Choosing Between Weekly and Monthly Volatility Drivers Within a Double Asymmetric GARCH-MIDAS Model

Alessandra Amendola, Vincenzo Candila, and Giampiero M. Gallo

Abstract Volatility in financial markets has both low- and high-frequency components which determine its dynamic evolution. Previous modelling efforts in the GARCH context (e.g. the Spline-GARCH) were aimed at estimating the low-frequency component as a smooth function of time around which short-term dynamics evolves. Alternatively, recent literature has introduced the possibility of considering data sampled at different frequencies to estimate the influence of macro-variables on volatility. In this paper, we extend a recently developed model, here labelled Double Asymmetric GARCH-MIDAS model, where a market volatility variable (in our context, VIX) is inserted as a daily lagged variable, and monthly variations represent an additional channel through which market volatility can influence individual stocks. We want to convey the idea that such variations (separately) affect the short- and long-run components, possibly having a separate impact according to their sign.

Keywords Volatility · Asymmetry · GARCH-MIDAS · Forecasting

1 Introduction

Volatility modelling has been extensively studied: more than 30 years have gone by since the seminal contributions by [9, 14]. As they have about 25 K citations each (and some pertinent papers do not even mention them), it is clear that GARCH-type models are the standard among academicians and practitioners alike. These models

A. Amendola (✉)
Department of Economics and Statistics, University of Salerno, Via Giovanni Paolo II, Fisciano, Italy
e-mail: alamendola@unisa.it

V. Candila
MEMOTEF Department, Sapienza University of Rome, Rome, Italy
e-mail: vincenzo.candila@uniroma1.it

G. M. Gallo
Italian Court of Audits (Corte dei conti), and NYU in Florence, Florence, Italy
e-mail: giampiero.gallo@nyu.edu

M. La Rocca et al. (eds.), *Nonparametric Statistics*, Springer Proceedings in Mathematics & Statistics 339, https://doi.org/10.1007/978-3-030-57306-5_3

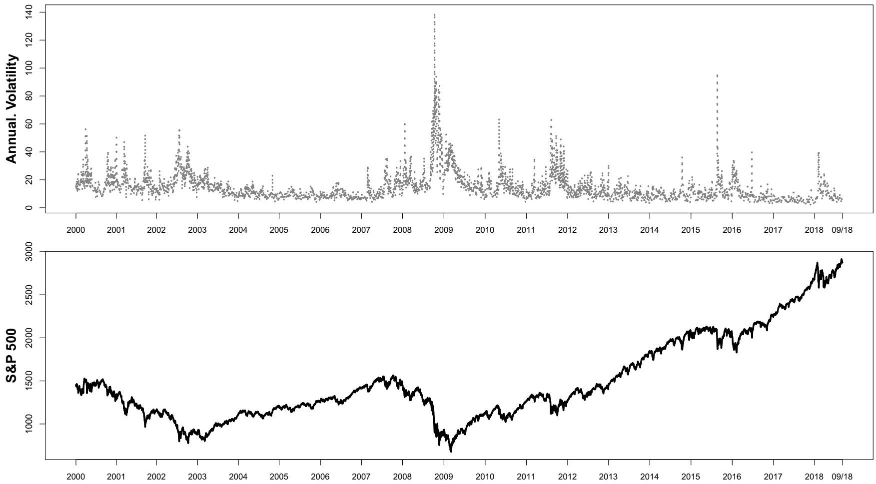

Fig. 1 S&P 500 Index and its realized volatility

build upon stylized facts of persistence in the conditional second moments (volatility clustering), an analysis made easier by the direct measurement of volatility, starting from the availability of ultra-high-frequency data (cf. [7]). Looking directly at the series of the Standard and Poor's (S&P) 500 Index and of its realized volatility, as illustrated in Fig. 1, one encounters two of such stylized facts in need of adequate modelling: the first is that volatility has a slow-moving/state-dependent *average local level* of volatility to be accounted for, and hence its dynamic evolution is driven by two components: a high-frequency and a low-frequency one. Another is that peaks of volatility are recorded in correspondence with streaks of downturns in the index, a sign of well-documented asymmetry in the dynamics.

Various suggestions exist in the literature to model the first of these two stylized facts: in a Markov Switching approach, GARCH parameters are state-dependent ([10, 13, 19], among others). The resulting high-frequency dynamics varies across states and evolves around a constant average level of volatility within states as a low-frequency component. In other contributions, the two components are additive; [12, 15] specify a model in which higher persistence is an identifying condition for the long-run component. The most popular GARCH specification is one in which long- and short-run components combine multiplicatively with the error term. Amado et al. [4] survey the contributions in this field: the common feature is that long run is a term which *smoothly* amplifies or dampens the short-run GARCH dynamics. The long-run term can be a deterministic function of time as in the Spline GARCH [16]; a logistic function of a forcing variable as in the Smooth Transition approach ([1–3], for instance); an exponential function of a one-sided filter of past values of a variable sampled at a lower frequency than the daily series of interest, as in the GARCH-MIDAS of [17]. In this paper, we take a modification of this latter model, called the Double Asymmetric GARCH-MIDAS (DAGM) introduced by [5], where a rate of change at a low frequency is allowed to have differentiated effects according to its sign, determining a local trend around which an asymmetric GARCH that describes the short-run dynamics. A market volatility variable (in our context, we choose the VIX index which is based on implied volatilities from at-the-money option prices) is inserted as a daily lagged variable, and monthly variations represent an additional channel through which market volatility can influence individual stocks.

The issue at stake in this empirically motivated paper is how this information about market-based volatility can help in shaping the MIDAS-GARCH dynamics. The idea we are pursuing is to illustrate

1. how a predetermined daily variable (in lagged levels) adds some significant influence to the short-run component (an asymmetric GARCH in the form of the GJR [18] model—this would be the first asymmetry considered); and, most importantly,
2. how the same variable observed at a lower frequency (in lagged first differences) can determine a useful combination (in the MIDAS sense seen in detail below) for the low-frequency component in the slowly moving level of local average volatility. In particular, it is of interest to explore what frequency (weekly or monthly), works best in this context, and what horizon is informative. In so doing, we maintain that positive changes (an increase in volatility) and negative ones should be treated differently in the model (this is the second asymmetry considered).

The results show that in characterizing the volatility dynamics, our model with monthly data and six months of lagged information works best, together with the contribution of the lagged VIX in the short-run component. Out-of-sample, the model behaves well, with a performance which is dependent on the sub-period considered.

The rest of the paper is organized as follows: Sect. 2 addresses the empirical question, illustrating first how the DAGM works and then we report the results of an application of various GARCH, GARCH-MIDAS and DAGM models on the S&P 500 volatility, both in- and out-of-sample perspectives. Finally, Sect. 3 concludes.

2 Modelling Volatility with the DAGM Model

Let us focus on the GARCH-MIDAS model, here synthetically labelled GM: the paper by [17] defines GARCH dynamics in the presence of mixed frequency variables. The short-run component varies with the same frequency as the dependent variable while the long-run component filters the lower frequency macro-Variable(s) (MV) observations. Recent contributions on (univariate) GARCH-MIDAS model are [6, 8, 11].

The paper by [5] proposes a DAGM where asymmetry in the short run is captured by a GJR-type [18] reaction to the sign of past returns, and positive and negative MV values have different impacts on the long-run.

2.1 The DAGM Framework

The DAGM-X model is defined as

$$r_{i,t} = \sqrt{\tau_t \times g_{i,t}}\varepsilon_{i,t}, \quad \text{with } i = 1, \ldots, N_t, \tag{1}$$

where

- $r_{i,t}$ represents the log-return for day i of the period t;
- N_t is the number of days for period t, with $t = 1, \ldots, T$;
- $\varepsilon_{i,t}|\Phi_{i-1,t} \sim N(0, 1)$, where $\Phi_{i-1,t}$ denotes the information set up to day $i-1$ of period t;
- $g_{i,t}$ follows a unit-mean reverting GARCH(1,1) process (short-run component);
- τ_t provides the slow-moving average level of volatility (long-run component).

The short-run component of the DAGM-X is given by

$$g_{i,t} = (1 - \alpha - \beta - \gamma/2) + \left(\alpha + \gamma \cdot \mathbb{1}_{\left(r_{i-1,t}<0\right)}\right) \frac{\left(r_{i-1,t}\right)^2}{\tau_t} + \beta g_{i-1,t} + z V_{i-1,t}, \tag{2}$$

where $\mathbb{1}_{(.)}$ is an indicator function and $V_{i,t}$ is an additional, positive volatility determinant, observed daily, whose importance on $g_{i,t}$ is given by z. In order to assure the positivity of $g_{i,t}$, we impose the constraint $z \geq 0$. In absence of $V_{i,t}$, the DAGM-X model becomes the DAGM specification.

The long-run component of the DAGM-X and DAGM is defined as

$$\tau_t = \exp\left(m + \theta^+ \sum_{k=1}^{K} \delta_k(\omega)^+ X_{t-k} \mathbb{1}_{(X_{t-k}\geq 0)} + \theta^- \sum_{k=1}^{K} \delta_k(\omega)^- X_{t-k} \mathbb{1}_{(X_{t-k}<0)}\right), \tag{3}$$

where
- m plays the role of an intercept;
- θ^+, θ^- represent the asymmetric responses to the one-sided filter;
- $\delta_k(\omega)^+$ and $\delta_k(\omega)^-$ are suitable functions weighing the past K realizations of the additional stationary variable X_t. As in the related literature, we opt for the Beta function, that is

$$\delta_k(\omega) = \frac{(k/K)^{\omega_1-1}(1-k/K)^{\omega_2-1}}{\sum_{j=1}^{K}(j/K)^{\omega_1-1}(1-j/K)^{\omega_2-1}}. \tag{4}$$

Given that we are only interested in the case of larger weights put on the most recent observations, we set $\omega_1 = 1$ and $\omega_2 \geq 1$. Note that the Beta function represented in (4) is readily applicable for both the GM and the DAGM. In this latter case, it is sufficient to replace $\delta_k(\omega)$ with $\delta_k(\omega)^+$ and $\delta_k(\omega)^-$.

Thus, the short-run component includes a term related to negative returns ("bad news" increasing volatility, the well-known *leverage* effect) and potentially a term associated with an additional MV observed with the same frequency of the dependent variable. The long-run component avoids positive and negative compensations within the one-sided filter, separating the positive MV variations from the negative ones.

Typically, MVs can only be observed at low frequency, but here we move out of the classic MV framework where observations are available only at low frequency. Thus we take a variable which is observable daily, but can be sampled at lower frequencies, e.g. weekly or monthly. We take the DAGM to the empirical evaluation of how different frequencies of observations in the MV may change the results both in estimation and forecasting. Besides that, we include the same variable at high frequency in the short-run component ("–X" specifications).

Assuming a conditional normal distribution for the error term $\varepsilon_{i,t}$ allows to apply the standard statistical inference (for details on the asymptotic properties of the GARCH-MIDAS class of models, see [22]) according to the maximization of the following log-likelihood:

$$\ln \mathscr{L} = -\frac{1}{2}\sum_{t=1}^{T}\left[\sum_{i=1}^{N_t}\left[\log(2\pi) + \log(g_{i,t}\tau_t) + \frac{\left(r_{i,t}\right)^2}{g_{i,t}\tau_t}\right]\right]. \quad (5)$$

2.2 The Role of VIX in the S&P 500 Volatility Dynamics

The returns of interest are daily log-differences of the S&P 500 Index (also examined on a different sample and context in [5]), annualized on a sample period: 7 January 2000–7 September 2018 (number of daily observations: 4686, collected from Yahoo Finance).

The MV in this paper is VIX (an implied volatility-based index built on the same index, cf. [23]) which in our setup will appear: (i) lagged daily as a predetermined variable in the short-run component g_i of the GARCH-X; (ii) lagged variations—end-of-month or end-of-week (with various choices of K) in the long-run component. All the observations concerning VIX have been collected from the Thomson Reuters Eikon provider. The distance between the estimated volatility, labelled as $\hat{h}_i$, and the chosen volatility proxy, the realized volatility at five minutes, labelled as σ_i and collected from the realized library of the Oxford-Man Institute, are investigated through three loss functions[1]: QLIKE, Root Mean Squared Error (RMSE) and Mean Absolute Errors (MAE), defined as follows:

[1]For ease of notation and because we are only interested in daily estimates, here the suffix t identifying the period at lower frequency has been suppressed.

$$\text{QLIKE}: E\left(\sigma_i/\hat{h}_i + \log(\hat{h}_i)\right);$$
$$\text{RMSE}: \sqrt{E\left(\left(\sigma_i - \hat{h}_i\right)^2\right)}; \quad (6)$$
$$\text{MAE}: E\left(|\sigma_i - \hat{h}_i|\right).$$

2.2.1 Estimation and Diagnostics

The estimation and diagnostics results are shown in Table 1, where we report the coefficients with their standard errors in parenthesis and their significance. GARCH is the standard (1, 1) model; the GJR allows for an asymmetric response to the lagged negative returns; the GARCH-X and GJR-GARCH-X and DAGM-X contain an extra predetermined variable, the lagged daily VIX. The GM and DAGM are built on a one-sided filter for the monthly VIX, while in the DAGM-W we consider the weekly VIX. The last six months of VIX have been used in GM, DAGM, DAGM-X, and DAGM-W, i.e. $K = 6$ and $K_{\text{week}} = 24$. The choice of the adopted MIDAS lags derives from some preliminary estimations aiming at finding the best values according to the Bayesian information criterion (BIC). The number of "*" indicates the significance (10%, 5%, 1%, respectively) of the estimated coefficients heteroscedasticity and autocorrelation consistent ([21], HAC) standard errors in parenthesis). LB_l and LM_l report the p-value of the Ljung-Box and ARCH-LM tests on the squared standardized residuals at lag l, respectively. RMSE is in percentage terms.

A few comments are in order: the GARCH models (first four columns) exhibit customary results, with the possible surprise of the non-significance of the lagged VIX in the X specifications. The GM has non-significant coefficient on the one-sided filter and the wrong sign: as a matter of fact, the information criteria and the QLIKE signal a worse fit of this model, relative to the standard models. When we introduce our DAGM, the signs of the impact coefficients θ^+ and θ^- are the right ones (positive, negative, respectively), and significant. The information criteria and the QLIKE report a marked improvement over the models previously considered, with the *best* model being the DAGM-X model where the significant coefficients on the low-frequency component are, besides the constant, those pertaining to the positive changes (inducing an increase in volatility). Generally, the residual diagnostics show a good fit of the models. In particular, almost all the models fail to reject the null hypotheses of the Ljung-Box and ARCH-LM tests, independently of the order of lags considered. The only model whose p-values are below the significance level of 5% is the DAGM-X, for $l = 12$, for both the Ljung-Box and ARCH-LM tests. Despite this, the conclusion is that the DAGM-X provides the most convincing performance with VIX contributing to a marked improvement over other models. The result can be appraised graphically as in Fig. 2 where we show the close proximity of the fitted values to the realized volatility.

Table 1 DAGM and GARCH estimates

	GARCH	GARCH–X	GJR	GJR–X	GM	DAGM	DAGM–X	DAGM–W
α	0.105^{***}	0.116^{***}	0.001	0.001	0.001	0.001	0.001	0.001
	(0.013)	(0.019)	(0.01)	(0.011)	(0.011)	(0.01)	(0.014)	(0.011)
β	0.884^{***}	0.876^{***}	0.889^{***}	0.878^{***}	0.94^{***}	0.874^{***}	0.852^{***}	0.884^{***}
	(0.015)	(0.018)	(0.015)	(0.02)	(0.013)	(0.015)	(0.018)	(0.015)
γ			0.192^{***}	0.225^{***}	0.11^{***}	0.194^{***}	0.198^{***}	0.19^{***}
			(0.023)	(0.037)	(0.023)	(0.022)	(0.022)	(0.023)
z		0.117		0.165			0.257^{***}	
		(0.091)		(0.104)			(0.039)	
m					5.169^{***}	4.956^{***}	0.686^{***}	5.123^{***}
					(0.315)	(0.192)	(0.121)	(0.205)
θ					−0.004			
					(0.005)			
ω					1.36			
					(1.385)			
θ^+						0.164^{***}	0.101^{***}	0.096^{***}
						(0.042)	(0.027)	(0.028)
ω_2^+						1.372^{***}	1.681^{***}	13.876^{***}
						(0.368)	(0.546)	(0.693)
θ^-						-0.192^{***}	−0.078	-0.431^{***}
						(0.065)	(0.07)	(0.11)
ω_2^-						1.017	1.124	1.455^{***}
						(0.883)	(0.765)	(0.492)
BIC	37586.899	37590.534	37393.477	37394.136	37546.527	37404.797	37367.454	37397.298
QLIKE	−3.867	−3.865	−3.876	−3.873	−3.882	−3.882	−3.882	−3.879
RMSE	0.418	0.433	0.395	0.418	0.402	0.376	0.364	0.382
LB_{12}	0.274	0.388	0.329	0.123	0.506	0.129	0.048	0.186
LB_{24}	0.322	0.361	0.518	0.239	0.416	0.384	0.229	0.384
LB_{36}	0.362	0.383	0.626	0.37	0.278	0.474	0.37	0.482
LM_{12}	0.26	0.381	0.311	0.092	0.526	0.118	0.037	0.17
LM_{24}	0.318	0.349	0.485	0.159	0.41	0.366	0.203	0.345
LM_{36}	0.411	0.391	0.614	0.253	0.366	0.482	0.371	0.479

Notes Annualized scale. Sample period: 7 January 2000–7 September 2018. Number of daily observations: 4686. Ticker: S&P 500. Comparison of the DAGM with other GARCH models. Model definitions and comments in the text. HAC standard errors in parentheses. *, ** and *** denote significance at the 10%, 5% and 1% levels, respectively

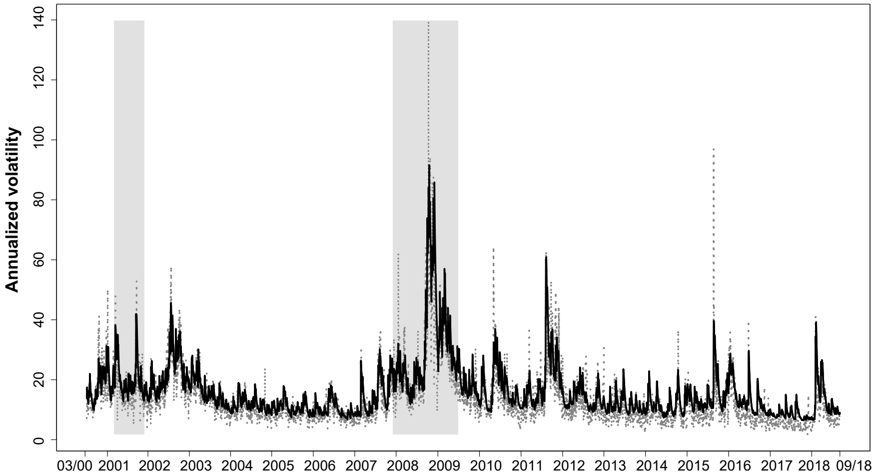

Fig. 2 Realized and DAGM-X volatilities. *Notes* The figure plots the DAGM-X volatility (solid black line) and the S&P 500 realized volatility (dashed grey line). Shaded areas represent NBER recession periods. Annualized scale

2.2.2 Forecasting

Further insights can be had moving to an out-of-sample exercise where we estimate the model over a 10-year period and project one-step ahead for one year and then move forward the estimation and forecasting window. The results are summarized in Fig. 3 where we report the presence in a Model Confidence Set as proposed by [20]. The results (at $\alpha = 10\%$) show that while the DAGM-X has a satisfactory performance, at the same time the standard GARCH or GJR models enter the set.

3 Wrapping Up

The slow-moving feature of conditional volatility can be addressed within a *Double Asymmetric* GARCH-MIDAS framework [5] where the low-frequency variable here is a volatility measure (variations in VIX). The main novelty in this approach is that the same variable can be inserted as a forcing variable (-X in levels) in the short-run component, and we can explore which frequency is the most suitable for the long-run component (in first differences). The fitting capabilities of this approach are comforting, with monthly movements in volatility providing the best in-sample results. In out-of-sample forecasting, though, the model is less satisfactory, in that it gives a performance very similar to a standard GARCH model.

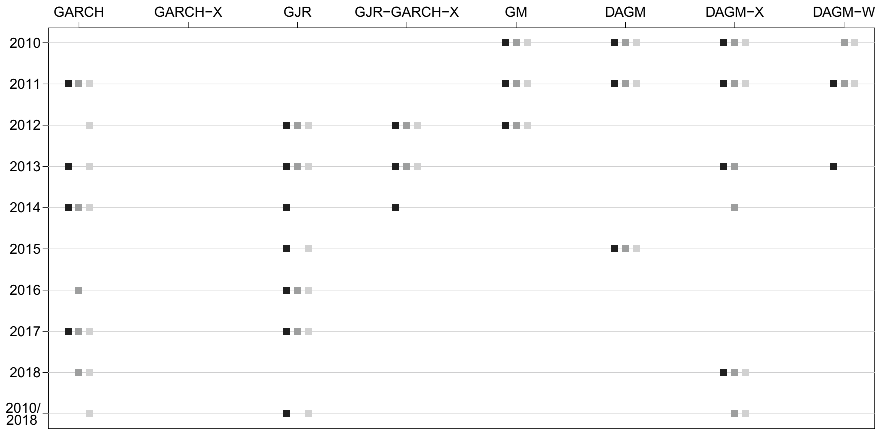

Fig. 3 MCS composition. *Notes* The figure plots the composition of the Model Confidence Set (MCS). For different loss functions, dark (QLIKE), medium-dark (MSE) and light (MAE) shades of grey indicate that a given model is included in the MCS, at a significance level of $\alpha = 0.1$.

References

1. Amado, C., Teräsvirta, T.: Modelling conditional and unconditional heteroskedasticity with smoothly time-varying structure. CREATES Res. Paper **8**, 1–53 (2008)
2. Amado, C., Teräsvirta, T.: Modelling volatility by variance decomposition. J. Econ. **175**(2), 142–153 (2013)
3. Amado, C., Teräsvirta, T.: Modelling changes in the unconditional variance of long stock return series. J. Empir. Finan. **25**, 15–35 (2014)
4. Amado, C., Silvennoinen, A., Teräsvirta, T.: Models with multiplicative decomposition of conditional variances and correlations. CREATES Res. Paper **14**, 1–46 (2018)
5. Amendola, A., Candila, V., Gallo, G.M.: On the asymmetric impact of macro-variables on volatility. Econ. Model. **76**, 135–152 (2019)
6. Amendola, A., Candila, V., Scognamillo, A.: On the influence of US monetary policy on crude oil price volatility. Empir. Econ. **52**(1), 155–178 (2017)
7. Andersen, T.G., Bollerslev, T., Christoffersen, P.F., Diebold, F.X.: Volatility and Correlation Forecasting. In: Elliott, G., Granger, C.W.J., Timmermann, A. (eds.) Handbook of Economic Forecasting, pp. 777–878. North-Holland, Amsterdam (2006)
8. Asgharian, H., Hou, A.J., Javed. F.: The importance of the macroeconomic variables in forecasting stock return variance: a GARCH-MIDAS approach. J. Forecast. **32**(7), 600–612 (2013)
9. Bollerslev, T.: Generalized Autoregressive Conditional Heteroskedasticity. J. Econ. **31**, 307–327 (1986)
10. Brunetti, C., Scotti, C., Mariano, R.S., Tan, A.H.H.: Markov switching GARCH models of currency turmoil in Southeast Asia. Emerg. Mark. Rev. **9**(2), 104–128 (2008)
11. Conrad, C., Loch, K.: Anticipating long-term stock market volatility. J. Appl. Econ. **30**(7), 1090–1114 (2015)
12. Ding, Z., Granger, C.W.J.: Modeling volatility persistence of speculative returns: a new approach. J. Econ. **73**(1), 185–215 (1996)
13. Dueker, M.: Markov switching in GARCH processes and mean-reverting stock-market volatility. J. Bus. Econ. Stat. **15**(1), 26–34 (1997)
14. Engle, R.F.: Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. Econometrica **50**, 987–1007 (1982)

15. Engle, R.F., Lee, G.J.: A permanent and transitory component model of stock return volatility. In: Engle, R.F., White, H. (eds.) Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W. J. Granger, pp. 475–497. Oxford University Press, Oxford (1999)
16. Engle, R.F., Rangel, J.G.: The Spline-GARCH model for low frequency volatility and its global macroeconomic causes. Rev. Finan. Stud. **21**, 1187–1222 (2008)
17. Engle, R.F., Ghysels, E., Sohn, B.: Stock market volatility and macroeconomic fundamentals. Rev. Econ. Stat. **95**(3), 776–797 (2013)
18. Glosten, L.R., Jagannathan, R., Runkle, D.E.: On the relation between the expected value and the volatility of the nominal excess return on stocks. J. Finan. **48**(5), 1779–1801 (1993)
19. Hamilton, J., Susmel, R.: Autoregressive conditional heteroskedasticity and changes in regime. J. Econ. **64**(1–2), 307–333 (1994)
20. Hansen, P.R., Lunde, A., Nason, J.M.: The model confidence set. Econometrica **79**(2), 453–497 (2011)
21. Newey, W.K., West, K.D.: A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica **55**(3), 703–708 (1987)
22. Wang, F., Ghysels, E.: Econometric analysis of volatility component models. Econ. Theory. **31**(2), 362–393 (2015)
23. Whaley, R.E.: Understanding the VIX. J. Portfolio Manage. **35**, 98–105 (2009)