

# The Influence of the Few

## Relationships in Corporate Power Structures

Project Report: CS224w

*Antonio Castro, Josh Neland, Kevin McKenzie*

### Abstract

Previous studies focusing on corporate power relationships have discovered that corporation-to-corporation and director-to-director networks exhibit properties of the Small World network model, with large central banks anchoring the network as the most central components. Most of these studies focus on the network as a whole and the static structure of the network at a given period with little emphasis placed on the evolution of the underlying community structures of the graph. Additionally, the effects of these relationships on corporate performance are unclear.

Our analysis confirms the Small World network structure and trends discovered in previous research. Furthermore, it validates the structure of the evolution of underlying collaboration networks follow similar characteristics of other collaboration networks such as the telephone call network and that of collaborating authors in the research community. Our analysis uses a community weight function to provide insight into the nature of the relationships within the company boards. Finally, we examine the link between shareholder value and the corporation's relationship to the overall network and to the communities embedded within the network.

### Introduction and Summary of Prior Work

Directors serving on the boards of powerful companies often sit on multiple boards in other industries. Companies fill the positions on their boards of directors with varying strategies, resulting in a network of relationships between companies. Companies are linked by shared board members, and directors themselves are linked to each other by sitting on the same boards as other directors.

There have been several studies that focus on these corporate power relationships to try to identify key characteristics about the network. The papers explored in this study investigate the measurement and impact of connectedness, communities and evolution within a network. Additionally, each can be applied to an undirected network of corporate relationships.

Insight into the structure of corporate power relationship networks helps to understand the power and influence held by groups of corporations. By exploring the structure and communities in the network, we hope to validate previously hypothesized trends and discover new properties of corporate power relationships.

In *The Small World of the Corporate Elite [1]*, authors Davis, Yoo, and Baker analyzed the network structure of US companies and their boards of directors to identify network changes as a response to political, economic, and social changes in the past 20 years. The authors gathered data from public SEC filings to develop two network graphs: one for director-to-director relationships and one for company-to-company. In the former graph, an edge connected two directors if both directors served on the same board at the time the data was collected. In the latter, an edge connected two companies if they shared a single director. Within each of these networks, they compared data from 1999 to similar data from previous years to determine whether changes in corporate culture and governance over the prior 20 years had affected the structure of the graphs. Prior research into corporate relationships showed that company relationship graphs were similar to the Small World model, with large US banks at the center of the graph.

The authors determined that many individual components of the network had changed, but the overall structure of the graphs and their key attributes had not changed significantly. The network structure was constant despite the major changes to corporate structure and governance over the prior 20 years. Banks maintained their central position in the graph as some of the most connected components, but their ranks and connectedness decreased as in the prior data sets. In addition, boards decreased in size over time, but the average path length and clustering coefficients remained close to those calculated over 20 years prior. This indicated that the network structure was more robust than previously imagined and that the Small World model continued to apply to both the director and company networks.

There are several other aspects of the corporate relationships networks that remain unexplored. *Quantifying Social Group Evolution [2]*, by Palla, Barabási, and Vicsek, proposes a set of measures of the evolution of communities over time within a social network and relates basic properties of these communities to the likelihood of their evolution.

The researchers use the Clique Percolation Method to determine communities at various points in time as the basis for the study. They propose a metric for determining the level of change in each community between two points of times using the number of nodes in the intersection of the two time stages divided by the number of nodes in their union. They further define a metric for stationarity as a function of the aggregation of this metric across time periods divided by the number of steps in the life of the community.

The study compares the behaviors of large and small communities within a collaboration network of scientists and a phone-call network of phone users. Based on the observations of these metrics in these two networks, the authors propose a set of rules that the small and large communities seem to follow. Specifically they observe that smaller networks' lifetimes are longer when stationarity is close to one while large networks with low stationarity metrics tend to have longer lifespans.

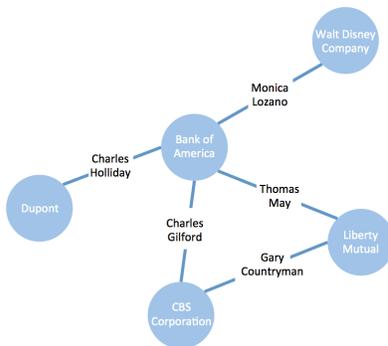
The paper continues to evaluate the structure of the relationships in the network as a predictor of the future evolution of the communities. Specifically, they measure the commitment of a node by the number of connections within the community compared to the ratio of connections outside a community as a measure of likelihood to leave a community. The commitment correlates to the likelihood to stay in a community and further confirms the relationship of the aggregate commitments of the members of a community to the lifetime of the community.

## Data and Approach

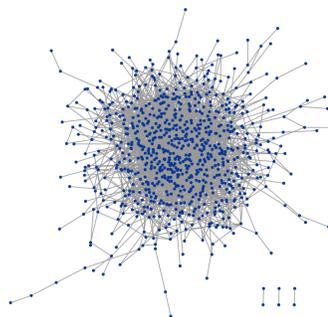
For our analysis, we developed two network graphs with the most recent data pulled from publicly accessible SEC filings on the EDGAR database, as collected by the company LittleSis.org [4]: one for director-to-director relationships between board of director members who sit on the same board, and one for corporation-to-corporation for companies that have shared board of director members.

We elected to use data from the top 896 largest US companies as provided by LittleSis.org, as the data for those companies were the largest (by market capitalization) and most powerful companies in the data set. The data set included industry information and the members of the boards of directors for the 896 companies. In addition to the LittleSis.org data set for SEC filings, we also collected adjusted stock price information that takes into account stock splits and dividends. We acquired this information using the Quandl platform [8] to create measures of the relative performance for each company and the volatility of that performance. During performance analysis, we included only the 625 companies that had been public for the entire 2002-2010 timespan.

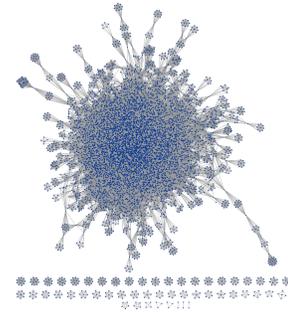
For example, a small subset of the corporation-to-corporation graph appears below, along with the visualizations of the entire graphs:



**Expanded Subset of Corporation Network**



**Corporation Network**



**Director Network**

Examination of a more holistic network of board-level relationships could reveal new insights about the effects of corporate relationships. Of the communities that exist in the director-to-director and corporation-to-corporation networks, the correlation to corporate performance has yet to be investigated.

## Overall Approach

At a high level, the following approach was used to analyze our data set:

1. Analyze the network structure for each graph and compare it to the findings of the previous studies performed by Davis et al for companies from 1982 to 2001.

- a. Determine the network structure and compare it to the Small World model described by Watts and Strogatz in 1998 [5] to determine whether the networks still exhibited the same phenomena.
  - b. Find the most connected and central nodes to determine if large national banks are still the most connected corporations in the network.
2. Identify communities of various sizes using the algorithm proposed in *Uncovering the overlapping community structure of complex networks in nature and society* [7] and explore the concepts of stationarity over time defined in *Quantifying Social Group Evolution* [2] for each community.
  - a. Compare the properties of these communities in the board relationships to those of the telephone network and research collaboration networks in [2]. Use the information from this analysis as inputs to determine whether there are any links between stability and market growth of the companies.
  - b. Further details about our community analysis can be found in the section below.
3. Attempt to determine the influence of powerful board members and companies to see how they affect the performance of other companies in their community. Corporate performance can be measured by calculating a moving average of normalized stock price for each stock and then determining if a company was above, below or near the normalized average stock price of all companies in the study.
  - a. Further details about our performance analysis can be found in the section below.
4. Evaluate the success of our analysis by comparing to the previous study by Davis et al., as well as similar studies performed on different data sets.
5. Draw conclusions from our analysis to describe the corporate power networks.

### Community Analysis Details

Utilizing CPM, we used the active dates for each board member to company relationship to extract bipartite graphs corresponding to each year from 2001 to 2013. Next, both the company and person collaboration networks were projected as a single weighted graph where the weight of each edge is proportional to the number of collaborations between members of the network.

The k-clique percolation method extracted the communities from the data sets. There were two non-trivial challenges to using the k-clique method: determining the value to choose for k, and matching communities from one year to the next given that the CPM method yields overlapping communities, as simple similarity matching can result in match "collisions" as explained in detail by Palla et. al in the supplemental information for [2].

For choosing k, we used a similar approach to the method described in the supplemental information for the research published by Palla et al. Various values of k between 2 and 5 were evaluated to balance between average network size, number of networks and coverage of both the network as a whole and of the strongly connected component of the graph. The differences in results between values 2, 3, 4, or 5 were trivial. A value of k=4 created consistency with the prior research and provided a higher resolution of communities to analyze.

For our analysis of company networks, we evaluated the company-company network communities. Testing the same four values of k, the communities covered around 6% of the network or 3% of the largest connected component with the exception of k=3 which returned 25-28% of the LCC and of the overall network depending on the year with between 190 and 230 communities. For the company network, we also explored various algorithms for using the link weights. Given the sparsity of weights above one, the network and community structures virtually disappear when using weights to prune edges.

As an alternative metric, we explored a cohesion metric of each community using the method described by J.-P. Onnela, J. Saramäki, J. Kertész, K. Kask in their paper on *Intensity and coherence of motifs in weighted complex networks* [10]. This metric essentially uses the sum of the edge weights divided by the max number of edges possible in the community subgraph:  $n*(n-1)/2$ . This metric provided a community weight that we use in our analysis of correlation to performance.

Evaluating this metric and the corresponding networks, networks that had a high weight and were large typically represented a combination of demographic characteristics such as region and industry. For example, one of the top communities in terms of size and weight contained many of the oil and gas companies in the rocky mountain area. While communities with a smaller weight were less homogenous and with a smaller size tended toward a merge formed by a common board member.

Lastly, we generated the stationarity metric for each community to be used in our analysis of correlation to performance. We use the algorithm in the Palla et. al research on the company network:

$$\zeta \equiv \frac{\sum_{t=t_0}^{t_{\max}-1} C(t, t+1)}{t_{\max} - t_0 - 1},$$

Where  $t_{\max}$  is the number of years evaluated,  $t_0$  is the first time step that the network appeared and the function  $C$  as defined below in the explanation of linking the communities generated for each year.

For matching the communities between steps of the graph, we use the fact that if you take the union of the network edges from the graph at time  $t$  ( $D$ ) and at the time  $t+1$  ( $E$ ), you will produce a graph  $V$ . Since the new combined graph contains all of the edges from  $D$  and  $E$ , the communities percolated from  $V$  match to exactly 1 community in  $D$  and 1 community in  $E$  as the joint graph can only grow, join or remain unchanged as pointed out by I. Derenyi. Matches are determined by taking all of the communities in  $E$  that are a subset of each community in  $V$  and ranking by the auto-correlation formula. This process is repeated for each subsequent time period  $t$  and form the community links including when a community is created, dissolved, and in many cases the entire “lifespan” of communities can be quantified.

To explain a bit further, we generate the communities of  $V$  as  $V_k$  then find each community of  $D$  and  $E$  where  $E_i^k \subseteq V_k$  and  $D_i^k \subseteq V_k$ . We then take the overlap of each of these sets using the correlation function:

$$C_{ij}^k = \frac{|D_i^k \cap E_j^k|}{|D_i^k \cup E_j^k|}$$

Finally, we take each of the pairs, sort them in descending order, and assign a linkage between the community in  $D$  and the community in  $E$ . If no candidate exists in  $D$ , we consider it a new community. If no candidate exists in  $E$  we consider it a community that has dissolved.

More information about these properties and algorithms can be found in the Palla et, al research [2].

For the community structure metrics in our analysis of performance, we computed the following for each company:

- The number of communities they were part of
- The maximum weight of the communities they were part of
- The min, max, and average stationarity of the networks they were part of
- The stationarity of the max weight community they were part of

### Performance Analysis Details

To introduce the concept of performance into the analysis, two measures are defined based upon the weekly-adjusted closing price for the stock of each company. The adjusted stock price takes into account stock splits and dividends, allowing for an easy comparison across stocks regardless of the time period being used.

To measure performance, two metrics were introduced. *Decade Performance* measures the percentage return of a given stock over the decade from 2002-2012, relative to the mean percentage return of the cohort. The *Annual Performance* metric answers the question, “did this stock perform above average or below average for a given year, relative to other stocks in the same year?”, and it is defined as *alpha* below:

$\phi_{y,w}(c)$  is the closing price of stock  $c$  in week  $w$  of year  $y$

$$\delta_{y,a,b}(c) = \frac{\phi_{y,b}(c) - \phi_{y,a}(c)}{\phi_{y,a}(c)}$$

$$\alpha_y(c) = \delta_{y,1,52}(c) - \text{median}_{i \in C} \{\delta_{y,1,52}(i)\}$$

$$\bar{\delta}_y(c) = \frac{1}{52} \sum_{i=1..52} \delta_{y,i,i+1}(c)$$

$$v_y(c) = \left( \frac{1}{52} \sum_{i=1..52} (\delta_{y,i,i+1}(c) - \bar{\delta}_y(c))^2 \right)^{\frac{1}{2}}$$

$$\beta_y(c) = v_y(c) - \text{median}_{i \in C} \{v_y(i)\}$$

While the overall and annual stock price percentage increase is a valuable measure of relative annual performance, it does not take into account the journey that stock price took to arrive at that performance. The volatility of the weekly stock price can reveal information about large changes in perception or inter-year drama that might be hidden by our performance measurement. The *Volatility* metric (beta) captures this information.

Five models were trained: logistic regression, random forest, SVM with rbf kernel, SVM with linear kernel and SVM with a polynomial kernel. Because of the relatively small sample size of 625 companies, a 10% holdout sample was used and 10-fold cross-validation was used to select model parameters.

## Results and Analysis

### Small World Properties of the Corporate Relationship Networks

#### Results

The work of Watts and Strogatz has shown that networks exhibiting properties of the Small World network will have a low "Average Shortest Path" metric and a high clustering coefficient while still maintaining an average node degree  $K$  that is significantly less than the number of nodes  $N$ .

Davis, et al. found that both the corporation-to-corporation and director-to-director networks both strongly match the Small World network properties. Using the newly collected data from 2012 for both networks, we calculated the same metrics.

**Table 1:** Small World Graph Properties

Network	N	K (Avg Degree)	L (Avg Shortest Path)	C (Clustering Coeff)
1982 Boards	648	10	3.38	0.24
1999 Boards	600	8.6	3.46	0.22
<b>2012 Boards</b>	<b>600*</b>	<b>6.1</b>	<b>4.6</b>	<b>0.15</b>
1982 Directors	6505	19	4.27	0.88
1999 Directors	5311	16	4.33	0.87
<b>2012 Directors</b>	<b>5429*</b>	<b>15</b>	<b>5.5</b>	<b>0.89</b>

\* Note: only the top 600 companies were used for this calculation only in the new data set to best compare to the previous data.

#### Analysis

There are several important characteristics to note from the data. First, each structure metric continued the trends from 1982 - 1999. The average degree of each fell, most likely because boards of directors have fallen in size continuously from the early 1980s to the present. In addition, the average number of boards served on by each board member has decreased each year since 1998 [9]. This leads to slightly lower average node degrees and longer average path lengths.

Even with the significantly smaller board sizes, the average shortest path has only moderately increased and the clustering coefficient has remained largely unchanged. Each of these metrics strongly indicates that the structure of the graph has not deviated from the Small World model of Watts and Strogatz. It also shows that the structure of the graph has not changed significantly from the networks using earlier data sets. This corroborates the findings of Davis, et al. that the structure of the graph remains remarkably stable despite the changes to the identities of nodes.

### Network Linchpins

#### Results

We analyzed the top linchpin companies to determine the nodes most central to the network. The highest ranked linchpin is defined as the node that best connects the rest of the graph to form shortcuts. The ranking was generated using the betweenness ranking of each node.

**Table 2:** Top Linchpin Companies

Rank	1982	1999	2012
1	Mellon Bank *	Bell Atlantic	<b>Cigna</b>
2	J P Morgan Co *	Bank of America *	<b>Mellon Bank *</b>
3	First Interstate Bank *	CSX	<b>Electronic Data Systems</b>
4	Bank America *	Chase Manhattan *	<b>Northern Trust *</b>
5	Continental Illinois *	Sara Lee	<b>Archer Daniels Midland</b>

Davis, et al. also ranked the most central directors in the directors-to-directors relationship network. We collected this information as well for the 2012 data set and compared to the findings in prior studies, as shown in the table below:

**Table 3:** Top Linchpin Directors

Rank	1982	1990	1999	2012
1	Richard T Baker	Rawleigh Warner, Jr	Jesse H Amelle	<b>Donna Zarcone</b>
2	Gene K Beare	William M Ellinghaus	Martin D Walker	<b>Arthur Martinez</b>
3	Samuel A Casey	Juanita M Kreps	Vernon Jordan, Jr	<b>J Whisler</b>
4	Kenneth C Foster	Vernon Jordan, Jr	Claudine B Malone	<b>R Martin</b>
5	C Jackson Grayson, Jr	Thomas P Stafford	Drew Lewis	<b>Marshall Larsen</b>

### Analysis

As Davis, et al indicated, the most likely reason to explain the similarities to the Small World model is that there are linchpin nodes that create shortcuts between disparate areas of the networks. Previous studies indicated that large US banks were the linchpins that held the network together, but later research proved that US banks had declined in centrality, but the overall network shape had remained unchanged.

In 1982, the top five linchpin companies were large US banks. In 1999, the number dropped to three, and then to two using our board membership data from 2012. Each large US bank is denoted by a \* in Table 2.

It has become clear that large US banks no longer dominate the center of the corporate relationship graphs, so we determined which industries have filled the void. We collected the primary industry of each company as filed with the SEC, then found the average betweenness for companies in each industry. We filtered the data to include only industries that included at least 4 companies in the dataset to reduce the effect of outliers.

**Table 4:** Top Linchpin Industries (2012) (Minimum of 4 Companies)

Rank	Industry	Avg Betweenness	No. of Companies
1	Motor Vehicles	3852	4
2	Paper Mills	2956	4
3	State Commercial Banks	2908	7
4	Food Production	2606	5
5	Hospital & Medical Service Plans	2544	8
...			
8	National Commercial Banks	2328	18

From this table, it is evident that commercial banks still play a crucial role in the network. State Commercial Banks and National Commercial banks are the third and eighth ranked industries by centrality respectively. If we were to only include industries with 10 or more companies, National Commercial Banks would again be ranked highest.

We can also see from Table 3 that there are very few directors who stay top linchpins in more than one study. It may be surprising that the most critical nodes to the structure of the network change in each study, considering that the key graph properties tabulated in Table 1 have remained very similar.

The most likely explanation for the turnover in the top linchpin directors list is that top directors are often near the age of retirement. In fact, the average age of each of the top linchpin directors in 2012 is 65.8, similar to the average of 65.4 from the 1999 list. With the average of 10 years between studies, it is likely that most of the directors will have

retired or limited the scope of their work due to their age by the time the next study is completed. According to the Spencer Stuart Board Index 2013 [9], the average tenure of a board member is 8.6 years on a top 500 company.

Still, it is interesting that the key graph properties remain so similar despite the changes to the core director nodes. However, it makes intuitive sense that when a board member retires, the company may seek to replace him or her with a director with similar connections. As concluded by Pfeffer and Salancik in 1978, "firms might seek directors that are executives of important competitors, buyers or suppliers (including banks) in order to co-opt them".[6] This indicates that the company would try and hire a board member with very similar connections as the retiring board member, thus preserving the network structure.

## Community Structure in Corporate Power Networks

### Results

Using the methods described in Palla, et al., and the data provided by LittleSis, we produced the community structures in years 2001 through 2013.

Some basic facts about the community structure evolution comparing macro community effects provide some early information about the stationarity of the communities in the network.

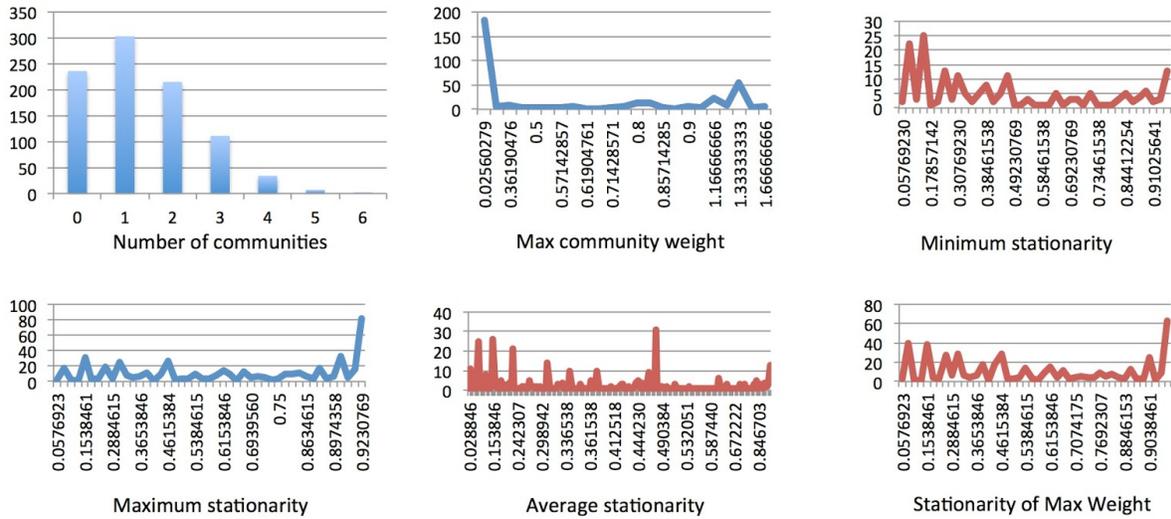
**Table 5:** Aggregate Community Properties of Board Members

	2001 vs. 2002	2001 vs. 2013
Number of new communities (splits or net new)	21	184
Number of communities dissolved (combined or disappeared)	2	9
% Of communities with a change in size	8.90%	45.10%
% Of communities with an increase in size	8.80%	42.80%
Number of communities with a decrease in size	0.10%	2.30%

**Table 6:** Aggregate Community Properties of Companies

	2001 vs. 2002	2001 vs. 2013
Number of new communities (splits or net new)	14	157
Number of communities dissolved (combined or disappeared)	0	4
% Of communities with a change in size	8.50%	25%
% Of communities with an increase in size	8.49%	24.50%
Number of communities with a decrease in size	0.01%	0.5%

The distributions of the generated descriptive variables generated for correlation to performance are below:

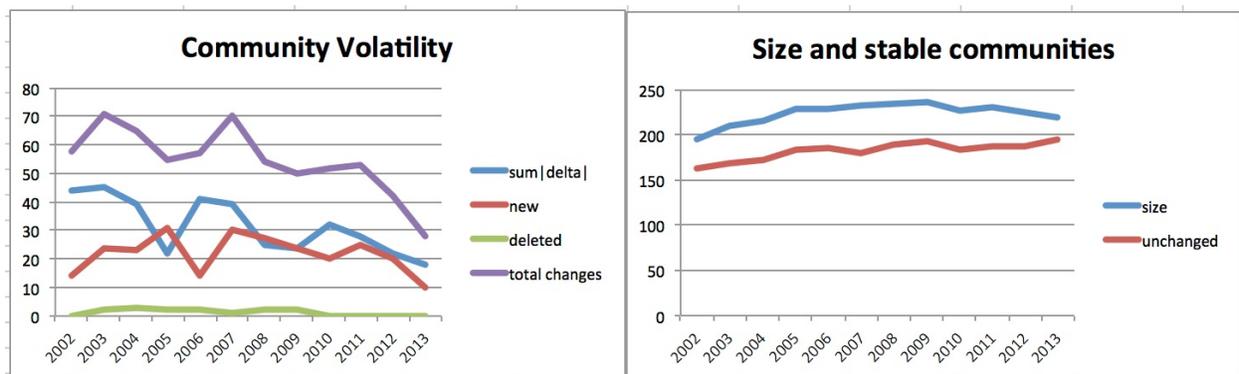


### Analysis

General observations of these properties for the company networks:

- For community weight, there are a large number of communities that are “weak” with a spike around 1.3 where we have smaller communities that are highly connected.
- For stationarity, no real pattern and the distributions seem relatively random.
- For stationarity of the max weight community, the distributions seem to cluster toward low stationarity and extremely high stationarity. This corresponds to the research by Palla, et al around the stability of smaller and larger communities.

Evaluating the volatility of the communities before and after the great recession in 2008, you can see that the communities were more volatile before the recession than they are after. Our hypothesis for this behavior is that the governance structures of the corporations have had a higher level of focus on stability after the recession. The charts below illustrate.



## The Influence of Board Members on Corporate Performance

### Results

The algorithm for generating the results below is found in the Performance Analysis Details section above. To test the influence of the corporate network on each company’s relative performance and volatility, the centrality (betweenness, closeness, and degree) and community characteristics of each company were combined as features and used to train and test various models for predicting the annual performance, decade performance and annual volatility.

**Table 7:** Predictive Accuracy of Community and Centrality Measures on Stock Behavior

<b>Model</b>	<b>Decade returns</b> (Accuracy/Random)	<b>Annual returns</b> (Accuracy/Random)	<b>Annual volatility</b> (Accuracy/Random)
Logistic Regression	0.8/0.8	0.5/0.51	0.61/0.62
Random Forest	0.77/0.75	0.45/0.51	0.57/0.58
SVM (rbf)	0.8/0.8	0.48/0.49	0.62/0.62
SVM (linear)	0.8/0.8	0.5/0.51	0.62/0.62
SVM (poly)	0.8/0.8	0.51/0.52	0.62/0.62

## Analysis

To baseline the accuracy measurement on the test set, the analysis was repeated with a random permutation of the training and test labels. In Table 7, these results are reported side-by-side with the predictions on the real data and show that no predictive power was found in the metrics produced related to centrality or community membership. Additionally, the variance of these predictions across the cross-validation scores was approximately +/- 0.05, meaning that every model performed similarly across both the data and the randomly permuted data.

Because multiple models, some linear and some non-linear with varying degrees of regularization, performed no better than random results (from a similar label distribution), we can eliminate the possibility that our model is missing a separation in the data. The reason accuracy is not 50%, even if the data is randomized, is because we create a binary label for our performance features based on a comparison to the mean value for the *decade returns* and *annual volatility* features. In the event of a highly skewed distribution (like decade returns), you can have many more companies that return above or below average.

## Conclusion

The corporate power networks for corporation-to-corporation and director-to-director relationships exhibited many of the same properties as those calculated by Davis et al. We concluded that the networks continue to be best described by the Small World network model. Furthermore, we discovered that large national banks continue to be among the most connected corporations in the network, but the trend of banks' decreasing centrality continued. Thus, we confirmed the network structure stability and network trends discovered by Davis et al. using updated corporate relationship data from 2012.

The community structure and the evolution of the communities resemble the characteristics of the telephone call and research collaboration research evaluated by Palla, et. al. An unexpected discovery of this analysis is how the community weight and community size combined provide insight into the value placed by companies in selecting board members within communities sharing region and industry focus. The evolution community structure also provided insight into the effects of the 2008 recession with a trend to more stability after the recession compared to before the recession.

While it seems intuitive that the relationships between corporate boards would facilitate the spread of corporate practices and that these practices would have a measurable effect on corporate performance, the analysis did not find such correlations. Neither the relationship of a company to the network as a whole, nor the communities to which the company belong seem to produce effects that matter to shareholders, assuming our metrics capture these characteristics effectively.

## Further Research Opportunities

While the analysis provided confirmation and additional information of the characteristics of the network, several opportunities exist to further the research and analyze the relationship between company performance and the boards that govern them. These further opportunities include:

- Use of other community characteristics:
  - Non-overlapping communities
  - Community size
  - Specific communities - are some communities more influential toward performance than others?
  - Community Centrality - does a community's position affect performance?
- Use of other company outcomes as targets: do any of these characteristics provide insight into

- Major Corporate events: Splits, Mergers, Acquisitions, Product launches
- Major Financial Events: Restatements, Capital Raises

## Citations

- [1] Davis, G. F., Yoo, M., & Baker, W. E. (2003). The small world of the American corporate elite, 1982-2001. *Strategic organization*, 1(3), 301-326.
- [2] Palla, G., Barabási, A. L., & Vicsek, T. (2007). Quantifying social group evolution. *Nature*, 446(7136), 664-667.
- [3] Newman, M. E., & Girvan, M. (2004). Finding and evaluating community structure in networks. *Physical review E*, 69(2), 026113.
- [4] *LittleSis.org Database*. Public Accountability Initiative. October 2014.
- [5] D. J. Watts and S. H. Strogatz. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393, 440-42.
- [6] Pfeffer, J. and Salancik, G.R. (1978) *The External Control of Organizations: A Resource Dependence Perspective*. New York: Harper & Row.
- [7] Gergely Palla, Imre Derényi, Illés Farkas, and Tamás Vicsek, *Uncovering the overlapping community structure of complex networks in nature and society* *Nature* 435, 814-818, 2005, doi:10.1038/nature03607
- [8] *Quandl.com*. November 2014.
- [9] Spencer Stuart. (2013) *Spencer Stuart Board Report 2013*. 19.
- [10] Intensity and coherence of motifs in weighted complex networks J.-P. Onnela, J. Saramäki, J. Kertész , K. Kaski *Phys. Rev. E* 71, 065103 (2005)

## Team Member Contributions

Each team member contributed equally to the research and report. The following sections were researched and written by the following team members:

- 1) Abstract - Kevin, Josh, and Antonio
- 2) Summary of Prior Work - Kevin, Josh, and Antonio
- 3) Data Model - Kevin, Josh, and Antonio
- 4) Algorithm - Kevin
- 5) Results and Findings
  - a) Data Model and Comparison to Previous Work - Kevin
  - b) Community Structure - Antonio
  - c) Performance Analysis - Josh
- 6) Conclusion - Kevin, Josh, and Antonio

The network graphs were collaboratively created and analyzed using NetworkX and Snap.py. We each contributed equally in designing the algorithms and coding our solutions. We collaborated using GitHub.