



EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Andres Azqueta-Gavaldon,
Dominik Hirschbühl, Luca Onorante,
Lorena Saiz

Nowcasting business cycle
turning points with stock networks
and machine learning

No 2494 / November 2020

Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Abstract

We propose a granular framework that makes use of advanced statistical methods to approximate developments in economy-wide expected corporate earnings. In particular, we evaluate the dynamic network structure of stock returns in the United States as a proxy for the transmission of shocks through the economy and identify node positions (firms) whose connectedness provides a signal for economic growth. The nowcasting exercise, with both the in-sample and the out-of-sample consistent feature selection, highlights which firms are contemporaneously exposed to aggregate downturns and provides a more complete narrative than is usually provided by more aggregate data. The two-state model for predicting periods of negative growth can remarkably well predict future states by using information derived from the node-positions of manufacturing, transportation and financial (particularly insurance) firms. The three-states model, which identifies high, low and negative growth, successfully predicts economic regimes by making use of information from the financial, insurance, and retail sectors.

JEL Classification: C45; C51; D85; E32; N1.

Keywords: real-time; turning point prediction; Granger-causality networks; early warning signal.

Non-technical summary

A real-time identification of recessions and economic “regimes” (i.e. phases of the business cycle) proves difficult, but it is of great importance for policy-makers. We propose a novel approach for identifying changes in economic regimes in real time. Our approach utilises various measures of connectedness in a stock returns network as a proxy for the transmission of shocks through the economy. We show that the dynamic network structure contains forward-looking cyclical information and is a good predictor of booms and recessions.

We emphasise the role of granular, firm-level information, and its connectedness to unveil aggregate shocks. When consumer and business spending softens in some part of the economy, corporate earnings start to decline and businesses seek ways to cut costs. They might downsize their workforce, put a freeze on hiring and delay investments. As these factors become more widespread, unemployment rises, aggregate wages fall and demand weakens further, tipping the economy into recession. Once demand for goods and services increases again, earnings rise supporting higher stock prices. Although every recession is different, there is a clear link between the state of the economy, corporate earnings and stock price developments. Our dynamic stock network captures the connectedness of shocks to expected corporate earnings and hence allows us to detect the propagation of shocks within the economic system.

The main findings can be summarised as follows. The baseline binary state model can predict upcoming recessions by using information from the node-positions of manufacturing, transportation and financial, particularly insurance firms. The ternary state model (featuring high, low growth, and recessions) successfully predicts economic regimes highlighting the role of information stemming from the financial, the insurance, and the retail sectors. Looking at the economic system as a whole, we highlight that during an expansion adverse shocks to firms are mostly idiosyncratic, while during contractions shocks become more widespread resulting in higher connectedness in certain parts of the network. Measures of centrality efficiently summarise economy-wide developments and allow us to monitor the state of the economy in real time.

1 Introduction

A real-time identification of recessions and economic “regimes” (i.e. phases of the business cycle) is of great importance for policy-makers. We propose a novel approach for identifying changes in economic regimes in real time. Our approach utilises measures of connectedness in a stock network as a proxy for the transmission of shocks through the economy and, hence, for widespread developments in corporate earnings. We show that the dynamic network structure contains forward-looking cyclical information, making it an optimal predictor of booms and recessions.

The identification of turning points or different phases of the business cycle in real time has proven rather difficult, despite a large literature on the subject and an ever increasing number of leading indicators. In this respect, the consensus in the literature is that the set of relevant indicators for predicting shifts in business cycle phases (e.g. recessions) changes over time. Therefore, indicators that are useful for predicting one recession do not necessarily serve to predict other recessions as “every cycle is different”. A general problem of forecasting economic regimes as shown by [Ng \(2014\)](#) is that only a few indicators are actually useful and the effectiveness of these indicators varies according to the forecast horizon. For example, financial variables (e.g. term and corporate spreads) are appropriate for forecasting 6 to 12 months ahead in view of their forward-looking nature.¹

One important reason for the delay in identifying shifts in business cycle phases is that the overall economy is normally reduced to a summation of its components and micro-level shocks are therefore assumed to offset each other at the aggregate level. By contrast, the approach adopted in this work emphasises the importance of granularity and the role of economy-wide return connectedness. When consumer and business spending softens in some parts of the economy, corporate earnings start to decline and businesses seek ways to cut costs. They might downsize their workforce, put a freeze on hiring and delay investments. As this situation becomes more widespread, unemployment rises, aggregate wages fall and demand weakens further,

¹ In our network approach, we focus on nowcasting. Investigating the forward-looking information content in the derived network measures is left for future work.

tipping the economy into recession. Although every recession is different, there is a clear link between the state of the economy, corporate earnings and stock price developments. Once demand for goods and services increases again, earnings rise supporting higher stock prices. Translating the company information contained in stock returns into a dynamic stock network allows us to proxy widespread developments in actual and expected corporate earnings and, hence, shock propagation within the economic system. Considering that changes in stock market returns can reflect both idiosyncratic and economy-wide shocks, we calculate pairwise Granger causality within companies' stock returns. A relation at a period t only exists if the daily returns of a company over a 12-month horizon provide statistically significant information for predicting the returns of another company and vice versa. During an expansion adverse shocks to firms are mostly idiosyncratic, while during contractions shocks become more widespread which results in higher connectedness in the network. We can monitor economy-wide developments by calculating measures of centrality that reflect increasing co-movement in the economy, thereby creating a “fat” data-set of roughly 5,002 time-series. Given that pure financial shocks are predominant in the network, we need to identify the characteristics of the dynamic network that are more likely to reflect economy-wide propagation leading to booms or recessions.

The work that is closest to ours is [Heiberger \(2018\)](#),² but we extend it in several ways. First, we construct the network on the basis of pairwise Granger-causalities rather than correlations. Limiting links to statistically relevant causality focuses on relationships that are economically relevant for shock transmission in the network, while reducing noise. The point is to identify the stocks/nodes that react earlier to shocks and play a role in their transmission (i.e. they are systemically relevant). We show that relevant economy-wide developments are only imperfectly reflected in the aggregate data, while firm-specific network characteristics are important early indicators of the state of the US economy. Second, we innovate by letting different algorithms such as support vector machines (SVM), naive Bayes (NB) and logistic

² The author creates a dynamic correlation-based network for stock market returns in a 12-month rolling window. It assumes a critical threshold of correlation of 0.7. The author's approach requires an optimal number of network features of around 2,200.

regression (LR) compete.

We analyse two models. The first tries to identify two economic regimes: “recession” (defined as negative growth in the quarter) and its complement. In the second model, we try to perform the more challenging task of distinguishing between three regimes: “negative growth” (defined as negative growth in the quarter), “low growth” (defined as positive, but below high growth) and “high growth” (defined as annualised growth higher than 3% in the quarter). Our findings suggest that our network approach can help to detect changes in economic regimes in real-time, giving policy-makers a tool to anticipate slowdowns and take measures against them. The main findings can be summarised as follows: (i) Applying recursive feature elimination (RFE) on the entire sample, we are able to determine ex-post which firms were exposed to economy-wide earnings shocks that turned into aggregate fluctuations, while remaining resilient to others. Looking at the sectors where these firms operate, a more complete narrative can be derived than is usually derived from more aggregate data. The two-state model for predicting periods of negative growth can remarkably well predict future states by using information from the node-positions of manufacturing, transportation and financial, particularly insurance firms. The ternary state model successfully predicts economic regimes by making use of information from a few node-positions from the financial and insurance sectors, as well as the retail sector. (ii) Using features (variables) that are selected to be optimal until the fourth quarter of 2006 (i.e. excluding the out-of-sample period), we validate our models in terms of true out-of-sample performance to verify that the in-sample results are not solely due to favorable feature selection. Unsurprisingly, the number of selected features is higher than in the in-sample exercise. While the predictions of the binary state model are weak, the ternary state captures the turning point of the great recession and performs overall well, while creating some false positives in the negative growth state in moments of weak growth momentum.

The rest of this paper proceeds as follows. Section 2 reviews differing approaches and sets them into the context of this study. Section 3 discusses the methodology of our approach. Section 4 illustrates the findings while Section 5 concludes.

2 Related literature

A large amount of literature has investigated models and economic and financial indicators that could be useful to predict or anticipate an economic downturn. The real-time predictive power of such models, however, has often proved unsatisfactory. In conventional frameworks, the inclusion of selected financial indicators has partially captured the effect of the financial cycle on recessions. The high costs that are associated with systemic risk and financial crises led to an extensive literature investigating the linkages between the financial sector and crises. Early warning models estimated with conventional techniques include [Alessi and Detken \(2011\)](#), [Rose and Spiegel \(2012\)](#), [Gourinchas and Obstfeld \(2012\)](#), [Duca and Peltonen \(2013\)](#) or [Drehmann and Juselius \(2014\)](#). This macro approach (based on macro aggregates) is partially successful, and has been able to identify some determinants of financial crises (e.g. the credit to GDP ratio). However, the problem remains that variables that have been found to be important, such as credit, are usually in the form of aggregates, which, by their nature, miss the tails and off-mean events that may trigger a recession.

In an attempt to overcome this, [Adrian et al. \(2019\)](#) estimate the predictive US GDP growth conditional distribution based on a synthetic index of financial conditions. This index aggregates variables that cover financial risk, leverage and credit quality. The analysis finds that the lower quantiles of GDP growth are more sensitive to financial conditions than upper quantiles. This finding suggests an asymmetric and non-linear relationship between financial and real variables. [Plagborg-Møller et al. \(2020\)](#) find that financial variables have very limited predictive power for the distribution of GDP growth at short horizons, especially, but not limited to, the tail risk. This is due to the fact that moments other than the mean are estimated imprecisely and information in monthly financial variables is highly correlated.

However, there are still several shortcomings of macro-based analysis. First, using aggregate data averages out shocks. That is, aggregation in the data “averages out” micro shocks because interactions are neglected in favour of averaging. Therefore, a macro approach (using aggregate data) is only suitable for identifying aggregate

shocks that result from macroeconomic policies. To the extent that systemic risk is determined by the complex interactions of micro shocks, macro aggregates would not be able to account for network effects. Second, aggregate data are usually not available on a timely basis. In the United States, the Bureau of Economic Analysis (BEA) releases the first estimate of the gross domestic product one month after the end of the reference quarter. A second estimate is published two months after the end of the reference quarter. And finally, a third estimate is released three months after the end of the reference quarter. A breakdown of GDP by industry is available later, with an average delay of 120 days.³ Third, there are substantial non-linearities that standard regression models struggle to capture. Attempts to identify non-linearities in economic processes mostly involve Markov switching models and multiple economic regimes (for instance, [Chauvet and Hamilton \(2006\)](#) and [Camacho et al. \(2018\)](#)), or the estimation of probit and logit regressions.

Recently, researchers have started to explore unconventional “big data” and/or techniques such as machine learning methods to forecast the business cycle, obtaining some promising results. For instance, [Qi \(2001\)](#) predicts US recessions with neural network models and a set of leading financial and economic indicators, and [Giusto and Piger \(2017\)](#) use Learning Vector Quantization (LVQ) to identify US business cycle turning points in real time.

In this paper, we follow this strand of research and propose to evaluate network information generated from stock returns to proxy the propagation of economic shocks that result in recessions. For example, [Gabaix \(2011\)](#) shows that idiosyncratic firm-level shocks can explain a large part of aggregate movements and provide a micro-foundation for aggregate shocks. In particular, the idiosyncratic movements of the largest 100 firms in the United States appear to explain about one-third of variations in output growth. [Acemoglu et al. \(2012\)](#) argue that in the presence of intersectoral input and output linkages microeconomic idiosyncratic shocks may lead to aggregate fluctuations. [Næs et al. \(2011\)](#) find that the liquidity in stock markets dries up prior to a crisis in the real economy while [Li \(2017\)](#) has shown that micro uncertainty (as

³ See the official release schedule published by the BEA on <https://www.bea.gov/news/schedule>

measured by the idiosyncratic component in stock prices), translates into macro uncertainty via credit frictions in line with [Bernanke et al. \(1999\)](#). [Ferreira \(2018\)](#) shows that financial skewness, measured by comparing cross-sectional upside and downside risks of the distribution of stock market returns of financial firms, is a powerful predictor of business cycle fluctuations. [Adam and Merkel \(2019\)](#) present a model of the business cycle with extrapolative belief formation in which price effects of technology shocks are amplified such that large and persistent boom and bust cycles occur as observed in stock prices.

Nonetheless, there are two main challenges when incorporating the information provided by a network into fore- and nowcasting models: i) how to accommodate a network structure in an econometric framework (i.e. selection of informative features of the nodes in the network); and ii) whether and how to deal with the curse of dimensionality and select what is relevant. Given that networks usually produce a large amount of information (e.g. complex interactions between the groups that form the network), standard econometric models would not be sufficient for this task (for example, our baseline specification is described by approximately 5,002 time-series), and would, therefore, go beyond what macroeconometricians define as big data (see [Giannone et al. \(2018\)](#)). To deal with this “curse of dimensionality”, we use machine learning techniques that perform variable selection, in particular RFE. In addition, to evaluate the remaining information we use competing machine learning methods: support vector machines, naive Bayes and a benchmark logistic regression.

We build our network on financial data. More precisely, we use stock returns of listed companies in the Standard and Poor’s (S&P) 500 index. Financial data have the advantage of being non-proprietary and available in real time. From an economist’s point of view, there are two main reasons why quotations on financial markets can help to explain the business cycle. First, financial markets process information efficiently, and quotations may be seen as informative summary statistics about a variety of firm characteristics such as the current status of company fundamentals and future expected performance as judged by the markets. Second, financial markets may additionally be the cause of sudden stops in the economy by turning

small, idiosyncratic shocks into big aggregate shocks. Moreover, the presence of a representative sample of insurers and banks covers financial and credit links. Most importantly, stock prices are not subject to meaningful errors or revisions, which makes them particularly useful for nowcasting and forecasting.

Notably, we recognise the difficulty of being robust to the “every cycle is different” criticism. To achieve as much robustness as possible across different cycles, we exploit different types of information. We compute our networks using rolling 12-months windows for each quarter ending on first working day of the second month in the quarter. Using those networks, we calculate node and aggregate network measures. Finally, we obtain for each measure in each network one metric that we combine over time to obtain one time-series. We combine the time-series for all measures in a matrix. Finally, we let machine learning techniques discriminate between useful information and noise.

3 Data and methodology

In this section we explain how we can model the relationships among companies using network theory. A network represents pair-wise relationships between nodes (firms in our case). Depending on the different attributes and relationships within the network, it will tell us relevant aggregate and sector-specific information about the way companies interact in the economy over time. Thus, representing the economy as a network has several potential advantages over looking at aggregate data (i.e. overall stock market indices). We can use a wide range of information embedded in the network, such as the key firm or sector in the economy, and, more importantly, we can account for non-linearities in the relationship between economic players and the economic cycle. The exact procedure is summarised in the following steps:

Steps of the methodology

1. We have 500 series of daily returns data for the period from the first quarter of 1980 to the first quarter of 2019.
2. For each quarter:

- (a) For each possible pair-wise relationship between stock returns:
 - i. Pick two firms' daily returns series over a 12-month (rolling) window ending on the first working day of the second month in the quarter (e.g. 1 February for the first quarter of the year).
 - ii. Test Granger causality of the two series using an heteroscedasticity and autocorrelation consistent (HAC) regression and 1% significance level.
 - iii. If a Granger-causal relationship fill the network's adjacency matrix $[500 \times (500-1)]$ with a 1 at the respective position, if not fill a 0.
 - (b) From the first quarter of 1982 to the first quarter 2019, we calculate the proposed network measures and combine them along the time dimension to track the evolution of these measures over time. In the feature matrix we consider solely the contemporary observation, hence we do not account for lagged values. We discard series that show too little variation over time.
3. We standardise all features to have a zero mean and a standard deviation of 1.
 4. We use RFE in-sample to find contemporary highly correlated network metrics and choose the combination that performs best assuming k-fold cross-validation with $k=5^4$ if we treat the entire data as in-sample and $k=3$ in case we consider an out-of-sample period. This unveils a reduced amount of features that have been independently highly correlated with GDP growth rates in different sub-periods in the past.
 5. We train the logistic, SVM and naive Bayes classifiers on ranges of GDP growth rates making use of the selected features.
 6. We use the classification to nowcast economic "regimes" using stock network data in the selected period.

3.1 Data

We build our explanatory variables using the daily stock return index (which accounts for dividends and reinvestment within the underlying company). We obtain the data

⁴ 3 in case of the ternary state model.

from Datastream and include all companies listed in the S&P 500 in the third quarter of 2018. The complete list of companies with identifiers can be found in Table A.3.3 in the Appendix. Table 1 illustrates the sectoral composition of the companies in the selected sample of firms as in the third quarter of 2018.⁵ As we can see, out of the 500 companies in the sample, the biggest share comprises manufacturing companies (191), the second biggest share comprise financial companies (97) and transportation and services companies are tied for the third largest share (both with 71 companies). Note that we obtain sector-specific information using the Standard Industrial Classification (SIC).⁶

Table 1 Sectoral composition of the listed companies in the S&P 500 index

Sector	#Nr. of companies
Primary and construction	24
Manufacturing	191
Transportation	71
Wholesale	11
Retail	35
Finance	97
Services	71
Total	500

Notes: The table displays the sectoral composition of the S&P 500 index according to the Standard Industrial Classification as described in Footnote 6.

The total time span is from 2 January 1980 to 31 March 2019. Following the standard procedure in the literature, if there are no observations for the return of a given company (e.g. a company did not exist in the 1980s) information about the company is retained, but as we will see later on, no causality is assumed. Using the constituents as at the third quarter of 2018 allows us to ignore changing compositions, but requires us to acknowledge that the network is imbalanced in the 1980s as it is gradually being compiled (see Figure A.16 in the Appendix). We obtained stock returns for each company i and time t by computing the difference in the logarithm of the return index at t and $t - 1$: $r_{it} = \log(RI_{it}) - \log(RI_{it-1})$.


⁵ The choice constitutes a trade-off between keeping the number of companies limited, on the one hand, and focusing on firms that are important at the present time, on the other hand. This results in imbalanced data as companies that were included in the third quarter of 2018 might not have existed in the 1980s.

⁶ The sectors are: Primary and construction (0100-1799), manufacturing (2999-3999), wholesale (5000-5199), transportation, electricity and gas (4000-4999), retail (5200-5999), financial, insurance and real-estate (6000-6799), and services (7000-8999). The classification we propose can be applied in a more granular way. This information is solely used in the statistical review of features.

3.2 Pair-wise Granger-causality network

In order to build networks using firm stock returns, we compute pairwise Granger causality tests among each of the companies. To illustrate how we construct the Granger causality network, consider Figure 1 which shows how a network can represent the interactions of three different companies. It is assumed that the returns of each company are available as a time series (top panel). In the left panel, the returns of companies B and C co-move whereas neither of them co-move with the returns of company A. These relationships are displayed as a network by connecting the nodes of only company B and company C. Moreover, these connections are bi-directional, in the sense that company B's returns cause company C's and vice versa. In the third panel, however, we see that company A's returns causes company B's returns, since they move earlier in the same direction, but the reverse is not true (i.e. company B's returns do not affect company A's returns). For this reason, as we will see later on, we will have asymmetric relationships among companies.

Figure 1 Construction of a Granger causality network



Notes: The graph illustrates how the Granger causality stock network is constructed. If returns of company B help to forecast returns of company C, and vice versa then a bi-directional edge is created between the two companies. If returns of company A Granger-cause returns of company B, but not vice versa a one-directional edge is assumed.

To account for causal interactions between stock returns rather than simply co-movements or correlation, we make use of Granger-causality networks. Granger-

causality networks have gained prominence by outperforming other approaches in displaying an extremely volatile degree of connectedness between financial institutions in times of financial crisis (see [Billio et al. \(2012\)](#)).⁷ Moreover, by applying them to a large and representative set of companies in the economy, it is also possible to capture patterns of economic developments over the business cycle.

In order to calculate the Granger-causality networks, we follow [Granger \(1969\)](#), [Granger \(1979\)](#) and [Billio et al. \(2012\)](#) and estimate linear Granger-causality networks on stock return indexes over rolling windows with a 12 month horizon. We do so in two steps: first, we estimate bivariate vector autoregressions (VAR) using the Bayesian Information Criterion (BIC) to specify the lags to be chosen.⁸ Second, we compute pairwise Granger-causality tests between the stock returns corrected for heteroscedasticity and autocorrelation (HAC). If the hypothesis of no causality is rejected, we assume a link between both stocks in the network. For both steps, we consider a significance level of 1% and a rolling window size of 12 months, as those parameters constitute a good choice for reflecting more persistent shocks. For instance, a 12-month rolling window is not solely focused on the short term, but also on somewhat more persistent shocks, which will be discussed in the next section. Specifically, we represent each network in each period t via an adjacency matrix $A_{ij}^{(t)}$:

$$A_{ij}^{(t)} = \begin{cases} 1 & \text{if } i \text{ Granger causes } j \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where we impose $A_{ii}^{(t)} \equiv 0$. It is worth noting that the adjacency matrix, unlike correlation matrices as in [Heiberger \(2018\)](#), is generally not symmetric, as variable i can Granger-cause j without the opposite being true.

⁷ In many applications, Granger-causality networks outperform networks based on dynamic Bayesian inference or simple correlation. In particular, if the data are not short they outperform dynamic Bayesian networks as shown by [Zou and Feng \(2009\)](#). In contrast to [Billio et al. \(2012\)](#), we apply heteroscedasticity and autocorrelation corrected Granger causality tests.

⁸ Assuming an informationally efficient financial market, short-term asset price changes should not be related to other lagged variables and no return series would be expected to Granger-cause another. However, the presence of value of risk constraints, costs of gathering and processing information, and institutional restrictions creates Granger causalities among price changes in financial assets over time.

3.3 Network topology

Nonetheless, there are many other rich attributes that we can extract from a network beyond how compact or loose it is (as displayed in Figure 3). In this section we describe aggregate and firm/sector-specific network measures that we use throughout the analysis. We will rely on a number of centrality measures, which aim to identify the most important vertices within a network.

- **Dynamic causality index (economy and sectors)**

$$DCI = \frac{1}{N(N-1)} \sum_{i \neq j}^N \sum_{j \neq i}^N A_{ij} \quad (2)$$

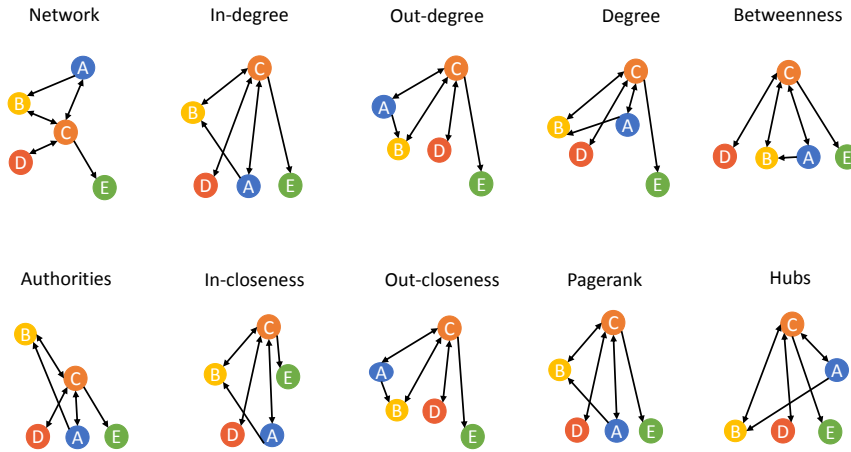
The dynamic causality index (DCI) is described by the number of causal relationships over a given period divided by the number of total possible causal relationships. The DCI describes the total number of Granger-causal relationships in the system and hence is a measure of overall connectedness. We consider the DCI of the whole network and of sector-specific sub-networks.

- **Firm-specific or node-specific measures**

In addition, we calculate various measures of node-level centrality which are illustrated in Figure 2. Measures of centrality attempt to identify the most important vertices within a network, e.g. the most influential firms or sectors. As we will see, there are many different ways of doing so, such as using the number of overall connections, connections within influential groups or flow across the network. In our set up, we consider five broad categories (and many more subcategories) of centrality which may characterise a network:

1. *Degree centrality.* Degree centrality is a basic metric of connectedness which measures the exposure of a company to the entire economy. More formally, it is the fraction of statistically significant Granger-causality relationships

Figure 2 Actual network and node-level measures of centrality



Notes: The graph illustrates the same network at a single period with nodes ordered and pictured according to different network measures. Higher nodes are assigned higher values.

among all $N(N-1)$ pairs of N companies:

$$DGC_i = \frac{1}{N(N-1)} \sum_{i=1}^N \sum_{j \neq i} (j \rightarrow i). \quad (3)$$

Subcategories of degree centrality include the in-degree and out-degree measures. These measures account for all in-going and out-going relationships respectively. A high in-degree measure identifies a receiver of shocks whereas a high out-degree connectedness identifies those nodes that are either originators of shocks or are exposed earlier to shocks in the network.

2. *Closeness centrality.* Closeness centrality defines node centrality in terms of the shortest paths. The distance between nodes i and j is given by the number of edges in the shortest path connecting them. A central node is therefore close to all other nodes in the network in terms of the average distance between this node and all others. Nodes with a high closeness reach the rest of the network in just a few steps and are able to quickly spread shocks through the network. Since our network is directed (i.e. there are Granger-causality directions among firms), it is better to distinguish be-

tween in-closeness and out-closeness centrality. In-closeness centrality measures the degree to which a node can be easily reached from other nodes (i.e. using the shortest path). Out-closeness centrality measures the degree to which a node can easily reach other nodes.

3. *Betweenness centrality.* Betweenness centrality is a measure related to closeness centrality since it is also based on shortest paths. Nonetheless, betweenness centrality does not require the network to be fully connected but can be calculated over multiple unconnected components (groups of nodes that are all connected to each other).

$$B_i = \sum_{i \neq j \neq k} \frac{\sigma_{j,k}(i)}{\sigma_{j,k}} \quad (4)$$

where $\sigma_{j,k}$ is the path between j and k , and $\sigma_{j,k}(i)$ is the number of paths a node i lies on. Firms with high betweenness centrality serve as a bridge to otherwise weakly connected nodes, and have a high influence over the flow and contagion of shocks across parts of the network (e.g. between different sectors). An example of firms that serve as bridges is banks or financial institutions. They are connected to different firms in the economy as they provide them with credit and therefore can easily spread negative shocks throughout the whole economy.

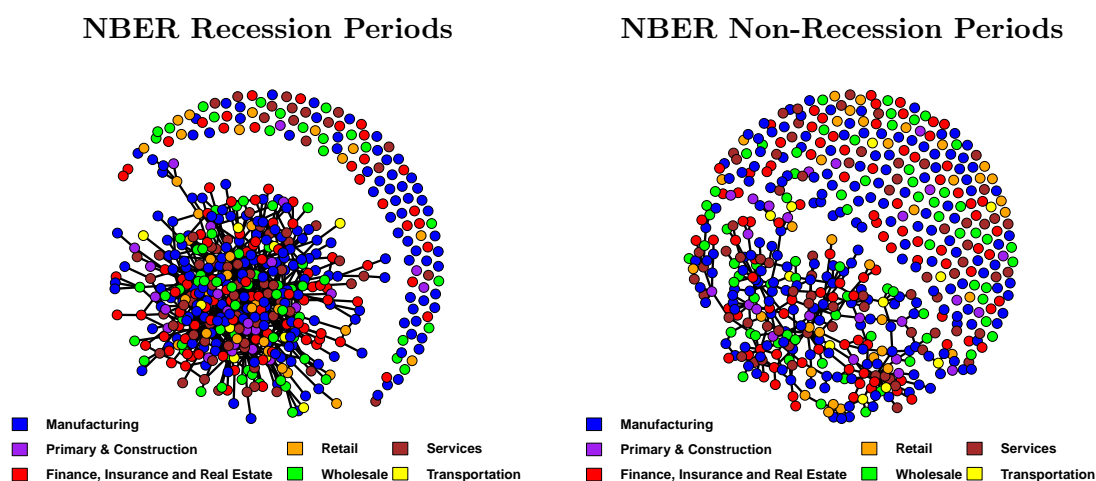
4. *PageRank centrality.* PageRank is another centrality measure developed by Google search in order to rank web pages according to how influential they are. PageRank works by counting the number and quality of links to a page to determine its importance. It does so by representing the network as a Markov chain in which each node is a state. The adjacency matrix can be transformed such that its elements represent the probability of transition between a pair of nodes. This measure computes the probability of arriving at the node i after a large number of steps following a random walk navigation through the network. Therefore, nodes with a higher probability are more central.

5. *Hubs and authorities centrality.* Hubs and authorities centrality are measures that give high scores to nodes which are so-called "hubs" or "authorities". A hub node is one that points to many authority nodes, and an authority node is one that is pointed to by many hub nodes.⁹

3.4 Mapping stock networks to the business cycle

Figure 3 shows the average structure of the stock network during periods of recession and non-recession as defined by the US National Bureau of Economic Research (NBER). The recession period chart displays all relationships that hold true in at least five of the nine NBER recession periods, while the non-recession period chart displays all other periods. It appears that the network displays a high degree of connectedness during economic contractions, while during expansions shocks and, hence, stock returns behave more idiosyncratically, leading to lower connectedness. This information is summarised in the histogram in Figure 4.

Figure 3 Stylised network during phases of the business cycle




Notes: The figures display averaged quarterly stock networks of the S&P 500 based on pairwise Granger causal relationships from 2000:Q1 to 2019:Q1 on a 12 month rolling window bandwidth and a 5% significance level. 68 non-recession periods and 9 recession periods. The threshold for the averaging is 50% causalities per relation for the non-recession and the recession periods. This means that at connection holds at least in 5 out of 9 periods for the NBER recession case. As the network for the non-recession periods would display no connectedness at a threshold of 34 out of 68 periods, we can lower the threshold to 7 relations per 68 periods.

There is evidence that the financial system has increased in size and complex-

⁹ The easiest way to understand this is by using the example of building a network of journal articles using citations. An article that is referenced by many other articles has a high authority score, and an article that cites many authority articles has a high hub score.

ity and its potential for causing economic disruption has correspondingly increased (see, for example Zingales (2015)). The strong decoupling between money and credit aggregates shows that the leverage of the financial sector has increased strongly in recent decades. The transformations in financial intermediation are poised to have far-reaching implications for real-financial interactions and ultimately for business cycle dynamics. Not only has the size of the financial sector been steadily increasing in comparison with the size of the real economy; but the influence of the sector has also become proportionally greater as the result of increasing interdependence between financial and non-financial firms. From the perspective of systemic risk, the financial sector can be thought of as a network of connected institutions that may benefit from having commercial relationships with each other but that can also translate micro shocks (e.g. firm-specific shocks) into systemic shock and cause financial crises.

Figure 4 Normalised stable node-relations during phases of the business cycle

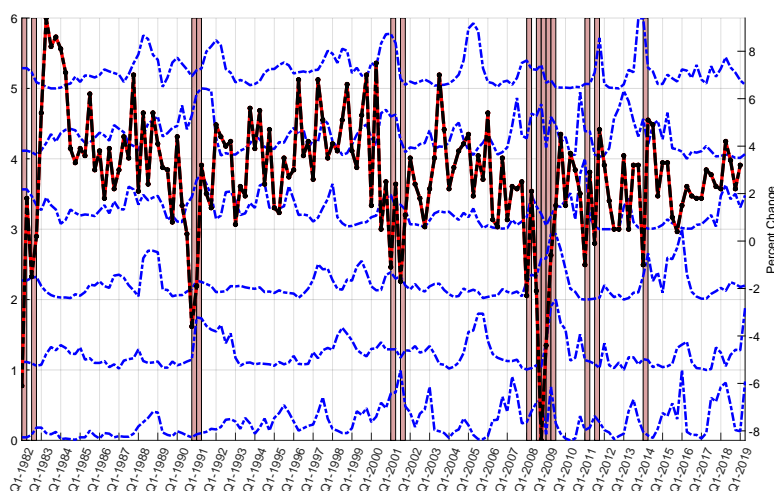


Notes: The graph illustrates normalised quarterly connectedness from the first quarter of 2000 to the first quarter of 2019 on a 12-month rolling window bandwidth with a 5% significance level. Connectedness in specific nodes in NBER recession periods are more persistent than in non-recession periods where connectedness might be more idiosyncratic. A histogram using data since the first quarter of 1982 is available in Figure A.17 in the Appendix.

To nowcast the business cycle using the wide range of information embedded in this network, we must convert network information into “useful” time series data. We do so in two steps. First, we describe the network using a wide set of network and node metrics. By observing the evolution of these measures over time, we can easily convert them into time-series. In the second step, the vast number of network measures has to be reduced. Our network is described by 5,002 time series, which

contain a considerable amount of noise. We apply machine learning techniques as standard econometrics cannot deal with such a sizeable dataset, owing to the “curse” of dimensionality (see, for example, [Giannone et al. \(2018\)](#)). We first eliminate uninformative signals that have characteristics that do not change over time (constant or semi-constant), then we allow for further statistical data selection making use of RFE. Dimension reduction techniques help to isolate relationships that are highly cyclical and that can be applied to predict recessions or periods of high and low growth.

Figure 5 Node-specific network measures vs GDP growth



Notes: The graph illustrates various illustrative firm-specific stock network measures derived from daily returns (dashed blue lines) and the annualised quarterly growth rate of the economy (dashed black-red line). The red bars indicate periods of negative growth.

After the “translation”, and as illustrated in Figure 5, a set of best-performing measures is used to predict the cycle. The business cycle data consist of seasonally adjusted quarterly real GDP growth figures (annualised rate) for the United States from the FRED database.¹⁰ We consider both aggregate and disaggregate (firm-specific) characteristics of the network and let the algorithm decide which are most helpful for classification. During expansions we expect shocks to returns to be firm-specific and less correlated across firms. Since the economy tends to be a shock absorber in normal times, our network should appear less aggregated. However, during downturns shocks might propagate across firms, thereby hitting the economy broadly and in a persistent manner. While aggregate measures can hardly reflect these granular developments,


¹⁰ The series identifier in the Federal Reserve Economic Data (FRED) database is A191RL1Q225SBEA.

we postulate that pooling node-specific patterns in the networks can capture those dynamics well. Furthermore, economic downturns can be detected earlier in specific firms and sectors.

Using the aforesaid metrics, we describe in this section the evolution of some aspects of the network. In particular, we discuss measures of the behaviour of the aggregate network and of single nodes or firms.

3.4.1 Aggregate measures

Figure 6 Real GDP growth rate vs dynamic causality index for the S&P 500



Notes: The red bars indicate periods of negative GDP growth. The different series display dynamic causality indexes for different rolling window bandwidths (i.e. 3, 6, 12, 18 and 24 months) and at a significance level of 1%. The dashed black line displays the seasonally adjusted annualised quarterly real GDP growth rate.

Figure 6 displays the dynamic aggregate behavior of the networks at the 1% confidence level for different rolling window bandwidths. The red bars represent periods of negative growth. Focusing on bust periods, it is possible to observe differences in the dynamic causality indexes (DCI) for different bandwidths. The measure computed with 24 months bandwidth seems to react early to big changes, but it is also very persistent which is problematic in an environment of increased volatility of connectedness. On the other hand, the measure with the shortest bandwidth (3 months)

misses the downturns. For this reason, we could expect dynamic causality indexes using intermediate bandwidths (i.e. 6 and 12 months) to have higher explanatory power. Besides peaking in periods with below zero growth (recessions) mostly, all measures tend to also peak in boom periods. For example, all dynamic causality indices peak in 1987:Q4 (when the stock market crashed), but this didn't lead to a slowdown in growth. The same happens for the peaks in 2000 and 2005. For this reason, we argue that the DCI computed at the level of the whole network is too noisy to predict economic activity.

3.4.2 Firm- or node-specific measures

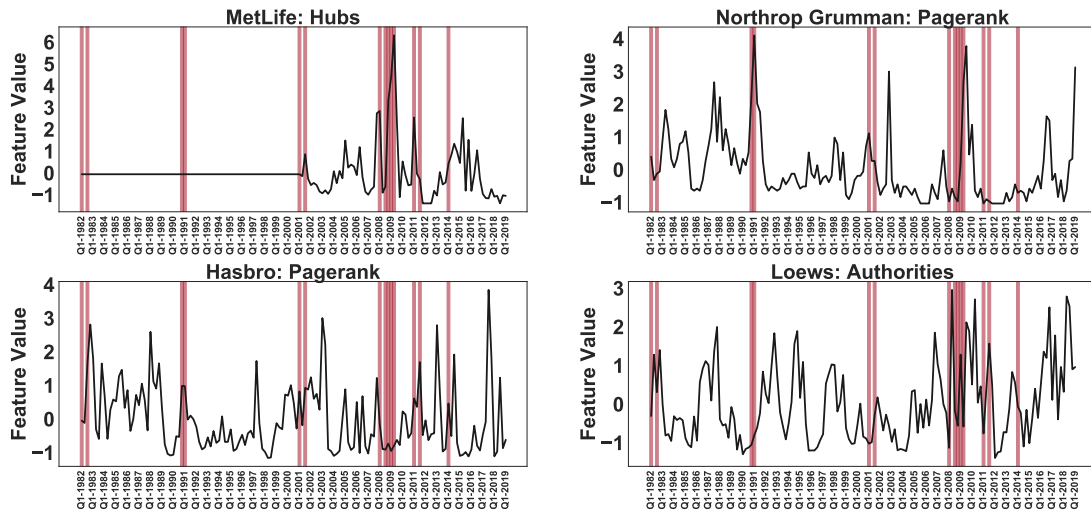
The most granular way to display information from the network is to evaluate the position of the individual nodes. Displaying connectedness at the node level over time is a rather difficult exercise owing to noise. Moreover, to isolate cyclical connectedness, we need to look at network topologies over time. For illustrative purposes Figure 7 displays the hubs centrality measure for MetLife, a large insurer, as well as the authorities centrality measure for Loews, a conglomerate dealing with insurance, pipeline transport, oil drilling and hotels. Furthermore, we show the PageRank measure for Northrop Grumman, which is an aerospace and defense technology company, and the PageRank measure for Hasbro, the largest toy maker in the world. These companies tend to have a central position in the stock network during recessions or periods of negative growth (e.g. Great Recession).

3.5 Feature normalisation and recursive feature elimination

It is important to have some flexibility when choosing from the wide range of centrality measures; while all measures are intended to reveal important node characteristics, some might be a better fit when trying to forecast business cycles at different horizons. As we have seen, all the centrality measures try to identify the key company or sector occupying pivotal structural positions within the network which might help us understand business cycle dynamics.¹¹ For this reason, it is of key importance to try to reveal the centrality measures that are more informative on the business cycle. To

¹¹ For a critical review on how centrality measures apply to social networks, see [Landherr et al. \(2010\)](#).

Figure 7 Firm-specific network measures



Notes: The black lines display firm-specific measures of centrality as outlined in the subtitles. The red bars indicate periods of negative growth. The network measures are derived from a 12-month rolling window bandwidth and a 1% level of significance.

do so, we employ a feature selection method to discard all uninformative measures in the context of our nowcasting exercise.¹²

We summarise the network measures in the vector of explanatory variables X , the so called feature matrix which comprises 5,002 features in total.¹³ In a first step, we manually remove all features exhibiting too low variance. After eliminating uninformative features, the X vector in our model contains about 4,104 features, each providing 149 time observations from the first quarter of 1982 until the first quarter of 2019. We standardise the features to have a zero mean and a standard deviation of 1.

Then we let algorithms, so called dimension reduction techniques, decide which features can improve predictive performance. In particular, we apply RFE as shown by Guyon et al. (2002) with a k-fold cross-validation as in Yan and Zhang (2015).

¹² Please note that we do not only discard less robust metrics, but also discard useful information that would be helpful at different forecasting horizons.

¹³ Note that richer specifications might improve model performance. Among others Davig and Hall (2019) including a rich lag structure in machine learning set-ups. This could be a promising avenue for future work as predictive information might be available several quarters before, which are in the current nowcasting setup only inadequately reflected. In alternative specifications, we have used a multitude of rolling window sizes and significance levels of the same network. While this can clearly enhance model performance it is computationally very costly. Therefore, we decided to discuss our findings based on reduced data-sets with 5,002 features. This number stems from 500 companies times ten node-specific measures (betweenness, degree, in-degree, out-degree, hubs, authorities, pagerank, closeness, in-closeness, out-closeness) and two aggregate measures (dynamic causality index, cluster coefficient).

RFE is a classical algorithm in machine learning used to select the most relevant set of features (variables) in a model. Unlike dimension reduction methods (e.g. principal components analysis), feature selection methods do not transform the features into a lower dimension, but rather remove those features discovered by the model to be of least importance. This step is very important because we are not imposing any structure on our model (e.g. conditioning the behaviour of the dependent variable into a set of variables that we have previously selected) but rather we are letting the model choose for us the hidden or unknown behavior that produces higher accuracy (that predicts GDP growth rates in our case). To achieve this, we first need to remove all variables that are irrelevant, insignificant or unimportant to the model. In particular, the algorithm removes weak features in each iteration until it has found the optimal combination according to the F1 score.¹⁴ This method performs well for our task as our features are (noisy) proxies of economy-wide shocks which require to be evaluated in pools. The algorithm is designed to use at least five features. RFE is an established and rigorous method for identifying relevant features before feeding them into a machine learning algorithm which reduces the problem of overfitting (or underfitting).¹⁵ To overcome the criticism of introducing information which was not known in that period, in order to select signals that are deemed optimal today, apply k-fold cross-validation to show that in each combination of sub-periods the measures must be optimal.

3.5.1 Model specification

The vector Y_k is given by (i) a binary model where negative annualised quarterly GDP growth rates take the value 1, while all other periods take the value 0:

$$Y_k = \begin{cases} 1 & \text{negative growth} & \text{if } QGDP_k < 0 \\ 0 & \text{normal} & \text{if } QGDP_k \geq 0 \end{cases} \quad (5)$$

¹⁴ The F1 score is defined as the harmonic mean of precision and recall. Precision is defined as the number of true positives over the number of true positives plus the number of false positives. Recall is defined as the number of true positives over the number of true positives plus the number of false negatives.

¹⁵ While our Y_k can take the binary form (negative growth and its complement) or the ternary form (high, low and negative growth periods). All cases will be defined in the next subsection.

or (ii) high, low and negative growth (HLN) as specified in Equation 6 for low growth where the rate is less than 0%, normal (between 0% and 3%), and high (more than 3%)¹⁶:

$$Y_k = \begin{cases} 2 & \text{high} & \text{if } QGDP_k > 3.0 \\ 1 & \text{negative} & \text{if } QGDP_k < 0 \\ 0 & \text{low} & \text{if } QGDP_k \leq 3.0 \text{ and } QGDP_k \geq 0 \end{cases} \quad (6)$$

We use data from the first quarter of 1982 to the fourth quarter of 2006 as training data and nowcast, using contemporaneous information from the current quarter, on an expanding window all periods from the first quarter 2007 to the first quarter of 2019.

3.5.2 The classifiers

In this subsection, we briefly describe the classifiers used in this study.

Naive Bayes

The naive Bayes (NB) classifier is a widespread method in machine learning applications.¹⁷ According to Bayes' theorem:

$$P(Y_k|X) = \frac{P(Y_k)P(X|Y_k)}{P(X)} \quad (7)$$

with Y_k representing the k-th class, in our case periods with below zero growth or recession periods, while $X=x_1, \dots, x_n$ are n observed variables, which are our features. Assuming that the explanatory variables are independent, $P(Y_k)P(X|Y_k)$ can be expressed as:

$$P(Y_k)P(x_1|Y_k)\dots P(x_n|Y_k) = P(Y_k) \prod P(x_i|Y_k) \quad (8)$$

The approach is called naive because the observations are assumed to be independent,

¹⁶ It is to be noted that we abstract from NBER recessions in this study as we are using non-standard data at a higher frequency for which this definition would appear somewhat artificial with conventional setups. We performed the RFE exercises with NBER data also and we identified a stronger impact on manufacturing as compared with the negative growth periods case.

¹⁷ See [Jordan and Mitchell \(2015\)](#), [Hagenau et al. \(2013\)](#) and [Heiberger \(2018\)](#).

while in reality they are not. However, the violation of this assumption would have only minimal consequences for the predictive ability of this Bayesian approach, and on the contrary could even contribute to improving it as shown by [Rish et al. \(2001\)](#). From Equation 1 and 7, we can derive the decision rule on how to assign a class \hat{Y} to a set of observations X with

$$\hat{Y} = \underset{k \in (1, \dots, k)}{\operatorname{argmax}} P(Y_k) \prod P(x_i | Y_k) \quad (9)$$

[Davig and Hall \(2019\)](#) showed that this simple technique can outperform logistic regressions in recession forecasting. This is due to the fact that it converges faster to its asymptotic error rate than, for example, logistic regression. This is an important advantage in applications with a short data span.

Support vector machines

Support vector machine (SVM) analysis is a popular machine learning tool for classification and regression, developed in 1992 by [Boser et al. \(1992\)](#). SVMs belong to the family of generalised linear classifiers. They are a prediction tool that uses machine learning theory to maximise predictive accuracy while automatically avoiding over-fitting to the data. SVM can be defined as systems which use a hypothesis space of linear functions in a high dimensional feature space, trained with a learning algorithm from optimisation theory that implements a learning bias derived from statistical learning theory. In other words, classes are separated by hyperplanes where the distance to the nearest elements of each class is the largest.

The main advantage of using SVM is that as non-parametric technique, it does not require to assume certain conditions or parameters in the data (e.g. linear combinations or lack of heteroscedasticity in our sample). SVM uses the principle of maximal margin, meaning that we are not so concerned about the prediction as long as the error term (ϵ) is less than a certain value. Put differently, maximal margin allows SVM to be viewed as a convex optimisation problem. Moreover, the regression can also be penalised using a cost parameter (as explained in more detail later), which helps to avoid over-fitting. Nevertheless, a shortcoming of the SVM model is that it

places observations above and below a classifying hyperplane and there is therefore no direct probabilistic interpretation.¹⁸ It should be noted that we do not rely on approximated probabilities, for this reason we calculate our test statistics and plot the predicted outcome.

To keep things simple, we use a linear kernel function to map lower dimensional data into a higher dimensional space: $u'v$, where u' and v are the vectors representing the inputs in the vector space. In addition, we set the cost of constraints violation to 0.5. This is the ' C '-constant of the regularisation term in the Lagrange formulation, or in other words, the extent to which misclassifying is to be avoided in each training example. For large values of C , the optimisation will choose a smaller-margin hyperplane of that hyperplane which does a better job.

Logistic regression

Logistic regression (LR) (see Equation 10) is commonly used to model the probability of a certain event and is a natural benchmark. Contrary to a linear regression, a logistic regression is bounded between 0 and 1. It thus provides a probability score that reflects the probability of the occurrence of an event.

$$P(Y_k|X) = E(Y_k|X) = \frac{e^z}{1 + e^z} = \frac{e^{\alpha+\beta_i X}}{1 + e^{\alpha+\beta_i X}} \quad (10)$$

The nowcasting of GDP turning points with our approach faces two major difficulties. First, a major challenge is to identify the correct signals in a very noisy environment. This step is managed by RFE. Given the shortness of the time-series of macroeconomic data, this approach is often criticised for using information which is not known in the past, e.g. cyclical node positions. However, by applying k-fold cross-validation, the criticism of selecting the best model using information available today but not known at the time is reduced, since the correct variables must also be valid when applied to prior subsets of data. Moreover, in an additional exercise we apply RFE in an out-of-sample consistent fashion. Second, given the availability of highly qualitative signals, it is still difficult to detect events in real-time that propagate slightly differently.

¹⁸Nevertheless, approximations of probabilities of outcomes exist, e.g. Platt (1999).

4 Nowcasting US business cycles

The following subsections present the results (in-sample and out-of-sample) for the binary and the ternary state model for nowcasting turning points in GDP.


4.1 In-sample feature selection and prediction

This subsection selects optimal features taking into account the whole sample including the out-of-sample period and splitting it into k-folds.

4.1.1 Two-state model: detecting negative growth

In this subsection, we nowcast periods of negative growth as specified in Equation 5. We apply a battery of models such as SVM, NB and LR on RFE-selected data. Figure 8 displays the outcome of an out-of-sample nowcasting exercise starting from

Figure 8 Real-time forecast, negative growth, three methods




Note: The graph displays probabilities of negative growth periods as calculated in an expanding window. For predicting t (from the first quarter of 2007 to the first quarter of 2019) we take into account network information until t , while training data ranges from the first quarter of 1982 to period $t-4$. The red bars indicate historical periods of negative growth. Features are selected with RFE until the first quarter of 2019 with five folds.

the first quarter of 2007 based on RFE feature selection over the entire sample until the first quarter of 2019.¹⁹ We make use of RFE to select the most informative set

¹⁹ The features are selected by using a method that is well-known in machine learning: recursive feature elimination, with k-fold cross-validation with $k=5$. The RFE algorithm divides the data into groups and evaluates the performance of the model on subgroups of data, while predicting 1 out of k folds. The k-fold cross-validation ensures that features would have also been regarded as optimal in different sub-periods prior the first quarter of 2007. In addition, it also provides the optimal number of features with specific properties dealing with the trade-off of adding information, while reducing noise, information redundancy,

of features for a given matrix of features X and a response variable Y . The SVM algorithm captures the turning points and other periods very well, while the naive Bayes produces few false positives and the logistic regression does not capture the initial period of the contraction during the financial crisis. Parts of the remaining inaccuracies stem from the fact that the data-set is large, containing 4,100 features, and from the standardisation of features, which make it hard for naive Bayes and logistic regression to process the information correctly.

Figure 9 Properties most informative features, negative growth



Notes: The graph displays the properties of the 16 features with the highest explanatory power for predicting periods with negative growth. Each subplot is to be interpreted independently. They summarise information on the type of measure and to which sectors the features relate to.

The conditional probabilities $P(Y_{kt}|X_{it})$ are ranked by their explanatory power. In the case of negative growth periods, the algorithm favours 16 features (see Figure A.19 in the Appendix). In an ex-post evaluation, we summarise the properties of the selected features for the real-time forecast (nowcast) model. Figure 9 at the top shows the properties of the selected measures. The features used for prediction are mostly selected from among the hubs, in-degree and PageRank metrics. All these measures emphasise the high importance of centrality and they assign a high weight to pivotal companies in the economy, whose returns are linked directly or indirectly to many

and overfitting. Using RFE in each period would add noise and make it harder to capture turning points, while at the same time it is also less computationally efficient.

and influential companies and which hence provide early signals for economy-wide developments. The algorithm selects features from the finance and manufacturing sectors.

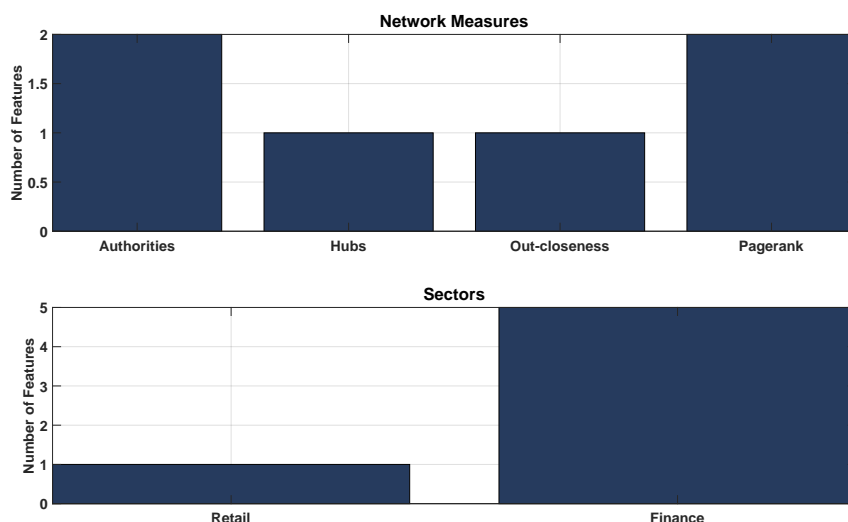
Table A.3 in the Appendix displays the individual properties of the features such as the measure, significance and bandwidth levels, the company and the sector it belongs to. For example, one of the selected features is associated with MetLife, one of the largest global providers of insurance, annuities and employee benefit programmes. This is very intuitive as insurers like MetLife are exposed to economy-wide developments owing to their shareholdings in a wide variety of companies. Similarly, with regard to services, Akamai Technologies is one of the largest content delivery network (CDN), cyber-security and cloud service providers, which are exposed to cyclical demand. On the manufacturing side, we select Dentsply Sirona, one of the world's largest dental equipment makers and dental consumables producers, and Amgen, a multinational bio-pharmaceutical company. Both are part of the health-care sector, which is very sensitive to developments in economic activity. The selection of hubs, PageRank and authorities measures as being more informative for predicting negative growth periods suggests that developments in returns of these companies influence the returns of many other companies that are not necessarily connected to each other.²⁰

4.1.2 Three-state model: high, low and negative growth

In this subsection, we extend the model to a ternary outcome as specified by Equation 6. The RFE algorithm proposes six features to be optimal. Figure 10 reveals the properties of the selected features. There is a strong emphasis on companies related to finance, which includes banks and insurers. With respect to the measures, the algorithm again favours hubs, PageRank and authorities measures.

²⁰ Additional analysis reveals that with varying significance levels cyclical information might be better approximated by different measures of centrality, while the node-position of a company remains unaltered. For example MetLife out-degree is regarded as cyclical by this algorithm at a significance level of 1% with the feature being standardised, while at 5% with the feature being normalised the algorithm detects cyclical information in the hubs measure associated with this company. The latter case is illustrated in Figure 7.

Figure 10 Properties most informative features, high, low and negative growth

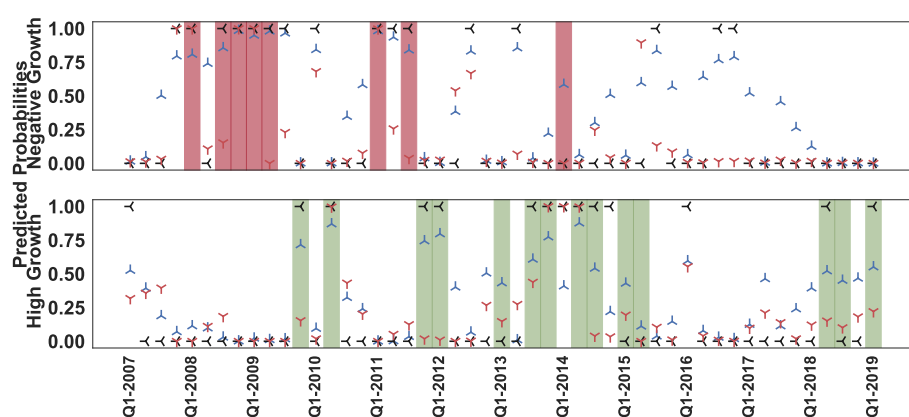


Notes: The graph displays the properties of the six selected features with the highest explanatory power for predicting periods with HLN growth. Each subplot is to be interpreted independently. They summarise information on the type of measure and to which sectors the features relate to.

Figure 11 displays the predicted probabilities for two of the three states. All three algorithms perform well despite the fact that there are a few false positives in the high growth probability case. Furthermore, logistic regressions have a harder time to interpret the data, providing elevated estimates in the post-crisis period. Surprisingly, the naive Bayes model captures the turning points well and produces few false positives.

Table A.5 in the Appendix displays the details of the selected features. As in the negative growth model, the MetLife hubs measure is selected. The Unum Group is another insurer, SVB Financial is a commercial bank that has specialised in lending to high-tech companies and UnitedHealth Group offers health care products and insurance services. UDR is a publicly traded real estate investment trust that invests in apartments, while Home Depot is the largest home improvement retailer in the United States.

Figure 11 High, low and negative growth probability, real-time forecast, three methods



Notes: The graph displays probabilities of high and negative growth periods as calculated in an expanding window. For predicting t (from the first quarter of 2007 to the first quarter of 2019) we take into account network information until t , while the training data range is from the first quarter of 1982 to period $t-4$. The red bars indicate historical periods of negative growth and the green bars indicate high growth (above 3%). For the SVM model, we plot the predicted outcome. Features are selected with RFE until the first quarter of 2019 with three folds.

4.2 Out-of-sample consistent feature selection and prediction


Although k -fold cross-validation requires features to be optimal in subsets of the entire time span - including sub-periods that occur only before the start of the nowcasting exercise - we acknowledge the criticism of potentially working with information that would not have been known at this time. Therefore, in this exercise we only employ features that would have been regarded as optimal based solely on information that was available up the fourth quarter of 2006 and we conduct the out-of-sample nowcasting exercise starting in the first quarter of 2007.

4.2.1 Two-state model: detecting negative growth

Figure 12 displays the outcome of a nowcasting exercise starting from the first quarter of 2007 with features being selected before the fourth quarter of 2006. Apparently, this version of the binary model would not immediately capture the onset of the Great Recession, but would otherwise perform relatively well. Therefore, we can assume that the selected features have been stable in their provision of signals regarding economic growth. As expected, the overall performance of the measures is worse relative to the performance of the complete sample RFE based binary model. Figure 13 summarises the properties of the 32 features selected as illustrated in Table A.4 in the Appendix.

Features from the manufacturing industry (including pharmaceutical companies and durable goods manufacturing companies) and features from the transportation, retail and financial sectors are chosen. The features selected from the manufacturing sector are mostly measures of degree centrality, while the features selected from retail, transportation and financial firms are chosen from the PageRank measure.

Figure 12 Real-time forecast, negative growth, three methods




Note: The graph displays probabilities of negative growth periods as calculated in an expanding window. For predicting t (from the first quarter of 2007 to the first quarter of 2019) we take into account network information until t , while the training data range is from the first quarter of 1982 to period $t-2$. Features are selected with RFE until the fourth quarter of 2006. The red bars indicate historical periods of negative growth. Features are selected with RFE until the fourth quarter of 2006 with three folds.

4.2.2 Three-state model: high, low and negative growth

Figure 14 displays the predicted probabilities for two of the three states at the out-of-sample horizon, employing features that are optimal until the fourth quarter of 2006. We observe a significant loss in efficiency, although the turning point of the financial crisis is still nowcasted. We find many false positives for the negative growth state, however, they correspond to very low (although non-negative) growth. The same holds true for cases of positive growth.

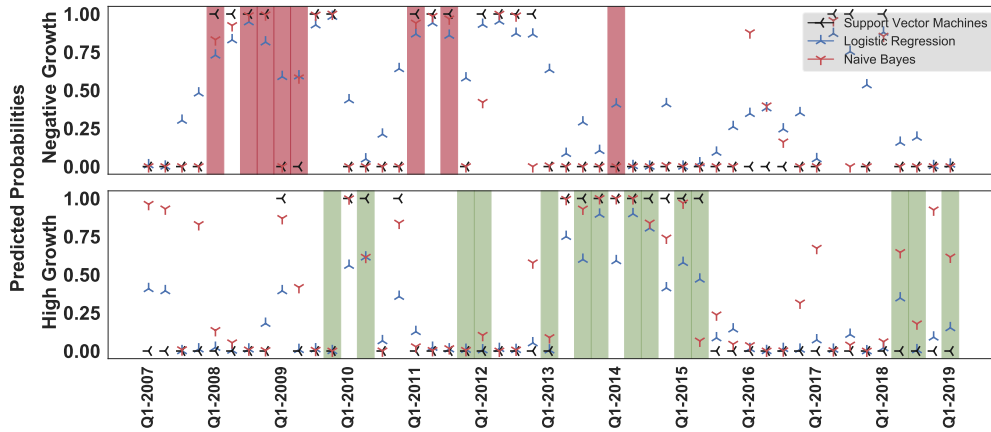
Figure 15 summarises the properties of the 12 features selected as illustrated in Table A.6 in the Appendix. Most features relate to the manufacturing, retail and services sectors, while the network measures are broad-based with the PageRank

Figure 13 Properties most informative features, negative growth



Notes: The graph displays the properties of the 32 features with the highest explanatory power for predicting periods of negative growth. Each subplot is to be interpreted independently. They summarise information on the type of measure and to which sectors the features relate to.


Figure 14 High, low and negative growth probability, real-time forecast, three methods



Notes: The graph displays probabilities of high and negative growth periods as calculated in an expanding window. For predicting t (from the first quarter of 2007 to the first quarter of 2019) we take into account network information until t , while the training data ranges from the first quarter of 1982 to period $t-2$. The red bars indicate historical periods of negative growth and the green bars indicate high growth (above 3%). For the SVM model, we plot the predicted outcome. Features are selected with RFE until the fourth quarter of 2006 with three folds.

measure again being predominant. The home improvement retailer Home Depot is once again among the selected features.

Figure 15 Properties most informative features, high, low and negative growth



Notes: The graph displays the properties of the 12 features with the highest explanatory power for predicting periods with HLN growth. Each subplot is to be interpreted independently. They summarise information on the type of measure and to which sectors the features relate to.

4.3 Algorithm performance comparison

In this subsection, we evaluate the performance of the different algorithms in an out-of-sample nowcast exercise. As correct probabilities for the SVM classification do not exist, we use the predicted outcome which restricts somewhat comparability with the other algorithms.

We apply four different metrics: i) the average precision score that summarises a precision-recall curve as the weighted mean of precision achieved at each threshold (this is done by using the increase in recall from the previous threshold as a weight); ii) the area under the receiver operator curve (AUROC): the higher the area, the better the model is at separating classes; iii) the mean squared error (MSE), which is a risk metric that corresponds to the value of the squared (quadratic) error or loss (it is always non-negative, and values closer to zero are better); and iv) the accuracy score, which is the fraction of correct predictions.

Table 2 illustrates the performance of the different algorithms for the two models where (i) RFE was conducted on the entire sample and (ii) where RFE was run on the data up to the fourth quarter 2006. Overall, case (ii) - the true out-of-sample

Table 2 REAL-TIME FORECAST, MODEL EVALUATION.

Metric/Algorithm	Negative growth		HLN growth	
	in-sample RFE	out-of-sample RFE	in-sample RFE	out-of-sample RFE
Average Precision Score				
SVM	1.00*	0.41*	0.57*	0.44*
NB	0.90	0.33	0.56	0.50
LR	0.98	0.42	0.76	0.55
AUROC				
SVM	1.00*	0.76*	0.77*	0.66*
NB	0.97	0.62	0.65	0.69
LR	0.99	0.76	0.83	0.68
Mean Squared Error				
SVM	0.00	0.26	0.72	0.98
NB	0.09	0.21	1.04	1.43
LR	0.02	0.19	0.91	0.85
Accuracy Score				
SVM	1.00	0.74	0.66	0.53
NB	0.91	0.78	0.60	0.47
LR	0.98	0.81	0.47	0.53

Notes: The average precision score and the AUROC are derived based on predicted probabilities. We decide to keep the predicted outcome for the SVM model and not to provide approximate probabilities. The measures affected are indicated with an asterisk *. The mean squared error and the accuracy score are based on the predicted outcomes, which preserves full comparability between the measures for all algorithms.

exercise - displays a weaker performance. For case (i), in the negative growth case, SVM outperforms the naive Bayes and the logistic regression. As expected the logistic regression behaves very similarly to SVM, but might have some problems with the standardisation. Visual inspection reveals that both algorithms are very efficient, but SVM captures the turning points. In the ternary state model, the SVM algorithm and the logistic regression capture the negative growth periods, in particular the turning points very well. However, they also produce false positives which mostly relate to periods of a weaker growth momentum. In particular, the logistic regression displays a high persistency in probabilities of false positives in the negative growth probability case. On the other hand, the naive Bayes does not produce as many false positives, but also misses some true positives. For case (ii), in the negative growth model the logistic regression performs better in all measures, while SVM is slightly worse relative to the AUROC and the average precision score. In the three-states model, the logistic regression and the naive Bayes model have a much higher average precision score and a slightly higher AUROC than SVM. Overall, the results show that SVMs perform best in in-sample analysis, since this method is prone to overfitting. But in true out-of-sample forecasting, logistic regression and naive Bayes outperform SVM.

5 Concluding remarks

Employing machine learning methods, we exploit the forward-looking behaviour of financial market information to identify the current state of the economy. We construct a dynamic network based on pairwise Granger causalities of stock returns of the constituents of the Standard & Poors 500 index. By calculating the topology of the network we create a “fat” data-set of 5,002 time-series, which we evaluate with machine learning methods. Once cyclical patterns in the network are identified, different classifiers (naive Bayes, support vector machines and logistic regression) are used to predict different economic regimes in real time (e.g. expansions and recessions).

The advantages of the proposed method are the following: (i) the network structure is able to capture the non-linear nature of shock propagation in the economy; (ii) the network contains many relationships in the economy, including some that might not have received attention before; (iii) pooling node-level information provides a mechanism that is more robust to the “every cycle is different” criticism; and (iv) the advantage of using financial data is that they are forward looking, non-proprietary and provided in real-time. Most importantly, they are not subject to larger revisions and reflect market expectations about economic fundamentals as well as the short-term outlook, which makes them ideal for nowcasting and forecasting.

We innovate upon existing approaches by: (i) employing pairwise Granger-causality networks as opposed to mere correlations, thereby taking into account directionality and statistical significance; (ii) allowing for many measures of centrality at the aggregate and node-level, but finally selecting fewer features than in comparable studies, potentially reducing the problem of overfitting; (iii) comparing different algorithms; and (iv) performing an exercise in which we compare the in-sample and out-of-sample stability of features.

We analyse two models. The first tries to identify two economic regimes: “recession” (defined as negative growth in the quarter) and its complement. In the second model, we try to perform the more challenging task of distinguishing between three

regimes: “low growth” (defined as negative growth in the quarter), “low growth” (defined as non-negative, but lower than high growth) and “high growth” (defined as annualised growth of 3% in the quarter). Our findings suggest that our network approach can help to detect changes in economic regimes in real-time giving policy-makers a tool to anticipate slowdowns and take measures against them. The main findings can be summarised as follows: (i) Applying RFE on the entire sample, we are able to determine ex-post which firms were exposed to earnings shocks that turned into aggregate fluctuations. Looking at the sectors where these firms operate, a more complete narrative can be derived than is the usually derived from more aggregate data. The two-state model for predicting periods of negative growth can predict future states remarkably well by using information from the node-positions of manufacturing firms, transportation firms and financial firms, particularly insurance firms. The ternary state model successfully predicts economic regimes by making use of information from the financial sector, the insurance subsector, and the retail sector. (ii) Using features that are selected to be optimal until the fourth quarter of 2006 (i.e. excluding the out-of-sample period), we validate our models in terms of true out-of-sample nowcasting performance and verify that the in-sample results are not solely due to favorable feature selection. While the predictions of the binary state model are weaker, the ternary state model captures the turning point of the great recession and performs well overall, despite producing some false positives for the negative growth state in periods of weak growth.

There are many ways to improve upon and extend the current framework. Stock market returns networks provide timely signals and implicitly include forward-looking information. However, they provide a somewhat narrow view of economic developments. Including features that could better capture the financial cycle²¹ or a different measurement of spillovers, such as transfer entropy or directed volatility spillovers, as proposed by [Diebold and Yilmaz \(2014\)](#), might deliver stronger signals. Furthermore, robustness exercises have shown that different rolling window bandwidths of 3 to 18 months could be promising variants to use. As computational efficiency and feature

²¹ [Lang et al. \(1996\)](#) find a negative relationship between high leverage and future firm growth. [Berger and Udell \(1998\)](#) show that for small firms different capital structures are optimal at different points in the cycle.

engineering is vital in this setup, a clear taxonomy of dimension reduction techniques could further improve the models' performance.

References

- Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. 2012. The network origins of aggregate fluctuations. *Econometrica* Vol. 80, Issue 5: 1977–2016.
- Adam, K., and S. Merkel. 2019. Stock price cycles and business cycles. *ECB Working Paper Series* No 2316, ECB, Frankfurt am Main, September.
- Adrian, T., N. Boyarchenko, and D. Giannone. 2019. Vulnerable growth. *American Economic Review* Vol. 109, Issue 4: 1263–1289.
- Alessi, L., and C. Detken. 2011. Quasi real time early warning indicators for costly asset price boom/bust cycles: A role for global liquidity. *European Journal of Political Economy* Vol. 27, Issue 3: 520–533.
- Berger, A. N., and G. F. Udell. 1998. The economics of small business finance: The roles of private equity and debt markets in the financial growth cycle. *Journal of Banking & Finance* Vol. 22, Issue 6-8: 613–673.
- Bernanke, B. S., M. Gertler, and S. Gilchrist. 1999. The financial accelerator in a quantitative business cycle framework. *Handbook of Macroeconomics* Vol. 1, Part C: 1341–1393.
- Billio, M., M. Getmansky, A. W. Lo, and L. Pelizzon. 2012. Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics* Vol. 104, Issue 3:35–559.
- Boser, B. E., I. M. Guyon, and V. N. Vapnik. 1992. A training algorithm for optimal margin classifiers pp. 144–152.
- Camacho, M., G. Perez-Quiros, and P. Poncela. 2018. Markov-switching dynamic factor models in real time. *International Journal of Forecasting* Vol. 34, Issue 4: 598–611.
- Chauvet, M., and J. Hamilton. 2006. Dating business cycle turning points. *in Milas, C. P. Rothman and D. van Dijk (eds), Nonlinear Time Series Analysis of Business Cycles* 276.

- Davig, T., and A. S. Hall. 2019. Recession forecasting using Bayesian classification. *International Journal of Forecasting* Vol. 35, Issue 3: 848–867.
- Diebold, F. X., and K. Yilmaz. 2014. On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics* Vol. 182, Issue 1: 119–134.
- Drehmann, M., and M. Juselius. 2014. Evaluating early warning indicators of banking crises: Satisfying policy requirements. *International Journal of Forecasting* Vol. 30, Issue 3: 759–780.
- Duca, M. L., and T. A. Peltonen. 2013. Assessing systemic risks and predicting systemic events. *Journal of Banking & Finance* Vol. 37, Issue 7: 2183–2195.
- Ferreira, T. 2018. Stock market cross-sectional skewness and business cycle fluctuations. *FRB International Finance Discussion Papers* No. 1223.
- Gabaix, X. 2011. The granular origins of aggregate fluctuations. *Econometrica* Vol. 79, Issue 3: 733–772.
- Giannone, D., M. Lenza, and G. E. Primiceri. 2018. Economic predictions with big data: The illusion of sparsity. Tech. Rep. 847.
- Giusto, A., and J. Piger. 2017. Identifying business cycle turning points in real time with vector quantization. *International Journal of Forecasting* Vol. 33, Issue 1: 174–184.
- Gourinchas, P.-O., and M. Obstfeld. 2012. Stories of the twentieth century for the twenty-first. *American Economic Journal: Macroeconomics* Vol. 4, Issue 1: 226–65.
- Granger, C. W. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* Vol. 37, Issue 3: 424–438.
- Granger, C. W. 1979. Forecasting in business and economics. *Elsevier Monographs* .
- Guyon, I., J. Weston, S. Barnhill, and V. Vapnik. 2002. Gene selection for cancer classification using support vector machines. *Machine Learning* Vol. 46: 389–422.


- Hagenau, M., M. Liebmann, and D. Neumann. 2013. Automated news reading: Stock price prediction based on financial news using context-capturing features. *Decision Support Systems* Vol. 55, Issue 3: 685–697.
- Heiberger, R. H. 2018. Predicting economic growth with stock networks. *Physica A: Statistical Mechanics and its Applications* Vol. 489: 102–111.
- Jordan, M. I., and T. M. Mitchell. 2015. Machine learning: trends, perspectives, and prospects. *Science* Vol. 349, Issue 6245: 255–260.
- Landherr, A., B. Friedl, and J. Heidemann. 2010. A critical review of centrality measures in social networks. *Business & Information Systems Engineering* Vol. 2, Issue 6: 371–385.
- Lang, L., E. Ofek, and R. Stulz. 1996. Leverage, investment, and firm growth. *Journal of Financial Economics* Vol. 40, Issue 1: 3–29.
- Li, J. 2017. Credit market frictions and the linkage between micro and macro uncertainty.
- Næs, R., J. A. Skjeltorp, and B. A. Ødegaard. 2011. Stock market liquidity and the business cycle. *The Journal of Finance* Vol. 66, Issue 1:139–176.
- Ng, S. 2014. Viewpoint: Boosting recessions. *Canadian Journal of Economics* Vol. 47, Issue 1: 1–34.
- Plagborg-Møller, M., L. Reichlin, G. Ricco, and T. Hasenzagl. 2020. When is Growth at Risk? *Brookings Papers on Economic Activity* Spring.
- Platt, J. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. in *Smola, A., P. Bartlett, B. Schölkopf and D. Schuurmans (eds.) Advances in Large Margin Classifiers*, Cambridge.
- Qi, M. 2001. Predicting US recessions with leading indicators via neural network model. *International Journal of Forecasting* Vol. 17, Issue 3: 383–401.
- Rish, I., et al. 2001. An empirical study of the naive Bayes classifier. *IJCAI 2001 Workshop on Empirical Methods in Artificial Intelligence* Vol. 3, Issue 22: 41–46.

- Rose, A. K., and M. M. Spiegel. 2012. Cross-country causes and consequences of the 2008 crisis: Early warning. *Japan and the World Economy* Vol. 24, Issue 1: 1–16.
- Yan, K., and D. Zhang. 2015. Feature selection and analysis on correlated gas sensor data with recursive feature elimination. *Sensors and Actuators B: Chemical* Vol. 212: 353–363.
- Zingales, L. 2015. Presidential address: Does finance benefit society? *The Journal of Finance* Vol. 70, Issue 4: 1327–1363.
- Zou, C., and J. Feng. 2009. Granger causality vs. dynamic Bayesian network inference: a comparative study. *BMC Bioinformatics* Vol. 10, Issue 1: 122.

A Additional results


A.1 Data

Figure A.16 Data coverage



Notes: The graph displays the amount of NaN series per total amount of companies.

Figure A.17 Normalised stable node-relations during phases of the business cycle



Notes: The graph illustrates normalised quarterly connectedness from 1982:Q1 to 2019:Q1 on a 12 month rolling window bandwidth and a 5% significance level. Connectedness in specific nodes in NBER recession periods are more persistent than in non-recession periods, where connectedness might be more idiosyncratic.

A.2 Negative growth periods

A.2.1 In-sample RFE

Figure A.18 Optimal number of features (RFE)




Table A.3 Selected features United States, negative growth

Company	Sector	Measure	Significance	Window
ABBOTT LABORATORIES	Manufacturing	Betweenness	1%	12m
EXELON	Transportation	Degree Centrality	1%	12m
METLIFE	Finance, Insurance and Real Estate	Out-degree	1%	12m
BECTON DICKINSON	Manufacturing	In-degree	1%	12m
AKAMAI TECHS.	Services	In-degree	1%	12m
DENTSPLY SIRONA	Manufacturing	In-degree	1%	12m
METLIFE	Finance, Insurance and Real Estate	Hubs	1%	12m
HUMANA	Finance, Insurance and Real Estate	Hubs	1%	12m
STATE STREET	Finance, Insurance and Real Estate	Hubs	1%	12m
DOLLAR TREE	Retail	Hubs	1%	12m
KIMCO REALTY	Finance, Insurance and Real Estate	Hubs	1%	12m
EXELON	Transportation	Pagerank	1%	12m
WELLTOWER	Finance, Insurance and Real Estate	Pagerank	1%	12m
AKAMAI TECHS.	Services	Pagerank	1%	12m
UDR	Finance, Insurance and Real Estate	Out-closeness	1%	12m
LOEWS	Finance, Insurance and Real Estate	Authorities	1%	12m

A.2.2 Out-of-sample consistent RFE

Figure A.19 Optimal number of features (RFE)




Table A.4 Selected features, negative growth

Company	Sector	Measure	Significance	Window
PFIZER	Manufacturing	Betweenness	1%	12m
ABBOTT LABORATORIES	Manufacturing	Betweenness	1%	12m
HP	Manufacturing	Degree Centrality	1%	12m
DOLLAR TREE	Retail	Degree Centrality	1%	12m
PEPSICO	Manufacturing	Out-degree	1%	12m
NIKE B	Manufacturing	Out-degree	1%	12m
STATE STREET	Finance, Insurance and Real Estate	Out-degree	1%	12m
DOLLAR TREE	Retail	Out-degree	1%	12m
NEWMONT GOLDCORP	Primary & Construction	Out-degree	1%	12m
PROCTER & GAMBLE	Manufacturing	In-degree	1%	12m
CVS HEALTH	Retail	In-degree	1%	12m
CERNER	Services	In-degree	1%	12m
DENTSPLY SIRONA	Manufacturing	In-degree	1%	12m
NISOURCE	Transportation	In-degree	1%	12m
STATE STREET	Finance, Insurance and Real Estate	Hubs	1%	12m
DOLLAR TREE	Retail	Hubs	1%	12m
CVS HEALTH	Retail	Pagerank	1%	12m
TJX	Retail	Pagerank	1%	12m
DUKE ENERGY	Transportation	Pagerank	1%	12m
NORTHROP GRUMMAN	Manufacturing	Pagerank	1%	12m
COGNIZANT TECH.SLTN.A	Services	Pagerank	1%	12m
EXELON	Transportation	Pagerank	1%	12m
SYSCO	Wholesale	Pagerank	1%	12m
ONEOK	Transportation	Pagerank	1%	12m
EQUITY RESD.TST.PROPS. SHBI	Finance, Insurance and Real Estate	Pagerank	1%	12m
WELLTOWER	Finance, Insurance and Real Estate	Pagerank	1%	12m
EVERSOURCE ENERGY	Transportation	Pagerank	1%	12m
HASBRO	Manufacturing	Pagerank	1%	12m
EXPEDITOR INTL.OF WASH.	Transportation	Pagerank	1%	12m
TORCHMARK	Finance, Insurance and Real Estate	Pagerank	1%	12m
H&R BLOCK	Services	Pagerank	1%	12m
MCCORMICK & COMPANY NV.	Manufacturing	Authorities	1%	12m

A.3 High, low and negative growth periods

A.3.1 In-sample RFE

Figure A.20 Optimal number of features (RFE)




Table A.5 Features, high, low and negative growth periods

Company	Sector	Measure	Significance	Window
METLIFE	Finance, Insurance and Real Estate	Hubs	1%	12m
HOME DEPOT	Retail	Pagerank	1%	12m
UNUM GROUP	Finance, Insurance and Real Estate	Pagerank	1%	12m
UDR	Finance, Insurance and Real Estate	Out-closeness	1%	12m
UNITEDHEALTH GROUP	Finance, Insurance and Real Estate	Authorities	1%	12m
SVB FINANCIAL GROUP	Finance, Insurance and Real Estate	Authorities	1%	12m

A.3.2 Out-of-sample consistent RFE

Figure A.21 Optimal number of features (RFE)




Table A.6 Features, high, low and negative growth periods

Company	Sector	Measure	Significance	Window
CINTAS	Services	Betweenness	1%	12m
METTLER TOLEDO INTL.	Manufacturing	Betweenness	1%	12m
JEFFERIES FINANCIAL GROUP	Manufacturing	Degree Centrality	1%	12m
ECOLAB	Manufacturing	Out-degree	1%	12m
HOME DEPOT	Retail	Pagerank	1%	12m
STANLEY BLACK & DECKER	Manufacturing	Pagerank	1%	12m
HESS	Manufacturing	Pagerank	1%	12m
PINNACLE WEST CAP.	Transportation	Pagerank	1%	12m
ACTIVISION BLIZZARD	Services	Out-closeness	1%	12m
EQUIFAX	Services	In-closeness	1%	12m
TEXAS INSTRUMENTS	Manufacturing	Authorities	1%	12m
KROGER	Retail	Authorities	1%	12m

A.3.3 Included S&P 500 companies

MNEMONIC	NAME	MNEMONIC	NAME
@AAPL	APPLE	U:MS	MORGAN STANLEY
@AMZN	AMAZON.COM	U:SLB	SCHLUMBERGER
@MSFT	MICROSOFT	U:UPS	UNITED PARCEL SER.'B'
@FB	FACEBOOK CLASS A	U:CVS	CVS HEALTH
U:JPM	JP MORGAN CHASE & CO.	U:NEE	NEXTERA ENERGY
U:JNJ	JOHNSON & JOHNSON	U:BLK	BLACKROCK
U:XOM	EXXON MOBIL	@SBUX	STARBUCKS
@GOOGL	ALPHABET A	@CHTR	CHARTER COMMS.CL.A
U:BAC	BANK OF AMERICA	U:DHR	DANAHER
U:WMT	WALMART	@KHC	KRAFT HEINZ
U:WFC	WELLS FARGO & CO	@WBA	WALGREENS BOOTS ALLIANCE
U:V	VISA 'A'	U:ANTM	ANTHEM
U:PFE	PFIZER	U:BDX	BECTON DICKINSON
U:UNH	UNITEDHEALTH GROUP	@BIIB	BIOGEN
U:T	AT&T	U:SCHW	CHARLES SCHWAB
U:HD	HOME DEPOT	U:EOG	EOG RES.
U:CVX	CHEVRON	@ADP	AUTOMATIC DATA PROC.
U:MA	MASTERCARD	U:TJX	TJX
@CSCO	CISCO SYSTEMS	U:AET	AETNA
U:VZ	VERIZON COMMUNICATIONS	U:AGN	ALLERGAN
@INTC	INTEL	U:AMT	AMERICAN TOWER
U:BA	BOEING	U:FDX	FEDEX
U:PG	PROCTER & GAMBLE	@MDLZ	MONDELEZ INTERNATIONAL CL.A
U:KO	COCA COLA	U:PNC	PNC FINL.SVS.GP.
U:C	CITIGROUP	U:CB	CHUBB
U:ORCL	ORACLE	U:SYK	STRYKER
U:MRK	MERCK & COMPANY	@ATVI	ACTIVISION BLIZZARD
@CMCSA	COMCAST A	@CELG	CELGENE
U:DIS	WALT DISNEY	@CSX	CSX
@PEP	PEPSICO	@ISRG	INTUITIVE SURGICAL
@NFLX	NETFLIX	U:OXY	OCCIDENTAL PTL.
@NVDA	NVIDIA	U:CL	COLGATE-PALM.
U:DWDP	DOWDUPONT	U:GD	GENERAL DYNAMICS
U:IBM	INTERNATIONAL BUS.MCHS.	@CME	CME GROUP
U:ABBV	ABBVIE	U:RTN	RAYTHEON 'B'
@AMGN	AMGEN	U:DUK	DUKE ENERGY
@ADBE	ADOBE (NAS)	U:SPGI	S&P GLOBAL
U:MDT	MEDTRONIC	U:SPG	SIMON PROPERTY GROUP
U:MMM	3M	@INTU	INTUIT
U:MCD	MCDONALDS	@MU	MICRON TECHNOLOGY
U:PM	PHILIP MORRIS INTL.	U:BK	BANK OF NEW YORK MELLON
U:HON	HONEYWELL INTL.	U:BSX	BOSTON SCIENTIFIC
U:ABT	ABBOTT LABORATORIES	@ESRX	EXPRESS SCRIPTS HOLDING
U:UNP	UNION PACIFIC	U:GM	GENERAL MOTORS
U:ACN	ACCENTURE CLASS A	@ILMN	ILLUMINA
U:MO	ALTRIA GROUP	U:NOC	NORTHROP GRUMMAN
U:CRM	SALESFORCE.COM	U:PSX	PHILLIPS 67
U:UTX	UNITED TECHNOLOGIES	U:CI	CIGNA
U:LLY	ELI LILLY	U:EMR	EMERSON ELECTRIC
U:NKE	NIKE 'B'	U:NSC	NORFOLK SOUTHERN
U:GE	GENERAL ELECTRIC	U:AIG	AMERICAN INTL.GP.
@QCOM	QUALCOMM	U:DE	DEERE
@PYPL	PAYPAL HOLDINGS	U:ITW	ILLINOIS TOOL WORKS
@TXN	TEXAS INSTRUMENTS	U:MET	METLIFE
U:BMY	BRISTOL MYERS SQUIBB	U:PH	PARKER-HANNIFIN
@AVGO	BROADCOM	U:PRU	PRUDENTIAL FINL.
@COST	COSTCO WHOLESALE	@FOXA	FOX A
@GILD	GILEAD SCIENCES	U:VLO	VALERO ENERGY
U:TMO	THERMO FISHER SCIENTIFIC	U:COF	CAPITAL ONE FINL.
U:AXP	AMERICAN EXPRESS	U:CCI	CROWN CASTLE INTL.
U:LMT	LOCKHEED MARTIN	@MAR	MARRIOTT INTL.'A'
U:CAT	CATERPILLAR	U:PX	PRAXAIR
U:GS	GOLDMAN SACHS GP.	@CTSH	COGNIZANT TECH.SLTN.'A'
U:LOW	LOWE'S COMPANIES	U:D	DOMINION ENERGY
@BKNG	BOOKING HOLDINGS	U:ECL	ECOLAB
U:USB	US BANCORP	U:HCA	HCA HEALTHCARE
U:COP	CONOCOPHILLIPS	U:HUM	HUMANA

MNEMONIC	NAME	MNEMONIC	NAME
U:SO	SOUTHERN	U:DG	DOLLAR GENERAL
U:TGT	TARGET	U:GIS	GENERAL MILLS
@VRTX	VERTEX PHARMS.	U:MCK	MCKESSON
U:ICE	INTERCONTINENTAL EX.	@ORLY	O REILLY AUTOMOTIVE
U:MMC	MARSH & MCLENNAN	U:OKE	ONEOK
U:PLD	PROLOGIS	U:PPG	PPG INDUSTRIES
U:SHW	SHERWIN-WILLIAMS	U:RCL	ROYAL CARIBBEAN CRUISES
U:ZTS	ZOETIS	U:YUM	YUM! BRANDS
U:BAX	BAXTER INTL.	@ALXN	ALEXION PHARMS.
U:EXC	EXELON	U:AVB	AVALONBAY COMMNS.
U:F	FORD MOTOR	U:DXC	DXC TECHNOLOGY
U:HPQ	HP	U:EQR	EQUITY RESD.TST.PROPS. SHBI
@REGN	REGENERON PHARMS.	U:HPE	HEWLETT PACKARD ENTER.
@AMAT	APPLIED MATS.	U:HLT	HILTON WORLDWIDE HDG.
U:BBT	BB&T	U:IQV	IQVIA HOLDINGS
U:DAL	DELTA AIR LINES	U:KR	KROGER
U:KMB	KIMBERLY-CLARK	@LRCX	LAM RESEARCH
U:KMI	KINDER MORGAN	@PAYX	PAYCHEX
U:LYB	LYONDELLBASELL INDS.CL.A	U:PEG	PUB.SER.ENTER.GP.
U:PGR	PROGRESSIVE OHIO	@TROW	T ROWE PRICE GROUP
U:AON	AON CLASS A	U:ZBH	ZIMMER BIOMET HDG.
U:ETN	EATON	U:APTIV	APTIV
U:IR	INGERSOLL-RAND	U:BBY	BEST BUY
U:SYU	SYSCO	U:CTL	CENTURYLINK
U:WM	WASTE MANAGEMENT	U:ED	CONSOLIDATED EDISON
U:AFL	AFLAC	U:CMJ	CUMMINS
U:APD	AIR PRDS.& CHEMS.	U:DLR	DIGITAL REALTY TST.
@ADI	ANALOG DEVICES	U:FCX	FREEMPORT-MCMORAN
U:CCL	CARNIVAL	U:K	KELLOGG
U:STZ	CONSTELLATION BRANDS 'A'	U:MTB	M&T BANK
@EA	ELECTRONIC ARTS	@NTRS	NORTHERN TRUST
@EQIX	EQUINIX REIT	@PCAR	PACCAR
U:FIS	FIDELITY NAT.INFO.SVS.	U:PCG	PG&E
U:HAL	HALLIBURTON	U:RHT	RED HAT
U:MPC	MARATHON PETROLEUM	U:RSG	REPUBLIC SVS.'A'
@ROST	ROSS STORES	U:ROK	ROCKWELL AUTOMATION
U:LUV	SOUTHWEST AIRLINES	U:SWK	STANLEY BLACK & DECKER
U:VFC	V F	U:SYF	SYNCHRONY FINANCIAL
U:ALL	ALLSTATE	U:TWTR	TWITTER
U:AEP	AMER.ELEC.PWR.	@UAL	UNITED CONTINENTAL HOLDINGS
@ADSK	AUTODESK	U:WELL	WELLTOWER
@EBAY	EBAY	U:WY	WEYERHAEUSER
@FISV	FISERV	@XEL	XCEL ENERGY
U:PSA	PUBLIC STORAGE	U:A	AGILENT TECHS.
U:STT	STATE STREET	@CTAS	CINTAS
U:TRV	TRAVELERS COS.	U:EIX	EDISON INTL.
U:APC	ANADARKO PETROLEUM	U:IP	INTERNATIONAL PAPER
U:EL	ESTEE LAUDER COS.'A'	U:MSCI	MSCI
U:JCI	JOHNSON CONTROLS INTL.	@NDAQ	NASDAQ
@MNST	MONSTER BEVERAGE	@NTAP	NETAPP
U:MCO	MOODY'S	U:COL	ROCKWELL COLLINS
U:SRE	SEMPRA EN.	@AAL	AMERICAN AIRLINES GROUP
U:STI	SUNTRUST BANKS	U:AMP	AMERIPRISE FINL.
U:TEL	TE CONNECTIVITY	U:AZO	AUTOZONE
U:WMB	WILLIAMS	@CERN	CERNER
@AMD	ADVANCED MICRO DEVICES	U:DVN	DEVON ENERGY
U:EW	EDWARDS LIFESCIENCES	U:ES	EVERSOURCE ENERGY
U:ROP	ROPER TECHNOLOGIES	U:HES	HESS
@ALGN	ALIGN TECHNOLOGY	U:HRL	HORMEL FOODS
U:ADM	ARCHER DANIELS MIDLAND	@IDXX	IDEXX LABORATORIES
U:CNC	CENTENE	@INFO	IHS MARKIT
U:GLW	CORNING	U:KEY	KEYCORP
U:FTV	FORTIVE	U:MRO	MARATHON OIL
U:PXD	PIONEER NTRL.RES.	U:MSI	MOTOROLA SOLUTIONS
U:APH	AMPHENOL 'A'	@MYL	MYLAN
U:CXO	CONCHO RESOURCES	U:NUE	NUCOR
U:DFS	DISCOVER FINANCIAL SVS.	U:PPL	PPL

MNEMONIC	NAME	MNEMONIC	NAME
U:RF	REGIONS FINL.NEW	U:NOV	NATIONAL OILWELL VARCO
U:WEC	WEC ENERGY GROUP	U:OMC	OMNICOM GROUP
U:ABC	AMERISOURCEBERGEN	U:O	REALTY INCOME
U:AME	AMETEK	U:RMD	RESMED
U:ANET	ARISTA NETWORKS	@SWKS	SKYWORKS SOLUTIONS
U:BCP	BOSTON PROPERTIES	U:FTI	TECHNIPFC
U:CBS	CBS 'B'	@ULTA	ULTA BEAUTY
U:CFG	CITIZENS FINANCIAL GROUP	U:VMC	VULCAN MATERIALS
U:CLX	CLOROX	U:WAT	WATERS
@DLTR	DOLLAR TREE	U:WCG	WELLCARE HEALTH PLANS
U:DTE	DTE ENERGY	U:ADS	ALLIANCE DATA SYSTEMS
@FITB	FIFTH THIRD BANCORP	U:AJG	ARTHUR J GALLAGHER
U:FLT	FLEETCOR TECHNOLOGIES	U:BHGE	BAKER HUGHES A
U:GPN	GLOBAL PAYMENTS	U:KMX	CARMAX
U:HRS	HARRIS	U:CF	CF INDUSTRIES HDG.
@MCHP	MICROCHIP TECH.	@CHRW	CH ROBINSON WWD.
U:NEM	NEWMONT GOLDCORP	U:CHD	CHURCH & DWIGHT CO.
U:TDG	TRANSIDGM GROUP	@CTXS	CITRIX SYS.
U:VTR	VENTAS	U:CMS	CMS ENERGY
@VRSN	VERISIGN	U:CAG	CONAGRA BRANDS
@VRSK	VERISK ANALYTICS CL.A	U:ETFC	E TRADE FINANCIAL
@WDC	WESTERN DIGITAL	U:EMN	EASTMAN CHEMICAL
@WLTW	WILLIS TOWERS WATSON	U:ETR	ENTERGY
U:GWW	WW GRAINGER	@GRMN	GARMIN
@XLNX	XILINX	U:IT	GARTNER 'A'
@ABMD	ABIOMED	U:GPC	GENUINE PARTS
U:APA	APACHE	@HAS	HASBRO
@CA	CA	@INCY	INCYTE
U:CAH	CARDINAL HEALTH	U:LEN	LENNAR 'A'
@EXPE	EXPEDIA GROUP	U:MHK	MOHAWK INDUSTRIES
@FAST	FASTENAL	U:NBL	NOBLE ENERGY
U:FE	FIRSTENERGY	U:DGX	QUEST DIAGNOSTICS
U:BEN	FRANKLIN RESOURCES	U:RJF	RAYMOND JAMES FINL.
U:HIG	HARTFORD FINL.SVS.GP.	@STX	SEAGATE TECH.
@HBAN	HUNTINGTON BCSH.	U:AOS	SMITH (AO)
U:LH	LABORATORY CORP.OF AM. HDG.	@SNPS	SYNOPSIS
U:L	LOEWS	@TTWO	TAKE TWO INTACT.SFTW.
@PFG	PRINCIPAL FINL.GP.	U:TPR	TAPESTRY
@SBAC	SBA COMMS.	U:URI	UNITED RENTALS
@SIVB	SVB FINANCIAL GROUP	U:VNO	VORNADO REALTY TRUST
U:TXT	TEXTRON	U:WRK	WESTROCK
U:TIF	TIFFANY & CO	@WYNN	WYNN RESORTS
U:TSS	TOTAL SYSTEM SERVICES	U:XYL	XYLEM
U:TSN	TYSON FOODS 'A'	U:AAP	ADV.AUTO PARTS
U:ARE	ALEXANDRIA RLST.EQTIES.	@AKAM	AKAMAI TECHS.
U:AEE	AMEREN	U:ALB	ALBEMARLE
U:AWK	AMERICAN WATER WORKS	@CDNS	CADENCE DESIGN SYS.
@ANSS	ANSYS	U:CPB	CAMPBELL SOUP
U:BALL	BALL	U:CBOE	CBOE GLOBAL MARKETS
U:BR	BROADRIDGE FINL.SLTN.	U:CNP	CENTERPOINT EN.
U:BF.B	BROWN-FORMAN 'B'	U:CMG	CHIPOTLE MEXN.GRILL
U:CBRE	CBRE GROUP CLASS A	@CINF	CINCINNATI FINL.
U:CMA	COMERICA	U:COO	COOPER COS.
U:DHI	D R HORTON	@CPRT	COPART
U:DRI	DARDEN RESTAURANTS	U:COTY	COTY CLA
U:EFX	EQUIFAX	U:DVA	DAVITA
U:ESS	ESSEX PROPERTY TST.	U:DOV	DOVER
U:EVERG	EVERGY	U:EQT	EQT
U:FRT	FEDERAL REALTY INV.TST.	@EXPD	EXPEDITOR INTL.OF WASH.
U:HSY	HERSHEY	@FFIV	F5 NETWORKS
U:HST	HOST HOTELS & RESORTS	U:FMC	FMC
@KLAC	KLA TENCOR	U:HCP	HCP
U:LLL	L3 TECHNOLOGIES	@HSIC	HENRY SCHEIN
U:LNC	LINCOLN NATIONAL	U:HFC	HOLLYFRONTIER
U:MCK	MCCORMICK & COMPANY NV.	@JBHT	HUNT JB TRANSPORT SVS.
U:MTD	METTLER TOLEDO INTL.	U:IFF	INTL.FLAVORS & FRAG.
U:MGM	MGM RESORTS INTL.	U:SJM	J M SMUCKER

MNEMONIC	NAME	MNEMONIC	NAME
U:KSU	KANSAS CITY SOUTHERN	@FLIR	FLIR SYSTEMS
U:KSS	KOHL'S	U:FLR	FLUOR
U:MLM	MARTIN MRTA.MATS.	U:IPG	INTERPUBLIC GROUP
U:MAS	MASCO	@IPGP	IPG PHOTONICS
U:MAA	MID-AMER.APT COMMUNITIES	U:NI	NISOURCE
U:TAP	MOLSON COORS BREWING 'B'	U:PNW	PINNACLE WEST CAP.
U:MOS	MOSAIC	U:RL	RALPH LAUREN CL.A
U:NCLH	NORWEGIAN CRUISE LINE HDG.	U:RHI	ROBERT HALF INTL.
@SYMC	SYMANTEC	U:SLG	SL GREEN REALTY
U:ARNC	ARCONIC	U:UNM	UNUM GROUP
U:AVY	AVERY DENNISON	U:WU	WESTERN UNION
U:COG	CABOT OIL & GAS 'A'	U:WHR	WHIRLPOOL
U:DRE	DUKE REALTY	U:AMG	AFFILIATED MANAGERS
U:EXR	EXTRA SPACE STRG.	U:AIV	APARTMENT INV.& MAN.'A'
U:GPS	GAP	U:AIZ	ASSURANT
U:HII	HNTGTN.INGALLS INDS.	U:FLS	FLOWSERVE
@HOLX	HOLOGIC	U:FBHS	FORTUNE BNS.HM.& SCTY.
U:IVZ	INVESCO	@GT	GOODYEAR TIRE & RUB.
U:IRM	IRON MOUNTAIN	U:HBI	HANESBRANDS
U:JEC	JACOBS ENGR.	U:HOG	HARLEY-DAVIDSON
U:JNPR	JUNIPER NETWORKS	U:HP	HELMERICH & PAYNE
U:LB	L BRANDS	U:JEF	JEFFERIES FINANCIAL GROUP
@LKQ	LKQ	U:KIM	KIMCO REALTY
U:M	MACY'S	U:MAC	MACERICH
@NKTR	NEKTAR THERAPEUTICS	U:PNR	PENTAIR
U:NLSN	NIELSEN	U:PHM	PULTEGROUP
U:JWN	NORDSTROM	U:SEE	SEALED AIR
U:NRG	NRG ENERGY	U:XRX	XEROX
U:PKG	PACKAGING CORP.OF AM.	@BHF	BRIGHTHOUSE FINANCIAL
U:PKI	PERKINELMER	@DISCA	DISCOVERY SERIES A
U:PRGO	PERRIGO	U:EVHC	ENVISION HEALTHCARE
U:PVH	PVH	U:FL	FOOT LOCKER
@QRVO	QORVO	U:HRB	H&R BLOCK
U:SNA	SNAP-ON	U:LEG	LEGGETT&PLATT
U:TMK	TORCHMARK	@MAT	MATEL
@TSCO	TRACTOR SUPPLY	U:NFX	NEWFIELD EXPLORATION
U:UDR	UDR	@NWSA	NEWS 'A'
U:UHS	UNIVERSAL HEALTH SVS.'B'	@PBCT	PEOPLES UNITED FINANCIAL
U:VAR	VARIAN MEDICAL SYSTEMS	U:PWR	QUANTA SERVICES
@VIAB	VIACOM 'B'	U:SCG	SCANA
@ZION	ZIONS BANCORP.	@SRCL	STERICYCLE
U:AES	AES	@TRIP	TRIPADVISOR 'A'
U:ALK	ALASKA AIR GROUP	U:UAA	UNDER ARMOUR A
U:ALLE	ALLEGION	@GOOG	ALPHABET 'C'
U:BWA	BORGWARNER	U:BRK.B	BERKSHIRE HATHAWAY 'B'
U:XEC	CIMAREX EN.	@DISCK	DISCOVERY SERIES C
@XRAY	DENTSPLY SIRONA	@NWS	NEWS 'B'
@DISH	DISH NETWORK 'A'	@FOX	FOX B
U:RE	EVEREST RE GP.	U:UA	UNDER ARMOUR 'C'

Acknowledgements

This study profited from discussions with Gabe de Bondt, Roberto de Santis, Lorenzo Frattarolo, Mirco Rubin and several participants of the European Commission's Joint Research Centre workshop on Big Data and Economic Forecasting in Ispra in 2019, the Modelling with Big Data and Machine Learning conference hosted by the Bank of England and King's College in London in 2019 and the Big Data/Advanced analytics internal workshop at the European Central Bank in 2019. All remaining errors are our own. The views expressed in this paper are solely those of the authors and do not necessarily represent the views of the ECB or of the European Commission.

Andres Azqueta-Gavaldon

European Central Bank, Frankfurt am Main, Germany; University of Glasgow; email: andres.azqueta_gavaldon@ecb.europa.eu

Dominik Hirschbühl

European Central Bank, Frankfurt am Main, Germany; email: dominik.hirschbuhl@ecb.europa.eu

Luca Onorante

European Commission – Joint Research Centre, Ispra, Italy; email: luca.onorante@ec.europa.eu

Lorena Saiz

European Central Bank, Frankfurt am Main, Germany; email: lorena.saiz@ecb.europa.eu

© European Central Bank, 2020

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-4411-3

ISSN 1725-2806

doi:10.2866/23967

QB-AR-20-146-EN-N