Stability and Similarity in Financial Networks - How do They Change in Times of Turbulence?

Tristan Millington*, Mahesan Niranjan

University of Southampton, University Road, Southampton, SO17 1BJ

Abstract

Diversified portfolios are a key component of modern portfolio theory, based on the idea of choosing uncorrelated or unrelated stocks to minimize risk. With this in mind, we use networks to study the correlations between stocks and how this varies over time, using daily returns from the S&P500 (US), FTSE100 (UK) and DAX30 (Germany). We study both the full correlation networks and those filtered using the PMFG method. We conclude that stocks tend to become more similar in the full correlation networks during times of market disruption for the US and UK markets - implying that nodes that were once dissimilar (and therefore a good choice for a low risk portfolio) are no longer so, demonstrating the difficulties of choosing a diversified portfolio. Furthermore, these full networks are also more stable by certain measures during these periods of disruption, contrary to expectations. However, these apply less to the PMFGs and the German market.

Keywords: networks, finance, correlation, portfolio selection

1. Introduction

The dynamics of asset prices and their interplay are regarded as a very complex system [1]. By using assets as nodes and relationships between them as edges, we can represent a system of assets as a network [2]. Networks are a popular approach due to their ability to represent such complex systems in an abstract manner, which can then be analyzed using a common toolkit [3].

We study these financial networks in order to understand interconnectedness in the financial markets. In order to reduce risk, many investors wish to own stocks that are unrelated, or even better, negatively correlated. This referred to as a diversified portfolio, and ensures that even if a certain portion of their portfolio declines, the overall value remains the same. This interconnectedness is a double edged sword however, as high levels of interconnectedness can create high amounts of systematic risk, and cause a failure in one company to propagate to the rest of the market [4].

*Corresponding author

Email addresses: t.millington@soton.ac.uk (Tristan Millington), m.niranjan@soton.ac.uk (Mahesan Niranjan)

These networks are also dynamic, and change over time. This means that companies that were considered unrelated during a period of market calm may suddenly become related during a time of market stress [5], meaning that we must study how these networks vary in order to accurately quantify the amount of interconnectedness present.

A wide variety of financial networks exist [6]. We focus on correlation networks between assets as they are easy to construct and interpret. This is a popular method following the seminal work of Onnela [7], Boginski [8] and Mantegna [2]. Once inferred these networks have many potential applications, including portfolio optimization [9] [10] [11] and understanding sector dominance [12].

Our contributions in this paper are as follows. Firstly, we study the similarity between nodes in financial networks. Other authors have studied the similarities between financial networks, but to the best of our knowledge this has not been done for node similarity. Secondly, we study how a variety of similarity and stability measures vary for full and filtered correlation networks, and in particular, look at how these are correlated with the volatility of the index. Some of these measures have been studied in isolation, but to the best of our knowledge, these have not been brought together, linked to the volatility and an interpretation drawn.

The similarity of nodes in the network is important due to the concept of creating diversified portfolios. If the nodes in a correlation network are very similar, this implies that it is not possible to construct a diversified portfolio from said nodes. Furthermore if the similarity varies over time, assets that were once dissimilar can become similar, changing from a good choice for a diversified portfolio to a poor choice.

Node centrality also has relevance for portfolio construction. Pozzi et al. [10] show that companies with a low centrality have a higher Sharpe ratio than those with a high centrality. Therefore if the centralities are not stable, this also makes it challenging for portfolio construction.

The stability of the networks can also be related to portfolio construction. If the networks are unstable, this makes it challenging to predict the future performance of assets, and by extension, portfolios constructed from said assets. Therefore understanding the stability of the networks is also important for exploring the possibility of constructing portfolios.

2. Related Work

In this section, we provide a brief review into the analysis of how financial networks vary during times of market calm and stress. Correlations in the market tend to increase during times of market stress, with correlations between negative returns being stronger than those between positive returns [13]. This causes previously unrelated companies to become related [5]. Furthermore, this increase in correlation is not geographically isolated, meaning that diversification across countries might not be successful [14] [15].

These factors affect the networks. In the partial correlation based ‘dependency network’ [16] the clustering coefficient also increases during these times [17]. Some authors have found that the markets tend to have more edge changes during these times [18], however, others have found that the cluster structure can become more stable [19] [20]. This has also been noted when studying the entropy of the spectrum of the correlation matrix [21], with the entropy dropping (i.e. information increasing) during crashes.

Crises impact both individual markets and the global system. On this theme, Silva et al [22] study the interconnectedness of the global financial system using networks created from the exposure of the banks of one country to those of another. Of particular interest to this paper, they show how the 2008 financial crisis changes the structure of these networks, indicating how investors and lenders have changed their decision making. A notable conclusion is that banks have avoided taking long term risky positions in global markets after the financial crisis, in contrast to before.

With regards to correlation networks, Sandoval et al. [15] explore correlations between stock indices to see how countries are connected, and by extension, if geographic diversification can produce robustness to downturns. They study the effects of financial market crises on the correlations between international stock indices, showing that correlations also increase between them during times of market stress. This does mean that it is challenging to diversify a portfolio, even if an investor looks to multiple countries. Correlations between indices also decay far more slowly than those between assets [14].

Heiberger [23] studies correlation networks with a set threshold during times of stress and calm. Using the May-Wigner theory of network stability from ecology, they claim the markets become less stable during times of market stress, and that the modularity tends to drop too.

As well as studying the full correlation networks, authors have also studied how market crashes affect filtered networks. For instance, Onnela et al. [24] find that the length of the minimum spanning tree (MST) decreases during market disruption, and others have noted it has a different structure during times of crashes [25] [26], with Wilinski et al. [27] showing that MSTs transition from a scale free MST with a hierarchy of local hubs to being orientated around a single super hub.

Li and Pi [28] study the effects of the financial crisis on minimum spanning trees and threshold networks constructed from index returns from a selection of countries. The dataset is divided into three sections, pre-crisis, during crisis, and post-crisis. The MSTs show geographic clusters both before, during and after the crisis, but the threshold networks constructed have less modularity structure during the crisis than before or after.

As well as global indices, some exist for specific sectors. Zhang et al. [25] analyze correlations between index returns that represent various US economic sectors, giving a 'coarse-grained' view of the US economy as a whole. They also hope to understand the differences in MST structure between times of growth and times of recession. These index returns are clustered to separate them into times of market stress and market calm in a data driven way. They conclude that the topology of the MSTs have a star-like structure in times of market calm, with the industrial sector at the center, but during times of crisis, they become chain-like.

Other authors have analyzed the structure of planar maximally filtered graphs (PMFGs) with similar goals. Musmeci et al. [29] study the persistence of structure in correlation matrices from US and UK financial returns. They quantify this using two measures, meta-correlation (the correlation between the entries of adjacent correlation matrices) and edge persistence (the fraction of edges that are shared between adjacent PMFGs). The meta-correlation shows there are blocks of time with a similar correlation structure on either side of the financial crisis and the edge persistence drops rapidly during the financial crisis, indicating the financial crisis causes a significant change in the structure of the PMFGs. Furthermore, they develop a measure q, which is the ratio of the amount of volatility present in the current window vs the next, and look at predicting this from the edge persistence and the meta-correlation. They show that this can be predicted, with an accuracy usually above 50%, and in some cases up to 83%.

Zhao et al. [30] perform a detailed study of PMFG structure during two financial crises, the dot-com bubble (1999 - 2002) and the subprime crisis (2007 - 2009) using returns from the S&P500. Firstly they investigate the consequences of choosing different window sizes for the structure of the networks and show this choice has a significant effect. Next, they take a novel approach to understanding the changes by using the entire dataset up to the particular day in question, rather than using a window of constant size. With this method they find the average shortest path length decreasing, the average clustering coefficient increasing and the heterogeneity of the network (as defined by [31]) decreasing (indicating the companies become more similar). They also find that these PMFGs constructed from the full dataset show a greater degree of sector clustering compared to those constructed using windows, even when the windows are large. Using InfoMap [32] to perform community detection, they show the modularity increases during these crisis periods, and that the number of intersector edges drops during both crises and stays low once it has dropped, indicating a permanent increase in sector correlation. Finally, they study the edit distance between adjacent PMFGs in time. This spikes before a crisis, indicating a large change in structure.

Song et al. [33] study the returns from 57 market indices from countries across the world, using a variety of time windows, to show how different correlation dynamics are present at different periods. Shorter windows allow easier detection of crises, with a 3 month window allowing the detection of the Asian 1997 crisis, Russian 1998 crisis, and the 2008 global crisis, while longer windows cause these to be smeared out. The mean edge weight spikes during crises in these PMFGs as the correlation in the market increases. Mutual information is then used to measure edge changes between PMFGs adjacent in time. During the months in which shocks occur, the structure of the PMFG changes significantly, with drops then large spikes in the mutual information. Comparing the average correlation in the PMFG to that in an MST using Welch's t-test, they find the p-value spikes during times of stress, indicating that the two are detecting similar stress.

On the theme of clustering, Musmeci et al. [20] study the effectiveness of a variety of clustering methods in picking up sector structure in correlation networks. Furthermore, they study how this success varies during different market conditions. The method with the most success in picking up the sector structure is the directed bubble hierarchical tree [34]. The most relevant finding to this paper is that the methods have less success in uncovering sector structure during the financial crisis.

If the reader desires more information on the subject of correlation-based financial networks, we direct them to the review provided by Marti et al [35].

3. Methods

3.1. Network Construction

To construct this network, we firstly calculate the log return of the stock prices

$$r_i(t) = \log x_i(t) - \log x_i(t-1) \tag{1}$$

where $x_i(t)$ is the price of asset i and time t in our dataset containing p stocks with n samples. Next we calculate the Pearson correlation coefficient between assets i and j as

$$c_{ij} = \frac{\sum_{i=1}^{n}(r_i(t) - \bar{r_i})(r_j(t) - \bar{r_j})}{\sqrt{\sum_{i=1}^{n}(r_i(t) - \bar{r_i})^2(r_j(t) - \bar{r_j})^2}} \tag{2}$$

where $\bar{r_i}$ is the mean of r_i.

This correlation matrix is then used as the adjacency matrix for the correlation network, with the assets becoming nodes and the correlation between two assets being the weight on the edge between them.

While this method does not account for outliers or non-linearity, it is simple to calculate and understand. Other authors have proposed to use different methods, such as mutual information [36], partial correlation [16] [37] [18], rank correlation [38] [11] or estimators designed for fat tails [39] to gain a different set of relationships, but these lose the previously mentioned attractive properties of Pearson correlation.

3.2. Network Stability

To measure the stability of the correlation networks we use the L_2 difference between the eigenvectors that correspond to the largest eigenvalue (the leading eigenvectors) from networks adjacent in time. These are also normalized so the entries add to one. This is defined as:

$$||\frac{\mathbf{v_t}}{\sum_{i=1}^{p} v_{t,i}} - \frac{\mathbf{v_{t-1}}}{\sum_{i=1}^{p} v_{t-1,i}}||_2 \tag{3}$$

where where p is the number of companies in the network, $\mathbf{v_t}$ is the leading eigenvector at time t and $||\mathbf{x}||_2$ is the L_2 norm of vector $\mathbf{x}$.

The entry that corresponds to a node in this eigenvector can be interpreted as the centrality of said node in the network. This difference therefore measures how much the centrality of each node changes from one network to the next - so effectively we measure if corresponding networks agree on which nodes are regarded as important.

3.3. Network Similarity

Many methods have been proposed to measure node similarity in networks. Broadly these fall into three categories, structural equivalence, regular equivalence or automorphic equivalence [3].

Structural equivalence is perhaps the easiest to understand and calculate, we simply measure how many neighbors two nodes have in common [40]. This can be quantified by measuring the similarity between the appropriate rows of the adjacency matrix with a vector distance measure e.g. Pearson correlation. To quantify the amount of structural equivalence in our networks, we use two measures. The first is the mean L_2 distance between the rows of the correlation matrix, defined as:

$$\frac{1}{p(p-1)} \sum_{i} \sum_{j \neq i} ||\mathbf{C_i} - \mathbf{C_j}||_2 \tag{4}$$

where $\mathbf{C_i}$ corresponds to column i of the correlation matrix.

The second is the mean cosine distance, defined as:

$$\frac{1}{p(p-1)} \sum_{i} \sum_{j \neq i} 1 - \frac{\sum_{k=1}^{p} C_{ik} C_{jk}}{||\mathbf{C_i}||_2 ||\mathbf{C_j}||_2} \tag{5}$$

We choose two measures due to the varying amount of correlation present at different times. The cosine distance is normalized while the L_2 distance is not. This allows us to see if this normalization affects the results. In both of these measures we ignore the diagonal, as the value of all entries is 1.

We may also be interested in methods that measure nodes that hold similar positions in networks, despite not sharing any neighbours. This can be measured by *Regular equivalence*, defined by [41]. Here, nodes are similar if they are connected to other nodes that are themselves similar. To measure this we use a form of Katz similarity, defined as

$$(D - \alpha A)^{-1} D \tag{6}$$

where D is the degree matrix (a matrix with the degree of each node on the diagonal), A the adjacency matrix of the network, and α a 'dampening factor'. The Katz similarity allows paths of all lengths to contribute to the similarity. α controls how longer paths contribute, a small α (i.e. close to 0) means longer paths will contribute little, while an α close to 1 ensures all paths will be given the same weight. This particular form of Katz similarity is normalized so that nodes of a higher degree are not biased to be more similar. We choose α to be 0.05.

Automorphic equivalence is formally defined as "two nodes of a graph are automorphically equivalent if all the nodes can be re-labeled to form an isomorphic graph with the labels of u and v interchanged" [42]. Calculating this can be computationally expensive, as it requires computing graph automorphisms. Furthermore, it is difficult to define for correlation networks as it does not depend on the weight of the edge. Due to these issues, we do not use automorphic equivalence in this paper.

If the reader is interested in more detail, we direct them to the following reviews [43] [44], both of which contain sections reviewing node similarity measures.

An alternative approach to measuring the similarity in the network is to look at the distribution of node centrality. This again assume that nodes that share a similar structure are more similar. Generally, these are based on the degree distribution. The first piece of work on this topic is by Bell [45], who proposed to use the variance of node degrees as a measure. Intuitively, if this is large then the nodes have a large difference in degree, and therefore would be expected to be dissimilar. If it is small then the nodes all have similar degrees, and therefore are more similar. We use this approach to measure the similarity in the full correlation networks. As previously mentioned, the entries in the leading eigenvector correspond to the eigenvector centrality of the nodes, and so we measure the standard deviation in the entries of the normalized leading eigenvector.

We also use a formulation proposed by Estrada et al [31]. This defines heterogeneity as follows:

$$H = \frac{\sum_{i,j \in E} (k_i^{-1/2} - k_j^{-1/2})}{p - 2\sqrt{p-1}} \tag{7}$$

where k_i is the degree of node i and E is the edge set of the graph. If the network is completely uniform, with all nodes having the same degree, the heterogeneity will be 0. On the other hand if there is a great amount of variation in node degree (for instance a star graph), the heterogeneity will be close to 1. This cannot be used on the full correlation networks as it assumes the networks are unweighted, and so we must turn to a filtration method to process the correlation networks before they can be analyzed with this measure.

3.4. Filtered Network Construction

Correlation networks by default are complete and weighted. Furthermore, these weights can be negative. Both of these can cause problems when using some similarity measures, which tend to either implicitly or explicitly assume that the networks are sparse and only contain edges with positive weight. This is specifically the case for heterogeneity and Katz similarity. Various filtration methods have been proposed to prune the networks and only keep a subset of representative edges. These include thresholding [8], topological methods (e.g. Minimum Spanning Tree [2], Planar Maximally Filtered Graph [46] and the Triangular Maximally Filtered Graph [47]) or random matrix theory [48] [49] [50]. For our analysis we choose the Planar Maximally Filtered Graph (PMFG) [46] to filter the full networks, and then apply the similarity measures to the filtered networks. We also take the absolute values of the correlations as input into the PMFG procedure as the similarity method can fail to converge when negative edges are present.

Briefly we describe how the PMFG method works. Construction is done in a greedy manner as follows:

- Initialize an empty graph G with the same nodes as the correlation matrix
- Sort edges of the correlation matrix in descending order in a list
- For each item in the sorted list
    - If the edge does not increase the genus of G, add it to the graph

This results in a planar graph with $3(p - 2)$ edges. A planar graph is one that can be drawn on a 2d plane with no edges crossing.

3.5. Community Detection

We can also study the similarities and differences between nodes in a network using community detection. To do this we maximize the modularity of the network using the Louvain algorithm [51]. The modularity for a network with adjacency matrix A and a vector of community assignments $\mathbf{c}$ is [52]

$$Q = \frac{1}{m} \sum_i \sum_j (A_{ij} - \frac{k_i k_j}{m}) \delta(c_i, c_j) \tag{8}$$

where $m = \sum_i \sum_j A_{ij}$ and $\delta(c_i, c_j)$ is the Kronecker delta function, equaling 1 when $c_i = c_j$ (i.e. nodes i and j are in the same community and 0 otherwise).

The modularity is maximized in a greedy manner. Firstly each node is assigned to its own community. Next we iterate through the nodes, calculating the loss in modularity of removing it from said community, and the gain for adding it to a different community, as follows:

$$\delta Q = \frac{\sum_{\text{in}} + 2k_i}{2m} - (\frac{\sum_{\text{tot}} + k_i}{2m})^2 - (\frac{\sum_{\text{in}}}{2m} - (\frac{\sum_{\text{tot}}}{2m})^2 - \frac{k_i}{2m}^2) \tag{9}$$

where $\sum_{\text{in}}$ is the sum of weights of all the edges inside the community node i is being moved into, $\sum_{\text{tot}}$ is the sum of weights of the edges to the community. The node is assigned to the community that gives the largest gain in modularity. This process is repeated until the overall gain in modularity is no longer positive.

This formulation of modularity does not permit negative edges, and so we simply add one to every edge in the network. While there are more elegant methods (e.g. [53] [54]), we find this method gives comparable results with a simpler implementation.

4. Software and Data

The data we use is downloaded from Yahoo Finance. For the UK data we use the FTSE100 companies, for the US returns we use the S&P500 companies, and for the German data we use the DAX30 companies. We use returns from 2000/03/01 to 2019/10/21. For each dataset, any company missing more than 10% of its data is removed, and any missing values are filled forwards from the first good value. If the values are missing from the start we backfill from the first good value. We use a window of 252 trading days (1 trading year) for Germany and the UK, and 504 (2 trading years) for the US due to the larger sample size and slide it along 30 days at a time. This results in 5065 days of return data for 70 companies for the UK, 5068 days for 23 companies for Germany, and 4790 days of return data for 229 companies for the US. Each company is tagged with a sector from the GICS classification using information from Bloomberg. This places each company into 1 of 11 sectors, Information Technology, Real Estate, Materials, Telecommunication Services, Energy, Financials, Utilities, Industrials, Consumer Discretionary, Healthcare or Consumer Staples.

We make use of Python, NumPy and SciPy [55] for general scripting, pandas [56] for handing the data, statsmodels [57] for some of the statistical analysis, matplotlib [58] for

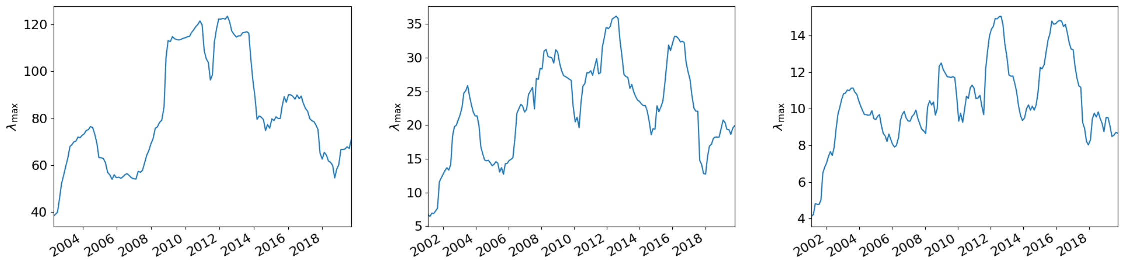

(a) US (b) UK (c) DE

Figure 1: Largest eigenvalue of the correlation matrices inferred from the market returns. 2009 and 2012 are noticeable peaks in the US markets probably due to the financial crisis from 2007 onwards, while the dot com crash of 2003 is also visible. The other markets seem to be noisier, although the UK has peaks in 2003, 2008, 2012 and 2016 - the latter of which could be related to the Brexit vote. Germany has peaks during 2003, 2012 and 2016 too, although these are smaller (note scale on axis) but this could be due to the smaller market

plotting, TopCorr (`https://github.com/shazzzm/topcorr`) for construction of the PMFGs, Cytoscape [59] for visualization of the networks and Networkx [60] for the network analysis.

5. Results

5.1. Periods of Disruption

Firstly we look at the largest eigenvalue of the correlation matrices to demonstrate when the networks detect a period of market disruption. The largest eigenvalue measures the intensity of correlation in that network and should be larger during periods of market disruption. This is shown in Figure 1.

The US market has a far larger eigenvalue than either the German or UK markets, but it is a far larger matrix. It also has more obvious peaks in intensity - notably in 2003, 2009, and 2012. These relate to the dot com crash, financial crash, and European sovereign debt crisis respectively [61]. The UK has similar peaks in 2003, 2008, 2012 and 2016 which contains periods of disruption during the financial crisis and the Brexit vote of 2016. Germany is noisier, but with noticeable peaks in 2003, 2012, and 2016. The two largest peaks are in 2012 and 2016, which could relate to the Eurozone sovereign debt crisis and the Brexit vote, both of which have deep ties into the German economy [62].

5.2. Stability

Next, we look at the difference between the leading eigenvectors of the networks. The results of this are shown in Figure 2. For the US, we see there is an obvious dip from 2009 to 2011, showing that surprisingly the system seems more stable during this time of disruption. If we look at values from times of growth and stability (e.g. 2006) the difference between adjacent windows is around 0.45, indicating most of the leading eigenvector changes during these times. The UK has a slightly different story, with the start of the dataset showing a

(a) US (b) UK (c) DE

Figure 2: L_2 difference between the eigenvector centrality of adjacent networks of the US market. Here our goal is to measure the stability of the underlying structure of the correlation matrices over time. Particularly for the US it is clear than the differences drop during 2009 and 2012, which are both times of market stress. With the UK we can see that there are minimums in 2009 and in 2016, which again are times of market stress for the UK, while Germany has low points during 2008, 2012 and 2016.

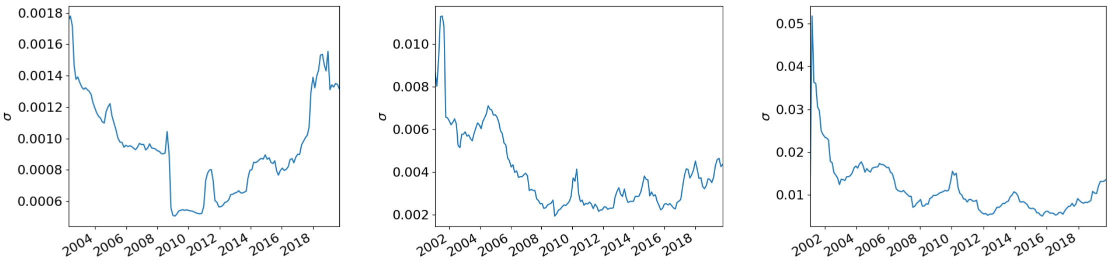

(a) US (b) UK (c) DE

Figure 3: Standard deviation of the normalized leading eigenvector for the networks over time. This is a measure of the similarity in the networks - if the standard deviation is low then the nodes all have a more similar centrality and therefore should be more similar. We again note the drops in 2009 and 2012 for the US market, and in 2012 and 2016 for the UK markets.

large change, following by the difference decreasing until 2008. From this year it is much more stable. There are small peaks in 2004, 2010, 2015 and 2017. It is interesting to note how the stability increases after 2004 until 2008, leveling off at around 0.3. Germany is again distinctive. At the start of the dataset there are large changes in the entries of the leading eigenvector, but it becomes far more stable as time goes on, with minimums in 2008, 2012, and from 2015 to 2017.

5.3. Full Network Similarity

In this section, we study the similarity present in the full correlation networks. The first measure we use is the standard deviation of the entries of the normalized leading eigenvector - which as previously mentioned relates to the eigenvector centrality of a node in the network. If this value is large then there is a big spread in the importance of the nodes and therefore they could be dissimilar. If it is small then they all have similar centralities and therefore could be similar. A plot of this over time is shown in Figure 3.

Starting with the US, we note that the standard deviation is very high in 2002 and 2018.

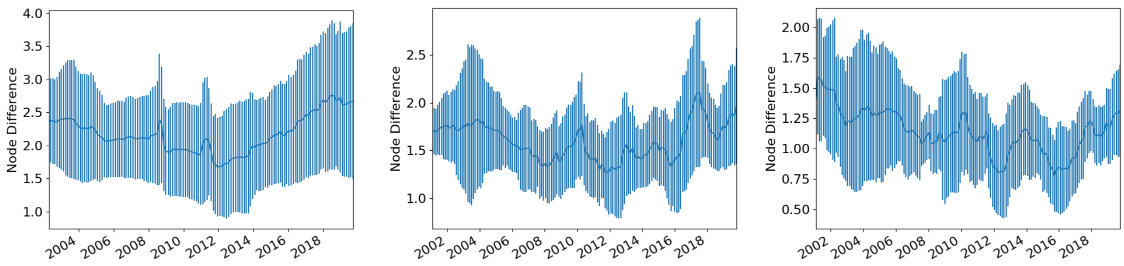

(a) US (b) UK (c) DE

Figure 4: Mean L_2 difference between the nodes in the network, defined in equation 4. Error bars are the standard deviation. For the US markets the financial crisis is particularly noteworthy with there being a significant drop in difference between 2009 and 2010, but peaks in 2008 and 2011. With the UK there is a peak of difference during 2017 when there was a strong bull market. Germany has minimums during 2008, 2012, and 2016.

Again we see this drop from 2009 - 2010, during the financial crisis. For the UK, we see drops during 2009, 2012, and 2016, which as previously mentioned correspond to large negative macroeconomic effects. Germany has a maximum in 2002 like the others, and minimums during 2008, 2012, and 2016, but the movement seems to be more subtle.

Next, we look at how the structural similarity of the networks varies over time. To start with we look at the L_2 distance between nodes. A plot of this over time is shown in Figure 4. For the US market, we can see the difference drops from 2008 to 2009 and from 2010 to 2012, which is consistent with our other networks measures of the perception of a crisis - implying the nodes do become more similar during times of disruption. With the UK market, we can see a small drop during 2012 and 2016, which as mentioned are times of difficulty. Interestingly, there is a large rise during 2017. This time was a period of growth for the FTSE100 due to the weakened pound making UK company currency holdings worth more. This does line up with our theory that bull markets correspond to times of greater difference between nodes, but since this is a general large macroeconomic effect it is worth mentioning that the gains do not seem to be as evenly spread out. Finally, we look at the results for the German market. Again the differences here are smaller than for the UK and US markets, but we do notice a drop during 2012 and 2016.

We then move to our other measure of structural similarity, the cosine distance. Plots of how this varies over time are shown in Figure 5. The effects here are even starker than for the L_2 distance, with both the mean distance and the standard deviation dropping significantly during the financial crisis. This larger drop is due to the effects of normalization. Since correlations are larger during times of stress, normalization reduces the relative distance between them, even if the absolute difference is larger.

5.4. PMFG Analysis

In this section, we study the properties of the PMFGs constructed from the stock returns of the three countries. Firstly we show example networks from the first window of data for all three countries. These are shown in Figure 6.

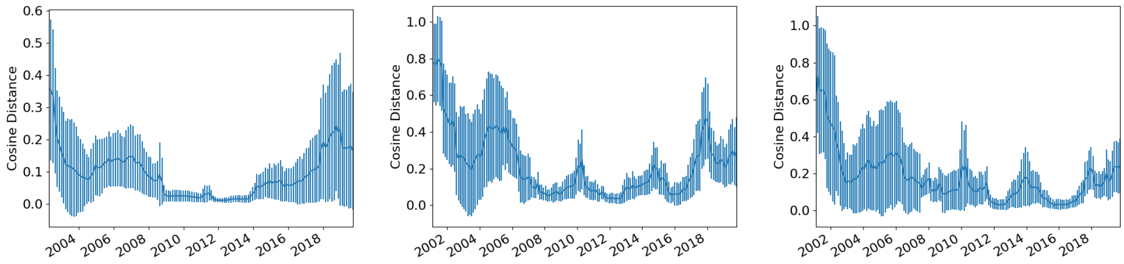

(a) US (b) UK (c) DE

Figure 5: Mean cosine difference between the nodes in the network. Error bars are the standard deviation. We see an even greater change in cosine distance between nodes during the financial crisis than L_2 distance, particularly for the US.

Next, we look at the degree heterogeneity. This is shown in Figure 7. From this we first note that the PMFGs are not particularly heterogeneous, indicating a relatively uniform degree distribution in general. This could be due to the construction procedure, which does enforce a certain structure on the network. For the US the heterogeneity is stable, being around 0.14 for most of the dataset. There is an increase from 2012 - 2014 and minimums in 2003 and 2019. This is not particularly easy to connect to any market crashes. For the UK there is significantly more variation. There are peaks in 2003, 2007, 2011, and 2016 and a noticeable drop from 2009 - 2011, which could be connected to the financial crisis. Germany is noisier still, with a great deal of variation in the heterogeneity. The increase in variation as we move from the US to the UK to Germany could be due to the decreasing size of the markets, making smaller changes more visible.

With this in mind, we next study Katz similarity in the PMFGs. A plot of its mean value over time is shown in Figure 8. This shows a very similar trend to the heterogeneity for all of the countries, indicating the two are picking up similar trends. Again this makes it difficult to connect to specific times for the US and Germany, but the UK can be linked to market states.

In terms of which nodes are regarded as similar to which, we found that if a node had a high similarity to another, they tended to be in the same sector, and in general the mean similarity was higher between nodes in the same sector. However, the full sector structure is not immediately visible. To quantify this we look at the point biserial correlation between the similarity between two nodes and whether they are in the same sector. Across the full dataset, this is 0.257 for the US, 0.245 for the UK, and 0.108 for Germany, indicating that the sector structure is not particularly strong. In fact, if we do the same for the full correlation network, we get 0.417 for the US, 0.393 for the UK, and 0.484 for Germany. Therefore the sector structure is stronger in the correlation matrices than the similarity matrices.

Finally, we look at how the stability of the PMFGs varies over time. To measure this we use the number of edge changes between networks adjacent in time. The results of this are shown in Figure 9. All three graphs show a sharp drop in edge changes in 2009, and for the US this low is maintained until 2011. The US also shows a drop in 2003, but for the other two countries there is no discernible pattern.

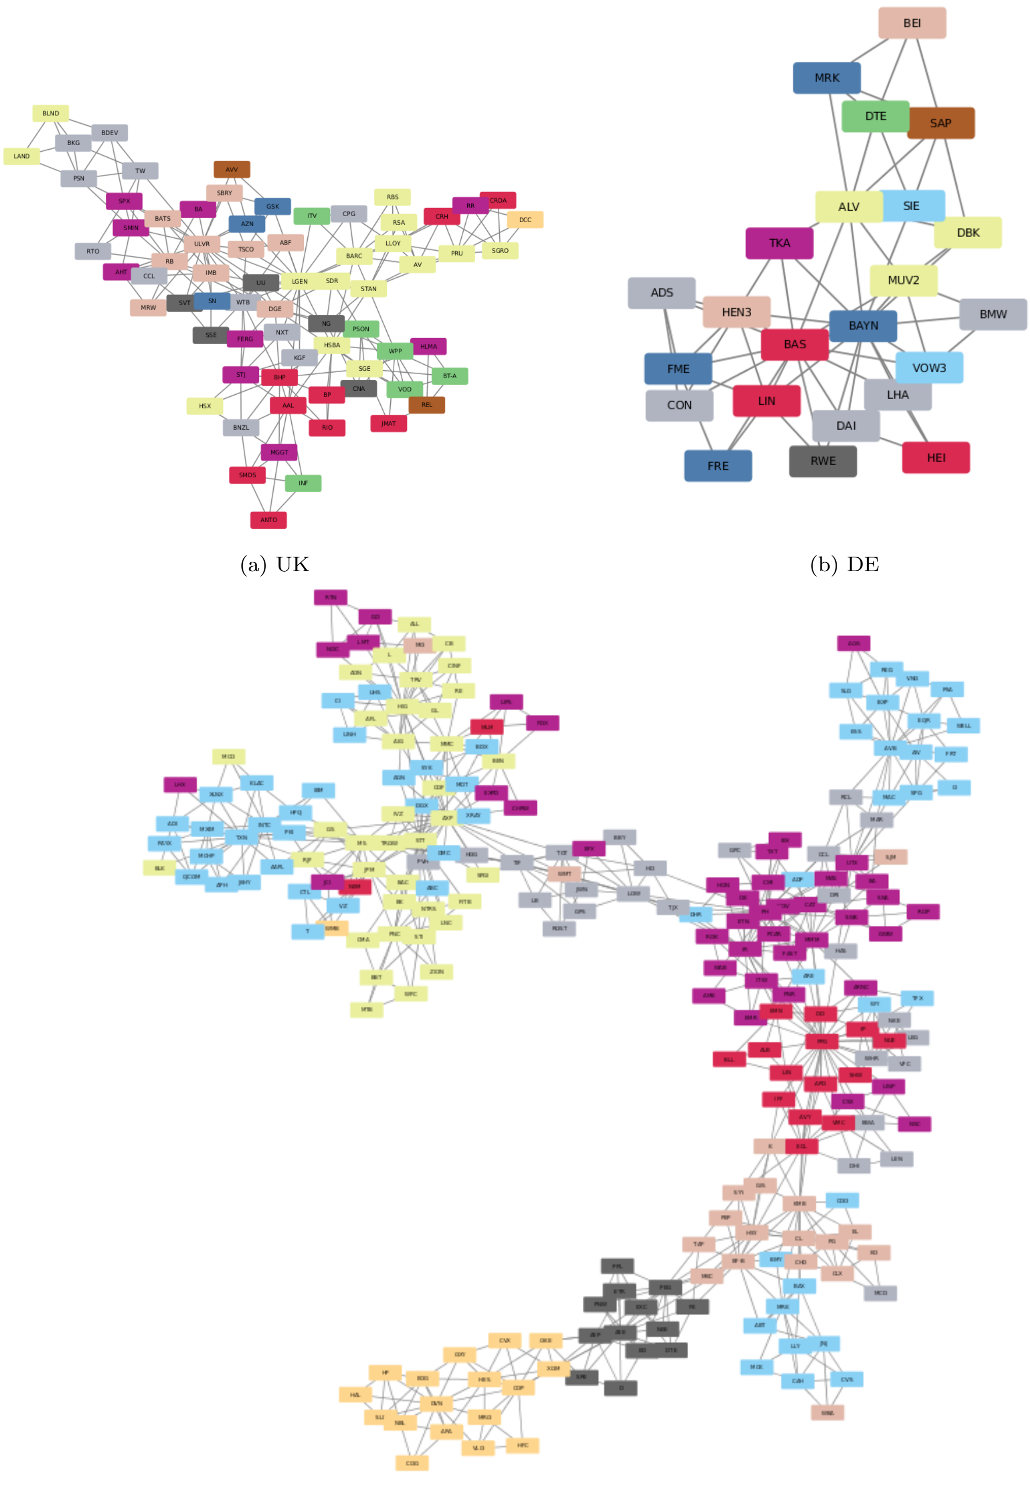

(a) UK

(b) DE

(c) US

Figure 6: Example PMFGs constructed from the first window of stock returns for all three countries. Nodes are colored according to sector membership.

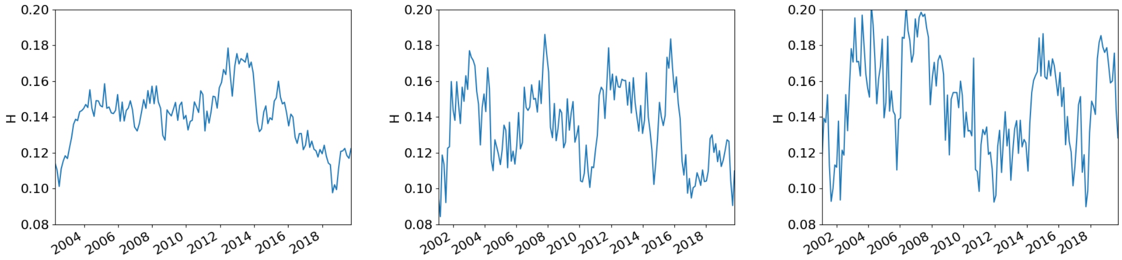

(a) US (b) UK (c) DE

Figure 7: Degree Heterogeneity of the PMFGs over time (this is defined in equation 7). For the US this is relatively stable in time, making it difficult to connect to a market state. The UK and Germany show much more variation.

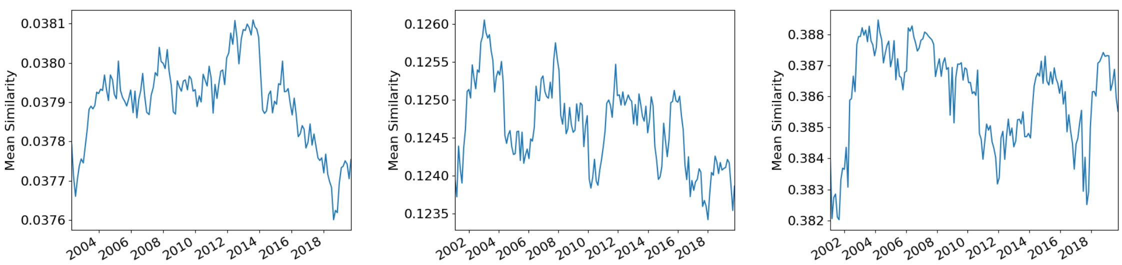

(a) US (b) UK (c) DE

Figure 8: Mean Katz Similarity in the PMFGs over time. This shows a very similar trend to the heterogeneity (see Figure 7), indicating the two are picking up on similar trends

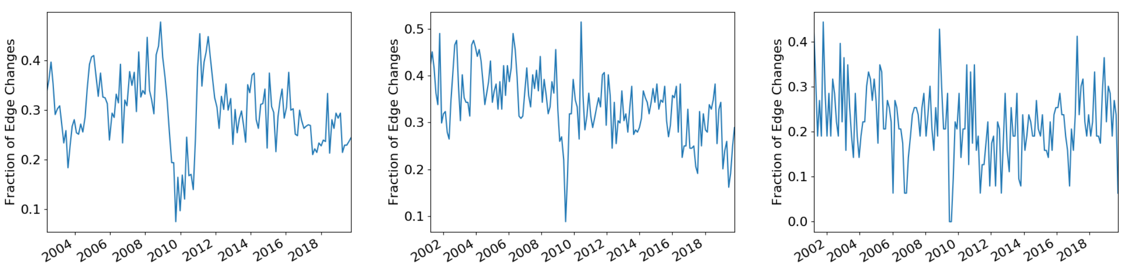

(a) US (b) UK (c) DE

Figure 9: Number of edge changes between PMFGs adjacent in time. All three show a sharp drop during 2009, indicating an increase in stability during this time, however beyond this there does not seem to be any other patterns.

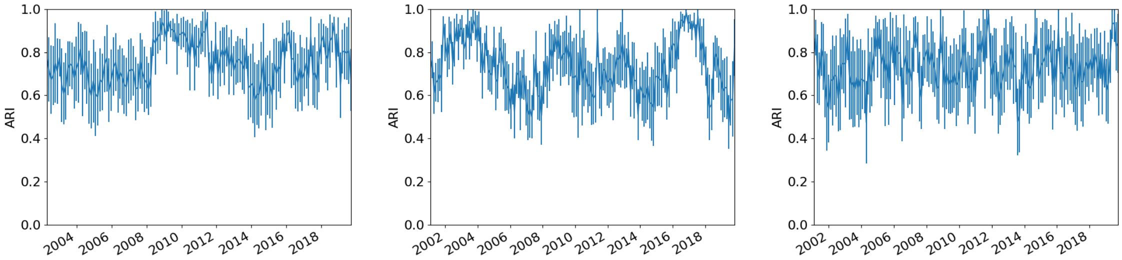

(a) US (b) UK (c) DE

Figure 10: Mean Adjusted Rand Index from the 10 runs on the same network. Since the algorithm is greedy a different result will be reached each time, but if there is structure to be discovered we would expect similarities between the different runs. In general the rand score is high and therefore the communities produced are quite similar. It does also vary over time, with periods of market stress seeming to have higher consistencies (see 2008-2010 in US and UK, 2016 in the UK)

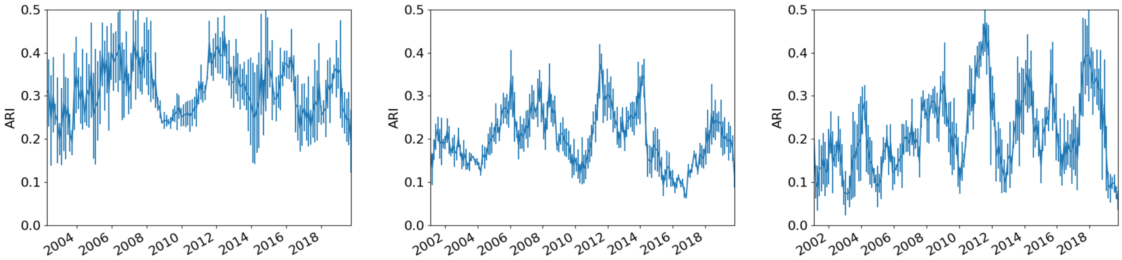

(a) US (b) UK (c) DE

Figure 11: Adjusted Rand Index of the community assignments when compared to the known GICS underlying sector assignments. In general there is not a great level of success in uncovering this underlying sector structure, perhaps indicating macroeconomic effects are more important than sector specific effects. There does also seem to be dips during times of market stress, with both the mean and standard deviation decreasing during these times in the US and UK markets.

5.5. Community Detection

In this section, we use the Louvain modularity maximization algorithm on the full correlation networks to perform community detection. We are interested in four measures, the number of communities in the network, the adjusted rand score between the 10 runs on each network (referred to as the community consistency), the adjusted rand score when comparing the community assignments to the underlying sector assignments (using the GICS assignments) and the adjusted rand score when comparing the assignments from one network to the next in time - referred to as the community stability over time.

Firstly, we look at how consistent the community detection is when run on the same network. This is shown in Figure 10. The communities produced seem to be consistent, with the US mean being at 0.751, UK mean at 0.732 and German mean at 0.730. This consistency is also time-varying - for both the US and the UK there are peaks during the financial crisis and the UK also has other peaks in 2004 and 2016. This would indicate that there is more signal during these times for the community detection procedure to uncover.

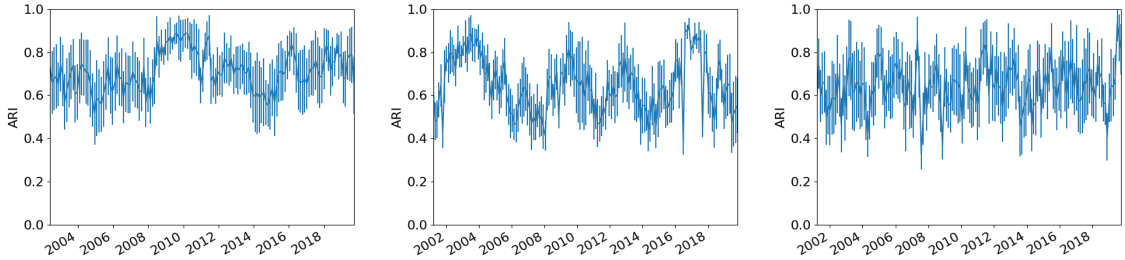

(a) US (b) UK (c) DE

Figure 12: Stability of the communities over time (measured as the adjusted rand index of the networks from one period to the next). For all the markets this is a noisy measure, but the communities do seem to be relatively consistent. In particular it increases for the US during 2003 and from 2008 - 2010, which fits in with our hypothesis that the markets are actually more similar during times of disruption

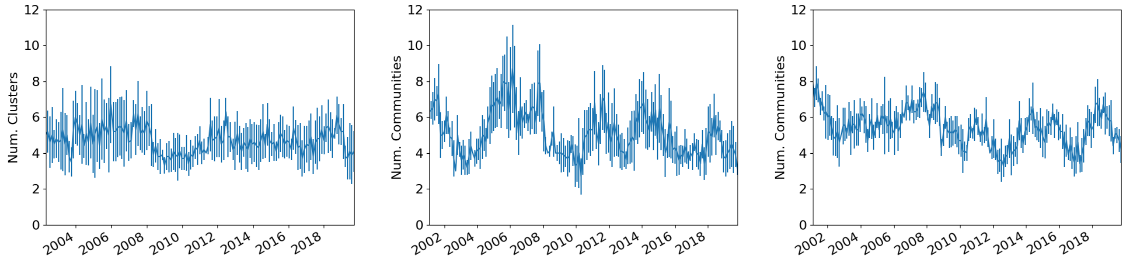

(a) US (b) UK (c) DE

Figure 13: Number of communities detected for each of the markets over the dataset. There does seem to be fewer communities during times of market stress, with perceptible lows for the US and UK between 2008 and 2010, and for the UK during 2016. Germany has drops during 2002, 2012 and 2016.

The adjusted rand score of the networks over time is shown in Figure 11. Interestingly, the level of success in uncovering this underlying sector structure is relatively small, showing that perhaps macroeconomic effects are more important than sector relationships. For the US markets, both the mean and standard deviation of the rand score decrease during 2003 and from 2008 - 2010, indicating that the market distress causes companies to behave more similarly irrelevant of whether they are in the same sector. The UK shows a similar trend, with a decrease in the rand index from 2008 - 2010. There are also lows in 2004 and 2017. Germany shows a great deal of variation compared to the other two countries, with maximums in 2012 and 2018, and a large number of lows. It is difficult to connect this to any market conditions.

The community stability over time is shown in Figure 12. We see peaks for the US markets from 2008 - 2010 indicating that surprisingly, there is more signal during times of disruption than calmer times. The UK has a similar trend, showing an increase in community stability over the financial crisis, but this is not quite as clear as it is for the US. There is generally more variation compared to the US. The German stability is very noisy, with no discernible time varying structure.

The number of communities in each network is shown in Figure 13. If the markets are

Network Measure	Correlation with Market Volatility		
Country	US	UK	DE
Largest Eigenvalue	0.593	0.494	0.483
Leading Eigenvector Difference	-0.404	-0.304	*0.067*
L_2 Distance	-0.311	*-0.122*	*-0.059*
Cosine Distance	-0.500	-0.289	0.202
Leading Eigenvector Standard Deviation	-0.411	-0.291	*0.069*
Adjusted Rand Index	-0.444	-0.230	-0.228
Community Stability over Time	0.400	0.416	*-0.097*
Number of Communities	-0.502	-0.513	-0.299
Community Consistency	0.378	0.430	-0.270
Heterogeneity	*0.103*	0.335	*-0.152*
Mean Katz Similarity	*0.111*	0.583	*-0.148*
PMFG Edge Changes	*-0.139*	*-0.024*	*-0.030*

Table 1: Spearman correlation of the various network measures with the volatility of the index. The US and UK markets clearly have strong relationships between the volatility of the index and the network measures, while for Germany this is not the case.

becoming more similar we would expect fewer communities, so we should see a drop during times of disruption. This does seem to be the case, with drops for the US and UK between 2008 and 2010, and for the UK and Germany during 2012 and 2016.

5.6. Correlation Between Network Measures and Volatility

Finally, we look at the Spearman correlation between the various market measures and the volatility of the index. We also conduct a 2 sided hypothesis test to ensure the correlation is significant. To measure the volatility we use the standard deviation of the log returns of the index over the time window the network is inferred from. The results for all three countries are shown in Table 1. Non-significant correlations ($p > 0.05$) are in italics.

Since the UK and US have similar results we focus on them first. Both have a strong correlation between the volatility of the index and the largest eigenvalue, showing this is a reasonable measure of the disruption present in the market. Both also show a negative correlation between the L_2 distance, cosine distance, eigenvector centrality difference, and the market volatility indicating that the nodes do become more similar during times of disruption. The leading eigenvector also seems to be more stable during these times of disruption. However, it is important to state the correlation between the L_2 distance and volatility is only significant for the US market. The German picture is different, having only two significant correlations between the volatility and the network measures. In this case, this is the largest eigenvalue and the cosine distance, both of which share a positive correlation with the market volatility.

For the community detection, we have a negative correlation for all markets between volatility and adjusted rand index, which is something that fits in with our hypothesis - if the nodes become more similar then the underlying sector structure matters less. The

US and the UK have positive correlation between volatility and community stability and consistency, and all markets have negative correlation between the number of communities and the volatility.

With the PMFG measures (heterogeneity, Mean Katz Similarity, and the edge changes) we see that only the UK has any significant correlations, with heterogeneity and Mean Katz Similarity having positive correlations with the market volatility. It is somewhat surprising that only the UK shows any relationships between volatility and network measures from the PMFGs. This could be due to the filtration procedure, as it discards the majority of the correlations present, and it could be some of these smaller correlations that are causing these effects in the full networks.

Throughout this investigation, it appears the US and UK markets show relatively similar effects, while Germany deviates significantly. Part of this could simply be due to the significantly smaller size of the German market (23 companies vs 70 and 229) but it is also known the German economy has a significantly different set up the US and the UK. Both the UK and the US have much larger financial sectors than Germany, which places greater emphasis on their stock markets.

6. Discussion and Conclusion

In this paper, we have inferred correlation networks using returns data from the US, UK, and German markets, and investigated how the stability and similarity between companies changes over time. To this end, we have used both full correlation networks and those filtered using the PMFG. We have also performed community detection on the full correlation networks.

Firstly we study the full correlation networks. To measure stability we use the L_2 norm of the difference between leading eigenvectors and the stability in the communities detected from adjacent networks. We find that the leading eigenvectors tend to change less during times of market disruption for the UK and US markets, and that the community consistency and stability increases. This is perhaps surprising, as we would expect the networks to change more during these times. Next, we explore how the similarity between nodes changes, measuring this using the standard deviation of the leading eigenvector, the L_2 and cosine distances between the rows of the correlation matrix, the number of clusters in the community detection, and the success of recovering the underlying sector structure. We find that all these measures drop during times of distress for the US and UK markets indicating that the nodes do become more similar during times of market disruption, but that for the German market the only measures with significant relationships with the market volatility are the cosine similarity and the rand index.

For the PMFGs the results are different. Using the mean Katz similarity and the heterogeneity to measure similarity in these filtered graphs, we firstly find that these measures show very similar trends. For Germany and the US, these measures are difficult to link to particular market states, and this is reflected in their low correlation with the index volatility. This is not the case with the UK, which shows a strong correlation with market volatility. Measuring the fraction of edge changes, all three show a sharp drop in the number of edge changes in 2009, but beyond that, there is little relation with market volatility. This is perhaps surprising, although this could be due to the PMFG construction procedure discarding many of the correlations that are changing. While nodes are more likely to be considered similar to those in the same sector, the relationship is not strong.

There are of course limitations to the approach taken in this paper. Firstly we have used Pearson correlation to quantify the relationships between assets. The advantage of Pearson correlation is that it is simple to interpret and calculate, properties which make it ideal for the initial exportation of complex systems. However, it assumes linearity and normality, and is sensitive to outliers. Financial returns tend to be neither linear nor normal and contain outliers. Potential future work could involve using rank correlations or mutual information to overcome these issues. A second limitation is that we have a set size for the window of data from which we infer the correlation matrices from. This does not take into account changes that could occur in the middle of windows, leading to a 'blurring' of different correlation matrices. Recent work in change point detection [63] [64] has shown that it is possible to dynamically select an appropriate window size, and hopefully avoid this issue. A future paper could exploit this to more accurately study the changes in correlation structure between times of market calm and stress.

References

[1] G. Bonanno, F. Lillo, R. N. Mantegna, Levels of complexity in financial markets, Physica A: Statistical Mechanics and its Applications 299 (1) (2001) 16–27, application of Physics in Economic Modelling. doi:https://doi.org/10.1016/S0378-4371(01)00279-5.
URL `https://www.sciencedirect.com/science/article/pii/S0378437101002795`

[2] R. N. Mantegna, Hierarchical structure in financial markets, The European Physical Journal B - Condensed Matter and Complex Systems 11 (1) (1999) 193–197.

[3] M. Newman, Networks: an introduction, Oxford university press, 2010. doi:10.1017/CBO9780511761942.

[4] M. J. Bricco, M. T. Xu, Interconnectedness and contagion analysis: a practical framework, International Monetary Fund, 2019.

[5] T. Preis, D. Y. Kenett, H. E. Stanley, D. Helbing, E. Ben-Jacob, Quantifying the behavior of stock correlations under market stress, Scientific Reports 2 (2012) 752 EP –, article. doi:10.1038/srep00752.

[6] A. Nagurney, J. Dong, D. Zhang, A supply chain network equilibrium model, Transportation Research Part E: Logistics and Transportation Review 38 (5) (2002) 281 – 303. doi:10.1016/S1366-5545(01)00020-5.

[7] J.-P. Onnela, K. Kaski, J. Kertész, Clustering and information in correlation based financial networks, The European Physical Journal B 38 (2) (2004) 353–362.

[8] V. Boginski, S. Butenko, P. M. Pardalos, Statistical analysis of financial networks, Computational Statistics & Data Analysis 48 (2) (2005) 431 – 443. doi:10.1016/j.csda.2004.02.004.

[9] G. Peralta, Network-based measures as leading indicators of market instability: the case of the spanish stock market, Journal of Network Theory in Financedoi:10.2139/ssrn.3400673.

[10] F. Pozzi, T. Di Matteo, T. Aste, Spread of risk across financial markets: better to invest in the peripheries, Scientific reports 3 (2013) 1665. doi:10.1038/srep01665.

[11] V. Boginski, S. Butenko, O. Shirokikh, S. Trukhanov, J. G. Lafuente, A network-based data mining approach to portfolio selection via weighted clique relaxations, Annals of Operations Research 216 (1) (2014) 23–34.

[12] L. Uechi, T. Akutsu, H. E. Stanley, A. J. Marcus, D. Y. Kenett, Sector dominance ratio analysis of financial markets, Physica A: Statistical Mechanics and its Applications 421 (2015) 488 – 509. doi:https://doi.org/10.1016/j.physa.2014.11.055.
URL `http://www.sciencedirect.com/science/article/pii/S0378437114010267`

[13] A. Ang, J. Chen, Asymmetric correlations of equity portfolios, Journal of Financial Economics 63 (3) (2002) 443 – 494. doi:https://doi.org/10.1016/S0304-405X(02)00068-5.
URL `http://www.sciencedirect.com/science/article/pii/S0304405X02000685`

[14] D. Wang, B. Podobnik, D. Horvatić, H. E. Stanley, Quantifying and modeling long-range cross correlations in multiple time series with applications to world stock indices, Phys. Rev. E 83 (2011) 046121. doi:10.1103/PhysRevE.83.046121.
URL `https://link.aps.org/doi/10.1103/PhysRevE.83.046121`

[15] L. Sandoval, I. D. P. Franca, Correlation of financial markets in times of crisis, Physica A: Statistical Mechanics and its Applications 391 (1) (2012) 187 – 208. doi:https://doi.org/10.1016/j.physa.2011.07.023.
URL `http://www.sciencedirect.com/science/article/pii/S037843711100570X`

[16] D. Y. Kenett, M. Tumminello, A. Madi, G. Gur-Gershgoren, R. N. Mantegna, E. Ben-Jacob, Dominating clasp of the financial sector revealed by partial correlation analysis of the stock market, PLOS ONE 5 (12) (2010) 1–14. doi:10.1371/journal.pone.0015032.

[17] D. Y. KENETT, T. PREIS, G. GUR-GERSHGOREN, E. BEN-JACOB, Dependency network and node influence: Application to the study of financial markets, International Journal of Bifurcation and Chaos 22 (07) (2012) 1250181. arXiv:https://doi.org/10.1142/S0218127412501817, doi:10.1142/S0218127412501817.
URL `https://doi.org/10.1142/S0218127412501817`

[18] T. Millington, M. Niranjan, Quantifying influence in financial markets via partial correlation network inference, in: 2019 11th International Symposium on Image and Signal Processing and Analysis (ISPA), 2019, pp. 306–311. doi:10.1109/ISPA.2019.8868437.

[19] A. Kocheturov, M. Batsyn, P. M. Pardalos, Dynamics of cluster structures in a financial market network, Physica A: Statistical Mechanics and its Applications 413 (2014) 523 – 533. doi:https://doi.org/10.1016/j.physa.2014.06.077.
URL `http://www.sciencedirect.com/science/article/pii/S0378437114005585`

[20] N. Musmeci, T. Aste, T. Di Matteo, Relation between financial market structure and the real economy: Comparison between clustering methods, PLOS ONE 10 (3) (2015) 1–24. doi:10.1371/journal.pone.0116201.
URL `https://doi.org/10.1371/journal.pone.0116201`

[21] S. Drożdż, F. Grümmer, A. Górski, F. Ruf, J. Speth, Dynamics of competition between collectivity and noise in the stock market, Physica A: Statistical Mechanics and its Applications 287 (3) (2000) 440 – 449. doi:https://doi.org/10.1016/S0378-4371(00)00383-6.
URL `http://www.sciencedirect.com/science/article/pii/S0378437100003836`

[22] T. C. Silva, S. R. S. de Souza, B. M. Tabak, Structure and dynamics of the global financial network, Chaos, Solitons & Fractals 88 (2016) 218–234.

[23] R. H. Heiberger, Stock network stability in times of crisis, Physica A: Statistical Mechanics and its Applications 393 (C) (2014) 376–381. doi:10.1016/j.physa.2013.08.0.
URL `https://ideas.repec.org/a/eee/phsmap/v393y2014icp376-381.html`

[24] J.-P. Onnela, A. Chakraborti, K. Kaski, J. Kertész, Dynamic asset trees and black monday, Physica A: Statistical Mechanics and its Applications 324 (1) (2003) 247 – 252, proceedings of the International Econophysics Conference. doi:10.1016/S0378-4371(02)01882-4.

[25] Y. Zhang, G. H. T. Lee, J. C. Wong, J. L. Kok, M. Prusty, S. A. Cheong, Will the us economy recover in 2010? a minimal spanning tree study, Physica A: Statistical Mechanics and its Applications 390 (11) (2011) 2020 – 2050. doi:10.1016/j.physa.2011.01.020.

[26] J.-P. Onnela, A. Chakraborti, K. Kaski, J. Kertesz, A. Kanto, Dynamics of market correlations: Taxonomy and portfolio analysis, Physical Review E 68 (5) (2003) 056110.

[27] M. Wiliński, A. Sienkiewicz, T. Gubiec, R. Kutner, Z. Struzik, Structural and topological phase tran-

sitions on the german stock exchange, Physica A: Statistical Mechanics and its Applications 392 (23) (2013) 5963 – 5973. doi:https://doi.org/10.1016/j.physa.2013.07.064.
URL http://www.sciencedirect.com/science/article/pii/S037843711300695X

[28] B. Li, D. Pi, Analysis of global stock index data during crisis period via complex network approach, PLOS ONE 13 (7) (2018) 1–16. doi:10.1371/journal.pone.0200600.
URL https://doi.org/10.1371/journal.pone.0200600

[29] N. Musmeci, T. Aste, T. Di Matteo, Interplay between past market correlation structure changes and future volatility outbursts, Scientific reports 6 (2016) 36320.

[30] L. Zhao, W. Li, X. Cai, Structure and dynamics of stock market in times of crisis, Physics Letters A 380 (5) (2016) 654 – 666. doi:https://doi.org/10.1016/j.physleta.2015.11.015.
URL http://www.sciencedirect.com/science/article/pii/S0375960115009937

[31] E. Estrada, Quantifying network heterogeneity, Phys. Rev. E 82 (2010) 066102. doi:10.1103/PhysRevE.82.066102.
URL https://link.aps.org/doi/10.1103/PhysRevE.82.066102

[32] Maps of random walks on complex networks reveal community structure, Proceedings of the National Academy of Sciences 105 (4) (2007) 1118–1123.
URL http://dx.doi.org/10.1073/pnas.0706851105

[33] D.-M. Song, M. Tumminello, W.-X. Zhou, R. N. Mantegna, Evolution of worldwide stock markets, correlation structure, and correlation-based graphs, Phys. Rev. E 84 (2011) 026108. doi:10.1103/PhysRevE.84.026108.
URL https://link.aps.org/doi/10.1103/PhysRevE.84.026108

[34] W.-M. Song, T. Di Matteo, T. Aste, Hierarchical information clustering by means of topologically embedded graphs, PLOS ONE 7 (3) (2012) 1–14. doi:10.1371/journal.pone.0031929.
URL https://doi.org/10.1371/journal.pone.0031929

[35] G. Marti, F. Nielsen, M. Bińkowski, P. Donnat, A review of two decades of correlations, hierarchies, networks and clustering in financial markets, arXiv preprint arXiv:1703.00485.

[36] G.-J. Wang, C. Xie, H. E. Stanley, Correlation structure and evolution of world stock markets: Evidence from pearson and partial correlation-based networks, Computational Economics 51 (3) (2018) 607–635. doi:10.1007/s10614-016-9627-7.

[37] T. Millington, M. Niranjan, Robust Portfolio Risk Minimization Using the Graphical Lasso, Springer International Publishing, Cham, 2017, pp. 863–872.

[38] T. Millington, M. Niranjan, Construction of minimum spanning trees from financial returns using rank correlation, Physica A: Statistical Mechanics and its Applications (2020) 125605doi:https://doi.org/10.1016/j.physa.2020.125605.
URL http://www.sciencedirect.com/science/article/pii/S0378437120309031

[39] T. Isogai, Dynamic correlation network analysis of financial asset returns with network clustering, Applied Network Science 2 (1) (2017) 8. doi:10.1007/s41109-017-0031-6.
URL https://doi.org/10.1007/s41109-017-0031-6

[40] F. Lorrain, H. C. White, Structural equivalence of individuals in social networks, The Journal of mathematical sociology 1 (1) (1971) 49–80.

[41] D. R. White, K. P. Reitz, Graph and semigroup homomorphisms on networks of relations, Social Networks 5 (2) (1983) 193–234.

[42] R. Hanneman, Introduction to social network methods, 2001.

[43] L. Lü, T. Zhou, Link prediction in complex networks: A survey, Physica A: statistical mechanics and its applications 390 (6) (2011) 1150–1170.

[44] S. Ontañón, An overview of distance and similarity functions for structured data, Artificial Intelligence Review (2020) 1–43.

[45] F. Bell, A note on the irregularity of graphs, Linear Algebra and its Applications 161 (1992) 45 – 54. doi:https://doi.org/10.1016/0024-3795(92)90004-T.
URL http://www.sciencedirect.com/science/article/pii/002437959290004T

[46] M. Tumminello, T. Aste, T. Di Matteo, R. N. Mantegna, A tool for filtering information in complex

systems, Proceedings of the National Academy of Sciences 102 (30) (2005) 10421–10426.
[47] G. P. Massara, T. Di Matteo, T. Aste, Network filtering for big data: Triangulated maximally filtered graph, Journal of complex Networks 5 (2) (2016) 161–178.
[48] J.-P. Bouchaud, M. Potters, Financial applications of random matrix theory: a short review, arXiv preprint arXiv:0910.1205.
[49] M. Potters, J.-P. Bouchaud, L. Laloux, Financial applications of random matrix theory: Old laces and new pieces, arXiv preprint physics/0507111.
[50] V. Plerou, P. Gopikrishnan, B. Rosenow, L. A. N. Amaral, T. Guhr, H. E. Stanley, Random matrix approach to cross correlations in financial data, Physical Review E 65 (6) (2002) 066126.
[51] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks, Journal of Statistical Mechanics: Theory and Experiment 2008 (10) (2008) P10008. doi:10.1088/1742-5468/2008/10/p10008.
[52] M. E. J. Newman, Modularity and community structure in networks, Proceedings of the National Academy of Sciences 103 (23) (2006) 8577–8582. arXiv:https://www.pnas.org/content/103/23/8577.full.pdf, doi:10.1073/pnas.0601602103.
URL `https://www.pnas.org/content/103/23/8577`
[53] S. Gómez, P. Jensen, A. Arenas, Analysis of community structure in networks of correlated data, Physical Review E 80 (1) (2009) 016114.
[54] M. MacMahon, D. Garlaschelli, Community detection for correlation matrices, Phys. Rev. X 5 (2015) 021006. doi:10.1103/PhysRevX.5.021006.
URL `https://link.aps.org/doi/10.1103/PhysRevX.5.021006`
[55] T. E. Oliphant, A guide to NumPy, Vol. 1, 2006.
[56] W. McKinney, Data structures for statistical computing in python, in: Proceedings of the 9th Python in Science Conference, 2010, pp. 51 – 56.
[57] S. Seabold, J. Perktold, Statsmodels: Econometric and statistical modeling with python, in: 9th Python in Science Conference, 2010.
[58] J. D. Hunter, Matplotlib: A 2d graphics environment, Computing in Science & Engineering 9 (3) (2007) 90–95. doi:10.1109/MCSE.2007.55.
[59] P. Shannon, A. Markiel, O. Ozier, N. S. Baliga, J. T. Wang, D. Ramage, N. Amin, B. Schwikowski, T. Ideker, Cytoscape: a software environment for integrated models of biomolecular interaction networks, Genome research 13 (11) (2003) 2498–2504.
[60] A. A. Hagberg, D. A. Schult, P. J. Swart, Exploring network structure, dynamics, and function using networkx, in: Proceedings of the 7th Python in Science Conference, Pasadena, CA USA, 2008, pp. 11 – 15. doi:10.25080/issn.2575-9752.
[61] Federal reserve bank of st. louis' financial crisis timeline.
URL `https://fraser.stlouisfed.org/timeline/financial-crisis`
[62] R. C. K. Burdekin, E. Hughson, J. Gu, A first look at brexit and global equity markets, Applied Economics Letters 25 (2) (2018) 136–140. arXiv:https://doi.org/10.1080/13504851.2017.1302057, doi:10.1080/13504851.2017.1302057.
URL `https://doi.org/10.1080/13504851.2017.1302057`
[63] I. Barnett, J.-P. Onnela, Change point detection in correlation networks, Scientific reports 6 (2016) 18893.
[64] S. Banerjee, K. Guhathakurta, Change-point analysis in financial networks, Stat 9 (1) (2020) e269, e269 sta4.269. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/sta4.269, doi:https://doi.org/10.1002/sta4.269.
URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/sta4.269`