Socio-Economic Planning Sciences 95 (2024) 102023

Contents lists available at ScienceDirect

Socio-Economic Planning Sciences

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

Investors' attention and network spillover for commodity market forecasting

Roy Cerqueti [a,b], Valerio Ficcadenti [c], Raffaele Mattera [a,*]

[a] Department of Social and Economic Sciences, Sapienza University of Rome, Rome, Italy
[b] GRANEM, University of Angers, Angers, France
[c] School of Business, London South Bank University, London, United Kingdom

ARTICLE INFO

Keywords:
Dynamic network model
Google Trends
Factor model
Prediction
Principal components
Commodity returns

ABSTRACT

This paper explores the role of network spillovers in commodity market forecasting and proposes a novel factor-augmented dynamic network model. We focus on a novel network definition based on investors' attention to commodities, positing that commodities exhibit spillovers if they share a similar level of interest. To this aim, we employ Google Trends search data as an instrumental measure for attention. The results reveal that including attention-driven spillovers significantly enhances the forecasting accuracy of commodity returns.

1. Introduction

Commodities play a relevant role in various industries, offering crucial insights into economic trends and often being part of the production process, as either input, output, or both. Therefore, the need for accurate commodity forecasts goes beyond the risk management task, as forecasts have a significant impact on strategic planning, production optimisation, supply chain management, and regulatory compliance [1–3]. Businesses heavily reliant on commodities stand to benefit significantly from accurate forecasts, and the ability of financial services to deliver dependable commodity return forecasts becomes integral in ensuring a functioning operation of supply chains. Indeed, accurate forecasts are instrumental in aligning production strategies with anticipated market trends, contributing to more efficient production cycles. Therefore, the process evaluation, which is important in terms of organizational efficiency, is intricately connected to the accuracy of commodity's performance forecasting [4].

Therefore, viewing commodity forecasting through the lenses of business decision-making, production optimisation, and process evaluation highlights its indispensable role in the strategic toolkit of both production firms and financial services [5–7]. The capacity to provide accurate forecasts enhances decision-making efficiency and contributes significantly to the overall operational excellence of businesses across diverse sectors [8].

This paper contributes to the debate on forecasting methods and models of commodities' returns. Specifically, we propose a novel factor-augmented dynamic network model (Fa-DNm) by including the spillover effect of the investors' attention to commodities. In this respect, a crucial step of the forecasting procedure is the identification of the way investors' attention intervenes in connecting the considered commodities. This relevant aspect is faced through complex network theory which is a versatile instrument to describe disaggregated entities along with their interconnections [9,10].

Indeed, integrating complex networks in commodity return forecasting adds a layer of sophistication and reliability to the evaluation process. In our specific case, understanding the intricate interdependencies within commodity markets and related financial instruments is essential for providing comprehensive and accurate forecasts. Complex network analysis aids in identifying hidden patterns, correlations, and systemic risks, contributing to a more nuanced evaluation of financial services. The relevance of complex network structures and the impact of network spillovers on forecasting commodities have gained increasing attention in the literature [11].

The papers above acknowledge complex networks' relevance for improving forecasting procedures. However, to the best of our knowledge, no contributions fully leverage the potential of internet-derived data for networking commodities.

On this, we notice that data collected from the internet allows researchers to build on and sometimes replace traditional methods in forecasting practice. While the use of internet data has become a common practice for leveraging the accuracy of modern forecasting approaches, this information has predominantly been employed as an additional predictor in statistical models [e.g., see 12,13]. Differently from previous studies, this paper introduces internet data sources, specifically Google Trends search data, in the context of network modeling. Following previous literature [14–16], we employ Google Trends search data as an instrumental measure for investors'

* Corresponding author.
E-mail address: raffaele.mattera@uniroma1.it (R. Mattera).

https://doi.org/10.1016/j.seps.2024.102023
Received 30 January 2024; Received in revised form 12 June 2024; Accepted 15 July 2024
Available online 22 July 2024

attention [see, for example, 17,18, where Google searchers volume have been used similalrly]. Our main contribution lies in the integration of internet data, not as supplementary variable [19–22], but as a fundamental component for detecting and modeling spillovers within the commodity market. By using Google Trends search data, we depart from conventional practices and introduce a dynamic network framework that captures the evolving relationships among commodities based on attention-driven spillovers. This approach allows for a real-time assessment of the investor's interest in various commodities and provides a more nuanced understanding of the interconnectedness within the market.

A promising approach for improving forecast accuracy is the use of factor models. Factor models are commonly used by forecasters in both finance and macroeconomics [23,24], but their use in the commodity market is more recent [e.g., see 14,25,26]. Factor-augmented models have been proposed to account for possible (unobservable) factors that may affect commodity returns. For instance, latent factors may be related to the fluctuations in global economic activity [27] or to the inventory levels and storage conditions [28]. Weaker global economic activity acts as a negative demand shock for commodities, but also storage demand or other commodity-specific demand shocks can affect commodity prices. In this regard, the traditional theory of storage assumes that holders of inventories receive implicit benefits that decline as inventory increases. Therefore, information about the disruptions to storage conditions and inventory levels can be included in the estimated latent factors. For this reason, in this paper, we propose the use of a factor-augmented model with network interactions for forecasting commodity returns.

Our findings reveal the existence of substitution effects between commodities since the returns of commodities with similar levels of attention are negatively correlated. Moreover, the inclusion of attention-driven spillovers significantly enhances the forecasting capabilities of the model. By adopting a dynamic network perspective, the developed model not only accommodates the presence of spillovers but also demonstrates their successful use in the forecasting task.

The paper is structured as follows. Section 2 presents a brief overview of the previous papers related to the use of factor models and internet data to instrument investors' attention in the context of commodity market forecasting. Section 3 discusses the forecasting method adopted in the paper, that is a Fa-DNm, and the network definition implied by Google Trends data. Section 4 discusses the data adopted for the empirical study, that is commodity returns and Google Trends volume searches, as well as the in-sample estimation of the model. Section 5 presents the forecasting experiment and the main results, comparing the accuracy of the forecasts obtained with the proposed approach and those based on different benchmark models, while Section 6 concludes with final remarks and some future research direction.

2. Related studies on investors' attention and forecasting

Before the Internet era and open data policies, observing investors' decision-making processes posed challenges due to a lack of public information. However, with the widespread adoption of information technology and the Internet, an increasing number of investors now rely on online sources for being updated about the most recent news, exchanging ideas with other investors, and so on. As a result, also academic research has delved into questioning the usefulness of internet-derived data and both sentiment and investor attention for forecasting, where sentiment can be seen as an alternative measure of the investors' attention on the financial market.[1]

[1] For instance, Mbanga et al. [29] found that investor attention causes changes in sentiment but not vice versa.

Early literature used consumer confidence survey data as a proxy of investor sentiment as it was ready-to-use and easily available from national statistical institutes. To improve sentiment measurement, alternative indicators have been however proposed [e.g., see 30]. Wang et al. [31] shown that more accurate out-of-sample forecasts can be achieved using sentiment index as predictors for stock returns, albeit Chung et al. [32] shown that these measures perform better in out-of-sample compared to other predictors only in the expansion states. Many fewer studies focused on the commodity market. In this regard, Gao and Süss [33] demonstrated empirically that sentiment, computed using an approach similar to Baker and Wurgler [30], allows for improvements in out-of-sample forecasting of commodity returns if used as an additional predictor.

With these approaches, it is crucial to have high-frequency data (e.g., intra-daily) and asset-specific measures, elements not always easily accessible. A prevalent method adopted nowadays for solving both issues involves the analysis of textual content from various sources such as news articles and social media platforms. Since 2010, the surge in user-generated content on platforms like Twitter has piqued interest in real-time data mining within the news analytics community. Bollen et al. [34] have been the first to show that sentiment computed on tweets allows forecasting the stock market, while Elshendy et al. [35], Li et al. [36] are noticeable examples of papers using text mining measures computed on data derived on the internet for forecasting oil prices. Abreu et al. [8] provides similar evidence for gold using Twitter-derived data.

More recently, researchers have used the volume of online searches, such as Google Trends, as a leading indicator for financial forecasting. The idea is that changes in search volume for asset-related terms can reflect evolving investor interest and information gathering, which may precede price movements. Google Trend allows downloading indicators of the volume of online searches for specific words. When these are related to commodities, they can be used as a proxy of investors' interests in those. Notably, the online search behavior of investors has shown good explanatory power and forecasting accuracy in the stock market [15,37]. Using web search data, Audrino et al. [38], Lin et al. [39], achieved improved accuracy also in forecasting volatility. Considering the commodities, Yu et al. [40] used Google Trends data for forecasting crude oil consumption, while Zhao et al. [41] for forecasting crude oil inventories. Salisu et al. [42] shown that Google Trend can be used for improving the out-of-sample forecasts of precious metals returns.

As emerges from a brief discussion of previous studies, it is clear that the information about sentiment and investors' attention has been mainly used as additional variables in forecasting models. In this paper, we consider a different perspective and construct a network of commodities in terms of attention levels. This allows us to investigate whether the information on commodities with similar attention levels is useful in forecasting them. Said differently, we evaluate the existence of substitution effects in attention dynamics and explore if such effects are useful in forecasting.

3. Methodology

In this section, we outline the methodology employed in our study, focusing on the use of spatial panel data models and their estimation through maximum likelihood. Additionally, we discuss the crucial aspect of selecting an appropriate spatial weight matrix, which in this paper is based on the commodities' similarity in terms of investor's attention. In doing so, we introduce a network structure among the considered commodities.

Hereafter, we refer to N commodities whose returns are observed for T consecutive times.

Notational agreement. Given a matrix $\mathbf{A}$ with dimension $I \times J$, its generic element is denoted by a_{ij}. The ith row of matrix $\mathbf{A}$ is $\mathbf{a}_i$. With a reasonable abuse of notation that does not generate confusion, we denote the jth column of $\mathbf{A}$ by $\mathbf{a}_j$ for simplicity. Moreover, R-dimensional vectors $\mathbf{x}$ and $\mathbf{x}_t$ have generic rth component denoted by x_r and x_{rt}, respectively.

3.1. Factor-augmented dynamic network model

Spatial panel data models are a valuable tool for analysing data exhibiting both temporal and spatial effects. These models account for spatial dependency and heterogeneity, making them particularly suitable for studying phenomena with geographical relevance. However, as argued by many authors [43–45], the concept of space can be generalised to any source of similarity and, therefore, to models with network interactions. In this paper, we consider a model of the form

$$\mathbf{y}_t = \rho \mathbf{W}\mathbf{y}_t + \phi \mathbf{y}_{t-1} + \delta \mathbf{W}\mathbf{y}_{t-1} + \boldsymbol{\varepsilon}_t, \tag{1}$$

where $\mathbf{y}_t$ is the vector of the commodities' returns at time t, $\mathbf{W}$ is a suitably selected weighted adjacency matrix of a network whose nodes are the considered commodities, and $\boldsymbol{\varepsilon}_t$ is the vector of error terms. The (1) can be called the Dynamic Network model (DNm), where the instantaneous network effect is represented by the coefficient ρ, while the lagged network effect is introduced through δ coefficient. In the end, a dynamic effect is introduced by coefficient ϕ, which is restricted to be the same for all the N commodities in the network. The model parameters can be estimated by Maximum Likelihood as in the case of Dynamic Spatial Panel Data models [DSPDm, see 46,47]. Given the model (1), its $h = 1$ step ahead forecast can be computed as follows

$$\hat{\mathbf{y}}_{t+1} = (\mathbf{I} - \hat{\rho}\mathbf{W})^{-1} \left[\hat{\phi}\mathbf{y}_t + \hat{\delta}\mathbf{W}\mathbf{y}_t\right]. \tag{2}$$

To enhance forecasting accuracy, we propose to augment model (1) by K latent factors. According to Elhorst [48], there are several possibilities for including latent factors in the model (1). The first one is to consider $K = 2$ common factors, where the first equals the cross-sectional fixed effects, which are constant over time but with heterogeneous coefficients, and the second is the time-period fixed effects, which change over time but with homogeneous coefficients. The second approach, discussed for instance in [49], considers cross-sectional averages of the dependent variables. A third alternative consists of estimating the unobservable common factors by one or more principal components. The first two approaches, although useful in many applicative domains, appear to not be appropriate in the case of commodities (and more generally financial) markets, where factors must be dynamic to model appropriately the intrinsic turbulence in both the markets and macroeconomic aggregates.

A simple approximate factor model assumes that observed time series stored in the $T \times N$-dimensional matrix $\mathbf{Y}$ can be decomposed into common factors $\mathbf{F}$ and idiosyncratic components $\mathbf{Z}$. We therefore consider the following factorial structure

$$\mathbf{Y} = \mathbf{F}\boldsymbol{\Lambda}' + \mathbf{Z}, \tag{3}$$

where $\mathbf{F}$ has dimension $T \times K$, being K the number of latent factors, $\mathbf{Z}$ has dimension $T \times N$ and $\boldsymbol{\Lambda}$ has dimension $N \times K$. To estimate the factor model, we employ the principal component estimator [PCE, 23,50], which estimates factors and loading from the covariance matrix of the observed time series. In particular, the PCE involves the following steps. First, the data is demeaned and the sample covariance matrix $\hat{\mathbf{S}} = \mathbf{Y}'\mathbf{Y}/N$ is computed. Then, the eigenvalues and eigenvectors of $\hat{\mathbf{S}}$ are obtained and ordered in decreasing order. The top K eigenvectors to form the loading matrix $\boldsymbol{\Lambda}$ are chosen and the common factors $\mathbf{F}$ are estimated using the selected eigenvectors.

Assuming $T > N$, the K global factors can be consistently estimated with

$$\hat{\mathbf{F}} = \mathbf{Y}\hat{\boldsymbol{\Lambda}}/N, \tag{4}$$

given that $\hat{\boldsymbol{\Lambda}}$ is estimated by $\sqrt{N}$ times the eigenvectors of $\mathbf{Y}'\mathbf{Y}$ associated with the K largest eigenvalues. To choose the number of factors K, following a consistent strand of literature, we follow the approach highlighted in [50,51]. According to Bai and Ng [50], the number of latent factors K is chosen by minimising the following information criterion

$$IC_K = \ln\left(V_K\right) + K\left(\frac{N+T}{NT}\right)\ln\left(\frac{NT}{N+T}\right) \tag{5}$$

with

$$V_K = \min_{\boldsymbol{\Lambda},\mathbf{F}} (NT)^{-1}\left(\sum_{i=1}^{N}\sum_{t=1}^{T} y_{it} - \lambda_i'\mathbf{f}_t\right)^2, \tag{6}$$

where V_K is the sum of squared residuals, divided by NT, when K factors are estimated. Given the estimation of K latent factor, we can write the Fa-DNm as follows

$$\mathbf{y}_t = \rho \mathbf{W}\mathbf{y}_t + \phi \mathbf{y}_{t-1} + \delta \mathbf{W}\mathbf{y}_{t-1} + \boldsymbol{\beta}\mathbf{f}_t + \boldsymbol{\varepsilon}_t \tag{7}$$

where $\mathbf{f}_t$ is the K-dimensional vector of estimated latent factors at time t. To estimate the parameters, we employ the maximum likelihood estimation approach also in this case, as the model (7) can be seen as a DSPDm with exogenous covariates. In this paper, we therefore consider a factor-augmented model based on a two-step approach, where the network model equation – i.e. the second step – is augmented by the inclusion of the estimated latent factors obtained in the first step. Like (2), $h = 1$ step ahead forecast for the model (7) can be obtained as

$$\hat{\mathbf{y}}_{t+1} = (\mathbf{I} - \hat{\rho}\mathbf{W})^{-1}\left[\hat{\phi}\mathbf{y}_t + \hat{\delta}\mathbf{W}\mathbf{y}_t + \hat{\boldsymbol{\beta}}\hat{\mathbf{f}}_{t+1}\right], \tag{8}$$

where $\hat{\mathbf{f}}_{t+1}$ is the vector of latent factors at time $t+1$ and $\mathbf{y}_t$ is the vector of commodities' returns at time t. In practice, we replace $\hat{\mathbf{f}}_{t+1}$ with $\mathbf{f}_t$, thus forecasting the latent factors with a random walk without drift.

3.2. Selection of the weighted adjacency matrix W

A critical aspect in the type of models discussed in Section 3.1 is the selection of the weighted adjacency matrix $\mathbf{W}$. The choice of $\mathbf{W}$ influences the model's performance and the interpretation of network relationships. Common weight matrix specifications include distance-based weights, contiguity weights, and k-nearest neighbor weights, among the others [52,53].

In this paper, we construct the weight matrix by considering the similarities in investors' attention to commodities. Indeed, similar patterns in Google Trends suggest that the two commodities experiment with an increase in web searches, which is our proxy of attention [14,16]. We investigate if this source of spillover is statistically significant, thus assessing if spillovers due to attention correlation exist and if the use of this information may generate a gain in forecasting terms.

Google Trends provides valuable insights into different search terms' popularity and relative interest over time. Data on search queries are anonymous and normalised to ensure comparability across diverse terms. As users input search queries, Google Trends compiles information on the number of searches for each term and in units of time. Then, search volumes are scaled between 0 (no searches) and 100 (associated with the time unit with the maximum number of searches over the considered period) to provide a relative measure of search interest over time. Therefore, the Google Trends values are not absolute search volumes but represent the popularity of a specific commodity's name search relative to the total search volume at a given time.

We collect Google Trends data for each commodity of interest and calculate the Euclidean distance between their respective search trends. The Euclidean distance serves as a measure of similarity, reflecting how closely the investors' attention to different commodities aligns over the study period. A similar approach is presented in [54], where the frequency of usage of economics-related words by US Presidents in their speeches is taken as a measure of the attention paid by them to economics and financial topics; then the series of these words are used to capture distance with the S&P 500 (trading volume, prices and returns) via a set of measures, including the Euclidean distance. Formally, let $\mathbf{s}_i$ and $\mathbf{s}_j$ be the Google Trends time series for commodities i and j, respectively. We define the distance between commodities i and j by

$$d_{ij} = \sqrt{\left(\mathbf{s}_i - \mathbf{s}_j\right)'\left(\mathbf{s}_i - \mathbf{s}_j\right)}. \tag{9}$$

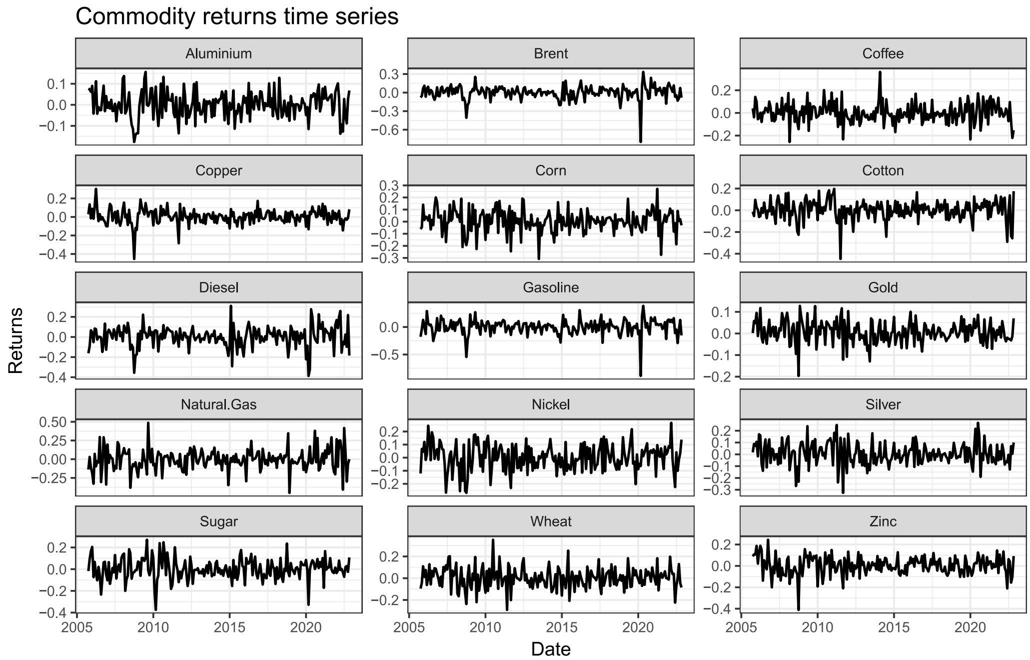

Fig. 1. Commodity returns time series.

The Euclidean distance in (9) is a simple but effective approach for measuring the similarity of the temporal patterns of two time series i and j. It is commonly employed in time series clustering literature for grouping time series of similar shapes, albeit also other dissimilarity definitions are possible. A low value of Euclidean distance d_{ij} suggests that the shape of the Google Trend time series for commodities i and j is very similar and, therefore, the considered commodities are characterised by similar patterns of investors' attention.

By reasonably assuming that the link between commodities i and j is stronger as i and j have smaller distance, one can define the weight between i and j by

$$w_{ij} = \frac{1}{d_{ij} + 1}. \tag{10}$$

The elements w's in (10) form the weighted adjacency matrix $\mathbf{W}$. Thus, the distances ds in (9) consider the similarity in the Google Trends time series shapes. According to (10), such distances ensure that commodities with higher investor attention similarity receive stronger weights in the weight matrix $\mathbf{W}$, reflecting their network interdependence.

4. Data

We consider monthly returns for various commodities, spanning from October 2005 to November 2022. The dataset includes the most frequently traded commodities, that is, aluminum, coffee, crude oil, copper, cotton, diesel, gasoline, gold, natural gas, nickel, silver, sugar, wheat, and zinc. Therefore, we have commodities obeying the standard taxonomy [55] that is, energy, metals, and agricultural commodities. Fig. 1 shows the temporal evolution of the commodity returns.

Fig. 1 highlights similarities in the patterns of commodities belonging to similar sectors. For example, energy-related commodities (e.g., Brent, diesel, gasoline) show a large negative peak in April 2020, when crude oil futures become negative for the first time in history [56]. In a forecasting setting, it is interesting to measure the ability of network models to forecast such events through spillovers from similar commodities. We notice that, while some commodities (e.g., copper, cotton, wheat, zinc) have not been exposed to such an event, it seems that some others (e.g., aluminum, sugar, nickel) belonging to a different sector experienced a shock on the same date. It seems, in other words, that spillovers do not always come from commodities of the same type. An intriguing source spillover channel that we explore in this paper, is the one associated with investor's attention.

To capture the impact of investor attention-driven spillovers, we incorporate Google Trends data into the problem at stake. In particular, we consider the volume of searches, for each commodity, during the period from October 2005 to November 2022. Fig. 2 shows the temporal patterns of the search volumes for each commodity. The years 2020 and 2022, which include COVID-19, negative oil futures prices, and the Russo-Ukranian war, are highlighted in gray color. The time series are normalised so that all the values fall in the range [0,100], where – as already stated in the previous section – 0 represents the date with the lowest search volume and 100 is the date with the largest one.

Considering the April 2020 shock to crude oil futures as a reference date, Fig. 2 clearly shows that on this date the investor's attention to crude oil reached its maximum value for the whole period under study. However, also the attention to other commodities, not energy-related, was maximum around the same period. Among the others, we can mention aluminum, sugar, and nickel which are not energy-related commodities, but experienced drops in returns similar to crude oil (see Fig. 1), and also an increase in the volume of Google searches. This simple evidence suggests that an attention-driven network could be a helpful tool for explaining spillovers across commodities of different categories, which cannot be explained by accounting for the standard taxonomy [55] alone.

To evaluate the usefulness of the model, we estimate the full model until the break of interest occurred in April 2020. First, we estimate the number and the value of the latent factors affecting the commodity returns. Following the Bai and Ng [50] approach, we estimate $K = 1$, which is then used as an additional predictor in the model. Given the network structure implied by the Google Trends time series is used to estimate the Fa-DNm in Table 1.

Given the results of Table 1, we get evidence of a statistically significant spillover coefficient ρ, which means that a spillover between commodities in terms of investors' attention exists. The spillover is negative, meaning that an increase in the returns of a commodity closer to the network space generates a decrease in the others. This result suggests the existence of a substitution effect between commodities

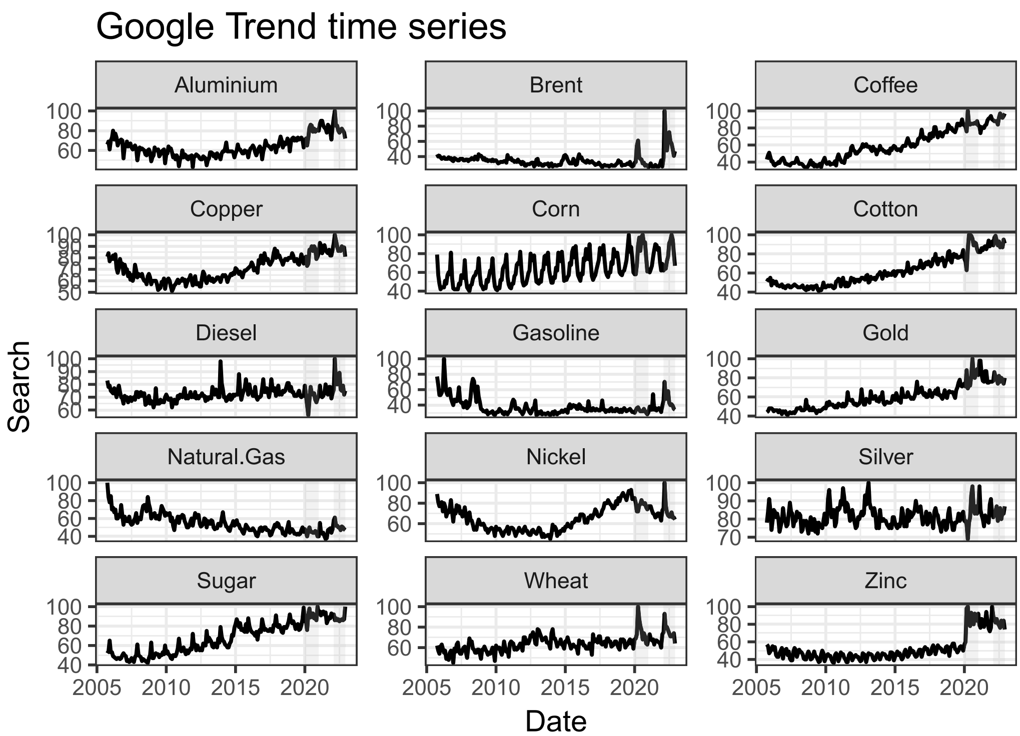

Fig. 2. Google Trends time series: temporal evolution of searching volumes for each commodity.

Table 1
Estimation results of the Fa-DNm.

Coefficient	Estimate	Std. Error	p-value
ρ	−0.11977	0.053780	0.02595
ϕ	0.026133	0.019473	0.17960
δ	−0.064658	0.035676	0.06993
β	−1.018410	0.056279	0.00000

that results close in the network, which is also associated with the idea of investors' diversification. In particular, a negative ρ coefficient indicates that investors hedge against negative returns of a commodity by investing in others with similar attention levels, being close in the network space. This result is also in line with the common finding that the investor's attention has a significant but negative relationship with returns; hence, periods of higher investor attention tend to be followed by lower returns for the aggregate market [e.g., see 37]. Moreover, we also find that the lagged temporal network lag coefficient, δ, is negative and of lower magnitude compared with ρ, suggesting that the instantaneous network effect is stronger than the lagged one. The estimated latent factor has also an overall negative relationship with returns, which suggests that it could be associated with some source of risk in the commodity market. This finding is also consistent with previous studies. For example, Guidolin and Pedio [26] show that macro factors estimated with principal components negatively affect the sample of commodities in both high and low-volatility periods. Furthermore, Daskalaki et al. [25] show that momentum and other macro-factors are also negatively correlated with commodity returns. From the theoretical perspective, the theory of storage predicts a negative correlation between inventories and commodity returns, that is lower inventories are associated with higher returns and higher volatility. Therefore, the negative sign could be explained if the estimated factors consider this information. Interestingly, the significance of all these quantities suggests that the simple autoregressive term is not statistically significant. This may also suggest that the estimated factors capture information included in the autocorrelation structure of commodities. We, however, prefer to still use past information to forecast commodity returns.

Table 2
Estimation results of restricted models.

Coefficient	Estimate	Std. Error	p-value
	Panel A: Purely spatial model		
ρ	−0.1137	0.0533	0.033
ϕ	–	–	–
δ	–	–	–
β	−1.007	0.0555	0.000
	Panel B: Purely temporal model		
ρ	–	–	–
ϕ	0.0040	0.0360	0.5437
δ	–	–	–
β	−0.9717	0.1061	0.0000

We then evaluate the robustness of these findings considering the estimation of restricted models, that is, purely temporal and purely "network" panel data models. In the case of a pure network model, we estimate a static version of (7) by restricting $\phi = \delta = 0$, while for the case of a purely temporal model, we restrict $\rho = \delta = 0$. We keep the latent factors in the models' specification, as we always find them to be statistically significant. The results are shown in Table 2.

Interestingly, the results are robust to alternative specifications of the model (7). Indeed, in the case of a purely spatial model, we find both ρ and β to be negative and statistically significant. Also, the magnitude of the coefficients is very similar between the two specifications. Moreover, in the case of a purely temporal model, we find that temporal lags are not significant while including the latent factors in the model and that the $K = 1$ latent factor has a negative relationship with the commodity returns. In this second case, the magnitude of the parameter is larger than that of the other two settings, namely −0.97 compared with -1, but the magnitude is still comparable. This suggests that the joint usage of the two provides a more comprehensive view of the commodities' returns dynamics, leaving to the factor a lower impact, and its almost doubled parameter's standard deviation increases its uncertainty when the spillover is not considered.

In sum, we find evidence that network structure in terms of Google Trends data provides important information for modeling both temporal and cross-sectional dependence of commodities. In other words, investors' attention, measured considering the Google Trends time series, is a suitable channel for return spillovers across commodities. In the next section, we evaluate the forecasting ability of such a piece of information.

5. Forecasting experiment

5.1. Experimental design

For the out-of-sample forecasting experiment, an expanding-window procedure is adopted. In doing so, we mimic the exercise of a commodity forecaster which starts making predictions at a given time t until the point in time T. To evaluate forecasting accuracy out-of-sample, the full sample is divided into a train and a test sample. We consider an initial window of $M = 100$ observations and leave the last $T - M$ observations for testing. In each step, a new point in time is added to the sample and the network structure $\mathbf{W}$ is updated following the same procedure discussed in practice in Section 3, and the K latent factors are estimated in each iteration. The model parameters are re-estimated in each iteration using the M observations. Therefore, the outcome of the forecasting exercise is a matrix time series containing out-of-sample forecasts $\widehat{\mathbf{Y}}$ of dimension $(T - M) \times N$ [57].

The out-of-sample performance of the Fa-DNm is compared with the predictions obtained from different benchmarks usually considered in commodity forecasting. First of all, we consider the following VAR model

$$\boldsymbol{y}_t = \boldsymbol{c} + \sum_{p=1}^{P} \boldsymbol{\Phi}_p \boldsymbol{y}_{t-p} + \boldsymbol{\varepsilon}_t \tag{11}$$

with $\mathbf{y}_t$ the N-dimensional vector of the commodities' returns at time t, $\mathbf{c}$ the vector of constant terms, $\boldsymbol{\Phi}_p$ the matrix of autoregressive coefficients of order p, and $\boldsymbol{\varepsilon}_t$ the vector of error terms. We also consider a factor-augmented version of (11), where the global factors $\mathbf{f}_t$ are included as an additional variable in the VAR model. We call this second model factor-augmented VAR, i.e. Fa-VAR. Comparing the Fa-DNm with these two alternatives allows us insight into the relevance of the network spillovers. Indeed, we get evidence favoring the relevance of the network information if the Fa-DNm performs better than the VAR. Moreover, we also compare the Fa-DNm with simpler well-known benchmarks such as the Random Walk (RW), that is $\widehat{\mathbf{y}}_{t+1} = \mathbf{y}_t$ and from those based on the latent factors (3) only (FM), that is $\widehat{\mathbf{y}}_{t+1} = \widehat{\boldsymbol{\beta}}\widehat{\mathbf{f}}_t$.

To evaluate the forecasting accuracy of a given model for each ith commodity, we rely on two commonly employed accuracy metrics, namely the Root Mean Squared Error (RMSE)

$$\mathrm{RMSE}_i = \sqrt{\frac{1}{T-M} \sum_{t=M+1}^{T} \left(\hat{y}_{it} - y_{it}\right)^2}, \tag{12}$$

and the Mean Absolute Error (MAE)

$$\mathrm{MAE}_i = \frac{1}{T-M} \sum_{t=M+1}^{T} \left|\hat{y}_{it} - y_{it}\right|. \tag{13}$$

To obtain an overall accuracy measure, we average both RMSE_i and MAE_i, so that $\mathrm{RMSE} = 1/N \sum_i \mathrm{RMSE}_i$ and $\mathrm{MAE} = 1/N \sum_i \mathrm{MAE}_i$.

Then, predictive accuracy tests of Diebold and Mariano [58], Diebold [59] are considered to evaluate the statistical significance of the forecasting error differences. Let us define $e_{1,t} = \hat{y}_t - y_t$ the forecasting error at time t for a generic unit and for the generic forecasting model 1, and $d_t = g\left(e_{1,t}\right) - g\left(e_{2,t}\right)$ the error differential between two forecasting approaches (say, forecasting models 1 and 2) up to some transformation g. In this paper, we consider the transformations $g(e) = e^2$ and $g(e) = |e|$. Assuming covariance stationarity of the loss T-dimensional differential series $\mathbf{d}$, Diebold and Mariano [58] show that the sample mean of the loss differential,

$$\bar{d} \equiv \frac{1}{T-M} \sum_{t=T-M+1}^{T} d_t, \tag{14}$$

follows an asymptotically standard Normal distribution. Therefore, testing the null hypothesis of equal forecast accuracy can be obtained by calculating the statistic

$$\mathrm{DM} = \bar{d}[V(\bar{d})]^{-1/2}, \tag{15}$$

where $V(\bar{d})$ is consistently estimated assuming a certain autocorrelation structure of the forecasting errors.

We notice that the [58] test is valid considering the entire out-of-sample testing period, while it could be interesting to evaluate the differences in the accuracy of the methods during different market conditions. For this aim, a weighted version of the [58] test can be adopted. Following Van Dijk and Franses [60], we can consider the weighted version of (14), that is

$$\bar{d}_w \equiv \frac{1}{T-M} \sum_{t=T-M+1}^{T} w(\omega_t) d_t, \tag{16}$$

where ω_t is the information set used for choosing the weighting function. Van Dijk and Franses [60] proposed the use of the function $w(\omega_t) = 1 - \kappa(y_t)/\max(\kappa(y_t))$, where $\kappa(y_t)$ denotes the density function of y_t estimated using a standard Nadaraya–Watson kernel estimator with a Gaussian kernel. This approach is used for weighting more the errors when the realised values are in tails of the distribution, while Mattera et al. [61] adopted the weighting function $w(\omega_t) = \kappa(y_t)/\max(\kappa(y_t))$ to provide less weight to the tails. We, therefore, also provide the results of both the weighted DM tests, as proposed in [60,61].

5.2. Results

In what follows, we assess the forecasting performance of the five models for commodity returns: the Fa-DNm, introducing a novel dynamic network perspective which is based on investor's attention, the VAR and its factor-augmented version (Fa-VAR), the naive approach RW and the factor model FM. While the FM and VAR represent two popular approaches for commodity forecasting [e.g., see 62], the RW is a common benchmark as it is difficult to beat in out-of-sample experiments [63].

Our evaluation relies on relative efficiency, measured in terms of Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE). Table 3 shows the forecasting accuracy of the models for each commodity as well as the averaged results. Panel A shows the results in terms of RMSE, while Panel B is in terms of MAE. The best model in terms of average loss is highlighted in bold font.

Table 3 reveals that the Fa-DNm consistently outperforms all the considered benchmarks in out-of-sample forecasting, as evidenced by its superior relative efficiency. In particular, consistent with previous studies, we find that both the VAR and the factor model provide more accurate than random walk for most commodities, but still, it provides less accurate forecasts compared to the proposed approach. Therefore, in terms of both RMSE and MAE, we find that the proposed Fa-DNm provides more accurate forecasts, which investors can use to make more informed decisions.

Fig. 3 shows the relative efficiency in terms of RMSE, while Fig. 4 in terms of MAE. Relative efficiency is computed, for each commodity, by the ratio of the benchmark and the proposed Fa-DNm. Values larger than 1 indicate better performance of the proposed approach compared with the benchmark. The dashed red lines in the figures highlight the thresholds of 1 where the Fa-DNm outperforms the benchmarks. Fig. 3(a) shows the relative efficiency of the Fa-DNm versus RW in terms of RMSE, while Fig. 3(b) compares Fa-DNm versus FM forecasts considering RMSE loss. Fig. 3(c) and Fig. 3(d) show the comparison with the VAR and Fa-VAR, respectively.

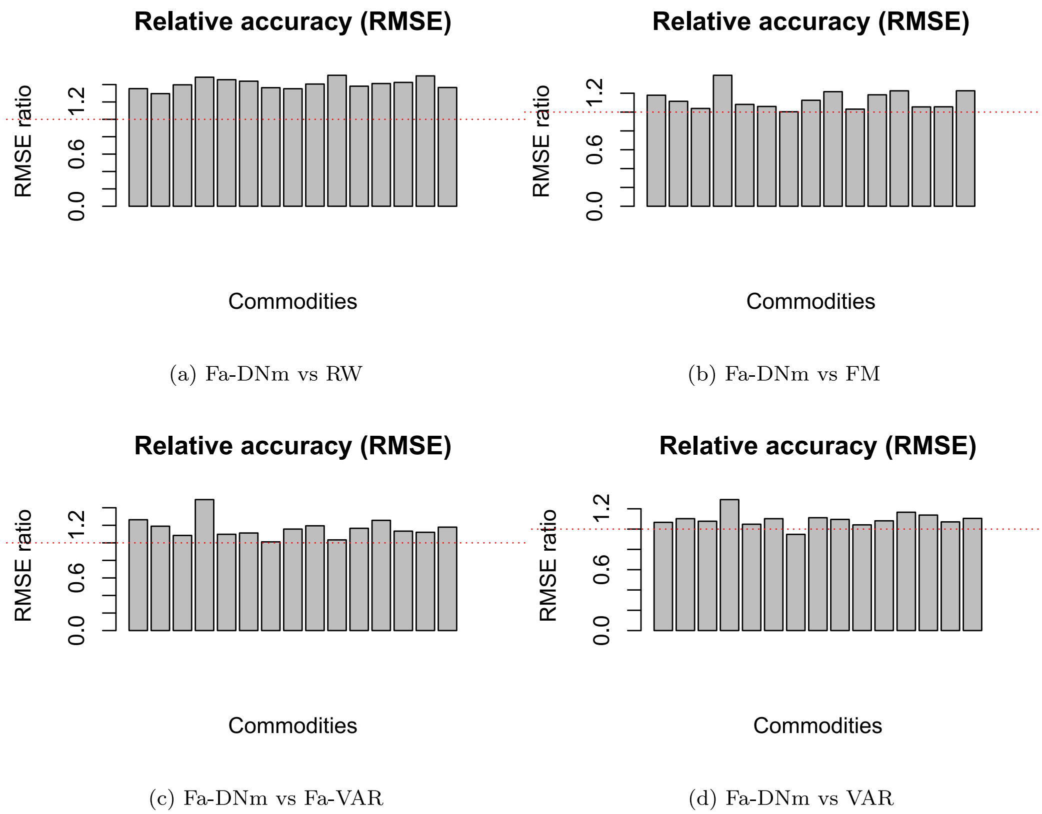

Fig. 3. Relative accuracy between the Fa-DNm (8) and the considered benchmarks in terms of RMSE. Values larger than 1 indicate better performance of the proposed approach. The dashed lines represent the thresholds of 1, associated with Fa-DNm outperforming the benchmark.

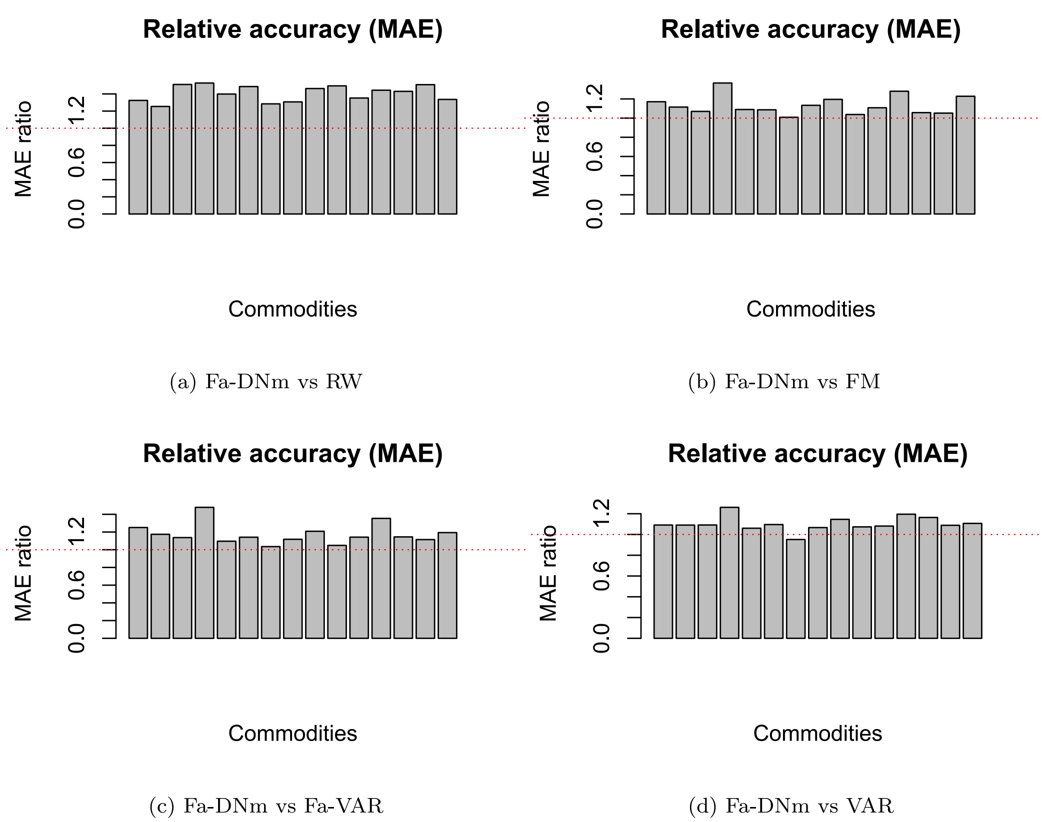

Fig. 4. Relative accuracy between the Fa-DNm (8) and the considered benchmarks in terms of RMSE. Values larger than 1 indicate better performance of the proposed approach.

Table 3
Predictive accuracy of the considered models for each commodity. Both RMSE and MAE loss functions are considered. The best model in terms of average loss is highlighted in bold font.

Commodity	Fa-DNM	RW	FM	Fa-VAR	VAR
Panel A: RMSE loss					
Aluminium	0.0552	0.0748	0.0651	0.0698	0.0589
Brent	0.1264	0.1638	0.1409	0.1504	0.1394
Coffee	0.0910	0.1272	0.0945	0.0987	0.0981
Copper	0.0596	0.0884	0.0828	0.0889	0.0769
Corn	0.0797	0.1160	0.0861	0.0874	0.0835
Cotton	0.0845	0.1216	0.0895	0.0940	0.0932
Diesel	0.1190	0.1624	0.1194	0.1204	0.1128
Gasoline	0.1485	0.2008	0.1670	0.1718	0.1651
Gold	0.0396	0.0557	0.0482	0.0473	0.0434
Natural Gas	0.1431	0.2156	0.1476	0.1479	0.1490
Nickel	0.0888	0.1227	0.1051	0.1035	0.0960
Silver	0.0758	0.1071	0.0929	0.0953	0.0884
Sugar	0.0802	0.1143	0.0846	0.0909	0.0913
Wheat	0.0860	0.1289	0.0908	0.0963	0.0920
Zinc	0.0666	0.0910	0.0816	0.0785	0.0736
Avg.	**0.0896**	0.1260	0.0998	0.1027	0.0974
Panel B: MAE loss					
Aluminium	0.0441	0.0584	0.0516	0.0552	0.0481
Brent	0.0866	0.1086	0.0966	0.1017	0.0945
Coffee	0.0683	0.1032	0.0731	0.0777	0.0746
Copper	0.0470	0.0718	0.0642	0.0695	0.0593
Corn	0.0593	0.0830	0.0647	0.0650	0.0629
Cotton	0.0636	0.0945	0.0691	0.0726	0.0696
Diesel	0.0887	0.1138	0.0894	0.0918	0.0844
Gasoline	0.1021	0.1334	0.1157	0.1141	0.1089
Gold	0.0312	0.0457	0.0373	0.0377	0.0358
Natural Gas	0.1004	0.1500	0.1042	0.1053	0.1078
Nickel	0.0717	0.0970	0.0795	0.0819	0.0776
Silver	0.0583	0.0841	0.0746	0.0789	0.0697
Sugar	0.0606	0.0866	0.0640	0.0693	0.0705
Wheat	0.0672	0.1013	0.0707	0.0749	0.0731
Zinc	0.0541	0.0722	0.0664	0.0645	0.0599
Avg.	**0.0669**	0.0936	0.0747	0.0773	0.0731

Given the relative efficiency, we observe that the ratio is larger than 1 for most commodities, considering all the benchmarks and both RMSE and MAE accuracy measures. The enhanced forecasting accuracy of the Fa-DNm can be attributed to several key factors. Firstly, the dynamic network structure of the Fa-DNm captures the evolving interdependencies among commodities over time. This adaptability is crucial in forecasting commodity returns, as market dynamics can rapidly change. The incorporation of spatial econometrics principles within the network framework allows the model to capture spillovers based on media-driven factors. Indeed, the Fa-DNm leverages Google Trends search data, a rich source of information on investor's attention. The use of this internet-derived data allows the model to respond in real time to shifts in investors' sentiment and interest, providing a timely representation of market conditions. In contrast, while the FM is a well-established method, it may not fully capture the evolving relationships and spillovers present in dynamic commodity markets. The RW, even if widely used as a benchmark hard to beat in out-of-sample, tends to underestimate the complexities of commodity returns by assuming a constant expected return.

Finally, we adopt the [58] predictive accuracy test to evaluate whether these differences are statistically significant. Considering average results, the comparison between Fa-DNm and RW leads to DM test statistics equal to −2.58, which suggests that the Fa-DNm provides more accurate forecasts on average than RW, with a p-value of 0.0112. Therefore, we reject the null hypothesis of equal predictive accuracy. Comparing Fa-DNm and FM leads to DM test statistics equal to −2.1627, which still suggests that the Fa-DNm provides more accurate forecasts on average than FM, with a p-value of 0.0328. Therefore, also in this case we reject the null hypothesis of equal predictive accuracy. The same findings are obtained by comparing the Fa-DNm with both VAR and Fa-VAR. Table 4 provides a detailed comparison of the pairwise tests for each commodity.

Table 4
p-values associated with the [58] predictive accuracy test. The proposed approach (Fa-DNm) is compared with RW, FM, Fa-VAR, and VAR. Under the null hypothesis, the two models under comparison provide equally accurate forecasts.

Commodity	Fa-DNm vs			
	RW	FM	Fa-VAR	VAR
Aluminium	0.00	0.01	0.00	0.18
Brent	0.18	0.25	0.19	0.37
Coffee	0.00	0.21	0.02	0.06
Copper	0.00	0.01	0.01	0.05
Corn	0.00	0.07	0.05	0.09
Cotton	0.00	0.39	0.12	0.06
Diesel	0.01	0.96	0.86	0.30
Gasoline	0.10	0.27	0.34	0.39
Gold	0.00	0.01	0.00	0.03
Natural Gas	0.00	0.41	0.33	0.17
Nickel	0.00	0.03	0.00	0.02
Silver	0.00	0.03	0.00	0.00
Sugar	0.00	0.13	0.00	0.00
Wheat	0.00	0.10	0.02	0.03
Zinc	0.00	0.01	0.00	0.02

In sum, only for Brent and Gasoline the difference in accuracy between Fa-DNm and the benchmarks are not statistically significant, while for all the other commodities the proposed approach provides significantly better forecasts. This result could be partly explained by the higher trading activity and efficiency of the Brent market compared to other commodities. Indeed, being a type of crude oil, Brent is one of the most widely traded commodities globally. The high liquidity and trading volume in the Brent market contribute to a more efficient price discovery process. In highly liquid markets, information is quickly reflected in prices [64], reducing the potential for predictive models to gain a significant edge over simpler models like RW. The case of Cotton is interesting as in this case all the models including the latent factors (i.e. Fa-DNm, FM and Fa-VAR) provide statistically the same forecasts, while models not using this additional information perform worse. Finally, given the comparison for all the other commodities, we reject the null hypothesis of the pairwise predictive accuracy tests. This result confirms that factor-augmented models are good tools for forecasting commodities, as highlighted by previous studies, and that investors' attention provides significantly relevant additional information for forecasting the returns.

In the end, Tables 5 and 6 show the results of the weighted predictive accuracy tests considering either the tails or the center of the distribution of the actual values, according to Van Dijk and Franses [60], Mattera et al. [61], respectively. In the first case, we evaluate the ability of the models to forecast the events in the tails, that is in turbulent times, while in the second case, we evaluate the comparison of the models during not turbulent periods. The results of the two tables are consistent with those of Table 4, in that the models provide statistically different forecasts. The main difference, however, is that with weighted approaches also commodities like Brent and Cotton showing some similar performance between the models, are now statistically different. Said differently, in this case, the Fa-DNm provides statistically more accurate predictions compared with the benchmarks. Therefore, the results are robust to the portion of the distribution considered for testing the equality of the predictive accuracy.

6. Conclusions

This paper introduces an innovative framework for enhancing the precision of commodity return forecasts, which has particular relevance in business decision-making and production optimisation. The proposed methodology involves a Fa-DNm, uniquely designed to incorporate the spillover effects originating from investors' attention to commodities.

Table 5
p-values associated with the [60] predictive accuracy test (tail-based). The proposed approach (Fa-DNm) is compared with RW, FM, Fa-VAR, and VAR. Under the null hypothesis, the two models under comparison provide equally accurate forecasts in the tails of the distribution.

Commodity	Fa-DNm vs			
	RW	FM	Fa-VAR	VAR
Aluminium	0.00	0.00	0.00	0.01
Brent	0.03	0.02	0.02	0.12
Coffee	0.00	0.00	0.00	0.00
Copper	0.00	0.00	0.00	0.02
Corn	0.00	0.00	0.00	0.01
Cotton	0.00	0.00	0.00	0.01
Diesel	0.00	0.00	0.00	0.02
Gasoline	0.00	0.00	0.00	0.05
Gold	0.00	0.00	0.00	0.00
Natural Gas	0.00	0.01	0.01	0.00
Nickel	0.00	0.05	0.00	0.00
Silver	0.00	0.00	0.00	0.00
Sugar	0.00	0.08	0.00	0.00
Wheat	0.00	0.01	0.00	0.00
Zinc	0.00	0.01	0.00	0.00

Table 6
p-values associated with the [60] predictive accuracy test (centrality-based). The proposed approach (Fa-DNm) is compared with RW, FM, Fa-VAR, and VAR. Under the null hypothesis, the two models under comparison provide equally accurate forecasts in the center of the distribution.

Commodity	Fa-DNm vs			
	RW	FM	Fa-VAR	VAR
Aluminium	0.00	0.00	0.00	0.01
Brent	0.02	0.01	0.01	0.11
Coffee	0.01	0.00	0.00	0.00
Copper	0.00	0.00	0.00	0.02
Corn	0.00	0.00	0.00	0.01
Cotton	0.00	0.00	0.00	0.01
Diesel	0.00	0.01	0.00	0.04
Gasoline	0.00	0.00	0.01	0.15
Gold	0.00	0.00	0.00	0.00
Natural Gas	0.00	0.01	0.01	0.01
Nickel	0.00	0.06	0.00	0.00
Silver	0.00	0.00	0.00	0.00
Sugar	0.00	0.07	0.00	0.00
Wheat	0.00	0.01	0.00	0.00
Zinc	0.00	0.01	0.00	0.00

Drawing upon the tenets of complex network theory, we explore the intricate interdependencies within commodity markets and associated financial instruments.

Our findings highlight that the use of complex network analysis helps us see hidden patterns, connections, and potential risks in financial services. This aligns with the growing interest in understanding how complex networks, especially those involving spillovers, can be crucial in predicting commodity patterns. An important aspect is that the proposed dynamic network model fully utilises internet data, specifically based on Google Trends search data. Instead of following the usual approaches, we introduce a Fa-DNm, where the network structure is developed in terms of investors' attention. In contrast to prior studies that treat internet-derived data merely as predictors, our novel approach views it as a fundamental component for shaping the network structure. This paradigm shift allows us to discern investor interest and dynamic connections between various commodities, offering novel insights into market dynamics.

The empirical findings highlight the substantive value of integrating attention-driven spillovers, elucidating their substantial enhancement of forecasting capabilities. The dynamic network perspective not only accommodates their presence but effectively leverages these spillovers in the forecasting task. The practical implications of this research extend to both academic researchers and industry practitioners within the financial domain. The ability to capture and quantify attention-driven spillovers through internet-derived data represents a pioneering approach, yielding insights that hold potential significance for understanding and predicting market dynamics. In an era of increasing global interconnectedness and rapid information dissemination, the integration of internet-derived data into forecasting models emerges as an invaluable tool for decision-makers grappling with the intricate dynamics of financial markets.

This paper primarily centers on empirically exploring the predictability of commodity market movements through the lens of investors' attention. From the theoretical perspective, the evidence of out-of-sample predictability naturally prompts inquiries into the efficiency of the commodity market. In particular, the presence of predictability enhanced by considering investors' attention spillovers, suggests that the commodity market may not be much efficient. However, from the theoretical viewpoint, the role of investors' attention and the significant spillover effect are the main aspects deserving of future investigations and deeper understanding. For example, the development of theoretical models like [15] for commodities could provide useful insights about investor behavior and how it relates to commodity markets. Our findings can corroborate the predictions of such a theoretical model.

In the end, some additional future research directions can be highlighted. From a modeling perspective, we recall that the proposed approach in this paper is based on a two-step procedure, where the factors are estimated considering a simple factor model in the first step. Then, the estimated factors are used in the second step for forecasting purposes. While, on one side, this approach is simple and effective in forecasting commodity returns in out-of-sample compared to different benchmarks, on the other side it does not account for the network interactions while estimating the latent factors. Other papers [e.g., see 65,66] consider more sophisticated approaches for dealing with latent factors in similar models. The use of these approaches for studying and forecasting commodity markets deserves future investigations. Moreover, a dynamic factor model may be considered instead of an approximate factor structure in future research to improve out-of-sample forecast accuracy.

From a different perspective, it could be interesting to apply the considered procedures to data with different periodicities. On this front, we are investigating some modifications of the proposed algorithms to deal with the computational challenges it offers. A promising strategy might be to reduce the cardinality of the dataset by preprocessing the series to cluster times and cross-sections to move to a dataset with reduced dimensionality. In this respect, some operational research techniques are leading to efficient dataset dimensionality reduction without losing too much information (see, e.g., the wide literature on data compression models, like the pioneering contribution by Rissanen [67] and, more recently, the handbook [68] and the monograph [69]).

Finally, we point out that we deal with a static network, that is the structure of the network does not evolve with time. The introduction of a dynamic network, where the spillover in terms of attention may be dynamic, albeit interesting, poses significant challenges from a modeling perspective, especially if latent factors are included.

CRediT authorship contribution statement

Roy Cerqueti: Writing – original draft, Visualisation, Supervision, Methodology, Investigation, Data curation, Conceptualisation. **Valerio Ficcadenti:** Writing – original draft, Visualisation, Software, Methodology, Data curation, Conceptualisation. **Raffaele Mattera:** Writing – original draft, Visualisation, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualisation.

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.

Data availability

Data will be made available on request.

References

[1] Edwards DR, Idoko FO, Vogiatzis C, Davis LB, Mirchandani P. Determining optimal fuel delivery strategies under uncertainty. Socio-Econ Plan Sci 2023;88:101612.
[2] Labys WC, Pollak PK. Commodity models for forecasting and policy analysis. Routledge; 1984.
[3] Labys WC. Modeling and forecasting primary commodity prices. Routledge; 2006.
[4] Yukseltan E, Yucekaya A, Bilge AH, Aktunc EA. Forecasting models for daily natural gas consumption considering periodic variations and demand segregation. Socio-Econ Plan Sci 2021;74:100937.
[5] Elzaki RM, Elrasheed MM, Elmulthum NA. Optimal crop combination under soaring oil and energy prices in the kingdom of Saudi Arabia. Socio-Econ Plan Sci 2022;83:101367.
[6] Rettinger M, Mandl C, Minner S. A data-driven approach for optimal operational and financial commodity hedging. European J Oper Res 2024. http://dx.doi.org/10.1016/j.ejor.2024.01.026.
[7] Vecchiato R. Environmental uncertainty, foresight and strategic decision making: An integrated study. Technol Forecast Soc Change 2012;79(3):436–47.
[8] Abreu P, Santos D, Barbosa-Povoa A. Data-driven forecasting for operational planning of emergency medical services. Socio-Econ Plan Sci 2023;86:101492.
[9] Barrat A, Barthelemy M, Vespignani A. Dynamical processes on complex networks. Cambridge University Press; 2008.
[10] Newman ME. The structure and function of complex networks. SIAM Rev 2003;45(2):167–256.
[11] Pagnottoni P, Spelta A. The motifs of risk transmission in multivariate time series: Application to commodity prices. Socio-Econ Plan Sci 2023;87:101459.
[12] Ardia D, Bluteau K, Boudt K. Questioning the news about economic growth: Sparse forecasting using thousands of news-based sentiment values. Int J Forecast 2019;35(4):1370–86.
[13] Thorsrud LA. Words are the new numbers: A newsy coincident index of the business cycle. J Bus Econom Statist 2020;38(2):393–409.
[14] Ahmed S, Tsvetanov D. The predictive performance of commodity futures risk factors. J Bank Financ 2016;71:20–36.
[15] Andrei D, Hasler M. Investor attention and stock market volatility. The Rev Financ Stud 2015;28(1):33–72.
[16] Da Z, Engelberg J, Gao P. In search of attention. J Finance 2011;66(5):1461–99.
[17] Cerqueti R, Cesarone F, Ficcadenti V. Portfolio decision analysis for pandemic sentiment assessment based on finance and web queries. Ann Oper Res 2024;1–31. http://dx.doi.org/10.1007/s10479-024-05966-x.
[18] Cerqueti R, Ficcadenti V. Anxiety about the pandemic and trust in financial markets. Ann Reg Sci 2024;72(4):1277–328.
[19] Bulut L. Google Trends and the forecasting performance of exchange rate models. J Forecast 2018;37(3):303–15.
[20] Huang MY, Rojas RR, Convery PD. Forecasting stock market movements using Google Trend searches. Empir Econ 2020;59:2821–39.
[21] Hu H, Tang L, Zhang S, Wang H. Predicting the direction of stock markets using optimized neural networks with Google Trends. Neurocomputing 2018;285:188–95.
[22] Salisu AA, Ogbonna AE, Adediran I. Stock-induced Google Trends and the predictability of sectoral stock returns. J Forecast 2021;40(2):327–45.
[23] Stock JH, Watson MW. Forecasting using principal components from a large number of predictors. J Am Stat Assoc 2002;97(460):1167–79.
[24] Rapach D, Zhou G. Forecasting stock returns. In: Handbook of economic forecasting. vol. 2, Elsevier; 2013, p. 328–83.
[25] Daskalaki C, Kostakis A, Skiadopoulos G. Are there common factors in individual commodity futures returns? J Bank Financ 2014;40:346–63.
[26] Guidolin M, Pedio M. Forecasting commodity futures returns with stepwise regressions: Do commodity-specific factors help? Ann Oper Res 2021;299:1317–56.
[27] Alquist R, Bhattarai S, Coibion O. Commodity-price comovement and global economic activity. J Monetary Econ 2020;112:41–56.
[28] Gorton GB, Hayashi F, Rouwenhorst KG. The fundamentals of commodity futures returns. Rev Finance 2013;17(1):35–105.
[29] Mbanga C, Darrat AF, Park JC. Investor sentiment and aggregate stock returns: the role of investor attention. Rev Quant Financ Account 2019;53:397–428.
[30] Baker M, Wurgler J. Investor sentiment in the stock market. J Econ Perspect 2007;21(2):129–51.
[31] Wang W, Su C, Duxbury D. Investor sentiment and stock returns: Global evidence. J Emp Financ 2021;63:365–91.
[32] Chung S-L, Hung C-H, Yeh C-Y. When does investor sentiment predict stock returns? J Emp Financ 2012;19(2):217–40.
[33] Gao L, Süss S. Market sentiment in commodity futures returns. J Emp Financ 2015;33:84–103.
[34] Bollen J, Mao H, Zeng X. Twitter mood predicts the stock market. J Comput Sci 2011;2(1):1–8.
[35] Elshendy M, Colladon AF, Battistoni E, Gloor PA. Using four different online media sources to forecast the crude oil price. J Inf Sci 2018;44(3):408–21.
[36] Li X, Shang W, Wang S. Text-based crude oil price forecasting: A deep learning approach. Int J Forecast 2019;35(4):1548–60.
[37] Chen T. Investor attention and global stock returns. J Behav Finance 2017;18(3):358–72.
[38] Audrino F, Sigrist F, Ballinari D. The impact of sentiment and attention measures on stock market volatility. Int J Forecast 2020;36(2):334–57.
[39] Lin P, Ma S, Fildes R. The extra value of online investor sentiment measures on forecasting stock return volatility: A large-scale longitudinal evaluation based on Chinese stock market. Expert Syst Appl 2024;238:121927.
[40] Yu L, Zhao Y, Tang L, Yang Z. Online big data-driven oil consumption forecasting with Google Trends. Int J Forecast 2019;35(1):213–23.
[41] Zhao L-T, Zheng Z-Y, Wei Y-M. Forecasting oil inventory changes with Google Trends: A hybrid wavelet decomposer and ARDL-SVR ensemble model. Energy Econ 2023;120:106603.
[42] Salisu AA, Ogbonna AE, Adewuyi A. Google trends and the predictability of precious metals. Resour Policy 2020;65:101542.
[43] Elhorst JP, et al. Spatial econometrics: from cross-sectional data to spatial panels. vol. 479, Springer; 2014.
[44] Fernandez-Aviles G, Montero J-M. Spatio-temporal modeling of financial maps from a joint multidimensional scaling-geostatistical perspective. Expert Syst Appl 2016;60:280–93.
[45] Kutzker T, Wied D. Testing the correct specification of a system of spatial dependence models for stock returns. Empir Econ 2023;1–21.
[46] Lee L-f, Yu J. A spatial dynamic panel data model with both time and individual fixed effects. Econometric Theory 2010;26(2):564–97.
[47] Billé AG, Tomelleri A, Ravazzolo F. Forecasting regional GDPs: a comparison with spatial dynamic panel data models. Spatial Econ Anal 2023;1–22.
[48] Elhorst JP. The dynamic general nesting spatial econometric model for spatial panels with common factors: Further raising the bar. Rev Reg Res 2022;42(3):249–67.
[49] Pesaran MH, Tosetti E. Large panels with common factors and spatial correlation. J Econometrics 2011;161(2):182–202.
[50] Bai J, Ng S. Determining the number of factors in approximate factor models. Econometrica 2002;70(1):191–221.
[51] Hallin M, Liška R. Determining the number of factors in the general dynamic factor model. J Amer Statist Assoc 2007;102(478):603–17.
[52] Fülle MJ, Otto P. Spatial GARCH models for unknown spatial locations–an application to financial stock returns. Spatial Econ Anal 2023;1–14.
[53] Mattera R, Otto P. Network log-ARCH models for forecasting stock market volatility. Int J Forecast 2024. http://dx.doi.org/10.1016/j.ijforecast.2024.01.002.
[54] Ficcadenti V, Cerqueti R. Economic keywords in political communications and financial markets. Ann Oper Res 2024;1–24. http://dx.doi.org/10.1007/s10479-024-05905-w.
[55] Tabak B, Serra T, Cajueiro D. Topological properties of commodities networks. Eur Phys J B 2010;74:243–9.
[56] Corbet S, Hou YG, Hu Y, Oxley L. An analysis of investor behaviour and information flows surrounding the negative WTI oil price futures event. Energy Econ 2021;104:105589.
[57] Hewamalage H, Ackermann K, Bergmeir C. Forecast evaluation for data scientists: common pitfalls and best practices. Data Min Knowl Discov 2023;37(2):788–832.
[58] Diebold FX, Mariano RS. Comparing predictive accuracy. J Bus Econ Stat 2002;20(1):134–44.
[59] Diebold FX. Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of Diebold–Mariano tests. J Bus Econom Statist 2015;33(1):1.
[60] Van Dijk D, Franses PH. Selecting a nonlinear time series model using weighted tests of equal forecast accuracy. Oxf Bull Econ Stat 2003;65:727–44.
[61] Mattera R, Athanasopoulos G, Hyndman RJ. Improving Out-of-Sample Forecasts of Stock Price Indexes with Forecast Reconciliation and Clustering. Monash University, Department of Econometrics and Business Statistics; 2023, URL https://www.monash.edu/business/ebs/research/publications/ebs/2023/wp17-2023.pdf.
[62] Zeng T, Swanson NR. Predictive evaluation of econometric forecasting models in commodity futures markets. Stud Nonlinear Dyn Econom 1998;2(4).
[63] Kwas M, Rubaszek M. Forecasting commodity prices: Looking for a benchmark. Forecasting 2021;3(2):447–59.
[64] Foucault T, Pagano M, Röell A. Market liquidity: theory, evidence, and policy. USA: Oxford University Press; 2013.
[65] Shi W, Lee L-f. Spatial dynamic panel data models with interactive fixed effects. J Econometrics 2017;197(2):323–47.
[66] Bai J, Li K. Dynamic spatial panel data models with common shocks. J Econometrics 2021;224(1):134–60.
[67] Rissanen J. Modeling by shortest data description. Automatica 1978;14(5):465–71.

[68] Salomon D, Motta G. Handbook of data compression. Springer Science & Business Media; 2010.
[69] Sayood K. Introduction to data compression. Morgan Kaufmann; 2017.

Roy Cerqueti is a Full Professor of Statistics at the Sapienza University of Rome, Italy and an affiliate of the GRANEM, University of Angers, France. He is a member of the faculty of the Doctoral School of Social and Economic Sciences. Previously, he served as University Researcher and Associate Professor at the University of Macerata, and as Associate Professor in Mathematical Methods for Economics, Finance and Actuarial Science at Sapienza University of Rome. He was also a review Panel expert in the COST Action Proposal Submission, Evaluation, Selection and Approval procedure (appointment by MIUR - Office of the National Coordinator of COST Projects, from 2017 to 2019). He has published more than 130 scientific papers, almost all of them in international journals. He regularly participates in international conferences also as an invited speaker. He has several editorial assignments and regularly serves as a referee for many journals and scientific institutions.

Valerio Ficcadenti, is an Associate Professor at London South Bank University – Business School, specializing in Research Methods. He earned a PhD in Quantitative Methods for Economic Policy from the University of Macerata. His career has included roles in academia and the private sector, including an assistant manager position at Deloitte, pricing center, in London. His research focuses on rank-size laws and text mining.

Raffaele Mattera holds a PhD in Econometrics at the University of Naples "Federico II" and is Assistant Professor of Statistics at the Sapienza University of Rome, Italy. His main research interest is in computational statistics and predictive analytics, with a focus on clustering and forecasting of temporal, spatio-temporal and network data. He is a member of the International Institute of Forecasters, the Italian Statistical Society and the Spatial Econometrics Association. He regularly serves as a referee for many scientific journals in the field.