2016/02

A Time-Varying Long Run HEAVY Model

Manuela Braione

50 YEARS OF
CORE
DISCUSSION PAPERS

CORE
Voie du Roman Pays 34, L1.03.01
Tel (32 10) 47 43 04
Fax (32 10) 47 43 01
Email: immaq-library@uclouvain.be
http://www.uclouvain.be/en-44508.html

A time-varying long run HEAVY model

Manuela Braione

Université catholique de Louvain, CORE, B-1348 Louvain-La-Neuve, Belgium
manuela.braione@uclouvain.be

February 1, 2016

Abstract

We propose a scalar variation of the multivariate HEAVY model of Noureldin et al. [1] which allows for a time-varying long run component in the specification of the daily conditional covariance matrix. Differently from the original model featuring a BEKK-type parameterization, ours extends it to allow for a separate modeling of the conditional volatilities and the conditional correlation matrix, in a DCC fashion. Estimation is performed in one step by QML and multi-step ahead forecasting is feasible applying the direct approach to the HEAVY-P equation. In an empirical application aiming at modeling and forecasting the conditional covariance matrix of a stock (BAC) and an index (S&P 500), we find that the new model statistically outperforms the original HEAVY model both in-sample and out-of-sample.

Keywords: HEAVY model, Long term models, Mixed Data Sampling, Direct forecasting.

1. Introduction

There is a vast consensus among practitioners that inclusion of high frequency information enables the development of more accurate forecasting models for the conditional covariance of daily returns. An outstanding example is represented by the class of multivariate High-frEquency-bAsed VolatilitY (HEAVY) models introduced by Noureldin et al. [1], which links the dynamics of the conditional covariance matrix to the realized measure using a system of two equations akin to the multivariate BEKK specification. The model has several advantages, the main ones that it is easy to estimate by MLE and able to provide closed-form forecasting formulas. Nevertheless, when the scalar version of the model is employed with *targeting* (as is often the case in financial applications), the conditional covariance dynamics are driven by only two parameters, thus strongly penalizing the flexibility of the model in times of significantly changing economic conditions. For this reason, the authors raise the interesting question of whether a more sophisticated parameterization could improve the forecasting ability of the model. We address this question by studying a new model, the Time Varying Long Run (TVLR) HEAVY, which extends the basic HEAVY specification to a component structure that decomposes the conditional covariance matrix into long-run (permanent) and short-run (transitory) components in a multivariate fashion, similarly to the approach adopted by Golosnoy et al. [2] and Bauwens et al. [3]. We model the trend component via a parametric Mixed Data Sampling (MIDAS)-type of filter and allow the short term dynamics to move according to a DCC specification, thus stepping away from the basic linear BEKK recursion. We compare the TVLR-HEAVY against the standard HEAVY model from both an in- and out-of-sample perspective, where in the latter case multi-step ahead forecasts are constructed using the *direct* approach which overcomes the difficulties created by the nonlinear structure of the model. In this way, the TVLR-HEAVY can still be feasibly estimated by MLE, thus keeping computational tractability in practical application.

Our set of results shows that introducing an additional component that captures the secular movements in the (co)volatility dynamics is well justified, as the new model is found to improve over the existing one both in the overall fit and predictive accuracy. Particularly, the forecast gains tend to be more pronounced at longer forecast horizons, when the impact of the time-varying trend component appears to be predominant.

The structure of the paper is as follows. Section 2 briefly recalls the multivariate general framework and formally introduces the new model and its estimation approach. Section 3 illustrates the aim of the empirical application and presents the results of both the in- and out-of-sample analysis. Section 4 concludes the paper with some final remarks.

2. General framework

Let $\boldsymbol{r}_t$ denote the $(n \times 1)$ vector of daily returns at time t and $P_t = \boldsymbol{r}_t\boldsymbol{r}_t'$ the $(n \times n)$ matrix obtained as the outer product of daily returns. The realized measure is denoted by C_t, and is a $(n \times n)$, symmetric and positive definite (PD) matrix. As in the original multivariate HEAVY paper, we use the realized covariance (RC) estimator obtained by summing up intra-daily returns at the 5 minute frequency, although any other consistent estimator could be used.

Conditionally on past information $\Im_{t-1}$ consisting of C_τ for $\tau \leq t-1$, C_t is assumed to follow a n-dimensional central Wishart distribution, i.e. $C_t \sim W_n(\nu, S_t/\nu)$ with $\nu > (n-1)$, while $P_t \sim SINGW_n(1, H_t)$, where $SINGW_n$ denotes a n-dimensional Singular Wishart distribution, given the assumption that $\boldsymbol{r}_t = H_t^{1/2}\boldsymbol{\epsilon}_t$ with $\boldsymbol{\epsilon}_t \sim N(\boldsymbol{0}, I_n)$. As already stressed in the paper by Noureldin et al. [1], the distinction between the Wishart and Singular Wishart densities is of no consequence to QML estimation.

Therefore, for the properties of the Wishart distribution, we have that

$$E(P_t|\Im_{t-1}) = E(\boldsymbol{r}_t\boldsymbol{r}_t'|\Im_{t-1}) = H_t \qquad (1)$$
$$E(C_t|\Im_{t-1}) = S_t \qquad (2)$$

where the PD matrices S_t and H_t are the conditional expectation of the realized measure and of the outer product of daily returns, respectively. Note that they both condition on the same high frequency information, hence they are assumed to be $\Im_{t-1}$ measurable.

The HEAVY model links the dynamics of H_t to the realized measure and is based on a system of two equations for H_t and S_t both akin to the multivariate BEKK specification.

Similarly to Noureldin et al. [1], we will refer to these equations as HEAVY-H and HEAVY-S, unless otherwise stated. Restricting to the scalar case, they are written as follows:

$$H_t = \Omega_H + \alpha_H^2 C_{t-1} + \beta_H^2 H_{t-1} \quad (3)$$
$$S_t = \Omega_S + \alpha_S^2 C_{t-1} + \beta_S^2 S_{t-1} \quad (4)$$

where the covariance stationarity condition requires that $\{\alpha, \beta\} > 0$ and $\{\alpha + \beta\} < 1$. In this case, the model can be expressed in its covariance targeting parameterization, where the intercept matrices Ω_H and Ω_S are expressed in terms of the unconditional first moments of H_t and S_t and model parameters, i.e. $B_H := E(P_t) = (1 - \alpha_H^2 - \beta_H^2)^{-1}\Omega_H$ and $B_S := E(C_t) = (1 - \alpha_S^2 - \beta_S^2)^{-1}\Omega_S$.

Equation (4) is not needed for computing one-step ahead forecasts of H_t, but is needed to achieve analytical multi-period ahead predictions due to the presence of C_{t-1} in Eq. (3).

As we will show in a moment, the TVLR-HEAVY model features a nonlinear parameterization for H_t due to the presence of the DCC structure that creates problems in constructing closed-form expressions for multi-step predictions. This issue can be solved by applying the *direct* forecast approach to the HEAVY-H equation, a method that has been extensively used in the empirical finance literature as an alternative to the *iterated* one, see for example Marcellino et al. [4], Ghysels et al. [5] and Proietti [6]. It entails to estimate a horizon-specific model of the (co)volatility, say weekly or monthly, which can then be used to form direct predictions over the next week or month. As only observed data are utilized to predict future periods, it is thought to yield reliable results. In this way, only a one dimensional system is needed to achieve direct multi-step ahead predictions of H_t, as those of C_t are directly taken into account in the same equation. We elaborate on this point in the following subsection, which formally introduces the proposed model and its estimation approach.

2.1. The model

The TVLR-HEAVY model features a multiplicative decomposition of the conditional covariance matrix of returns H_t into a secular component $M_t = G_t G_t'$ and a short term component $H_t^\star$, as follows:

$$H_t = G_t H_t^\star G_t' \quad (5)$$

with $H_t^\star$ a $(n \times n)$ PD matrix and G_t a lower triangular matrix obtained as a Cholesky factorization of M_t. M_t captures the long term movements around which (co)volatilities fluctuate from day to day while $H_t^\star$ represents the transitory component of the covariance dynamics. In order to identify the model, we impose $E(H_t^\star) = I_n$, with I_n the $(n \times n)$ identity matrix, as otherwise the two components could be interchangeable. This restriction allows the interpretation of $H_t^\star$ as an autocorrelated disturbance with respect to the long term level M_t.

As already mentioned, we focus on a horizon-specific model, for which horizons equal to $h = 1, 5, 10, 22$ days are considered.

Inspired by the recent work of Bauwens et al. [3], the secular component is specified parametrically and modeled using a Mixed Data Sampling (MIDAS)-type filter driven by a weighted sum of lagged realized covariance matrices over a long horizon of K[1] days:

$$M_t = \bar{\Lambda} + \theta \sum_{k=0}^{K} \phi_k(\omega) C_{t-h-k}, \tag{6}$$

where $\bar{\Lambda}$ is a symmetric matrix of constant parameters, θ is a positive scalar and $\phi_k(\cdot)$ is a weight function parametrized according to the restricted Beta polynomial.[2] The scalar parameter ω dictates the shape of the function and is constrained to be larger than one in order to achieve a time-decaying pattern of the weights, or in other words, to favor more recent over older observations. For identification it also holds $\sum_{k=1}^{K} \phi_k(\omega) = 1$. As it is specified, the long run volatility component is allowed to change over time as long as θ is positive, as in the case $\theta = 0$ it will be time invariant and limited to the constant intercept matrix $\bar{\Lambda}$. The parameterization in Eq. (6) guarantees that M_t (and consequently, H_t) is positive definite for all t assuming that $\bar{\Lambda}$ is a full rank matrix, which can be achieved estimating $\bar{\Lambda} = \Lambda\Lambda'$, where Λ is a lower triangular matrix with $(n(n+1)/2)$ free parameters.

The dynamics of the short term component $H_t^\star$ is modeled according to a scalar DCC parametrization that enables for a higher degree of flexibility and a more challenging structure than the corresponding scalar BEKK. Namely, $H_t^\star$ is further decomposed in the product of the diagonal matrix of short term conditional standard deviations $D_t^\star = \text{diag}\{H_t^\star\}^{1/2}$, and the PD short term conditional correlation matrix $R_t^\star$, such that $H_t = G_t D_t^\star R_t^\star D_t^\star G_t'$.

Assuming no spillover terms across the univariate series and a lag-one structure by ease of exposition, for each asset $i = 1, ..., n$ the volatility process is defined as follows:

$$H_{ii,t}^\star = (1 - \gamma_i - \delta_i) + \gamma_i \frac{C_{ii,t-h}}{M_{ii,t-h}} + \delta_i H_{ii,t-h}^\star, \tag{7}$$

where $\{\gamma_i, \delta_i\} > 0$ for every i. Thus, each short-term volatility component mean reverts to unity at a geometric rate of $\{\gamma + \delta\}$ if $0 < \{\gamma + \delta\} < 1$. Note that in Eq. (7) we use as regressor the asset specific short-term realized variance, i.e. $C_{ii,t-h}/M_{ii,t-h}$, as we found it to be a more precise factor to drive the volatility dynamics compared to $r_{i,t-h}^2/M_{ii,t-h}$.

At this stage the vector of standardized residuals is obtained as $\boldsymbol{\epsilon}_t = (D_t^\star G_t)^{-1} \boldsymbol{r}_t$, for which it holds

$$\begin{aligned} E\left[\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t' \middle| \Im_{t-h}\right] &= E_{t-h}\left[(D_t^\star G_t)^{-1} \boldsymbol{r}_t \boldsymbol{r}_t' (D_t^\star G_t)'^{-1}\right] \quad (8) \\ &= E\left[(D_t^\star G_t)^{-1} E_{t-h}\left(\boldsymbol{r}_t \boldsymbol{r}_t'\right) (D_t^\star G_t)'^{-1}\right] \\ &= D_t^{\star -1} G_t^{-1} G_t D_t^\star R_t^\star D_t^\star G_t' G_t'^{-1} D_t^{\star -1} \\ &= R_t^\star. \quad (9) \end{aligned}$$

[1]The number of K lags spanned in the MIDAS specification is set equal to 260 in order to minimize the trade-off between the highest in-sample likelihood value and the number of observations lost to initialize the filter. For a discussion about alternative MIDAS schemes we refer to Bauwens et al. [3].

[2]The Restricted Beta function is defined as $\phi_k(\omega) = \frac{\left(1-\frac{k}{K}\right)^{\omega-1}}{\sum_{j=1}^{K}\left(1-\frac{j}{K}\right)^{\omega-1}}$.

The dynamic equations for the conditional correlation matrix can finally be defined as follows

$$Q_t^\star = (1-\alpha-\beta)I_n + \alpha(\boldsymbol{\epsilon}_{t-h}\boldsymbol{\epsilon}_{t-h}') + \beta Q_{t-h}^\star, \tag{10}$$

$$R_t^\star = diag\{Q_t^\star\}^{-1/2} Q_t^\star diag\{Q_t^\star\}^{-1/2} \tag{11}$$

where, as before, $\{\alpha, \beta\}$ are positive scalars such that $\{\alpha+\beta\} < 1$ and I_n is the mean reverting $(n \times n)$ identity matrix. Due to the use of the outer product of standardized residuals, regularization of $Q_t^\star$ through Eq. (11) is necessary to obtain a well defined correlation matrix.

The TVLR-HEAVY model is parameterized with a finite-dimensional parameter vector $\boldsymbol{\phi} = \{vech(\Lambda)', \theta, \omega, \boldsymbol{\gamma}, \boldsymbol{\delta}, \alpha, \beta\}'$, where *vech* is the operator that stacks the lower triangular part including the diagonal of a matrix into a $(n(n+1)/2 \times 1)$ vector and $\boldsymbol{\gamma} = (\gamma_1, ...\gamma_n)'$, $\boldsymbol{\delta} = (\delta_1, ...\delta_n)'$ are $(n \times 1)$ vectors.

For reasonably small cross sectional systems, estimation can be performed by maximum likelihood (ML) in one step.[3] Given the Wishart assumption made on P_t, the log-likelihood function for T observations $\ell_T(\boldsymbol{\phi})$, net of constant terms, is expressed as follows:

$$\ell_T(\boldsymbol{\phi}) = -\frac{\nu}{2}\sum_{t=1}^{T}\left[log|H_t| + tr(H_t^{-1}P_t)\right]. \tag{12}$$

Note that the last two terms on the right hand side are linear in the parameter ν. Hence, the first order conditions for the estimation of the parameter vector do not depend on ν, implying that the shape parameter is of no consequence when estimating $\boldsymbol{\phi}$ by MLE. Furthermore, the estimator based on the maximization of the (Singular) Wishart log-likelihood function maintains a QML interpretation, i.e. if the conditional expectation of P_t is correctly specified, the score of the log-likelihood function in Eq. (12), evaluated at the true value of the parameters $\boldsymbol{\phi}_0$, is a martingale difference sequence (MDS) (see also Noureldin et al. [1]).

3. Empirical example

Our empirical application aims at exemplifying the advantages of the proposed model over the existing benchmarks from both an in- and out-of-sample perspective. As in Noureldin et al. [1], we focus on the modeling and forecasting of the conditional covariance matrix of a stock (BAC) and an index (S&P 500) using the scalar HEAVY model (with and without the targeting) and the TVLR-HEAVY. Figure 1 displays the annualized daily returns and realized volatility of the series over the period 1 February 2001 to 31 December 2009 (2242 daily observations) while Table 1 reports some useful summary statistics.

[3] The scalar parameterization adopted for the model serves the purpose of mitigating the proliferation of parameters typical of multivariate models. Nevertheless, for large cross-sections, the major source of burden in the TVLR-HEAVY is mostly represented by the $(n \times n)$ intercept matrix in the MIDAS filter, whose dimension grows at order $O(n^2)$. A similar problem has been recently discussed and addressed in Bauwens et al. [7].

Figure 1: SPY and BAC annualized daily returns and realized volatility over the sample period.

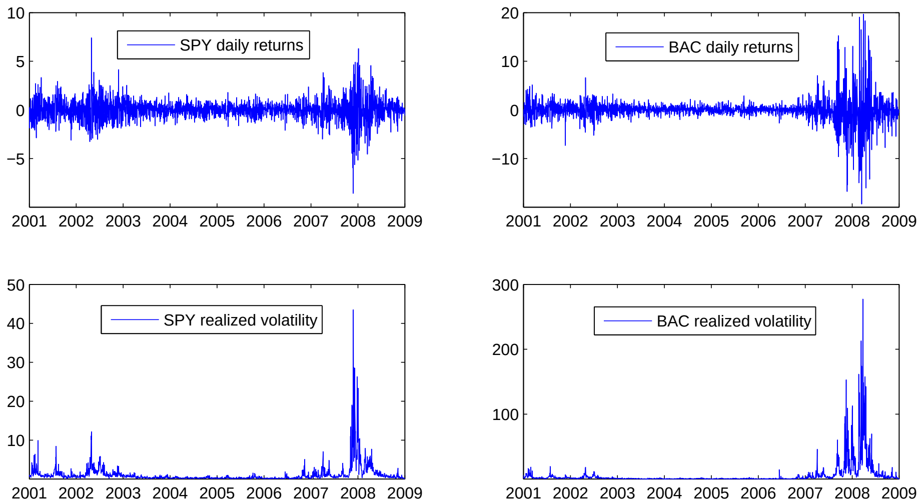

Table 1: Descriptive Statistics over the period 1 February 2001 to 31 December 2009 (2242 observations)

	Mean	Max.	Min.	Std.dev.	Skewness	Kurtosis
Panel A: Daily returns						
SPY	0.000	7.414	-8.574	1.050	-0.123	9.699
BAC	-0.000	19.687	-19.356	2.390	0.328	21.811
Panel B: Realized volatility						
SPY	1.123	43.463	0.041	2.248	8.104	102.035
BAC	5.458	277.308	0.074	16.811	7.178	72.570

3.1. Full sample results

Table 2 reports full sample estimation results for the three models. We also present parameter estimates of the HEAVY-S equations as they are used to compute multi-step ahead predictions of H_t in the HEAVY and HEAVY-CT (i.e. with covariance targeting) cases. At the outset, it should be noted that the TVLR-HEAVY, despite being more heavily parameterized, outperforms both competitors displaying lower AIC values which reflect an increase in the overall fit of the model. Parameter estimates for the HEAVY-CT and the HEAVY are fairly similar (with the latter slightly outperforming the former) and close to one in sum, thus suggesting a somewhat high level of persistence. This effect is mitigated for the TVLR-HEAVY ($\alpha + \beta = 0.91$) due to the role played by the additional (co)volatility component. The estimates of the MIDAS filter are indeed both significant and the relatively low value of the coefficient ω suggests a slow decaying pattern of the weights that leads to a pretty smooth temporal dynamics of the long term component.

Table 2: Panel A reports full sample parameter estimates and corresponding robust standard errors in brackets. HEAVY-CT and HEAVY refers to the scalar HEAVY model in (3)-(4) estimated with and without the targeting, respectively. Estimated intercept matrices are reported in Panel B, where the superscript * denotes significance at the 1% level. The sample is February 1, 2001 until December 31, 2009 (2242 observations). The Akaike information criterion in Panel C is computed as $AIC = -2Loglik + 2Np$, where Np is the number of model parameters.

	HEAVY-CT		**HEAVY**		**TVLR-HEAVY**	
	HEAVY-H	**HEAVY-S**	**HEAVY-H**	**HEAVY-S**	**HEAVY-H**	
Panel A: Estimated parameters						
α	0.277 (0.084)	0.399 (0.031)	0.241 (0.060)	0.424 (0.034)	0.033 (0.039)	
β	0.723 (0.085)	0.596 (0.032)	0.692 (0.077)	0.570 (0.034)	0.886 (0.206)	
θ					0.429 (0.068)	
ω					6.185 (2.027)	
γ_1					0.192 (0.050)	
γ_2					0.305 (0.066)	
δ_1					0.742 (0.072)	
δ_2					0.341 (0.122)	
Panel B: Unconditional covariance and estimated intercepts						
	$E(P_t)$		Ω_H		$\bar{\Lambda}$	
	1.105	1.692	0.033^*	0.035^*	0.240^*	0.237^*
	1.692	6.130	0.035^*	0.037	0.237^*	0.234
Panel C: Model selection criteria						
Loglik	–1161	–1411	–1110	–1409	–1079	
AIC	2326	2826	2231	2829	2180	

Figure 2 displays the innovations to the daily return of SPY and BAC[4] which appear to be centered around the identity matrix and, from unreported Ljung-box test results, not autocorrelated up to 100 lags.

3.2. Forecasting comparison

We assess the forecasting performance of the TVLR-HEAVY model by considering 1, 5, 10 and 22-step ahead forecasts of the covariance matrix. The predictions for the TVLR-HEAVY model are constructed using the direct approach mentioned in Section 2 while those for the benchmark models are based on the analytical formulas given by Noureldin et al. [1] (see Section 2.4 in their paper). The comparison is based on some consistent symmetric and asymmetric loss functions commonly used in practical applications (see Table 3 for their definition), for which we report averaged values over the forecasting period. Also, we evaluate the significance of the loss function differences by implementing the Model Confidence Set (MCS) approach of Hansen et al. [8], which identifies the single model or the set of models having the best forecasting performance at a given confidence

[4] They are obtained as first and second series of the vector $\boldsymbol{\epsilon}_t = \hat{H}_t^{-1/2}\boldsymbol{r}_t$.

Figure 2: SPY and BAC residuals

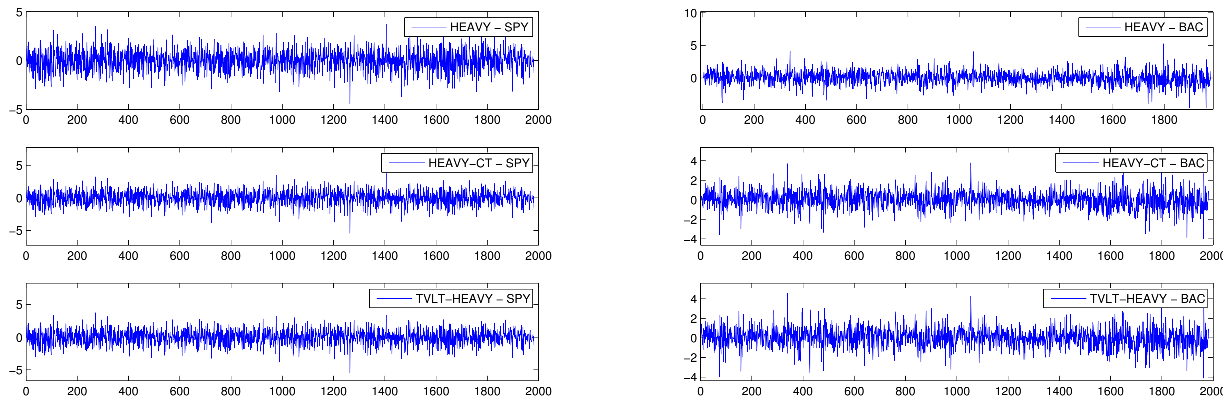

level.[5] Forecasts are obtained using an expanding-window scheme; specifically, the first 1942 trading days are taken as an in-sample period to estimate the model parameters and then each following estimation window is increased by one observation. Hence the out-of-sample period starts just before the heat of the financial crisis (October 2008) and covers the last part of the sample of roughly 300 observations. Table 4 shows the values of the loss functions over the forecasting period for the three models along with those included in the 75% MCS.

Results can be fairly summarized by horizon. The performance of the TVLR-HEAVY model is far from impressive if considering the first panel results (i.e. $h = 1$), as the HEAVY with covariance targeting is clearly the best performing model achieving the smallest values on the majority of the considered criteria. It also results as the only model included in the MCS. Nevertheless, as we move further in time the scenario drastically changes : the TVLR-HEAVY performs well and appears to deliver the most accurate predictions over longer horizons, with the biggest gains obtained for $h = 22$ in terms of RMSE and Stein reduction of respectively 10% and 17% with respect to the second best model. These findings are further confirmed by the dominance of the TVLR-HEAVY over the competitors in terms of frequency of occurrence in the MCS. This is somehow not surprising, considering that the (co)volatilities dynamics are driven by only two parameters in the HEAVY-CT ($n(n+1)/2+2$ in the HEAVY), while the TVLR-HEAVY allows for a major level of flexibility in the modeling of both the short and long term components. The effect of the latter on the overall model performance is especially valuable as the estimation interval increases: the bottom panel of Figure 3 shows the shift in value of the MIDAS parameters across the different horizons and suggests how the shape of the weight function adapts to the changing economic conditions (higher values of θ and lower values of ω imply a smoother

[5]We compute the MCS at the 75% confidence level, with block-length bootstrap parameter and number of bootstrap samples used to obtain the distribution under the null set equal to 2 and 10000, respectively.

Table 3: Implemented loss functions. Note: H_t is the predicted conditional covariance matrix while C_t is the realized measure; n denotes the number of assets.

		Matrix loss function
ST	Stein	$\mathrm{tr}(H_t^{-1}C_t) - \log \lvert H_t^{-1}C_t \rvert - n$
vND	von Neumann Divergence	$\mathrm{tr}\left(C_t \log C_t - C_t \log H_t - C_t + H_t\right)$
QL	QLIKE	$\log \lvert H_t \rvert + \mathrm{tr}\left(H_t^{-1}C_t\right)$
RMSE	Root mean square error	$\lvert\lvert C_t - H_t \rvert\rvert^{1/2}$

decaying pattern of the weights and thus a higher impact of the trend component on the estimated conditional covariance).

In light of this set of results, it appears that the introduction of an additional component that captures the secular movements in the (co)volatility dynamics is well justified. Indeed, even using a 'naive' forecasting approach as the direct one implemented in this paper, the improvements over existing benchmarks not accounting for time varying long term dynamics can be rather substantial.

Table 4: Out of sample evaluation. Panel A reports averaged values of the loss functions listed in Table 3 over the forecasting period, with the winning model for each horizon highlighted in bold; bottom panel entries are the p-values of the 75% MCS (included models in bold).

	Horizon 1			Horizon 5			Horizon 10			Horizon 22		
	HEAVY-CT	HEAVY	LT-HEAVY	HEAVY-CT	HEAVY	LT-HEAVY	HEAVY-CT	HEAVY	LT-HEAVY	HEAVY-CT	HEAVY	LT-HEAVY
Panel A: Loss function evaluation												
ST	**0.367**	0.455	0.475	0.760	0.850	**0.693**	1.029	1.199	**0.785**	1.230	1.512	**1.016**
vND	**6.017**	6.913	6.753	10.750	11.572	**9.945**	13.077	13.751	**11.584**	**17.669**	18.292	18.960
QL	**4.290**	4.378	4.398	4.643	4.734	**4.576**	4.888	5.058	**4.644**	4.978	5.260	**4.764**
RMSE	**2.381**	2.517	2.532	3.512	3.405	**3.257**	3.599	**3.469**	3.924	5.338	7.118	**4.799**
Panel B: 75% MCS p-values												
ST	**1.000**	0.000	0.003	**0.333**	**0.333**	**1.000**	0.035	0.035	**1.000**	0.042	0.042	**1.000**
vND	**1.000**	0.114	0.114	**0.433**	**0.433**	**1.000**	**0.552**	**0.552**	**1.000**	**0.699**	**0.713**	**1.000**
QL	**1.000**	0.001	0.005	**0.328**	**0.328**	**1.000**	0.035	0.035	**1.000**	0.045	0.045	**1.000**
RMSE	**1.000**	0.000	0.001	0.000	0.001	**1.000**	0.000	**1.000**	0.000	0.000	0.000	**1.000**

4. Conclusion

This paper introduces a new multivariate HEAVY model that combines daily return observations and realized covariance matrices to estimate and forecast the underlying conditional covariance of asset returns. The proposed model explicitly accounts for different components that capture the short and long run movements in the (co)volatility dynamics, thus enabling for a higher degree of flexibility at the cost of a handful of additional parameters. Despite the DCC-type structure, estimation can be easily performed by MLE in one step. Furthermore, using a horizon specific parameterization, the model can be reduced to an unidimensional system which proves convenient to achieve multi-step ahead covariance forecasts. In an application to a bivariate dataset we showed that the model adequately

Figure 3: Upper panel: HEAVY-CT parameter evolution across re-estimations. Lower panel: parameters driving the long term MIDAS component in the TVLR-HEAVY model across re-estimations and horizons.

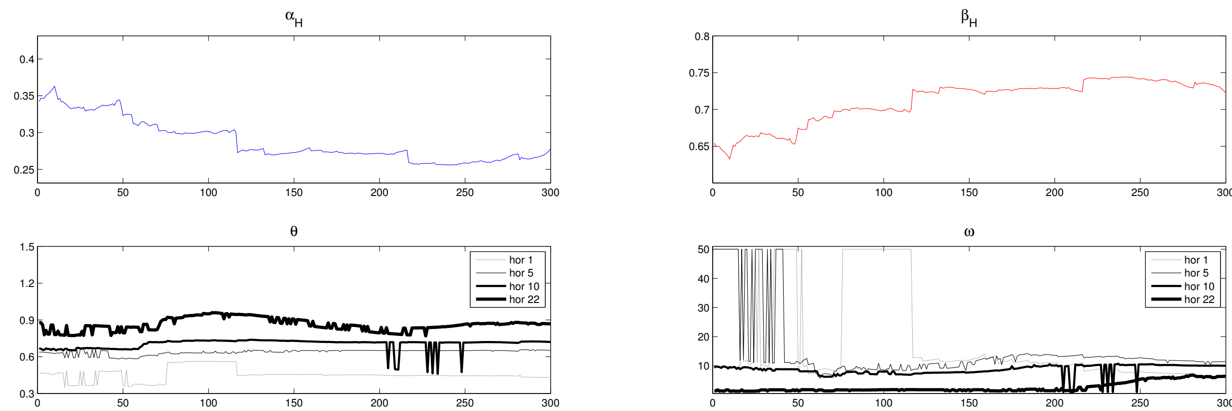

captures the (co)volatility dynamics over the sample period and that it substantially improves the out-of-sample fit of the covariance over longer horizons. The improvements include the period of the recent financial crisis of 2008-2012.

[1] D. Noureldin, N. Shephard, K. Sheppard, Multivariate High-Frequency-Based Volatility (HEAVY) Models, Journal of Applied Econometrics 27 (2012) 907–933.

[2] V. Golosnoy, B. Gribisch, R. Liesenfeld, The conditional autoregressive Wishart model for multivariate stock market volatility, Journal of Econometrics 167 (1) (2012) 211–223.

[3] L. Bauwens, M. Braione, G. Storti, Forecasting comparison of long term component dynamic models for realized covariance matrices, Annals of Economics and Statistics (Forthcoming) .

[4] M. Marcellino, J. H. Stock, M. W. Watson, A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series, Journal of econometrics 135 (1) (2006) 499–526.

[5] E. Ghysels, R. I. Valkanov, A. R. Serrano, Multi-period forecasts of volatility: Direct, iterated, and mixed-data approaches, in: EFA 2009 Bergen Meetings Paper, 2009.

[6] T. Proietti, Direct and iterated multistep AR methods for difference stationary processes, International Journal of Forecasting 27 (2) (2011) 266–280.

[7] L. Bauwens, M. Braione, G. Storti, A dynamic component model for forecasting high-dimensional realized covariance matrices Dises Working paper No. 3.234, University of Salerno.

[8] P. R. Hansen, A. Lunde, J. M. Nason, The model confidence set, Econometrica 79 (2) (2011) 453–497.