involved with implementing the minimum-variance portfolio strategy. More specifically, we consider the turnover (TO_t), the concentration (CO_t), and the total short positions (SP_t) for each competing model at time t . The turnover at time t is defined as

$$\text{TO}_t = \sum_{i=1}^k \left| w_{t+1}^{(i)} - w_t^{(i)} \frac{1 + y_t^{(i)}}{1 + w_t' y_t} \right| \quad (25)$$

where $w_t^{(i)}$ and $y_t^{(i)}$ the i -th element of the weight vector w_t and return vector y_t . A model that produces more stable covariance matrix forecasts implies in general less turnover and hence, less transaction costs.

Further, the portfolio concentration and total portfolio short position both measure the amount of extreme portfolio allocations. Again, more stable forecasts of V_t should result in less extreme portfolio weights. The portfolio concentration reads

$$\text{CO}_t = \left(\sum_{i=1}^k w_t^{(i)2} \right)^{1/2}, \quad (26)$$

while the total portfolio short positions SP_t is given by

$$\text{SP}_t = \sum_{i=1}^k w_t^{(i)} I[w_t^{(i)} < 0] \quad (27)$$

with $I[\cdot]$ an indicator function that takes the value 1 if the i -th element of the weight vector is lower than zero.

4.4 In-Sample Results

Table 3 shows parameter estimates and standard errors based on the sandwich (robust covariance matrix) estimator $A_0^{-1} B_0 A_0^{-1}$ with B_0 the inverse Hessian of the likelihood evaluated at the optimum (information matrix), and A_0 the expected value of the outer product of the gradients at the optimum. We show the results for a selection of $k=5$ randomly chosen stocks and for the complete set of all 15 equities.

The results in Table 3 show that the FIGAS model has the best fit to the data compared with the other models. In a preliminary analysis, the coefficient ϕ in the FIGAS specification turned out to be statistically insignificant, such that we estimate a FIGAS(0, d , 1) model as our preferred fractionally integrated score-driven model.

Based on the QLIK loss function, the FIGAS model has the best value, followed by the GAS tF, HAR, and the HEAVY model, respectively. Comparing the FIGAS, GAS, and HEAVY models, the QLIK values suggest that the largest gain is obtained by introducing the score-driven dynamics: the average QLIK drops from 7.89 (20.06) for the HEAVY model to 7.66 (18.96) for the GAS model. Hence, allowing for fat-tailedness in the realized covariance kernels improves the fit substantially. The further drop in QLIK when moving from GAS to FIGAS is more modest, but still sizable given the sample size. For $k=5$, the simple HAR model still does quite well. The HAR model's relative performance, however, quickly deteriorates in higher dimensions such as $k=15$, as is seen in Panel B of Table 3.

The likelihood and the BIC values for the different models underline that the FIGAS model provides a better fit to the data than the GAS and HEAVY models. Especially when $k=15$, the differences are large. In order to compare the FIGAS-type model with the GAS, HEAVY, and HAR models, we only consider the likelihood of these models associated with the realized kernel covariance matrix. For the GAS tF model, this implies that we only consider the likelihood of the matrix- F distribution, while for the HEAVY and HAR models we report the Wishart distribution with k degrees of freedom for the realized kernels. This makes the likelihood comparable to that of the FIGAS model, as the latter consists only of the matrix- F distribution.

We see a clear distinction between the likelihoods of the HEAVY and HAR model and the FIGAS model. This occurs mainly due to the likelihood contribution for the realized kernels. More specifically, this contribution equals 25,000 points ($k=5$) or 130,000 points

Table 3 Parameter estimates, likelihoods, and loss function

Panel A: AA/BA/CAT/GE/KO										
	d	α	β	β_2	β_3	ν_0	ν_1	ν_2	$\mathcal{L}^*$	QLIK
FIGAS	0.663 (0.012)		-0.086 (0.025)				45.642 (1.662)	36.698 (1.192)	-20,172	7.647
GAS		0.611 (0.025)	0.987 (0.001)			10.46 (0.730)	46.39 (1.669)	35.58 (1.112)	-20,585	7.664
HEAVY		0.284 (0.049)	0.589 (0.080)						-46,793	7.894
HAR			0.306 (0.016)	0.418 (0.023)	0.214 (0.020)				-45,223	7.684
Panel B: AA/AXP/BA/CAT/GE/HD/HON/IBM/JPM/KO/MCD/PFE/PG/WMT/XOM										
	d	α	β	β_2	β_3	ν_0	ν_1	ν_2	$\mathcal{L}^*$	QLIK
FIGAS	0.653 (0.009)		0.157 (0.016)				66.94 (1.209)	62.08 (0.984)	70,954	18.88
GAS		0.386 (0.012)	0.991 (0.001)			12.30 (0.571)	66.63 (1.161)	61.54 (0.975)	68,540	18.96
HEAVY		0.160 (0.019)	0.738 (0.033)						-65,977	20.06
HAR			0.204 (0.008)	0.419 (0.012)	0.296 (0.011)				-52,238	19.45

Notes: This table reports maximum-likelihood parameter estimates of the FIGAS, GAS, HEAVY, and the HAR model, applied to daily realized kernels of 5 and 15 assets. In case of the HAR model [see Equation (16)], we suppress the constant vector and report only β_1 , β_2 , and β_3 . Asset identifiers are explained in Table 2. Standard errors are provided in parenthesis and based on the (sandwich) robust covariance matrix estimator. The α column corresponds to α coefficient of the GAS model [see Equation (19)], and the α_V coefficient of the HEAVY model [see Equation (17)], respectively. Likewise, β corresponds with β of the (FIGAS) model, the β_V coefficient of the HEAVY model or the β_1 coefficient of the HAR model. The table reports the log-likelihood, the BIC criteria, as well as the mean of the QLIK loss function, which is defined in Equation (21). For comparative reasons, the likelihood consists of matrix- F , and the Wishart distribution for GAS, HEAVY, and HAR models. The sample is January 2, 2001–December 31, 2012 (3017 observations).

Figure 3 FIGAS implied correlograms for V_t for $k=5$ (left) and $k=15$ (right). This figure plots the implied correlograms of the conditional volatility of Alcoa Inc. (AA) corresponding with the FIGAS model based on parameter estimates in Table 3. The left panel shows the correlogram implied by the estimated model on AA/BA/CAT/GE/KO, while the right panel uses all ($k=15$) equities.

($k=15$) when going from the Wishart distribution (HEAVY and HAR model) to the matrix- F distribution (FIGAS model). Fat-tailedness of RK_t thus appears a prevalent feature in the data.

The likelihood values for the short-memory GAS specification and the FIGAS model are easier to compare. The likelihood increases by almost 400 points ($k=5$) or even more than 2400 points ($k=15$), with one parameter less to be estimated. This underlines that the long-memory features also play an important role in explaining the volatility and correlation dynamics.

Looking at the individual parameter estimates, we first note the positive and significant long-memory coefficient d . A similar estimate of d is found by, for example, Proietti (2016) in the univariate case. The value of d is highly robust across the dimensions considered and indicates that autocorrelations only die out very slowly.

The high degree of persistence in the FIGAS model is mirrored by the other models. For example, for the HAR model the estimated values of $\beta_1 + \beta_2 + \beta_3$ are also close to 1. Similarly, the estimate of the autoregressive coefficient for the short-memory GAS model is very close to 1 ($\beta \simeq 0.99$), indicating a strong persistence. The value of B of the FIGAS model changes from significantly negative for $k=5$ to significantly positive for the case of all equities. Based on Proposition (1), we empirically satisfy the constraint for positive definiteness of the resulting covariance matrices V_t for both dimensions. Although the value of B is negative for $k=5$ and positive for $k=15$, both values imply a highly similar set of autocovariance functions; see Figure 3. If anything, the increase of the dimension leads to a slightly stronger long-memory feature.

The degrees of freedom parameter ν_2 is estimated at around 35 and 65 for 5 and 15 dimensions, respectively. Despite that the value of $\hat{\nu}_2$ may appear high, such values already result in a substantial moderation of the effect of incidentally large observations RK_t in Equation (15) through the matrix weighting scheme. Also fat-tailedness for these values of ν_2 is considerably larger than that of the Wishart distribution (see also Opschoor et al., 2018).

Figure 4 plots some of the fitted volatilities and correlations. We show the results for Alcoa (AA) and Caterpillar (CAT) for the FIGAS model and the HAR model of Equation

Figure 4 Estimated volatilities and correlations. This figure plots the estimated volatilities of AA and CAT (see Table 2) in the upper-left and lower-right panels, and the pairwise covariances and correlations in the upper-right and lower-left panels, respectively. Time varying parameter paths are estimated using the FIGAS model and HAR model. The estimates are based on the full sample, January 2, 2001 until December 31, 2012 (3017 observations).

(16). The figure shows remarkable differences between the two models for both the volatility and the covariances and correlations. Focusing first on the volatilities and covariances, the robust transition scheme based on matrix- F score dynamics produces considerably fewer spikes: spikes of the HAR model stand out much more clearly than the spikes of the FIGAS model. Notable differences are apparent for both companies during the periods 2001–2003, 2007–2008, and 2010–2011. The FIGAS framework is able to mitigate the impact of temporary RK_t on the estimates of V_t .

We conclude that already for short horizons there is evidence that a combination of long-memory and fat-tailedness improves model performance. We expect the long-memory properties to become even more important once we move into longer forecasting horizons, as we do next.

4.5 Out-of-Sample Results

In our out-of-sample analysis, we assess both the short-term and long-term forecasting performance of the FIGAS model. We consider h -step ahead forecasts, with $h = 1, 5, 10$, and 22. In addition, we consider aggregated covariance forecasts for the next one or two trading

weeks and for the next month, that is, $V_{t:t+h} = V_{t+1} + V_{t+2} + \dots + V_{t+h}$ with $h = 5, 10$, and 22. Similar to the in-sample analysis of the previous subsection, we compare the FIGAS model with the HEAVY model, the GAS model, and the multivariate HAR model.

We test the predictive ability of the different models based on the loss-differences of the QLIK loss function (21) using the test-statistic defined in Equation (23). We use a moving window of 1500 observations and re-estimate the parameters after each 25 observations (≈ 1 month). The first in-sample period corresponds to the period January 2001–December 2006, which is well before the financial crisis of October 2008. This forecasting experiment therefore constitutes a major robustness test for all the models considered.

Table 4 contains the results for the whole out-of-sample period (Panel A). We present the results for the five-dimensional case, as well as for all assets ($k = 15$). Negative *t*-test statistics (in parentheses) indicate that the FIGAS model performs better. The overall significant negative values in Panel A for horizons $h = 1, 5, 10, 22$ as well as for the aggregated forecasts clearly show that the FIGAS model statistically outperforms the HEAVY and HAR models. The short-memory GAS model still does quite well for horizons of (up to) $h = 5$ or 10, though worse than the FIGAS specification. Also here, however, the FIGAS model does significantly better for longer horizons such as $h = 22$ (for both dimensions) and $h = 1 : 22$ (for $k = 5$). The improvements due to the long memory features of the FIGAS specification thus appear particularly pronounced in cases where they matter most, namely at long horizon forecasting.

To further investigate where the performance of the fractionally integrated specification comes from, we split the sample in two periods: the Financial Crisis period (July 2007–December 2009) and the non-crisis period (December 2006–July 2007 and January 2010–December 2012). Panels B and C of Table 4 show that discriminating between the crisis and non-crisis periods provides us two additional insights. First, the differences between the average QLIK values between the FIGAS model and the HAR and HEAVY models increase during the crisis compared with the non-crisis period. This holds especially for $k = 15$. The crisis period is characterized by more spikes, and the FIGAS specification can better deal with these due to the score dynamics and the fat-tailed distributional assumptions. The FIGAS model, therefore, performs particularly well in the crisis period compared with its HAR and HEAVY counterparts.

Second, the FIGAS model and the short-memory GAS specification perform similarly well during the crisis period, but not during the non-crisis period. In non-crisis years, the FIGAS model statistically outperforms its short-memory counterpart. Hence, the fractionally integrated dynamics appear particularly valuable during calm periods, while accounting for fat-tailedness of returns and realized covariances are more important during turbulent years. To summarize, taking account of both fat-tailedness (during turbulent times) and long-memory effects (during calm periods) provides the FIGAS model with its superior forecasting performance over longer time spans.

Table 5 illustrates the economic significance of the covariance matrix forecasts by showing the mean of the *ex post* conditional portfolio standard deviation, computed by implementing the period-by-period *ex ante* minimum variance portfolio weights obtained from Equation (24). In addition, we show three *ex post* portfolio statistics such as the turnover, concentration, and short positions. Panel A displays results for the 1, 5, 10, and 22 step ahead predictions, while panel B corresponds to the cumulative forecasts with a window of 5, 10, or 22 trading days. Both panels show the average out-of-sample portfolio standard

Table 4 Test-statistics on predictive ability (QLIK criterion)

	1	5	10	22	1:5	1:10	1:22
Panel A: Full out-of-sample							
A.1: AA/BA/CAT/GE/KO							
FIGAS	8.04 (—)	8.41 (—)	8.67 (—)	9.11 (—)	16.30 (—)	19.93 (—)	24.15 (—)
GAS	8.06 (-2.6)	8.44 (-1.1)	8.75 (-1.7)	9.32 (-1.9)	16.32 (-1.1)	19.96 (-1.4)	24.24 (-1.9)
HEAVY	8.29 (-14.9)	8.70 (-8.2)	9.03 (-3.6)	9.72 (-2.3)	16.56 (-11.3)	20.21 (-6.6)	24.52 (-3.5)
HAR	8.06 (-1.9)	8.51 (-3.2)	8.84 (-2.2)	9.50 (-1.6)	19.73 (-34.1)	26.14 (-34.8)	33.77 (-26.6)
A.2: All assets ($k = 15$)							
FIGAS	19.05 (—)	20.01 (—)	20.75 (—)	21.86 (—)	43.74 (—)	54.58 (—)	67.17 (—)
GAS	19.12 (-3.7)	20.06 (-0.8)	20.89 (-1.1)	22.23 (-1.8)	43.78 (-1.2)	54.65 (-1.0)	67.32 (-1.4)
HEAVY	20.63 (-19.9)	21.66 (-9.7)	22.50 (-5.4)	24.00 (-2.9)	45.34 (-13.6)	56.23 (-8.7)	68.92 (-5.0)
HAR	19.41 (-6.8)	20.65 (-5.1)	21.71 (-3.1)	23.48 (-1.9)	53.15 (-32.6)	72.22 (-34.4)	95.06 (-27.7)
Panel B: Crisis period							
B.1: AA/BA/CAT/GE/KO							
FIGAS	10.82 (—)	11.37 (—)	11.84 (—)	12.69 (—)	19.17 (—)	22.91 (—)	27.33 (—)
GAS	10.83 (-0.7)	11.36 (0.3)	11.94 (-0.9)	12.94 (-1.1)	19.16 (0.4)	22.92 (-0.3)	27.41 (-0.9)
HEAVY	11.05 (-7.0)	11.72 (-4.6)	12.43 (-2.6)	13.84 (-1.9)	19.43 (-5.6)	23.27 (-3.8)	27.93 (-2.5)
HAR	10.84 (-1.1)	11.52 (-2.4)	12.18 (-2.0)	13.53 (-1.5)	22.21 (-16.4)	28.49 (-15.9)	35.89 (-11.4)
B.2: All assets ($k = 15$)							
FIGAS	27.65 (—)	29.11 (—)	30.42 (—)	32.71 (—)	52.57 (—)	63.73 (—)	76.91 (—)
GAS	27.70 (-1.1)	29.02 (0.8)	30.43 (0.0)	32.91 (-0.5)	52.53 (0.6)	63.68 (0.3)	76.92 (0.0)
HEAVY	29.26 (-9.5)	31.17 (-5.3)	32.89 (-3.3)	36.28 (-2.1)	54.40 (-6.8)	65.77 (-4.7)	79.42 (-3.2)
HAR	28.14 (-6.1)	30.16 (-4.2)	32.25 (-2.8)	36.12 (-1.8)	60.72 (-15.2)	79.37 (-15.7)	101.5 (-12.0)

(continued)

Table 4 Continued

	1	5	10	22	1:5	1:10	1:22
Panel C: Non-crisis period							
C.1: AA/BA/CAT/GE/KO							
FIGAS	6.01 (—)	6.25 (—)	6.35 (—)	6.50 (—)	14.21 (—)	17.75 (—)	21.83 (—)
GAS	6.03 (-3.5)	6.30 (-2.2)	6.42 (-2.6)	6.67 (-2.8)	14.24 (-2.2)	17.79 (-2.4)	21.91 (-2.6)
HEAVY	6.27 (-16.5)	6.50 (-9.7)	6.54 (-5.7)	6.70 (-3.2)	14.46 (-13.1)	17.97 (-9.4)	22.03 (-5.3)
HAR	6.02 (-1.8)	6.31 (-2.1)	6.39 (-1.2)	6.56 (-0.8)	17.92 (-35.8)	24.41 (-43.5)	32.22 (-45.2)
C.2: All assets ($k = 15$)							
FIGAS	12.76 (—)	13.36 (—)	13.67 (—)	13.93 (—)	37.28 (—)	47.89 (—)	60.04 (—)
GAS	12.85 (-6.3)	13.50 (-4.2)	13.92 (-3.6)	14.41 (-3.4)	37.38 (-4.6)	48.04 (-3.8)	60.30 (-3.3)
HEAVY	14.32 (-29.1)	14.71 (-20.7)	14.91 (-15.1)	15.01 (-9.3)	38.72 (-25.3)	49.25 (-19.6)	61.24 (-13.9)
HAR	13.03 (-3.9)	13.69 (-3.0)	13.99 (-1.6)	14.23 (-0.9)	47.62 (-35.5)	66.99 (-42.5)	90.32 (-45.8)

Notes: This table shows test statistics on superior predictive ability between the FIGAS model and the GAS, HEAVY, or multivariate HAR model, respectively, based on the QLIK loss function defined in Equation (21). The test is based based on 1, 5, 10, and 22-step ahead predictions of the covariance matrix, applied to 5 and 15 (all) equities. In addition, we consider cumulative forecasts with a window of 5, 10, and 22 trading days. Panel A presents results for full out-of-sample (December 2006–December 2012), Panel B describes the results for the Financial Crisis period (July 2007–December 2009), and Panel C shows the results for the non-crisis period (December 2006–July 2007 and January 2010–December 2012). The subpanels 1 and 2 correspond to the number of considered assets, that is, $k = 5$ and $k = 15$. We report the average QLIK loss for each model with the associated DM-type of test statistic in parentheses. A negative test statistic indicates superior predictive ability of the FIGAS model. We use a moving window of 1500 observations. The prediction period contains 1495 observations.

deviation, the associated DM test statistics vis-à-vis the FIGAS model (in parentheses), and the *ex post* mean of the portfolio turnover, concentration, and short positions. For all pairs of assets and all forecasting horizons considered, the FIGAS model produces the lowest *ex post* portfolio standard deviation. This result also holds for the aggregated forecasts. The reductions in standard deviations are statistically significant compared with all of the benchmarks.

The second parts in all subpanels (containing TO, CO, SP) show that overall the FIGAS model shows also the best results with respect to the portfolio statistics related to the practical implementation of the minimum variance strategy. First of all, all subpanels of Table 5 indicate that short positions are always lowest for the FIGAS model. This implies that on average the stable forecast of the covariance matrix implies less extreme weights.

Table 5 Ex post minimum variance portfolio standard deviations

	AA/BA/CAT/GE/KO				All assets ($k = 15$)			
	FIGAS	GAS	HEAVY	HAR	FIGAS	GAS	HEAVY	HAR
Panel A.1: 1 step ahead								
$\bar{\sigma}_p$	0.925	0.926	0.949	0.927	0.688	0.689	0.737	0.693
t -stat (DM)	(—)	(−3.1)	(−21.1)	(−2.8)	(—)	(−3.9)	(−34.9)	(−6.9)
TO	0.146	0.123	0.117	0.197	0.249	0.201	0.202	0.382
CO	0.723	0.723	0.797	0.740	0.483	0.486	0.578	0.517
SP	−0.025	−0.026	−0.094	−0.034	−0.173	−0.175	−0.337	−0.223
Panel A.2: Five step ahead								
$\bar{\sigma}_p$	0.933	0.937	0.956	0.935	0.700	0.703	0.746	0.705
t -stat (DM)	(—)	(−6.4)	(−12.9)	(−3.0)	(—)	(−5.1)	(−21.7)	(−5.4)
TO	0.084	0.112	0.102	0.117	0.162	0.189	0.162	0.246
CO	0.723	0.722	0.793	0.737	0.477	0.485	0.569	0.512
SP	−0.020	−0.024	−0.087	−0.028	−0.166	−0.174	−0.325	−0.217
Panel A.3: 10-step ahead								
$\bar{\sigma}_p$	0.938	0.942	0.959	0.940	0.707	0.711	0.750	0.712
t -stat (DM)	(—)	(−4.6)	(−9.9)	(−3.0)	(—)	(−4.4)	(−16.3)	(−4.3)
TO	0.066	0.102	0.093	0.093	0.128	0.175	0.142	0.193
CO	0.723	0.720	0.789	0.735	0.476	0.484	0.562	0.509
SP	−0.018	−0.022	−0.081	−0.026	−0.164	−0.174	−0.316	−0.215
Panel A.4: 22 step ahead								
$\bar{\sigma}_p$	0.946	0.952	0.962	0.948	0.718	0.723	0.755	0.723
t -stat (DM)	(—)	(−3.9)	(−6.6)	(−2.0)	(—)	(−3.9)	(−12.0)	(−3.5)
TO	0.050	0.083	0.074	0.067	0.097	0.150	0.117	0.138
CO	0.724	0.718	0.781	0.731	0.475	0.483	0.552	0.507
SP	−0.016	−0.018	−0.072	−0.023	−0.163	−0.174	−0.302	−0.215
Panel B.1: 5-step ahead (cumulative)								
$\bar{\sigma}_p$	2.121	2.126	2.172	2.124	1.589	1.593	1.693	1.597
t -stat (DM)	(—)	(−5.5)	(−14.6)	(−3.5)	(—)	(−4.9)	(−22.4)	(−6.1)
TO	0.107	0.117	0.110	0.150	0.199	0.196	0.181	0.294
CO	0.723	0.722	0.795	0.738	0.479	0.485	0.573	0.514
SP	−0.021	−0.025	−0.090	−0.031	−0.168	−0.175	−0.330	−0.218
Panel B.2: 10-step ahead (cumulative)								
$\bar{\sigma}_p$	3.032	3.043	3.101	3.038	2.282	2.289	2.425	2.294
t -stat (DM)	(—)	(−5.7)	(−11.4)	(−3.5)	(—)	(−4.7)	(−16.9)	(−4.9)
TO	0.090	0.112	0.105	0.127	0.172	0.191	0.170	0.255
CO	0.723	0.721	0.792	0.737	0.478	0.485	0.568	0.512
SP	−0.020	−0.024	−0.086	−0.029	−0.166	−0.174	−0.324	−0.217

(continued)

Table 5 Continued

	AA/BA/CAT/GE/KO				All assets ($k = 15$)			
	FIGAS	GAS	HEAVY	HAR	FIGAS	GAS	HEAVY	HAR
Panel B.3: 22-step ahead (cumulative)								
$\bar{\sigma}_p$	4.567	4.588	4.657	4.577	3.461	3.477	3.656	3.480
t -stat (DM)	(—)	(−4.8)	(−8.4)	(−3.2)	(—)	(−4.5)	(−12.0)	(−3.4)
TO	0.077	0.104	0.098	0.108	0.147	0.183	0.160	0.215
CO	0.723	0.720	0.788	0.735	0.476	0.484	0.561	0.509
SP	−0.018	−0.021	−0.080	−0.026	−0.164	−0.174	−0.315	−0.215

Notes: This table shows results on a GMVP, based on 1-, 5-, 10-, and 22-step ahead predictions of the covariance matrix, according to the FIGAS, GAS, HEAVY, and the multivariate HAR model, applied to 5 and 15 equities. Panel A shows the results corresponding with 1-, 5-, 10-, and 22-step ahead predictions, while Panel B reports results on the cumulative forecasts with a window of 5, 10, and 22 trading days. For each model, the table shows the *ex post* mean of the daily portfolio volatility, portfolio turnover (TO_t), concentration (CO_t), and the total number of short positions (SP_t), as defined in Equations (25)–(27). The number between parentheses shows the test-statistic on equal portfolio volatility between the FIGAS model and the HEAVY, GAS, or HAR model. The (absolute) lowest number value of all statistics across the models are marked bold. We use a moving window of 1500 observations. The prediction period runs from December, 2006–December, 2012 (1495 observations).

Differences are particularly apparent for higher dimensional settings ($k = 15$), where short positions can be about twice as large for instance for the HEAVY model. Second, the FIGAS model produces the lowest portfolio concentration in case $k = 15$ across all models. For $k = 5$ the FIGAS still has less concentration than the HEAVY and HAR models, but is sometimes outperformed by the GAS specification, though only slightly so. Third, the portfolio turnover of the FIGAS models is the lowest across all models when the horizon increases. For example, for 10 and 22 step ahead point forecasts and the cumulative forecast of 22 trading days, the FIGAS portfolio turnover is considerably lower than the turnover of the benchmarks. All in all, the FIGAS specification also appears to outperform its benchmarks in economic terms.

5 Conclusions

We introduced a new multivariate fractionally integrated model with score-driven volatility dynamics (FIGAS) for matrix-variate realized covariance matrix observations. The proposed model explicitly acknowledges that realized (co)variances display long-memory behavior. It does so in a way that ensures positive definiteness of the covariance matrices by simple parameter restrictions in the model. In addition, the model takes into account that realized covariance matrices are typically fat-tailed. The score-driven matrix-valued dynamics automatically correct for influential observations in the realized covariances.

For a number of S&P500 equity returns over the period 2001–2012, we showed that both in-sample and out-of-sample and both statistically and economically the new model outperformed strong recent competitors such as the HEAVY model of Noureldin, Shephard, and Sheppard (2012) and the multivariate HAR model of Corsi (2009) and

Chiriac and Voev (2011). Interestingly, the fractionally integrated dynamics appear particularly valuable during calm periods. The outlier robust features of the model due to the score dynamics and the fat-tailed distributional assumptions, by contrast, are most useful during turbulent times. Combining the two, the FIGAS model shows the best overall performance over the entire sample.

Supplementary Data

Supplementary Data are available at *Journal of Financial Econometrics* online.

References

- Andersen T., T. Bollerslev, F. X. Diebold and P. Labys. 2001. The Distribution of Realized Exchange Rate Volatility. *Journal of the American Statistical Association* 96: 42–55.
- Asai M., M. McAleer and J. Yu. 2006. Multivariate Stochastic Volatility: A Review. *Econometric Reviews* 25: 145–175.
- Baillie R. T., T. Bollerslev and H. O. Mikkelsen. 1996. Forecasting Multivariate Realized Stock Market Volatility. *Journal of Econometrics* 74: 3–50.
- Barndorff-Nielsen O. E., and N. Shephard. 2002. Econometric Analysis of Realised Volatility and Its Use in Estimating Stochastic Volatility Models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64: 253–280.
- Barndorff-Nielsen O. E., P. R. Hansen, A. Lunde and N. Shephard. 2008. Designing Realised Kernels to Measure the Ex-Post Variation of Equity Prices in the Presence of Noise. *Econometrica* 76: 1481–1536.
- Barndorff-Nielsen O. E., P. R. Hansen, A. Lunde and N. Shephard. 2009. Realized Kernels in Practice: Trades and Quotes. *Econometrics Journal* 12: 1–32.
- Barndorff-Nielsen O. E., P. R. Hansen, A. Lunde and N. Shephard. 2011. Multivariate Realized Kernels: Consistent Positive Semi-Definite Estimators of the Covariation of Equity Prices with Noise and Non-Synchronous Trading. *Journal of Econometrics* 12: 1–32.
- Bauer G. H., and K. Vorkink. 2011. Forecasting Multivariate Realized Stock Market Volatility. *Journal of Econometrics* 160: 93–101.
- Blasques C., S. J. Koopman and A. Lucas. 2015. Information Theoretic Optimality of Observation Driven Time Series Models for Continuous Responses. *Biometrika* 102: 325–343.
- Bollerslev T.. 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics* 31: 307–327.
- Bollerslev T., A. J. Patton and R. Quaadvlieg. 2018. Modeling and Forecasting (Un)reliable Realized Covariances for More Reliable Financial Decisions. *Journal of Econometrics* 207: 71–91.
- Bollerslev T., and H. O. Mikkelsen. 1996. Modeling and Pricing Long Memory in Stock Market Volatility. *Journal of Econometrics* 73: 151–184.
- Brownlees C. T., and G. M. Gallo. 2006. Financial Econometric Analysis at Ultra-High Frequency: Data Handling Concerns. *Computational Statistics and Data Analysis* 51: 2232–2245.
- Chiriac R., and V. Voev. 2011. Modelling and Forecasting Multivariate Realized Volatility. *Journal of Applied Econometrics* 26: 922–947.
- Conrad C., and B. R. Haag. 2006. Inequality Constraints in the Fractionally Integrated GARCH Model. *Journal of Financial Econometrics* 4: 413–449.
- Corsi F. 2009. A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics* 7: 174–196.

- Cox D. R. 1981. Statistical Analysis of Time Series: Some Recent Developments. *Scandinavian Journal of Statistics* 8: 93–115.
- Creal D., S. J. Koopman and A. Lucas. 2011. A Dynamic Multivariate Heavy-Tailed Model for Time-Varying Volatilities and Correlations. *Journal of Business and Economic Statistics* 29: 552–563.
- Creal D., S. J. Koopman and A. Lucas. 2013. Generalized Autoregressive Score Models with Applications. *Journal of Applied Econometrics* 28: 777–795.
- Diebold F. X., and R. S. Mariano. 1995. Comparing Predictive Accuracy. *Journal of Business and Economic Statistics* 13: 253–263.
- Engle R., and B. Kelly. 2012. Dynamic Equicorrelation. *Journal of Business & Economic Statistics* 30: 212–228.
- Engle R. F., and G. M. Gallo. 2006. A Multiple Indicators Model for Volatility Using Intra-daily Data. *Journal of Econometrics* 131: 3–27.
- Giacomini R., and H. White. 2006. Tests of Conditional Predictive Ability. *Econometrica* 74: 1545–1578.
- Golosnoy V., B. Gribisch and R. Liesenfeld. 2012. The Conditional Autoregressive Wishart Model for Multivariate Stock Market Volatility. *Journal of Econometrics* 167: 211–223.
- Gourieroux C., J. Jasiak and R. Sufana. 2009. The Wishart Autoregressive Process of Multivariate Stochastic Volatility. *Journal of Econometrics* 150: 167–181.
- Harvey A. C. 2013. *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. New York: Cambridge University Press.
- Janus P., S. J. Koopman and A. Lucas. 2014. Long Memory Dynamics for Multivariate Dependence under Heavy Tails. *Journal of Empirical Finance* 29: 187–206.
- Karanasos M., Z. Psaradakis and M. Sola. 2004. On the Autocorrelation Properties of Long-Memory GARCH Processes. *Journal of Time Series Analysis* 25: 265–282.
- Konno Y. 1991. A Note on Estimating Eigenvalues of Scale Matrix of the Multivariate F-Distribution. *Annals of the Institute of Statistical Mathematics* 43: 157–165.
- Koopman S. J., B. Jungbacker and E. Hol. 2005. Forecasting Daily Variability of the S&P 100 Stock Index Using Historical, Realised and Implied Volatility Measurements. *Journal of Empirical Finance* 12: 445–475.
- Laurent S., J. Rombouts and F. Violante. 2013. On Loss Functions and Ranking Forecasting Performances of Multivariate Volatility Models. *Journal of Econometrics* 173: 1–10.
- Markowitz H. 1952. Portfolio Selection. *Journal of Finance* 7: 77–91.
- Noureldin D., N. Shephard and K. Sheppard. 2012. Multivariate High-Frequency-Based Volatility (HEAVY) Models. *Journal of Applied Econometrics* 27: 907–933.
- Opschoor A., P. Janus, A. Lucas and D. van Dijk. 2018. New HEAVY Models for Fat-Tailed Realized Covariances and Returns. *Journal of Business and Economic Statistics*, 643–657.
- Proietti T. 2016. Component-Wise Representations of Long-Memory Models and Volatility Prediction. *Journal of Financial Econometrics* 14: 668–692.
- Shephard N., and K. Sheppard. 2010. Realising the Future: Forecasting with High-Frequency-Based Volatility (HEAVY) Models. *Journal of Applied Econometrics* 25: 197–231.
- Silvennoinen A., and T. Teräsvirta. 2009. “Multivariate GARCH Models.” In T.G. Andersen, A. Davis, J.P. Kreib, and T. Mikosch. (eds.), *Handbook of Financial Time Series*, pp. 201–229. Berlin: Springer-Verlag.
- Zumbach G. 2006. “The RiskMetrics 2006 methodology.” Working paper.