ARTICLE IN PRESS

International Journal of Forecasting xxx (xxxx) xxx

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Network log-ARCH models for forecasting stock market volatility

Raffaele Mattera [a,*], Philipp Otto [b]

[a] Department of Social and Economic Sciences, Sapienza University of Rome, Italy
[b] Leibniz University Hannover, Institute of Cartography and Geoinformatics, Hannover, Germany

ARTICLE INFO

Keywords:
ARCH models
Network processes
Stock market volatility
Financial networks
Risk prediction
Spatial econometrics

ABSTRACT

This paper presents a dynamic network autoregressive conditional heteroscedasticity (ARCH) model suitable for high-dimensional cases where multivariate ARCH models are typically no longer applicable. We adopt the theoretical foundations from spatiotemporal statistics and transfer the dynamic ARCH model processes to networks. The model integrates temporally lagged volatility and information from adjacent nodes, which may instantaneously spill across the entire network. The model is used to forecast volatility in the US stock market, and the edges are determined based on various distance and correlation measures between the time series. The performance of alternative network definitions is compared with independent univariate log-ARCH models in terms of out-of-sample prediction accuracy. The results indicate that more accurate forecasts are obtained with network-based models and that accuracy can be improved by combining the forecasts of different network definitions. We emphasise the significance for practitioners to integrate network structure information when developing volatility forecasts.

1. Introduction

Forecasting volatility in the stock market is crucial, as it provides relevant information for investment and risk management purposes. Volatility forecasting usually employs generalised autoregressive conditional heteroscedasticity (GARCH) models and their extensions (Andersen & Bollerslev, 1998; Bollerslev, 1986; Francq & Zakoian, 2019). Geweke (1986) advocates using the log-ARCH model instead of the non-exponential GARCH when dealing with highly persistent volatility processes, large jumps in the data, outliers, and skewness. However, in times of worldwide trading systems and common international markets, interactions between stocks and their returns play a crucial role. If the volatility of one or more stocks changes, this change can immediately spill over to the entire financial network, where "close stocks" are affected by the largest degree. These interactions can also be exploited for volatility forecasting in financial networks. In this paper, we propose a network log-ARCH and develop a novel modelling framework inspired by spatiotemporal statistics to improve the out-of-sample volatility forecasting performance.

Considering volatility modelling, Otto et al. (2018) introduced spatial ARCH (spARCH) models, while Otto et al. (2021) analysed its properties in detail. Further, Sato and Matsuda (2017) suggested a logarithmic expression of the volatility equation, which is the equivalent of a log-ARCH model for spatial data. In Sato and Matsuda (2021), the spatial log-ARCH model was generalised to a spatial log-GARCH model. Otto and Schmid (2022) introduced a

* Corresponding author.
E-mail addresses: raffaele.mattera@uniroma1.it (R. Mattera), philipp.otto@ikg.uni-hannover.de (P. Otto).

https://doi.org/10.1016/j.ijforecast.2024.01.002

generalization of the spARCH model in a unified framework, allowing for a variety of possible spatial GARCH-type models. More recently, Otto et al. (2023) proposed a dynamic spatiotemporal log-ARCH approach for modelling house prices in Berlin. In this paper, we propose using the Otto et al. (2023) spatiotemporal modelling approach for forecasting stock market volatility. The out-of-sample forecasting performance is compared with the benchmark time-series log-ARCH models and ensemble forecasts.

Although spatial and spatiotemporal modelling is popular in many domains such as epidemiology (Mattera, 2022; Sahu & Böhning, 2022), environmental sciences (Cameletti et al., 2019; Fassò et al., 2022; Huang et al., 2011; Piter et al., 2022), and real estate economics (Baltagi & Li, 2014; Holly et al., 2010; Otto & Schmid, 2018), it is less explored in finance. By modelling volatility with a spatiotemporal approach, we assume that risk is influenced not only by the temporal fluctuations but also by dependencies with other stocks that are close in geographical space or, more generally, are in some sense similar.

According to Pirinsky and Wang (2006), two main explanations of why adjacent information can be helpful for modelling stock market data can be found in local investors' correlated trading activities and by the presence of locally correlated fundamentals. Nevertheless, many studies argue that geographical distance has quite limited relevance in explaining the correlations of financial returns, see e.g. Barker & Loughran, 2007; Eckel et al., 2011. Notice also that it is unclear which geographical information could and should be used to analyse stock market data, especially if country-specific stocks are considered. For example, using firms' headquarters to define spatial closeness can be reductive, considering that production plants can be located in many places. For this reason, previous studies propose alternative definitions of similarity across stocks measured in an attribute space rather than the geographical one, for example, the space spanned by financial indicators, balance sheet positions, or alternative indicator,s see e.g. Asgharian et al., 2013; Fernández-Avilés et al., 2012; Fülle & Otto, 2023. In this way, spillovers arise from stocks with similar characteristics, although not necessarily close in geographical space.

Therefore, the use of spatial econometric approaches in the financial domain requires a different definition of the proximity of locations and how spatial weights are determined. By substituting spatial contiguity with a broader concept of similarity, we can construct a financial network and model the volatility of a stock as a function of past values and the relationship with adjacent nodes in the network. Working with this kind of model is convenient because financial networks provide a suitable framework for understanding the propagation mechanisms of shocks occurring in the market. Specifically, in an empirical application, we consider several network definitions based on correlation metrics and closeness in the volatility dynamics. Moreover, various choices of edge weights are considered for each network definition, which are eventually compared concerning their out-of-sample prediction performance. In this way, we can identify the best financial network structure in terms of forecasting ability.

As shown by many authors, see e.g. Barigozzi & Hallin, 2017; Betancourt et al., 2020; Demiris et al., 2014; Diebold & Yılmaz, 2014; Liu et al., 2021; Vinciotti et al., 2019; Zhou et al., 2023, the financial market is well represented by networks where stocks are the nodes, and the edges reflect the degree of similarity across them. We introduce the spillovers from adjacent nodes in an ARCH-like manner. The information from adjacent network nodes can be successfully used to model relationships across stock returns and volatility. The idea that better predictions can be achieved by incorporating network information has attracted the interest of researchers in the field, see e.g. Huang et al., 2023; Wu et al., 2022. But we still know very little, especially about the form of interactions, and further studies are needed. Indeed, network modelling of returns and volatilities is a recent and flourishing research area in finance.

We contribute to this literature in three main directions. First, unlike previous papers (Billio et al., 2021; Caporin & Paruolo, 2015; Zhou et al., 2020), we propose the use of a spatiotemporal dynamic log-ARCH model for volatility forecasting in financial networks. Specifically, we consider instantaneous network interactions in an ARCH-like manner. Hence, the volatility may immediately spill over to adjacent/similar stocks, reflecting the simultaneity of investors' trading decisions. Furthermore, our proposed modelling approach shares the same relevant features of exponential volatility models. Second, we extend the dynamic spatiotemporal ARCH models of Otto et al. (2023) with homogeneous temporal ARCH effects by introducing stock-specific temporal ARCH parameters. Consequently, we directly extend univariate log-ARCH models for each stock by including interactions across the financial network. We assess the usefulness of the network log-ARCH model by a rigorous evaluation of its forecasting performance. To the best of our knowledge, this was not the focus of previous research. Third, under the assumption of unknown locations, we evaluate how the approach used for constructing the network affects the forecasting results. More precisely, we compare several network and edge weight definitions, and we select the network structure with the best forecasting performance. Previous papers proposed the construction of an adjacency based on objective criteria such as shared shareholders (Zhou et al., 2020), industry sectors (Billio et al., 2021; Caporin & Paruolo, 2015), or balance sheet data (Fülle & Otto, 2023). By contrast, we adopt three alternative definitions of similarity across stocks, taking inspiration from the time series clustering literature, for an overview see Maharaj et al., 2019. In doing so, we adopt an intuitive and fast data-driven approach for constructing the network that only requires the time series to be predicted for its construction. Then, we build the networks considering the inverse distance and k-nearest neighbours, obtaining 12 alternative network log-ARCH models. Ensemble forecasts to combine all model alternatives are considered to further improve the out-of-sample forecasting performance.

The empirical application is carried out on data about the stocks included in the Dow Jones Industrial Average Index during the period from 2010–2022. The forecasting exercise is conducted considering a rolling window approach, and the out-of-sample accuracy is evaluated in terms of the root mean squared forecast error (RMSFE) and mean absolute forecast error (MAFE). Moreover, the difference in the predictive performance of the network-based models (and their combinations) with respect to the benchmark is tested by means of the Diebold and Mariano (2002) and Clark and West (2007) tests. The results demonstrate that the proposed network-based log-ARCH approach provides more accurate out-of-sample forecasts than the traditional log-ARCH modelling relying solely on temporal information. Moreover, considering predictive accuracy tests, we find that both the similarity measure and the procedure employed in constructing the network model affect the out-of-sample forecasting accuracy. Although all the network approaches outperform the benchmark, we identify a superior subset (three out of 12) of network models with forecasts that are more statistically accurate. Finally, we show that a combination of forecasts obtained with alternative network definitions further improves forecasting accuracy. Given the good results obtained through network combinations, we suggest adopting combination procedures in practice when forecasting volatility with alternative network-based models.

The rest of the paper is structured as follows. Section 2 discusses the dynamic spatiotemporal log-ARCH model of Otto et al. (2023) and how spatial weights are determined. Section 3 presents the data used for the empirical analysis and explains how the out-of-sample forecasting methodology is conducted. Section 4 shows the main results, i.e. forecasting accuracy and predictive accuracy tests. Section 5 concludes with final remarks and some suggestions for future research.

2. Forecasting models

The following sections introduce the modelling framework used for forecasting comparisons. We consider a process of stock market returns on a network $G = (V, E)$, which consists of a set of nodes/vertices V (i.e., stocks) that are possibly connected by directed or undirected edges. These edges are contained in the set E. Furthermore, we observe a process $\{Y_t(s_i) : t = 1, \ldots, T, s_i \in V\}$ of nodal attributes across time (i.e., return series). In particular, our focus is on the volatility of this process, which should be forecasted for $T+1, T+2, \ldots$ using the dynamic network ARCH model. Moreover, let $V = \{s_1, \ldots, s_n\}$ be the set of n stocks and $\boldsymbol{Y}_t = (Y_t(s_1), \ldots, Y_t(s_n))'$ the n-dimensional vector of the observed process on the network. It is important to note that we do not consider dynamic networks (i.e., time-varying sets of nodes and/or edges), but the networks are assumed to be constant over time.

2.1. Univariate logarithmic ARCH models (baseline model)

We start our model comparison with univariate logarithmic ARCH (log-ARCH) models, which are fitted independently to each time series. In this way, we get a very flexible model, allowing for different dependence structures for each stock to get the forecast performance of the volatility series. The model was originally proposed by Geweke (1986). To be precise, the log-ARCH(P) for the ith stock can be written as follows:

$$Y_t(s_i) = \sqrt{h_t(s_i)}\varepsilon_t(s_i), \tag{1}$$

$$\ln h_t(s_i) = \omega_i + \sum_{p=1}^{P} \gamma_{ip} \ln Y_{t-p}^2(s_i)\,, \tag{2}$$

where ω_i is the constant term, $\gamma_{i1}, \ldots, \gamma_{iP}$ are the ARCH parameters for the ith stock, and P is the order of the log-ARCH process. We can fit n univariate log-ARCH models and predict stocks' volatilities by only using idiosyncratic temporal information. The $n(P+1)$ unknown parameters of the log-ARCH model can be consistently estimated via ARMA representation of the process, see e.g. Sucarrat et al., 2016. Applying a log-square transformation of (1) shows that $\ln Y_t^2(s_i) = \ln h_t(s_i) + \ln \varepsilon_t^2(s_i)$. As a consequence,

$$\ln Y_t^2(s_i) = \omega_i + \sum_{p=1}^{P} \gamma_{ip} \ln Y_{t-p}^2(s_i) + \ln \varepsilon_t^2(s_i)\,.$$

However, since the distribution of $\ln \varepsilon_t^2(s_i)$ does not have a mean of zero, $E\left(\ln \varepsilon_t^2(s_i)\right)$ is added and subtracted from sides of the equation, leading to

$$\ln Y_t^2(s_i) = \phi_{i0} + \sum_{p=1}^{P} \phi_{ip} \ln Y_{t-p}^2(s_i) + u_t(s_i) \tag{3}$$

with a new error term $u_t(s_i) = \ln \varepsilon_t^2(s_i) - E\left(\ln \varepsilon_t^2(s_i)\right)$ and the constant $\phi_{i0} = \omega_i + E\left(\ln \varepsilon_t^2(s_i)\right)$. Notice that $E\left(\ln \varepsilon_t^2(s_i)\right) = \mu_i^*$ for all t, so this transformation ensures that the error term $u_t(s_i)$ has a zero mean. Hence, the ARMA representation (3) allows for consistent estimations of the log-ARCH parameters. More precisely, we have that $\gamma_{ip} = \phi_{ip}$ and $\omega = \phi_{i0} - E\left(\ln \varepsilon_t^2(s_i)\right)$.

From the discussion highlighted so far, it is clear that we need an estimate of the term $E\left(\ln \varepsilon_t^2(s_i)\right)$ to estimate ω_i. Bauwens and Sucarrat (2010), Sucarrat and Escribano (2012), Sucarrat et al. (2016) propose ex post scale adjustments based on the estimated residuals $\{\hat{u}_t(s_i), t = 1, \ldots, T\}$. That is,

$$\hat{\mu}_i^* = -\ln\left[\frac{1}{T}\sum_{t=1}^{T} \exp(\hat{u}_t(s_i))\right]. \tag{4}$$

This is a consistent and asymptotic normal smearing estimator of the log-square transformed errors' mean, see Duan, 1983; Francq & Sucarrat, 2018; Sucarrat et al., 2016.

Lastly, the resulting one-step-ahead forecast at $T+1$ of the log-ARCH(1) model is given by

$$\ln \hat{h}_{t+1}(s_i) = \left[\hat{\phi}_{i0} + \ln\left(\frac{1}{T}\sum_{t=1}^{T} \exp(\hat{u}_t(s_i))\right)\right] + \hat{\phi}_{i1} \ln y_t^2(s_i), \tag{5}$$

with $y_t^2(s_i)$ being the observed values of $Y_t^2(s_i)$. For the comparison, we focus on one-step-ahead forecasts and a model order of $P = 1$, but the methods can easily be extended to multi-step-ahead forecasts and higher model orders.

Generally, network processes could also be represented as multivariate time series. Thus, a natural extension to account for dependence between the nodes would be multivariate log-ARCH models (Francq & Sucarrat, 2017). However, they are limited in the sense that (1) they do not account for the inherent network structure, and (2) the number of parameters grows quadratically when the number of nodes n increases. It is worth mentioning that the time series length (without any changes in the parameters or structural breaks) must be larger than n^2 to get unique and reasonably precise estimates. In this paper, we particularly focus on the case with large n (compared to the length of the time series T). Hence, we propose to include instantaneous ARCH effects across the network to describe the dependence between the stock returns. In other words, the log volatility $\ln h_t(s_i)$ of the ith stock is influenced by all other observations $y_t(s_j)$ for $j = 1, \ldots, n$ and $j \neq i$, whereby the dependence structure is determined by the edges E of the network. In this way, large volatilities (i.e., large values of $y_t(s_j)^2$) can spill over to the adjacent nodes and lead to an increase in $\ln h_t(s_i)$. As a consequence, volatility clusters can be observed across the network.

2.2. Dynamic network logarithmic ARCH model

The new dynamic network log-ARCH model is based on the dynamic spatiotemporal ARCH models proposed by Otto et al. (2023). As for the univariate log-ARCH models, the observed process is given by

$$Y_t(s_i) = \sqrt{h_t(s_i)}\varepsilon_t(s_i), \tag{6}$$

but now $h_t(s_i)$ is being influenced by past observations at the same node, $Y_{t-1}(s_i)$, and simultaneously by the adjacent observations at the same time point, $\{Y_t(s_j) : j \in E_i\}$, where E_i is the subset of edges with links to node s_i. Let $\boldsymbol{h}_t^* = (\ln h_t^2(s_1), \ldots, \ln h_t^2(s_n))'$ and $\boldsymbol{Y}_t^* = (\ln Y_t^2(s_1), \ldots, \ln Y_t^2(s_n))'$. Then, the network log-ARCH process of order one can be written as follows:

$$\boldsymbol{h}_t^* = \boldsymbol{\omega} + \boldsymbol{\Gamma}\boldsymbol{Y}_{t-1}^* + \rho\mathbf{W}\boldsymbol{Y}_t^*, \tag{7}$$

where $\mathbf{W} = (w_{ij})_{i,j=1,\ldots,n}$ is a matrix of edge weights which define the relative degree of the volatility spillovers, ρ is an unknown parameter for these instantaneous network interactions, $\boldsymbol{\Gamma} = \mathrm{diag}(\gamma_1, \ldots, \gamma_n)'$ is a diagonal matrix of stock-specific temporal ARCH effects, and $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_n)'$ is the constant term. The matrix of edge weights $\mathbf{W}$ is analogously specified as the spatial weight matrix in spatial econometrics. That is, the diagonal entries are supposed to be zero (i.e., no self-loops). The matrix is non-stochastic and uniformly bounded in row and column sums in absolute terms. The latter assumption is needed to limit the network interactions to a constant degree when the number of nodes n is increasing. Typical choices are inverse-distance matrices (e.g., for road networks) or k-nearest neighbour matrices, where the proximity is defined by any network characteristic. We discuss the definition of $\mathbf{W}$ in more detail in Section 2.3.

The log-volatility terms follow a process characterised by the presence of instantaneous network effects, which is the key difference to previously proposed network GARCH models e.g. Zhou et al., 2020 or GARCH models including artificial neural network structures, e.g. Donaldson & Kamstra, 1997; Kristjanpoller & Minutolo, 2015. These previous models include network interactions only at the first temporal lag. That is, network spillovers can only happen in the next time instance but not instantaneously. The network log-ARCH model allows for deriving an ARMA representation of the model, namely

$$\boldsymbol{Y}_t^* = \boldsymbol{\phi}_0 + \rho\mathbf{W}\boldsymbol{Y}_t^* + \boldsymbol{\Gamma}\boldsymbol{Y}_{t-1}^* + \boldsymbol{u}_t, \tag{8}$$

where $\boldsymbol{u}_t = (\ln\varepsilon_t^2(s_1) - E(\ln\varepsilon_t^2(s_1)), \ldots, \ln\varepsilon_t^2(s_n) - E(\ln\varepsilon_t^2(s_n)))'$ are the log-squared errors, and $\boldsymbol{\phi}_0 = \boldsymbol{\omega} + \boldsymbol{\mu}^*$ with $\boldsymbol{\mu}^* = \left(E(\ln\varepsilon_t^2(s_1)), \ldots, E(\ln\varepsilon_t^2(s_n))\right)'$. It is important to note that we allow for different means $E(\ln\varepsilon_t^2(s_i))$ for each stock, which is constant over time. To estimate $\boldsymbol{\phi}_0^*$, we propose using the smearing estimate proposed by Sucarrat et al. (2016) for time-series log-ARCH models. That is,

$$\hat{\boldsymbol{\phi}}_0^* = \frac{1}{T}\sum_{t=1}^{T}\hat{\boldsymbol{u}}_t. \tag{9}$$

From (8), we see that the instantaneous spillovers lead to endogeneity in the model (i.e., $\boldsymbol{Y}_t^*$ appears on both sides of the equation), which is not the case when the network interactions are restricted to the first temporal lag, like in Zhou et al. (2020). Hence, we apply the estimation proposed by Otto et al. (2023) based on orthonormal transformations and the generalised method of moments (GMM). The key idea of the GMM estimator is to instrument $\mathbf{W}\boldsymbol{Y}_t^*$ by higher-order network lags originally proposed for spatiotemporal autoregressive models by Lee (2007), Lee and Yu (2014). For further details, we refer the interested reader to Otto et al. (2023).

Finally, we obtain the one-step-ahead forecast of all stocks at time $T+1$ as

$$\boldsymbol{h}_{T+1}^* = \left(\mathbf{I}_n - \hat{\rho}\mathbf{W}\right)^{-1}\left[\hat{\boldsymbol{\Gamma}}\boldsymbol{Y}_T^* + \hat{\boldsymbol{\phi}}_0\right], \tag{10}$$

where $\mathbf{I}_n$ is the n-dimensional identity matrix. Notice that $\boldsymbol{\phi}_0$ includes $\boldsymbol{\omega}$ and $\boldsymbol{\mu}^*$, but these two quantities are jointly estimated from the residuals' process, as in (9), because the orthonormal transformation eliminates all cross-sectional fixed effects.

For this network log-ARCH model, finding a suitable edge weight matrix for the ARCH-type interactions across the network is crucial. The edges are typically unknown for financial networks or stock market interactions and therefore have to be estimated. Nevertheless, it is reasonable to assume that with the increasing similarity between the stocks, they are more likely to experience spillovers in the risks, i.e., conditional volatilities. In the following section, we discuss several options to estimate the similarity/dissimilarity in stock return series, which is a basis for the edge weights in $\mathbf{W}$.

2.3. Determining similarity across stocks

Measuring distance across time series can be done in many different ways. In this regard, one mainly distinguishes between raw-data-based, feature-based and model-based approaches (Maharaj et al., 2019). In the first class, we consider dissimilarities computed on raw data, such as using the standard Euclidean distance on temporal observations or dynamic time warping when stocks have different lengths. Following a future-based approach, dissimilarities across financial time series can be calculated based on asset correlations (Mantegna, 1999; Tumminello et al., 2010), auto-correlation structures (D'Urso & Maharaj, 2009), periodograms (Caiado et al., 2006, 2020), or Hurst exponents (Cerqueti & Mattera, 2023; Lahmiri, 2016). Further, model-based approaches measure dissimilarity across stocks by considering parameters estimated from statistical models, such as ARIMA (Piccolo, 1990), GARCH, see e.g. D'Urso et al., 2016; Otranto, 2008, the multiplicative error model, see e.g. Gallo et al., 2021, or score-driven models, see e.g. Cerqueti et al., 2022, 2021. The most commonly employed approaches are those based on volatility dynamics when dealing with stock returns. Indeed, volatility-based approaches can measure similarities by directly exploiting conditional heteroscedasticity.

In this paper, we consider three alternative configurations for the matrix **W**: the standard Euclidean distance across stock returns, the use of a correlation-based approach as suggested by Mantegna (1999), and a model-based approach based on the log-ARCH estimates. For the first approach, the dissimilarities over time are computed as follows:

$$d_{ij} = \sqrt{\sum_{t=1}^{T} \left(y_t(s_i) - y_t(s_j)\right)^2}\,. \tag{11}$$

Considering the correlation-based approach, instead, the generic entries of the dissimilarity matrix are given by

$$d_{ij} = \sqrt{2\left(1 - \rho_{ij}\right)}, \tag{12}$$

with ρ_{ij} being the estimated correlation coefficient between stocks i and j over the entire time horizon $t = 1, \ldots, T$. In this way, we assume that stocks are similar to each other if their correlation is high. Finally, we propose using a log-ARCH approach, accounting for the stocks' underlying log-volatility dynamics. According to Piccolo (1990), we can define the dissimilarity between two time series i and j by means of the AR(∞) of log-squared returns:

$$d_{ij} = \sqrt{\sum_{p=1}^{\infty} \left(\gamma_{ip} - \gamma_{jp}\right)^2}, \tag{13}$$

where γ_{ip} is the autoregressive parameter (of order p) log-squared returns series of the ith stock. In practice, the infinite sum in (13) has to be truncated at some order P. The selection can be made according to the Akaike or Bayesian information criterion. If two time series i and j have different orders, P_i and P_j, we take $P = \max(P_1, P_2)$ and let $\gamma_{ip} = 0$ for $P > P_1$ and, similarly, $\gamma_{jp} = 0$ for $P > P_2$ (Piccolo, 1990).

Based on the pairwise distances in (11) to (13), we consider two different strategies for constructing the weighting matrix **W**. First, **W** is defined as an inverse-distance weight matrix with

$$w_{ij} = d_{ij}^{-1} \qquad \text{for all } i, j = 1, \ldots, n\,. \tag{14}$$

Second, **W** is considered to have constant edge weights based on the k-nearest neighbours:

$$w_{ij} = \begin{cases} 1/k & \text{if } s_j \text{ is closer than the } k+1 \\ & \text{nearest neighbours of } s_i \\ 0 & \text{otherwise.} \end{cases} \tag{15}$$

The so-obtained weighting matrix **W** expresses distances in the (temporal) attribute space. Whereas the inverse-distance matrix is a symmetric and dense matrix by construction with equal weights for both directions (i.e., undirected graph, fully connected), the k-nearest-neighbour matrix is not necessarily symmetric (it will usually be asymmetric), but each neighbour has the same weight (i.e., a directed graph, homogeneously weighted). Alternatively, the weights could be obtained from estimated distances in balance sheet data, cf. Fülle & Otto, 2023.

3. Empirical analysis: Data and forecasting methodology

For our empirical exercise, we focused on volatility forecasting of stocks in the Dow Jones Industrial Average Index. To facilitate replication and further analysis, we have made our code, data, and examples available at philot789.github.io/Network_ARCH/.

3.1. Data

The time series of daily stock returns span from October 1, 2010, to October 31, 2022. We removed from our analysis any stocks with missing values. The list of considered stocks is shown in Table 1, along with their main descriptive statistics. The return series are displayed in Fig. 2 in terms of the median log returns of all stocks on each day (first row), and their median absolute returns as a measure of the stock market volatility (second row). In total, we observe the logarithmic returns of $n = 29$ nodes/stocks for $T = 3040$ days. To evaluate the benefit of including instantaneous network ARCH effects, we include the same information for all predictions. That is, we choose the temporal lag order $P = 1$ for the models described in Sections 2.1 and 2.2. We allow for different values of μ_i^* for each stock and different temporal ARCH parameters for each stock. That is, for $\rho = 0$, the network model is identical to the independent time-series log-ARCH models considered as the benchmark model. It is worth noting that, given the number of stocks n, the number of parameters of multivariate GARCH would exceed a reasonable degree, so it is useful to assume a certain structure of the covariance. More precisely, the number of parameters of a scalar Gaussian dynamic conditional correlation (DCC) model is 495, which corresponds to one

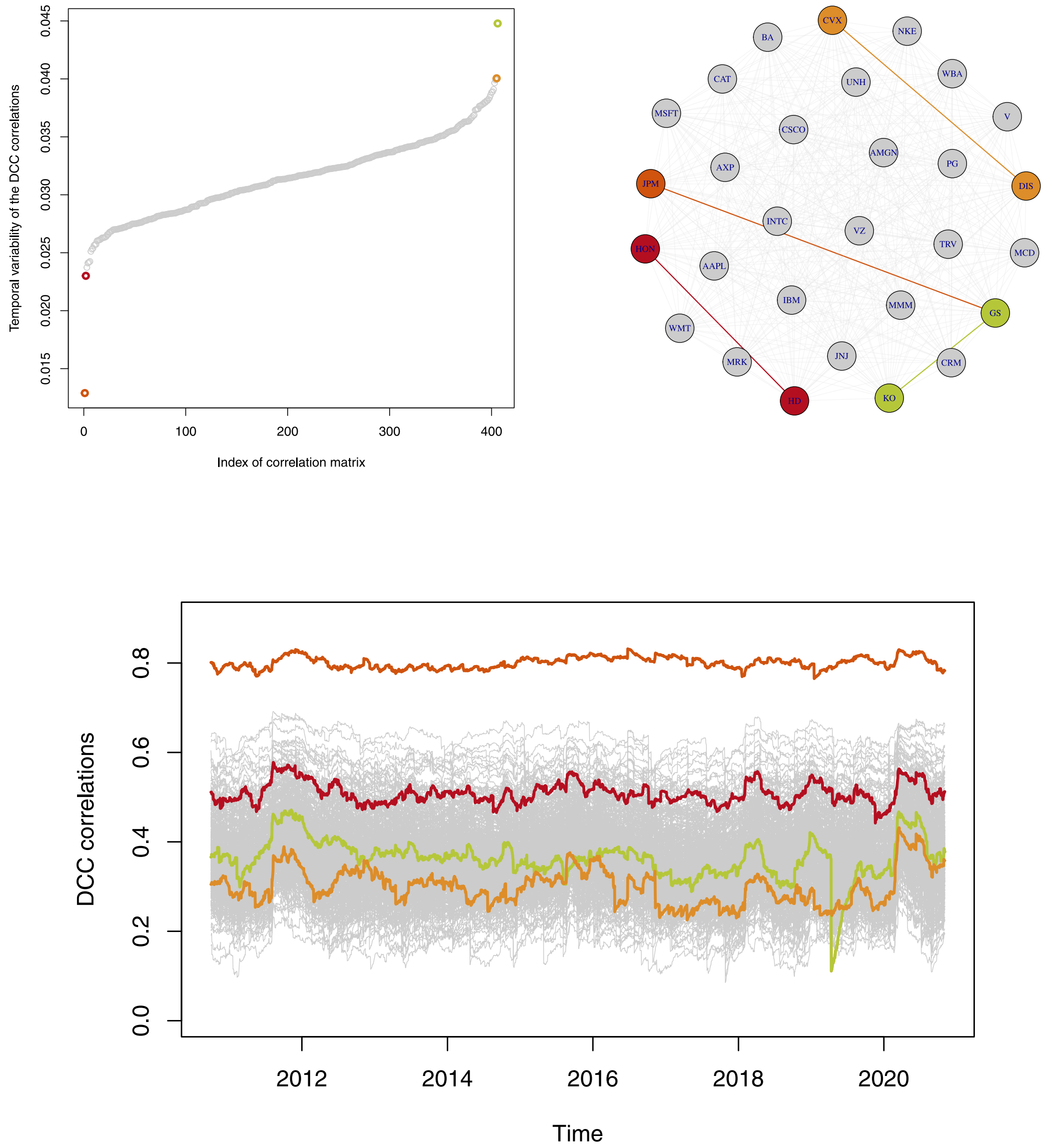

Fig. 1. Temporal stability of the network structure based on estimated correlations of a DCC model. Top left: Sorted temporal stability index (i.e., standard deviation of the estimated correlations across time for each pair of stocks). Top right: representation of four selected pairs of stocks with the highest/lowest temporal stability. Bottom: DCC correlations plotted across time, where the colour corresponds to the selected plots in the top panels.

parameter per roughly five time points in the estimation window (Caporin & McAleer, 2012). Nevertheless, we show the estimated DCC correlations in Fig. 1 to analyse the temporal stability of the network structure. The temporal correlation is largely stable across time, and the largest variations in the pairwise correlations are observed due to some peaks (light orange and green curves). To assume a structure in the covariance, we consider the proposed network ARCH structure. Thus, we allow for instantaneous volatility interactions across the network in an ARCH sense. Moreover, we compare the forecasting performance for various choices of edge weights, namely all three distance measures described in (11)–(13) and inverse-distance weighting (models A.1, A.2, and A.3) and k-nearest-neighbour weights with $k = 2, 3, 5$, and 10 (models B.k.1, B.k.2, and B.k.3).

The networks obtained considering the alternative approaches are shown in Fig. 3. The nodes in Fig. 3 are located according to their distances, and in the case of the inverse-distance approach, the edges are coloured according to their weights such that the higher the weight, the darker the edge. The network structures highlight interesting differences which can be exploited in the forecasting task.

In the case of Euclidean distances, as described in (11), we observe two outlying nodes, namely the CRM and BA stocks. When the weights are constructed based on the inverse distance (left-hand graphs), these two stocks have a minor influence on all other stocks. By contrast,

Table 1
Considered stocks in the empirical analysis and main descriptive statistics.

Company	Symbol	Mean	St. Dev.	Min	Max
Apple	AAPL	0.0010	0.0180	−0.1377	0.1132
Amgen	AMGN	0.0006	0.0153	−0.1008	0.1034
American Express	AXP	0.0005	0.0183	−0.1604	0.1979
Boeing	BA	0.0003	0.0232	−0.2724	0.2177
Caterpillar	CAT	0.0004	0.0183	−0.1541	0.0983
Salesforce	CRM	0.0006	0.0226	−0.1730	0.2315
Cisco	CSCO	0.0004	0.0169	−0.1769	0.1480
Chevron	CVX	0.0004	0.0177	−0.2501	0.2049
Dow	DOW	0.0004	0.0161	−0.1391	0.1346
Goldman Sachs	GS	0.0003	0.0182	−0.1359	0.1620
Home Depot	HD	0.0008	0.0148	−0.2206	0.1288
Honeywell	HON	0.0006	0.0147	−0.1288	0.1404
IBM	IBM	0.0002	0.0144	−0.1375	0.1071
Intel	INTC	0.0003	0.0187	−0.1990	0.1783
Johnson & Johnson	JNJ	0.0005	0.0107	−0.1058	0.0769
JPMorgan Chase	JPM	0.0005	0.0179	−0.1621	0.1656
Coca-Cola	KO	0.0004	0.0111	−0.1017	0.0628
McDonald’s	MCD	0.0005	0.0121	−0.1729	0.1666
3M	MMM	0.0002	0.0139	−0.1386	0.1187
Merck	MRK	0.0005	0.0131	−0.1038	0.0990
Microsoft	MSFT	0.0008	0.0164	−0.1595	0.1329
Nike	NKE	0.0006	0.0172	−0.1371	0.1444
Procter & Gamble	PG	0.0004	0.0111	−0.0914	0.1134
Travelers	TRV	0.0005	0.0144	−0.2332	0.1248
UnitedHealth	UNH	0.0010	0.0161	−0.1897	0.1204
Visa	V	0.0008	0.0161	−0.1456	0.1397
Verizon	VZ	0.0002	0.0112	−0.0697	0.0740
Walgreens Boots Alliance	WBA	0.0001	0.0178	−0.1548	0.1187
Walmart	WMT	0.0004	0.0124	−0.1208	0.1107

for the k-nearest-neighbour weights (right-hand graphs), their influence to adjacent stocks will be of the same degree. Interestingly, considering the main descriptive statistics in Table 1, BA and CRM have the highest variability in the sample. Specifically, BA has the lowest minimum value, while CRM has the highest maximum. Thus, we observe the highest Euclidean distance for these two stocks, which is not robust in the case of outlying observations. Furthermore, these two stocks show the most volatile patterns in terms of their returns. However, under this network structure and inverse-distance weights, the information included of these two stocks will have pretty low relevance for forecasting other stocks in the network. Conversely, the leading network structure is characterised by a central block (comprising, among the others, IBM, UNH, and MMM) affected mainly by the influence of the other stocks. Simultaneously, this group of stocks strongly affect those placed in the tails of the network. The central block of stocks has a stronger relationship with CRM, while the effect of BA is more pronounced for stocks in the right tail of the network.

The second row in Fig. 3 shows the networks constructed under the correlation distance defined in (12). From one side, under this scenario, the stock BA is not far away from the other stocks, as it is highly correlated with AXP, HON, JPM, and the other closer stocks. The same pattern was also observed for the Euclidean distance. On the other side, CRM is still the most distant from the others, although it shows an interestingly high correlation with IT-related stocks such as AAPL, MSFT, INTC, and CSCO. In light of this evidence, stocks belonging to the IT-related sectors would rely on information from stocks of the same or similar industry sector. In this way, the information included in highly correlated stocks is successfully employed to improve the volatility forecasts, due to the network structure. For the inverse-distance weights (left-hand graphs), this information spillover is represented by the edge colour, as the more relevant arrows have darker colours, and we observe a block of closely connected stocks in the bottom-right area of the graph. Contrary to these nodes, stocks placed on the left have lighter arrows, meaning that their information would scarcely be used for predicting stocks to the right of the structure and that their predictions rely mainly on idiosyncratic temporal information rather than the information coming from the network.

The last row in Fig. 3 shows the network structure constructed in terms of the volatility-based measure given by (14). Under this network structure, closer stocks are those sharing similar log-volatility dynamics. BA is again an outlying stock in this case, as it was under the Euclidean-based network. Thus, information from BA will scarcely be used for predicting the other stocks, and at the same time, BA forecasts will be mainly based on their temporally lagged values. The other two stocks placed far from the main cluster are AXP and TRV. However, the edges of these stocks show darker colours compared with BA, meaning that these stocks still have informative power for the volatility dynamics of the network.

3.2. Forecasting evaluation

The forecasting methodology is based on a rolling window procedure. For this reason, we first divide the sample into training and testing sets for each stock, leaving the

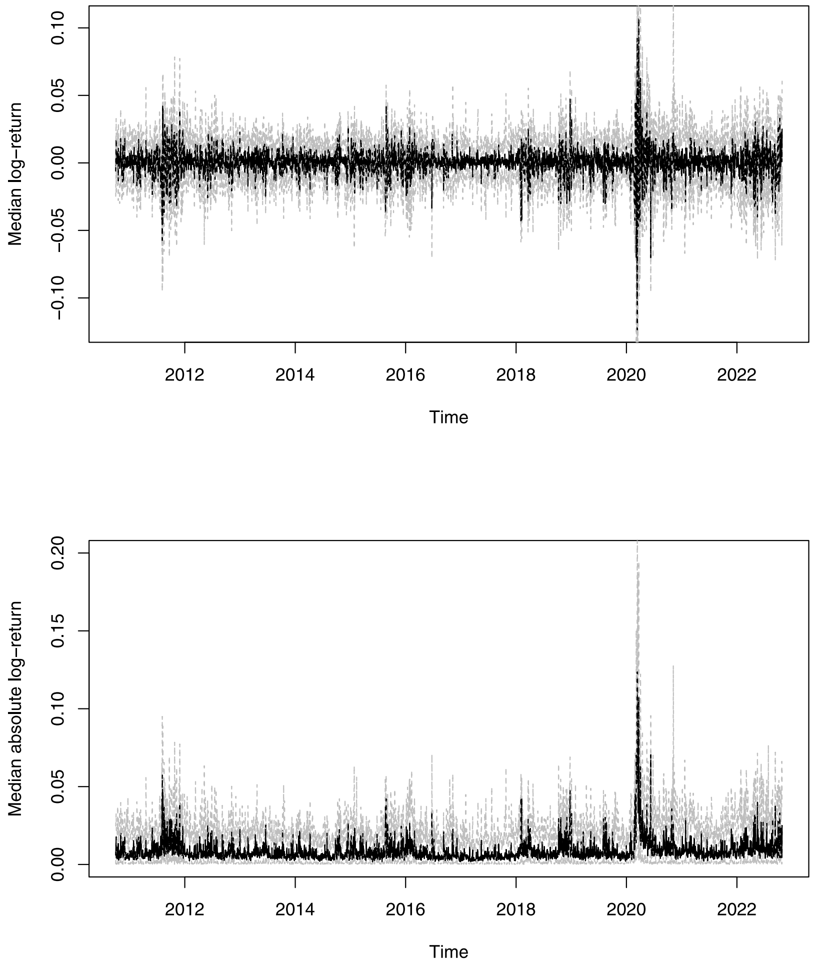

Fig. 2. Summary of $n = 29$ time series of log returns. Top: Median log returns of each day (solid line) and 5% and 95% quantiles (dashed lines). Bottom: Median (solid line) and 5% and 95% quantiles of the absolute log returns to depict the temporally varying volatility.

last 500 observations (i.e., the last two years) for out-of-sample testing. The first $M = 2540$ observations are used to obtain the edge weight matrix $\mathbf{W}$ and to estimate the model parameters. Then, the models are used for the one-step-ahead forecasts at $M+1$. Then, according to a rolling window procedure, the oldest observation is removed for the next step, and the new realised observation at $M + 1$ is included in the estimation sample. Parameters are re-estimated with the new data, and the forecasts are obtained for $M + 2$. This procedure is repeated until no new observation is available and all $T - M = 500$ volatilities are predicted for each stock. Thus, we always have an estimation window equal to 2540 observations at each recursion step.

To evaluate forecasting accuracy, we rely on two commonly employed accuracy metrics, namely the root mean squared forecast error (RMSFE), i.e.,

$$\text{RMSFE} = \sqrt{\frac{1}{T-M}\sum_{t=M+1}^{T}\left(\ln \hat{h}_{it} - \ln y_{it}^2\right)^2}, \tag{16}$$

and the mean absolute forecast error (MAFE), i.e.,

$$\text{MAFE} = \frac{1}{T-M}\sum_{t=M+1}^{T}\left|\ln \hat{h}_{it} - \ln y_{it}^2\right|. \tag{17}$$

Notice that we use realised squared log returns as the proxy of volatility for the out-of-sample accuracy evaluation. Furthermore, we use a predictive accuracy test to evaluate whether the forecasting errors of the competing statistical models are significantly different.

Let $d_t = g\left(e_{1,t}\right) - g\left(e_{2,t}\right)$ be the error differential between two forecasting approaches up to some transformation $g(\cdot)$, that in this paper are squaring $g(e_{1,t}) = e_{1,t}^2$ and absolute value $g(e_{1,t}) = |e_{1,t}|$. We perform the test independently for each stock; thus, we drop the index i in this notation. Assuming covariance stationarity of the loss differential series d_t, Diebold and Mariano (2002) show that the sample mean of the loss differential

$$\bar{d} \equiv \frac{1}{T-M}\sum_{t=M+1}^{T} d_t \tag{18}$$

asymptotically follows a standard normal distribution. Hence, a test decision about the null hypothesis of equal

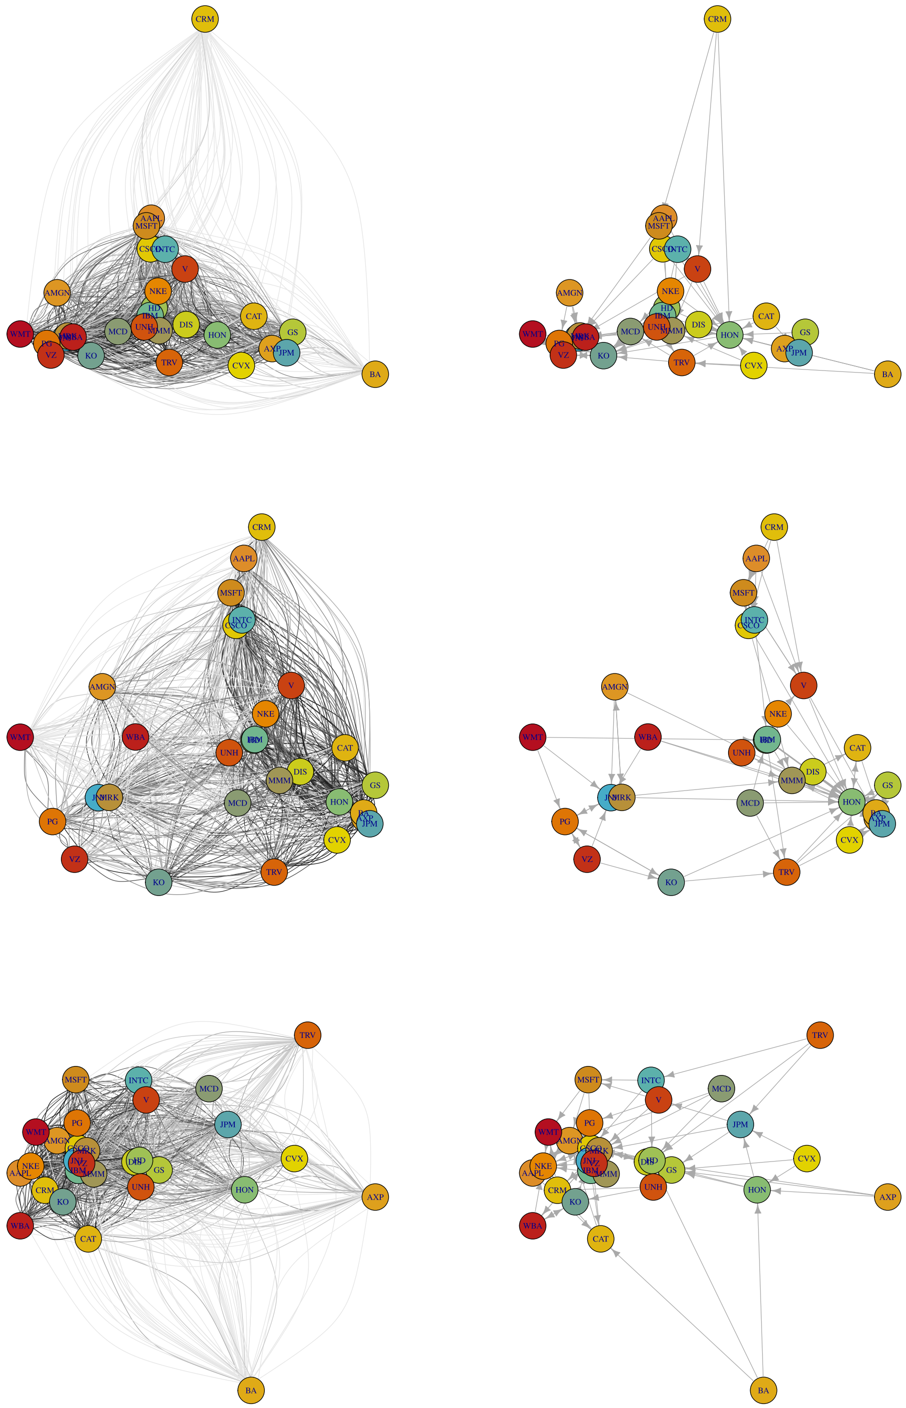

Fig. 3. Network of the considered stocks with different dissimilarity measures and weighting schemes. The nodes are located according to their distances. Top row: Euclidean dissimilarity (11). Centre row: correlation dissimilarity (12). Bottom row: log-ARCH dissimilarity (13). Left column: inverse-distance edge weights (14) with edges coloured according to their weights (i.e., the higher the weight, the darker the edge). Right column: $k = 3$-nearest-neighbour weights (15), where the arrows point towards the direction of the influence.

forecast accuracies can be obtained based on the following statistic:

$$\mathrm{DM} = \frac{\bar{d}}{\sqrt{\hat{V}(\bar{d})}}, \quad \mathcal{N} \sim \left(0, \hat{V}(\bar{d})\right), \tag{19}$$

where $V(\bar{d})$ can be consistently estimated assuming a particular autocorrelation structure of the forecasting errors. Moreover, we note that the network log-ARCH model nests the log-ARCH when $\rho = 0$. In this framework, under the null that the parsimonious model generates the data, it is known that forecasts obtained by the larger models are affected by noise introduced by the estimation of parameters that are zero in the population (Clark & McCracken, 2001). Therefore, we expect the forecasting errors of the log-ARCH to be smaller than those of the network log-ARCH. To test the predictive accuracy of nested models, we also include the results of the Clark and West (2007) test that is approximately normal under squared forecasting errors. In the end, to compare forecasts obtained from multiple models, we consider the model confidence set (MCS) procedure of Hansen et al. (2011) for sequential testing based on the statistics in (19).

4. Results of the out-of-sample forecasting exercise

Below, we discuss the results of the out-of-sample exercise. We aimed to predict the log volatility for the $n = 29$ stocks considered in Table 1. Forecasts obtained from 13 alternative models are compared. The time-series log-ARCH is the selected benchmark, as it exploits temporal information only. For the network log-ARCH models, we considered three different distance measures and four alternative weight definitions (i.e., inverse-distance, $k = 3$-nn, $k = 5$-nn, and $k = 10$-nn), resulting on 12 different network models.

4.1. Does the network-based approach improve the forecasting accuracy?

In what follows, we evaluate whether the additional information from the network nodes is useful for forecasting volatility. An overview of the results in terms of both the MAFE and RMSFE, computed considering the average forecasting errors for the n stocks, is shown in Table 2. The p-values associated with the tests of equal predictive accuracy against the benchmark are also reported. The first column shows the log-ARCH results, while the other columns report the forecasting results of different network structures. The stock-specific results are reported in Appendix. More precisely, Table A.1 shows the comparison in terms of predictive accuracy between the benchmark and the best network model, while Table A.2 compares the benchmark with the worst network model. Interestingly, all the network-based models provide more accurate out-of-sample predictions than the log-ARCH, because of their increased model flexibility.

Let us consider the RMSFE results first. The log-ARCH model provides an average RSME of 2.82, while the best network log-ARCH model ($k = 3$ nearest neighbours with Euclidean distance, i.e. model B.3.1) reaches an average RMSFE of 2.44, which is about 15% lower. The worst network model ($k = 3$ nearest neighbours with correlation-based distance, i.e. model B.3.2), instead, has an average RMSFE of 2.55 thus providing a not negligible improvement in the forecasting accuracy compared to the log-ARCH. Considering MAFE loss, the log-ARCH has an average value of 2, which is much larger than 1.85, which is the average MAFE obtained with the best network model ($k = 10$ nearest neighbours with volatility-based distance, i.e. model B.10.3). The worse network model in terms of MAFE ($k = 3$ nearest neighbours with correlation-based distance, i.e. model B.3.2) has an average value of 1.94 which is still lower than the log-ARCH. The superiority of the network-based log-ARCH models is supported by the predictive accuracy tests, as we reject the null hypothesis for all the considered network log-ARCH models and for both the considered forecasting losses.

From a first view, we can improve the forecasting accuracy of the log-ARCH by using any network structure. In this respect, it is worth noting that even if we include the market volatility factor as an explanatory variable in the log-ARCH model, we do not find significantly better forecasting results compared with the network log-ARCH. This confirms that the network interactions account for common market factors to a certain extent. However, not all network structures are the same regarding out-of-sample forecasting accuracy. First, a comparison across network models highlights that the k-nearest-neighbour approach leads to the construction of more effective financial networks from a financial point of view. In other words, fully connected networks obtained with the inverse-distance approach are not the best forecasting choice. To be precise, the fully connected network constructed with the volatility-based approach (13) (model A.3) provides the best forecasts compared to the other two inverse-distance approaches (models A.1 and A.2) in terms of both the RMSFE and MAFE. However, almost all k-nearest-neighbour networks provide more accurate forecasts in the validation set.

The stock-specific results, reported in Appendix (see Tables A.1 and A.2), suggest that even the worse network approach provides statistically more accurate forecasts out of sample than the benchmark. Interestingly, outlier stocks highlighted in some network structures, e.g. BA and CRM for both returns or AXP and BA for volatility distances, are better predicted considering network information. This result is not straightforward, because it is reasonable to assume that the more distant the adjacent nodes, the less relevant their information for forecasting.

4.2. Does the network structure matter?

The previous results showed that network-based log-ARCH models are useful for predicting volatilities. However, Table 2 highlights differences across the alternative network models in forecasting accuracy. For example, it is clear that the k-nearest-neighbour network provides more accurate forecasts on average than inverse-distance approaches. Thus, we may raise the question of whether the network structure matters. In other words, how can the best-fitting network be interpreted from a financial perspective?

To get more insights about the issue of finding the best network log-ARCH model, we apply the model confidence set (MCS) procedure of Hansen et al. (2011), which aims to find a smaller set of network models with the same statistical performance. The MCS procedure does not include in the superior sets network ARCH models with statistically lower forecasting performance. As in the previous assessment, we consider the results of the MCS procedure in terms of both squared and absolute

Table 2
Comparison of the forecasting results using different edge weight matrices. The best RMSFE/MAFE in each row is printed in bold. We report the average performance across stocks. The p-values of the Diebold and Mariano (2002) and Clark and West (2007) tests are reported. Under the null hypothesis of both tests, the benchmark log-ARCH model has better or equal performance than the network log-ARCH with alternative network definitions.

Model	Log-ARCH	Inverse-distance weights			k-nearest-neighbour weights								
	Benchmark	A.1	A.2	A.3	B.k.1			B.k.2			B.k.3		
k					3	5	10	3	5	10	3	5	10
Average RMSFE	2.8202	2.5324	2.5229	2.5194	**2.4457**	2.4904	2.4990	2.5476	2.4668	2.5017	2.4940	2.4571	2.4568
DM test (p-value)	–	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
CW test (p-value)	–	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001
Average MAFE	2.0018	1.9000	1.8898	1.8896	1.8669	1.8905	1.8960	1.9369	1.8733	1.8920	1.8953	1.8675	**1.8559**
DM test (p-value)	–	0.0076	0.0035	0.0038	0.0007	0.0005	0.0040	0.0873	0.0013	0.0015	0.0100	0.0008	0.0001

Table 3
Model confidence set: superior set of models – MSE and MAFE losses for the average errors. $e_{R,M}$ is the elimination rule, "p-val." is the MCS p-value, and "Loss" is the associated (MSE or MAFE) loss value. "Rank" provides the ranking of the models within the superior set in terms of the selected loss function.

Network structure	Rank	$e_{R,M}$	p-val	Loss
Panel A: Average squared error loss				
k-NN with (11) and $k = 3$ (B.3.1)	1	−1.75	1.00	6.033726
k-NN with (13) and $k = 3$ (B.5.3)	3	1.11	0.39	6.093567
k-NN with (13) and $k = 10$ (B.10.3)	2	0.79	0.62	6.088585
Panel B: Average absolute error loss				
k-NN with (11) and $k = 3$ (B.3.1)	2	0.85	0.56	1.866924
k-NN with (13) and $k = 5$ (B.5.3)	3	1.12	0.39	1.867474
k-NN with (13) and $k = 10$ (B.10.3)	1	−1.79	1.00	1.855948

forecasting errors. The results are shown in Table 3, where Panel A reports the results under the squared error loss, while Panel B reports them under the absolute error loss.

Interestingly, Table 3 highlights that the superior set composition is the same regardless of the adopted loss. Indeed, only three network-based ARCH models belong to the superior set, and all of them are based on networks constructed according to the k-nearest-neighbour procedure (15). In particular, the models included in the superior set are the Euclidean distance, which is based on returns, with $k = 3$, and the log-ARCH distance, which is based on volatilities, with $k = 5$ and $k = 10$.

However, it is interesting that the best model in the superior set differs according to squared and absolute forecasting errors. In the case of squared forecasting errors, the Euclidean distance (11) with $k = 3$ nearest neighbours provides the lowest loss, while in the case of absolute error, the best model is $k = 10$ nearest neighbours under log-ARCH distance (13). Therefore, we can conclude that we only need the information from a few adjacent stocks in forecasting under returns-based networks. By contrast, information from a higher number of nodes is required for volatility-based networks.

Overall, the results confirm that fully connected networks provide less accurate forecasts out of sample; thus, k-nearest-neighbour approaches should be preferred in practice. Moreover, the results suggest that although correlation-based approaches are the most widely used in the construction of financial networks, the correlation-based network ARCH is not included in the superior set. This means that information included in most correlated stocks is not as valuable for out-of-sample exercises as it appeared previously.

In summary, Table 3 shows that the network structure matters in terms of out-of-sample forecasting accuracy. Therefore, researchers and practitioners have to carefully specify the kind of network underlying the network ARCH model, even though the forecasting performance is good when not choosing the best network. Using a suitable network structure, it is possible to enhance the forecasting ability of the model further. The choice of the network can rely on the user's experience or on a cross-validation study. A third suitable approach that we evaluate next in Section 4.3 is the combination of alternative network-based models. In fact, if the prediction performance increases by combination, practitioners can be agnostic about the best network specification to adopt and combine the forecasts obtained by alternative models.

4.3. Can the prediction performance be increased by considering multiple network definitions?

Finally, we ask whether it is possible to improve the volatility forecasting accuracy with network log-ARCH models. A suitable idea is to use a combination of forecasts from the alternative models considered in the paper. Forecasting combination, also known as ensemble forecasting, is a technique used to improve the accuracy of predictions by combining multiple forecasts. The basic idea is that by combining the predictions of different models, the strengths of each can be leveraged to produce more accurate forecasting. By combining forecasts from multiple models, forecasters can reduce the risk of relying on a single model. The use of combination methods is nowadays widespread not only in economics see e.g. Proietti & Giovannelli, 2021 but also in other research areas such as sociology (Tollenaar & van der Heijden, 2013), epidemiology (Deb & Deb, 2022) and, meteorology (Di Narzo & Cocchi, 2010). In the context of volatility forecasting, ensemble techniques are also commonly considered (Becker & Clements, 2008).

Although there are many ways of combining forecasts, we consider the three most common approaches: the simple average, minimum-variance combination, and constrained OLS (COLS) for details, see Timmermann, 2006. For the simple average method, forecasts are obtained by averaging the predictions from the alternative models. Although straightforward, there is wide evidence supporting the superiority of simple averaging compared with optimal combination approaches (i.e. the so-called forecast-combination puzzle). In the case of the minimum-variance approach, combination weights are obtained by minimising the resulting forecasting error variance. In the COLS combination, the weights are obtained as the parameters of a linear regression, with a constraint on the parameters such that they sum up to one. In particular, the actual values of the time series to be predicted are regressed on the set of the alternative forecasts.

Below, we combine for each stock the forecasts obtained with both log-ARCH and the network log-ARCH, under the three aforementioned ensemble approaches. Then, we evaluate whether combination further enhances forecasting accuracy in the validation sample. Table 4 shows the comparison in terms of both the RMSFE and MAFE computed on average forecasting errors between the three competing combination approaches. Panel A of Table 4 shows the average accuracy metrics, while Panel B and Panel C show the relative accuracy of each combination approach with both the benchmark and the best network-based model. The ensemble forecast results can be compared with those of Table 2. Furthermore, Panel B and Panel C include the p-values associated with the Diebold and Mariano (2002) and Clark and West (2007) predictive accuracy tests.

In terms of the average RMSFE and MAFE, the best combination approach is represented by COLS. For example, the best network approach of Table 2 has an average RMSFE of 2.45. With the COLS combination, we reduce it to 2.35, which is about 5% lower. In terms of absolute errors, the best network achieves an average MAFE of 1.86, while with the COLS combination, we reduce the loss to 1.76, which is about 6% smaller. Therefore, the improvements in the forecasting accuracy with the combination are not negligible.

Interestingly, the simple average is the worst combination scheme to adopt. Contrary to this evidence, however, the best benchmark model performs (a bit) worse than this relatively easy combination approach. This suggests that we can improve forecasting accuracy with little effort or, more generally, by proficiently handling uncertainty about what model to use in a straightforward manner. The minimum-variance combination also improves forecasting accuracy, albeit by less than COLS.

Stock-specific results of the combination schemes are shown in Tables A.3 and A.4 in Appendix. In particular, Table A.3 shows the RMSFE and MAFE for each single stock and combination approach, while Table A.4 shows the results of stock-specific predictive accuracy tests, comparing the best combination approach with the best network log-ARCH model.

Based on the results presented in Table 4, combining forecasts obtained from the benchmark model and several network log-ARCH models appears more advantageous than relying solely on a single model. This approach is beneficial for addressing the uncertainty associated with selecting an appropriate network structure. As discussed in Section 4.2, the forecasts generated from different network structures are statistically different. While selecting the most suitable network structure is crucial, doing so ex ante can be challenging and complex. By combining forecasts obtained from different networks, we can enhance the accuracy of out-of-sample forecasts and alleviate concerns about selecting the appropriate network structure.

Table 4
Ensemble forecasting (average values). Results are reported in terms of the RMSFE and MAFE accuracy metrics. Relative performance in terms of the benchmark log-ARCH model and the best network log-ARCH model is also evaluated. "Min. var." indicates the minimum-variance ensemble, while "COLS" is the constrained OLS approach.

	Simple Average	Min. var.	COLS
Panel A – Accuracy metrics			
RMSFE	2.4596	2.3703	2.3572
MAFE	1.8319	1.7684	1.7634
Panel B – Relative accuracy (benchmark):			
Relative RMSFE	1.1465	1.1897	1.1964
DM test (*p*-value)	0.0000	0.0000	0.0000
CW test (*p*-value)	0.0000	0.0000	0.0000
Relative MAFE	1.0927	1.1319	1.1351
DM test (*p*-value)	0.0001	0.0000	0.0000
Panel C – Relative accuracy (best network model):			
Relative RMSFE	0.9943	1.0317	1.0375
DM test (*p*-value)	0.0000	0.0000	0.0000
CW test (*p*-value)	0.0000	0.0000	0.0000
Relative MAFE	1.0131	1.0494	1.0524
DM test (*p*-value)	0.0000	0.0000	0.0000

Table 5
Model confidence set, including forecast combination approaches: superior set of models – MSE and MAFE losses for the average errors. Here, $e_{R,M}$ is the elimination rule, "p-val" is the MCS *p*-value, "Loss" is the associated (MSE or MAFE) loss value, and "Rank" provides the ranking of the models within the superior set in terms of the selected loss function.

Network structure	Rank	$e_{R,M}$	p-val	Loss
Panel A: Squared error loss				
COLS combination	1	−5.52	1.00	5.56
Panel B: Absolute error loss				
COLS combination	1	−2.04	1.00	1.76

Finally, we provide a comparison across all the models considered in the paper—that is, the benchmark log-ARCH, all the network log-ARCH models under alternative network specifications, and their combination by means of the model confidence set (see Table 5). Interestingly, we find that the superior set includes only the COLS-based network combination. This finding is very important because we do not find any model competing with the best combination approach. Therefore, we recommend combining the results obtained from different network definitions.

In sum, the results highlighted in Tables 4 and 5 provide a suggestion to practitioners—that is, to combine different network log-ARCH models. This suggestion is

Table A.1
Diebold and Mariano (2002) and Clark and West (2007) predictive accuracy tests: benchmark vs. best network model. Under the null, the benchmark model (log-ARCH) has better or equal predictive accuracy than the best network log-ARCH approach.

Stock	Diebold and Mariano (2002)				Clark and West (2007)	
	Squared errors		Absolute errors			
	DM stat	p-value	DM stat	p-value	CW stat	p-value
AAPL	5.10	0.00	45.39	0.00	11.61	0.00
AMGN	10.20	0.00	42.37	0.00	16.97	0.00
AXP	3.24	0.00	41.08	0.00	9.72	0.00
BA	3.30	0.00	38.80	0.00	10.01	0.00
CAT	5.99	0.00	47.71	0.00	12.27	0.00
CRM	3.74	0.00	47.03	0.00	11.49	0.00
CSCO	7.73	0.00	50.06	0.00	13.72	0.00
CVX	4.89	0.00	41.18	0.00	10.48	0.00
DIS	3.86	0.00	41.18	0.00	9.04	0.00
GS	6.41	0.00	44.79	0.00	11.01	0.00
HD	4.03	0.00	43.63	0.00	10.42	0.00
HON	6.24	0.00	41.39	0.00	11.45	0.00
IBM	7.33	0.00	45.51	0.00	12.78	0.00
INTC	4.69	0.00	40.16	0.00	10.28	0.00
JNJ	5.59	0.00	42.88	0.00	10.52	0.00
JPM	4.78	0.00	47.30	0.00	11.53	0.00
KO	5.69	0.00	46.59	0.00	11.46	0.00
MCD	6.70	0.00	51.82	0.00	13.49	0.00
MMM	6.90	0.00	39.89	0.00	12.45	0.00
MRK	7.18	0.00	47.51	0.00	12.11	0.00
MSFT	4.03	0.00	52.28	0.00	10.20	0.00
NKE	3.14	0.00	43.75	0.00	9.85	0.00
PG	5.52	0.00	41.48	0.00	11.59	0.00
TRV	4.57	0.00	34.94	0.00	10.86	0.00
UNH	10.19	0.00	45.27	0.00	14.57	0.00
V	4.51	0.00	49.77	0.00	11.18	0.00
VZ	0.98	0.16	23.34	0.00	11.95	0.00
WBA	6.73	0.00	43.15	0.00	12.14	0.00
WMT	5.27	0.00	51.70	0.00	12.34	0.00

particularly relevant in the absence of any information about the most suitable network structure to adopt.

5. Conclusion

In this paper, we proposed a novel approach to forecasting volatility that extends the log-ARCH to incorporate the network structure of financial time series. The stock market is well represented by networks, where stocks are the nodes and the edges reflect the degree of similarity across them. By including the network connectives in the statistical model, we explicitly introduce the effect of instantaneous spillovers from adjacent nodes reflecting the simultaneity of investors' trading decisions. The information from the adjacent nodes of a financial network can be used for forecasting purposes.

There are many different ways of constructing financial networks. We evaluated the performance of 12 alternative network log-ARCH configurations. Inspired by the time series clustering literature, three alternative dissimilarity definitions were considered for constructing the networks, i.e. the Euclidean distance across returns, a correlation-based approach, and a volatility-based approach. In addition, networks were considered fully connected, employing an inverse-distance approach, and not fully connected, utilising k-nearest neighbours with $k = \{3, 5, 10\}$. In practice, a suitable underlying financial network structure could be chosen in a cross-validation study. However, given the good results obtained with the combination of alternative network log-ARCH models for forecasting out-of-sample volatility, we recommend using the combination in the absence of any a priori information. Finally, we used the proposed modelling approach to forecast the out-of-sample volatility of the stocks in the Dow Jones Index. The network structure would also allow modelling larger financial networks, even in cases where n is larger than T, which would be an interesting point for future research.

First, we found that the forecasting accuracy of log-ARCH models significantly increases when including network information. This means that the information on adjacent network nodes is helpful for forecasting volatility. Moreover, we showed that the network structure matters regarding out-of-sample forecasting accuracy. In particular, we found that networks constructed with inverse distance seem less effective at forecasting than those based on k-nearest neighbours (15). Thus, fully connected networks appear not to be the best forecasting choice. Moreover, we found three alternative network log-ARCH models belonging to the superior set, as suggested by Hansen et al. (2011). Interestingly, none of these models adopts a correlation-based network, although this is one of the most common choices for constructing financial networks.

We can highlight two possible future research directions. First, we included stock-specific constant terms in the volatility equation. However, future research could

Table A.2
Diebold and Mariano (2002) and Clark and West (2007) predictive accuracy tests: benchmark vs. worst network model. Under the null, the benchmark model (log-ARCH) has better or equal predictive accuracy than the worst network log-ARCH approach.

Stock	Diebold and Mariano (2002)				Clark and West (2007)	
	Squared errors		Absolute errors			
	DM stat	*p*-value	DM stat	*p*-value	CW stat	*p*-value
AAPL	5.12	0.00	72.30	0.00	12.37	0.00
AMGN	4.62	0.00	35.63	0.00	11.78	0.00
AXP	2.74	0.00	59.40	0.00	10.07	0.00
BA	1.60	0.05	49.00	0.00	8.66	0.00
CAT	5.77	0.00	58.68	0.00	12.77	0.00
CRM	4.22	0.00	51.47	0.00	12.02	0.00
CSCO	7.67	0.00	54.11	0.00	14.51	0.00
CVX	3.46	0.00	53.53	0.00	10.12	0.00
DIS	3.72	0.00	53.13	0.00	9.69	0.00
GS	5.34	0.00	64.53	0.00	10.84	0.00
HD	3.75	0.00	50.69	0.00	11.19	0.00
HON	5.21	0.00	51.50	0.00	11.79	0.00
IBM	6.94	0.00	52.73	0.00	13.59	0.00
INTC	4.20	0.00	56.41	0.00	11.25	0.00
JNJ	5.65	0.00	47.31	0.00	11.75	0.00
JPM	4.16	0.00	62.63	0.00	11.77	0.00
KO	5.13	0.00	39.35	0.00	12.08	0.00
MCD	6.93	0.00	48.82	0.00	13.50	0.00
MMM	6.92	0.00	49.24	0.00	13.21	0.00
MRK	5.01	0.00	35.67	0.00	11.25	0.00
MSFT	3.77	0.00	63.15	0.00	11.70	0.00
NKE	2.72	0.00	55.61	0.00	10.37	0.00
PG	5.20	0.00	38.07	0.00	11.90	0.00
TRV	3.30	0.00	45.66	0.00	11.06	0.00
UNH	9.66	0.00	45.74	0.00	14.75	0.00
V	3.62	0.00	46.29	0.00	10.55	0.00
VZ	−3.52	0.00	21.45	0.00	5.80	0.00
WBA	7.28	0.00	47.04	0.00	13.12	0.00
WMT	5.24	0.00	40.49	0.00	12.49	0.00

Table A.3
Ensemble forecasting for each stock. Results are reported in terms of both the RMSFE and MAFE accuracy metrics. “Min. var.” indicates the minimum-variance ensemble, while “COLS” is the constrained OLS approach.

Stock	Simple average		Min. var.		COLS	
	RMSFE	MAFE	RMSFE	MAFE	RMSFE	MAFE
AAPL	2.3522	1.8440	2.3263	1.8136	2.3209	1.8129
AMGN	2.4915	1.8271	2.2754	1.6764	2.2399	1.6625
AXP	2.3537	1.8207	2.2533	1.7024	2.2465	1.7019
BA	2.5508	1.9933	2.4392	1.8734	2.4171	1.8550
CAT	2.4031	1.8318	2.3514	1.8150	2.3381	1.8056
CRM	2.5308	1.9261	2.4094	1.8271	2.4021	1.8318
CSCO	2.2689	1.7101	2.2476	1.7247	2.2419	1.7226
CVX	2.3810	1.8424	2.2831	1.7204	2.2431	1.6966
DIS	2.6537	1.9728	2.6170	1.8982	2.5961	1.8808
GS	2.3094	1.7037	2.2547	1.6654	2.2500	1.6718
HD	2.2341	1.7274	2.2090	1.6892	2.2041	1.6922
HON	2.3814	1.8023	2.3249	1.7485	2.3156	1.7451
IBM	2.2139	1.6232	2.1901	1.6131	2.1803	1.6070
INTC	2.4565	1.8827	2.4006	1.8101	2.3908	1.8050
JNJ	2.6467	1.7668	2.6123	1.7642	2.5844	1.7545
JPM	2.3275	1.7554	2.2674	1.7184	2.2550	1.7090
KO	2.3686	1.7281	2.3178	1.6926	2.3154	1.6940
MCD	2.3091	1.7454	2.2816	1.7295	2.2680	1.7209
MMM	2.5672	1.9107	2.5072	1.9084	2.4964	1.9110
MRK	3.2397	2.0732	3.1407	2.0694	3.1319	2.0739
MSFT	2.6221	1.9454	2.5923	1.8978	2.5889	1.8992
NKE	2.2772	1.7594	2.2148	1.6787	2.1956	1.6806
PG	2.3088	1.7356	2.2350	1.7131	2.2317	1.7125

(*continued on next page*)

Table A.3 (*continued*).

Stock	Simple average		Min. var.		COLS	
	RMSFE	MAFE	RMSFE	MAFE	RMSFE	MAFE
TRV	2.1022	1.6435	2.0094	1.5689	1.9966	1.5610
UNH	2.4257	1.7304	2.3590	1.7597	2.3481	1.7681
V	2.1860	1.6969	2.0752	1.6361	2.0660	1.6341
VZ	3.1516	2.4379	2.5011	1.9102	2.4768	1.8846
WBA	2.4308	1.8502	2.3965	1.8413	2.3841	1.8278
WMT	2.4345	1.8394	2.3940	1.8205	2.3793	1.8184

Table A.4
Diebold and Mariano (2002) and Clark and West (2007) predictive accuracy tests: best combination approach vs. best network model. Under the null, the best combination approach has better or equal predictive accuracy than the best network log-ARCH approach.

Stock	Diebold and Mariano (2002)				Clark and West (2007)	
	Squared errors		Absolute errors			
	DM stat	p-value	DM stat	p-value	CW stat	p-value
AAPL	2.40	0.01	2.99	0.00	4.57	0.00
AMGN	4.05	0.00	2.65	0.00	6.86	0.00
AXP	3.88	0.00	3.54	0.00	6.48	0.00
BA	4.30	0.00	3.03	0.00	6.22	0.00
CAT	2.67	0.00	2.93	0.00	4.55	0.00
CRM	3.63	0.00	2.81	0.00	5.85	0.00
CSCO	2.09	0.02	2.90	0.00	3.49	0.00
CVX	4.98	0.00	2.89	0.00	7.36	0.00
DIS	3.60	0.00	3.09	0.00	6.16	0.00
GS	2.73	0.00	3.03	0.00	4.53	0.00
HD	2.54	0.01	2.85	0.00	4.93	0.00
HON	3.02	0.00	2.93	0.00	5.37	0.00
IBM	1.85	0.03	2.75	0.00	3.87	0.00
INTC	2.70	0.00	2.95	0.00	5.00	0.00
JNJ	2.51	0.01	2.81	0.00	4.38	0.00
JPM	2.83	0.00	3.09	0.00	5.42	0.00
KO	3.40	0.00	2.66	0.00	5.74	0.00
MCD	3.03	0.00	2.78	0.00	5.48	0.00
MMM	2.60	0.00	2.99	0.00	5.09	0.00
MRK	2.14	0.02	2.92	0.00	3.88	0.00
MSFT	2.32	0.01	3.00	0.00	4.90	0.00
NKE	3.38	0.00	2.84	0.00	6.60	0.00
PG	2.33	0.01	2.72	0.00	4.92	0.00
TRV	4.16	0.00	2.66	0.00	6.83	0.00
UNH	1.99	0.02	2.76	0.00	3.66	0.00
V	4.53	0.00	2.78	0.00	7.41	0.00
VZ	4.71	0.00	2.78	0.00	8.81	0.00
WBA	2.74	0.00	2.78	0.00	5.16	0.00
WMT	2.78	0.00	2.89	0.00	4.76	0.00

extend these stock-specific constant terms to a common-factor representation. Second, whereas we assumed the network structure to be constant over time, time-dynamic adjacency matrices or regime-switching network log-ARCH models could also be developed in future studies. In this case, it would be interesting to consider the estimated correlation structure from DCC models.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix. Additional results for single stocks

See Tables A.1–A.4.

References

Andersen, T., & Bollerslev, T. (1998). 'Answering the skeptics: Yes, standard volatility models do provide accurate forecasts'. *International Economic Review*, *88*, 5–905.

Asgharian, H., Hess, W., & Liu, L. (2013). 'A spatial analysis of international stock market linkages'. *Journal of Banking & Finance*, *37*(12), 4738–4754.

Baltagi, B., & Li, J. (2014). 'Further evidence on the spatio-temporal model of house prices in the United States'. *Journal of Applied Econometrics*, *29*(3), 515–522.

Barigozzi, M., & Hallin, M. (2017). 'A network analysis of the volatility of high dimensional financial series'. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, *66*(3), 581–605.

Barker, D., & Loughran, T. (2007). 'The geography of S & P 500 stock returns'. *Journal of Behavioral Finance*, *8*(4), 177–190.

Bauwens, L., & Sucarrat, G. (2010). 'General-to-specific modelling of exchange rate volatility: A forecast evaluation'. *International Journal of Forecasting*, *26*(4), 885–907.

Becker, R., & Clements, A. (2008). 'Are combination forecasts of S & P 500 volatility statistically superior?'. *International Journal of Forecasting 24*(1), 122–133.

Betancourt, B., Rodríguez, A., & Boyd, N. (2020). 'Modelling and prediction of financial trading networks: An application to the New York Mercantile Exchange natural gas futures market'. *Journal of the Royal Statistical Society. Series C. Applied Statistics*, *69*(1), 195–218.

Billio, M., Caporin, M., Frattarolo, L., & Pelizzon, L. (2021). 'Networks in risk spillovers: A multivariate GARCH perspective'. *Econometrics and Statistics*, *28*, 1–29.

Bollerslev, T. (1986). 'Generalized autoregressive conditional heteroskedasticity'. *Journal of Econometrics*, *31*(3), 307–327.

Caiado, J., Crato, N., & Peña, D. (2006). 'A periodogram-based metric for time series classification'. *Computational Statistics & Data Analysis*, *50*(10), 2668–2684.

Caiado, J., Crato, N., & Poncela, P. (2020). 'A fragmented-periodogram approach for clustering big data time series'. *Advances in Data Analysis and Classification*, *14*, 117–146.

Cameletti, M., Gomez-Rubio, V., & Blangiardo, M. (2019). 'Bayesian modelling for spatially misaligned health and air pollution data through the INLA-SPDE approach'. *Spatial Statistics*, *31*, Article 100353.

Caporin, M., & McAleer, M. (2012). 'Do we really need both BEKK and DCC? A tale of two multivariate GARCH models'. *Journal of Economic Surveys*, *26*(4), 736–751.

Caporin, M., & Paruolo, P. (2015). 'Proximity-structured multivariate volatility models'. *Econometric Reviews*, *34*(5), 559–593.

Cerqueti, R., D'Urso, P., De Giovanni, L., Giacalone, M., & Mattera, R. (2022). 'Weighted score-driven fuzzy clustering of time series with a financial application'. *Expert Systems with Applications*, *198*, Article 116752.

Cerqueti, R., Giacalone, M., & Mattera, R. (2021). 'Model-based fuzzy time series clustering of conditional higher moments'. *International Journal of Approximate Reasoning*, *134*, 34–52.

Cerqueti, R., & Mattera, R. (2023). 'Fuzzy clustering of time series with time-varying memory'. *International Journal of Approximate Reasoning*, *153*, 193–218.

Clark, T., & McCracken, M. (2001). 'Tests of equal forecast accuracy and encompassing for nested models'. *Journal of Econometrics*, *105*(1), 85–110.

Clark, T., & West, K. (2007). 'Approximately normal tests for equal predictive accuracy in nested models'. *Journal of Econometrics*, *138*(1), 291–311.

Deb, S., & Deb, S. (2022). 'An ensemble method for early prediction of dengue outbreak'. *Journal of the Royal Statistical Society Series A: Statistics in Society*, *185*(1), 84–101.

Demiris, N., Kypraios, T., & Smith, L. (2014). 'On the epidemic of financial crises'. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, *177*(3), 697–723.

Di Narzo, A., & Cocchi, D. (2010). 'A Bayesian hierarchical approach to ensemble weather forecasting'. *Journal of the Royal Statistical Society. Series C. Applied Statistics*, *59*(3), 405–422.

Diebold, F., & Mariano, R. (2002). 'Comparing predictive accuracy'. *Journal of Business & Economic Statistics*, *20*(1), 134–144.

Diebold, F., & Yılmaz, K. (2014). 'On the network topology of variance decompositions: Measuring the connectedness of financial firms'. *Journal of Econometrics*, *182*(1), 119–134.

Donaldson, R., & Kamstra, M. (1997). 'An artificial neural network-GARCH model for international stock return volatility'. *Journal of Empirical Finance*, *4*(1), 17–46.

Duan, N. (1983). 'Smearing estimate: A nonparametric retransformation method'. *Journal of the American Statistical Association*, *78*(383), 605–610.

D'Urso, P., De Giovanni, L., & Massari, R. (2016). 'GARCH-based robust clustering of time series'. *Fuzzy Sets and Systems*, *305*, 1–28.

D'Urso, P., & Maharaj, E. (2009). 'Autocorrelation-based fuzzy clustering of time series'. *Fuzzy Sets and Systems*, *160*(24), 3565–3589.

Eckel, S., Löffler, G., Maurer, A., & Schmidt, V. (2011). 'Measuring the effects of geographical distance on stock market correlation'. *Journal of Empirical Finance*, *18*(2), 237–247.

Fassò, A., Maranzano, P., & Otto, P. (2022). 'Spatiotemporal variable selection and air quality impact assessment of COVID-19 lockdown'. *Spatial Statistics*, *49*, Article 100549.

Fernández-Avilés, G., Montero, J.-M., & Orlov, A. (2012). 'Spatial modeling of stock market comovements'. *Finance Research Letters*, *9*(4), 202–212.

Francq, C., & Sucarrat, G. (2017). 'An equation-by-equation estimator of a multivariate log-GARCH-X model of financial returns'. *Journal of Multivariate Analysis*, *153*, 16–32.

Francq, C., & Sucarrat, G. (2018). 'An exponential chi-squared QMLE for log-GARCH models via the ARMA representation'. *Journal of Financial Econometrics*, *16*(1), 129–154.

Francq, C., & Zakoian, J.-M. (2019). *GARCH models: Structure, statistical inference and financial applications*. John Wiley & Sons.

Fülle, M., & Otto, P. (2023). 'Spatial GARCH models for unknown spatial locations – An application to financial stock returns'. *Spatial Economic Analysis*, *19*(1), 92–105.

Gallo, G., Lacava, D., & Otranto, E. (2021). 'On classifying the effects of policy announcements on volatility'. *International Journal of Approximate Reasoning*, *134*, 23–33.

Geweke, J. (1986). 'Modeling the persistence of conditional variances: A comment'. *Econometric Reviews*, *5*(1), 57–61.

Hansen, P., Lunde, A., & Nason, J. (2011). 'The model confidence set'. *Econometrica*, *79*(2), 453–497.

Holly, S., Pesaran, M., & Yamagata, T. (2010). 'A spatio-temporal model of house prices in the USA'. *Journal of Econometrics*, *158*(1), 160–173.

Huang, D., Hu, W., Jing, B., & Zhang, B. (2023). 'Grouped spatial autoregressive model'. *Computational Statistics & Data Analysis*, *178*, Article 107601.

Huang, W., Wang, K., Jay Breidt, F., & Davis, R. (2011). 'A class of stochastic volatility models for environmental applications'. *Journal of Time Series Analysis*, *32*(4), 364–377.

Kristjanpoller, W., & Minutolo, M. (2015). 'Gold price volatility: A forecasting approach using the artificial neural network–GARCH model'. *Expert Systems with Applications*, *42*(20), 7245–7251.

Lahmiri, S. (2016). 'Clustering of Casablanca stock market based on Hurst exponent estimates'. *Physica A. Statistical Mechanics and its Applications*, *456*, 310–318.

Lee, L.-f. (2007). 'GMM and 2SLS estimation of mixed regressive spatial autoregressive models'. *Journal of Econometrics*, *137*(2), 489–514.

Lee, L.-f., & Yu, J. (2014). 'Efficient GMM estimation of spatial dynamic panel data models with fixed effects'. *Journal of Econometrics*, *180*(2), 174–197.

Liu, S., Caporin, M., & Paterlini, S. (2021). 'Dynamic network analysis of North American financial institutions'. *Finance Research Letters*, *42*, Article 101921.

Maharaj, E., D'Urso, P., & Caiado, J. (2019). *Time series clustering and classification*. CRC Press.

Mantegna, R. (1999). 'Hierarchical structure in financial markets'. *The European Physical Journal B*, *11*(1), 193–197.

Mattera, R. (2022). 'A weighted approach for spatio-temporal clustering of COVID-19 spread in Italy'. *Spatial and Spatio-temporal Epidemiology*, *41*, Article 100500.

Otranto, E. (2008). 'Clustering heteroskedastic time series by model-based procedures'. *Computational Statistics & Data Analysis*, *52*(10), 4685–4698.

Otto, P., Doğan, O., & Taşpınar, S. (2023). 'Dynamic spatiotemporal ARCH models'. In *Spatial economic analysis*. http://dx.doi.org/10.1080/17421772.2023.2254817.

Otto, P., & Schmid, W. (2018). 'Spatiotemporal analysis of German real-estate prices'. *The Annals of Regional Science*, *60*(1), 41–72.

Otto, P., & Schmid, W. (2022). 'A general framework for spatial GARCH models'. In *Statistical papers* (pp. 1–27).

Otto, P., Schmid, W., & Garthoff, R. (2018). 'Generalised spatial and spatiotemporal autoregressive conditional heteroscedasticity'. *Spatial Statistics*, *26*, 125–145.

Otto, P., Schmid, W., & Garthoff, R. (2021). 'Stochastic properties of spatial and spatiotemporal ARCH models'. *Statistical Papers*, *62*, 623–638.

Piccolo, D. (1990). 'A distance measure for classifying ARIMA models'. *Journal of Time Series Analysis*, *11*(2), 153–164.

Pirinsky, C., & Wang, Q. (2006). 'Does corporate headquarters location matter for stock returns?'. *The Journal of Finance*, *61*(4), 1991–2015.

Piter, A., Otto, P., & Alkhatib, H. (2022). 'The Helsinki bike-sharing system—Insights gained from a spatiotemporal functional model'. *Journal of the Royal Statistical Society Series A: Statistics in Society*, *185*(3), 1294–1318.

Proietti, T., & Giovannelli, A. (2021). 'Nowcasting monthly GDP with big data: A model averaging approach'. *Journal of the Royal Statistical Society Series A: Statistics in Society*, *184*(2), 683–706.

Sahu, S., & Böhning, D. (2022). 'Bayesian spatio-temporal joint disease mapping of Covid-19 cases and deaths in local authorities of England'. *Spatial Statistics*, *49*, Article 100519.

Sato, T., & Matsuda, Y. (2017). 'Spatial autoregressive conditional heteroskedasticity models'. *Journal of the Japan Statistical Society*, *47*(2), 221–236.

Sato, T., & Matsuda, Y. (2021). 'Spatial extension of generalized autoregressive conditional heteroskedasticity models'. *Spatial Economic Analysis*, *16*(2), 148–160.

Sucarrat, G., & Escribano, A. (2012). 'Automated model selection in finance: General-to-specific modelling of the mean and volatility specifications'. *Oxford Bulletin of Economics and Statistics*, *74*(5), 716–735.

Sucarrat, G., Grønneberg, S., & Escribano, A. (2016). 'Estimation and inference in univariate and multivariate log-GARCH-X models when the conditional density is unknown'. *Computational Statistics & Data Analysis*, *100*, 582–594.

Timmermann, A. (2006). 'Forecast combinations'. *Handbook of Economic Forecasting*, *1*, 135–196.

Tollenaar, N., & van der Heijden, P. (2013). 'Which method predicts recidivism best?: A comparison of statistical machine learning and data mining predictive models'. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, *176*(2), 565–584.

Tumminello, M., Lillo, F., & Mantegna, R. (2010). 'Correlation, hierarchies, and networks in financial markets'. *Journal of Economic Behaviour and Organization*, *75*(1), 40–58.

Vinciotti, V., Tosetti, E., Moscone, F., & Lycett, M. (2019). 'The effect of interfirm financial transactions on the credit risk of small and medium-sized enterprises'. *Journal of the Royal Statistical Society Series A*, *182*(4), 1205–1226.

Wu, J., Xu, K., Chen, X., Li, S., & Zhao, J. (2022). 'Price graphs: Utilizing the structural information of financial time series for stock prediction'. *Information Sciences*, *588*, 405–424.

Zhou, Y., Chen, Z., & Liu, Z. (2023). 'Dynamic analysis and community recognition of stock price based on a complex network perspective'. *Expert Systems with Applications*, *213*, Article 118944.

Zhou, J., Li, D., Pan, R., & Wang, H. (2020). 'Network GARCH model'. *Statistica Sinica*, *30*(4), 1723–1740.