

Mining Longitudinal Network for Predicting Company Value

Yingzi Jin

The University of Tokyo, Japan
eiko-kin@biz-model.t.u-tokyo.ac.jp

Yutaka Matsuo

The University of Tokyo, Japan
matsuo@biz-model.t.u-tokyo.ac.jp

Ching-Yung Lin

IBM T. J. Watson Research Center
chingyung@us.ibm.com

Mitsuru Ishizuka

The University of Tokyo, Japan
ishizuka@i.u-tokyo.ac.jp

Abstract

Real-world social networks are dynamic in nature. Companies continue to collaborate, align strategically, acquire, and merge over time, and receive positive/negative impact from other companies. Consequently, their performance changes with time. If one can understand what types of network changes affect a company's value, he/she can predict the future value of the company, grasp industry innovations, and make business more successful. However, it often requires continuous records of relational changes, which are often difficult to track for companies, and the models of mining longitudinal network are quite complicated. In this study, we developed algorithms and a system to infer large-scale evolutionary company networks from public news during 1981–2009. Then, based on how networks change over time, as well as the financial information of the companies, we predicted company profit growth. This is the first study of longitudinal network-mining-based company performance analysis in the literature.

1 Introduction

Social Network Analysis (SNA) examines relations and structures of actors in a network to measure an actor's behavior and performance [Wasserman and Faust, 1994]. For example, some companies form numerous alliances with many small companies, whereas others make few collaborations but with big companies only; the partners of some companies are mutually well connected, whereas those of others' are connected with the target company only. These relations and structural embeddedness influence the behavior and performance growth of companies. For instance, degree, betweenness, and closeness centralities have been used to analyze actors' performance and value in the literature [Uzzi, 1997].

Relationships among companies are dynamic in nature. Companies continually establish and eliminate relationships with others, and continually receive positive or negative impacts from other companies. As a result, their performance changes with time. If one can understand useful effects of longitudinal networks of companies as well as the mechanism responsible for changing the company value, he/she

can infer a company's future value, decide strategic relational management for the company, and analyze the entire company's growth. Longitudinal networks are studied by sociologists to understand network evolution, belief formation, friendship formation, and so on [Wasserman and Faust, 1994; Doreian and Stokman, 1997; McCulloh and Carley, 2009].

As described in this paper, we explore a new analytical paradigm that uses large-scale representation of a longitudinal network of companies to predict a company's performance. Our goal is to find answers to the following research questions: Is it possible to predict a company's value (such as revenue and profit) based on dynamic (i.e. longitudinal) company networks? How can we infer evolutionary company networks? Two challenging algorithms are proposed in this paper: extraction of a longitudinal inter-company network from public news and mining longitudinal networks for future value prediction. For the first question, we develop a simple algorithm for temporal company network mining from public news. We specifically examine determination of the *impact* of relationships that a company shares with other companies via news articles, and use document and sentence co-occurrence to extract impactors for each target company to construct valued directed intercompany networks over years. Regarding the second question, we propose to investigate network effects related with company value from longitudinal impact relational networks. We generate network effects from local and global relationships, historical relationships, and delta-change in the relationships for each target company. We integrate those effects from the networks with a company's financial information to predict the company value. This paper is the first in the literature to describe performance analysis of a longitudinal network-mining-based company.

Experimentally obtained results show that our prediction model captures the trend of profit changes in the group companies or in an individual company's profit over the years. Results show that company profit prediction by a joint network and financial analysis outperforms network-only analysis by 150% and financial-only analysis by 34%. We can track the evolution of company networks over the years with different structural characteristics. The proposed algorithms are applicable not only to company domains, but also to people, products, and Web document ranking (or value) prediction from their dynamic networks.

This paper is organized as follows. The following section

presents a description of some related studies. Section 3 proposes a method for extracting relationships among companies from public news and shows the constructed networks. Section 4 proposes a method for mining longitudinal networks for developing a value prediction model and shows the predicted results. Section 5 concludes the study.

2 Related Work

In the literature, most prior approaches that have been attempted for the prediction of company values fall into three categories. The first approach (designated as a financial approach) is based on a company's financial statement (e.g., return on assets, capital ratio, number of employees) to measure the company's future earnings and performance [Xiao and Dasgupta, 2009]. The second approach (named a technical approach) is applying historical trends to identify price patterns and trends and to exploit those patterns to predict the direction of company values [Wang and Chan, 2007]. The third approach (named Social network Analysis, SNA) views relational and structural embeddedness of companies on inter-company networks from positional characteristics [Wasserman and Faust, 1994; Uzzi, 1997]. This study employs the third approach, but we combined both historical and financial information.

Analysis of over time network data has been presented in the social sciences literature [Katz and Proctor, 1959; Doreian and Stokman, 1997; Snijders, 1997]. Longitudinal network analysis is used to elucidate network evolution, belief formation, friendship formation, etc. [Feld, 1997; Snijders, 1997; Xiao and Dasgupta, 2009; Ben-Zvi, 2009]. However, few studies of longitudinal network analysis have addressed intercompany networks. The reason is that the relationships among companies are complex and unspecific, and it is difficult to track companies' network changing over time. Some studies focused on a specific relation (e.g. alliance) only, or, using self-report data, simulation data with time. A common complaint is that scalability suffers and incomplete information problems arise [Xiao and Dasgupta, 2009; Ben-Zvi, 2009].

News articles contain titles, content, and publishing time. Therefore, they are good resources for longitudinal network mining. Several studies specifically examined large-scale of public news articles to extract valuable information such as risk statements, future earnings etc. of companies [Tetlock *et al.*, 2008]. Bao *et al.* reported that a company is more likely to co-occur with its competitors on Web pages (i.e., document-level co-occurrence) than with non-competitors. In this study, we extract longitudinal *impact* networks among companies by mining *New York Times* articles published during 1981–2009. Companies receive different degrees of impact from different companies during the year. Therefore, the network is a directed valued network. This makes our task unique.

3 Extraction of Longitudinal Networks from Public News

We developed algorithms and a system to infer large-scale evolutionary company networks from public news during 1981–2009. We have a company name list from the New

York Times (7,594 companies can be indexed)¹. Because we need continuous records of company activities during years, we select only those large companies that have appeared on the Fortune list² at least three times and indexable form NYT articles as target companies, so we can match the network period with the obtainable company values.

To evaluate a company from a long-term view, one must be able to collect and analyze relevant information from a broad range of news stories to capture the company's activity [Bernstein *et al.*, 2002]. For example, *IBM* appeared in about 300 news articles in New York Times in 2009 (277 articles as *I.B.M.* and 84 articles as *International Business Machines*). If one can read and remember all the news stories, general knowledge about the company would become clear—which companies made an *impact* on *IBM*, and how much? Our assumption for the *impact* relation is that if a company *frequently* has co-appeared in the target company's *important* news articles and has been *frequently* described together with the target company in *important* sentences over a period of time, the company will make a large impact on the target company in that period. Therefore, they are regarded as having strong relationships. We propose the use of document-level and sentence-level co-occurrences to measure the *frequency*, and to assign weight to each document and sentence to measure the *importance*.

For each target company x , we score candidate companies³ Y by their impact to x in a period t . First, for each candidate company $y \in Y$, we collect a document set $D_{x,y}^t$ and a sentence set $S_{x,y}^t$, in which it has co-occurred with the target company x in the period t . Then, we sum up each of those document-weight $w_d(i)$ and sentence-weight $w_s(j)$ to calculate the final relational score for each y related to x as follows:

$$score_x(y) = a \cdot \sum_{i \in D_{x,y}^t} w_d(i) + b \cdot \sum_{j \in S_{x,y}^t} w_s(j). \quad (1)$$

As described in this paper, we use the following equations to assign *importance* in terms of weight to each co-occurring document and sentence.

$$w_d(i) = \log\left(1 + \frac{1}{|Y'|} + \frac{tf_x}{\sum_{y \in \{x,Y\}} tf_y}\right) \quad (2)$$


$$w_s(j) = \log\left(1 + \frac{1}{|Y''|}\right) \quad (3)$$

In those equations, Y' and Y'' denote the company names from document i and sentence j , respectively. $|Y'|$ and $|Y''|$ are counts of those names, and tf_y is the frequency of name y appearing in a document. Intuitively, if a document includes many company names, then it will be less important for those two companies than a document that mentions only few companies. In addition, the sentence weight is high for company x and y if it mentions only two companies, and low if it lists many companies. Constants a and b represent a tradeoff be-


¹For the name alias (e.g., *I.B.M.* and *International Business Machines*), we use a unified name.

²<http://money.cnn.com/magazines/fortune/fortune500/2010/>

³Candidate companies are those companies which appear in news articles about x for that year



(a) IBM 2003.



(b) IBM 2009.

Figure 1: Evolution of Networks in Different Years.

Table 1: Example of generic relation extraction for IBM in 2009.

r	name	score	Examples of documents and sentences.
1	Microsoft	85.85	<i>I.B.M. – Microsoft (55 articles, 264 sentences, score=85.85455)</i>
2	Oracle	65.49	http://www.nytimes.com/2009/03/06/business/06layoffs.html
3	Google	57.75	– Two days after I.B.M.’s report, <u>Microsoft</u> said that its quarterly profits were disappointing.
4	HP	50.70	http://www.nytimes.com/2009/01/31/business/31nocera.html
5	Intel	48.52	– Caterpillar, <u>Kodak</u> , Home Depot, <u>I.B.M.</u> , even mighty <u>Microsoft</u> : they are all cutting jobs.
6	Dell	32.75	<i>I.B.M. – SPSS (1 articles, 9 sentences, score =13.675)</i>
7	Sun	29.45	http://www.nytimes.com/2009/07/29/technology/companies/29ibm.html
8	EMC	15.16	– I.B.M. to Buy <u>SPSS</u> , a Maker of Business Software
9	Apple	14.65	– I.B.M.’s \$50-a-share cash offer is a premium of more than 40 percent over <u>SPSS</u> ’s closing stock price...
10	SPSS	13.67	
11	GM	13.18	<i>I.B.M. – Nike (4 articles, 9 sentences, score =8.212)</i>
12	Xerox	12.13	http://www.nytimes.com/2009/01/22/business/22pepsi.html
13	Nokia	8.95	– The list of companies that have taken steps to reduce carbon emissions includes <u>I.B.M.</u> ,
14	Nike	8.21	<u>Nike</u> , <u>Coca-Cola</u> and <u>BP</u> , the oil giant.
15	Thomson Reuters	7.78	http://www.nytimes.com/2009/11/01/business/01proto.html
			– Others are water-based shoe adhesives from <u>Nike</u> and a packing insert from <u>I.B.M.</u>

tween the document weight and sentence weight. Heuristically, we set $a=1$, $b=5$ ⁴

Finally, we obtain longitudinal inter-company networks year by year. Figure 1 compares the evolution of ego networks in different years. From *IBM* evolution networks between 2003 and 2009, some companies such as *Motorola*, *Novell*, and *NEC* were listed on the top in 2003, but disappeared from the network in 2009. Instead of them, *Google*, *SPSS*, and *Xerox* newly appeared on the network of *IBM*. *Microsoft*, *HP*, and *Sun* remained on the network, but slightly changed their relational strength with *IBM*. From *Microsoft* evolution networks in 1995, 2003, and 2009, we found that their top-related companies changed from *Intuit Inc.*, *IBM* to *Google*, and the relational strengths for top companies also changed slightly during these years.

⁴ Actually, we use different sets of a and b , and compare the results by user questionnaire. Consequently, $a=1$ and $b=5$ receive the best votes.

Here we give representative examples for *IBM* in 2009. From Table 1, it is apparent that *Microsoft* had the most impact on *IBM* in 2009. They co-occurred in 55 articles and were described together in 264 sentences. From these sentences, we can learn that they are direct competitors. Many top-ranked companies are competitors of *IBM*, which made a big impact on *IBM* in 2009. Sometimes impact relationships are not described in many articles. For example, *SPSS* and *IBM* are not competitors. They co-occurred in only 1 article and were described together in 3 sentences, but their relationship is important. The reason for this is that, *SPSS* and *IBM* co-appeared in an article in a high-weight document (which describes only *SPSS* and *IBM*’s acquisition relation in the entire article), and that they are described together in high-weight sentences (in which they are closely described, and no other companies appear). *Nike* and *IBM* are not competitors, and have no specific relationships, but they were described together because they took similar action or came together

for one product. Consequently, they might also exert impact on each other. Therefore, our algorithm can extract competitors, specific relationships, and other relationships that might make impact on the target company, i.e. *impact* relationships.

As described in this paper, we do not categorize relationships as negative or positive ones because a company might have many competitors because it has a high demonstrated performance. Our intercompany networks are extracted based on a statistic count obtained from news articles about companies over a period. Therefore, they are suitable for predicting a company's long-term value change. Prediction of short-term changes (e.g. market price) might require the use of another algorithm. We will work to develop such an algorithm in the future.

4 Longitudinal Network Mining for Company Value Prediction

After constructing a dynamic impact relational intercompany network, we measure how networks change over time as well as the value changing mechanism, to predict the company's future value.

4.1 Network Effect Generation

First, we calculate network effects for each target node x by its embeddedness in longitudinal *impact* networks $\mathcal{G}^T$. We use a vector $\mathbf{F}_x^T$ to indicate multi-dimensional network effects for x , which includes current network effects, historical network effects, and the delta-change of effects.

The current network effect (denoted as $\mathbf{F}_x^t$) for the target node x is generated based on the idea from [Karamon *et al.*, 2008] as follows. First, we define a node set N_x for x that might exert impact on x directly or indirectly. Then we define node pairs of three types among N_x : $\langle x, i \rangle$ (in which $i \in N_x$), $\langle i, j \rangle$ (in which $i \in N_x, j \in N_x, i \neq j$), and $\langle i, k \rangle$ (in which $i \in N_x, k \in V$). We conduct basic operations⁵—*connectivity* $\beta(i, j)$ (returns 1 if i and j are reachable; 0, otherwise), *distance* $\mu(i, j)$ (returns distance between i and j), and *betweenness* $\zeta^x(i, j)$ (returns 1 if the shortest path between i and j includes x ; 0, otherwise)—for these node pairs, and take the sum and standardize those values by the network size $|V|$, to compare effect values across networks. Finally, we obtain six types of basic network effects of x as the following list.

- $\sum_{i \in N_x} \beta(x, i) / (|V| - 1)$, which means the number of connections that x has.
- $\sum_{i \in N_x} \mu(x, i) / (|V| - 1)$, which signifies the distance between x and its related nodes.
- $\sum_{k \in V} \beta(i, k) / (|V| - 1)$, which means the number of connections that nodes related to x have.
- $\sum_{i, j \in N_x} \beta(i, j) / (|V| - 1)(|V| - 2)$, which means the number of connections among x 's related nodes.
- $\sum_{i, j \in N_x} \mu(i, j) / (|V| - 1)(|V| - 2)$, which means the distance between x 's related nodes.

⁵the basic operators can be extended to consider triad relations [Wasserman and Faust, 1994] and more to expand the basic effect vector in the future.

- $\sum_{i, j \in N_x} \zeta(i, j) / (|V| - 1)(|V| - 2)$, which means the number of node pairs having x on the shortest path.

We consider node set N_x , which might exert an impact on x by a *neighboring node set* L_x (i.e. directly connected) and a *reachable node set* G_x (i.e. indirectly connected) of x . In addition, the difference of impact from local and global node sets is also important. For example, the ratio of connections with x between L_x and G_x sets indicates the degree to which companies are directly related with x rather than indirectly. Furthermore, from the constructed valued directed network, we reduce the direction and weight information to generate networks of different types. For example, if we retain only the direction and ignore the weights, then the binary-directed network will represent who exerts an impact on whom but it will not show how much impact is exerted. This is similar to a friendship network (e.g., *Facebook* or *LinkedIn*), we only consider who treats whom as friends, but do not know how strong the friendship is. Therefore, by considering local and global impacts and their ratio, as well as networks of four different types (i.e., valued / unvalued, directed / undirected), we can generate a $72 = (3 \times 4 \times 6)$ -dimensional network effect vector for each target node x in the current network, i.e. $\mathbf{F}_x^t = F(N_x, d, v, t)$, where $N_x \in \{L_x, G_x, L_x/G_x\}$, $(v, d) \in \{(0, 1)\} \times \{0, 1\}$, and $t \in T$.

After we have generated network effects from the current network, we further generate historical network effects $\mathbf{F}_x^H$ by considering temporal information, i.e., $\mathbf{F}_x^H = \{\mathbf{F}_x^{t-1}, \mathbf{F}_x^{t-2}, \dots, \mathbf{F}_x^{t-w}\}$, where $\mathbf{F}_x^{t-w}$ indicates the network effects from the network that existed w years ago, which implies a historical network impact exerted by other companies. In addition, the amount of change over time is also considered, that is $\Delta \mathbf{F}_x^H = \{\Delta \mathbf{F}_x^{t-1}, \Delta \mathbf{F}_x^{t-2}, \dots, \Delta \mathbf{F}_x^{t-w}\}$. For example, we can examine the delta-change in neighboring nodes from last year to this year, or delta-changes from three years ago to this year.

As a result, for each company for each year, we generate 72-dimensional current-year network effects $\mathbf{F}_x^t$, plus $72 \times$ window size-dimensional historical network effects $\mathbf{F}_x^H$, and plus $72 \times$ delta size-dimensional network effects $\Delta \mathbf{F}_x^H$.

$$\mathbf{F}_x^T = \{\{\mathbf{F}_x^t\}, \{\mathbf{F}_x^H\}, \{\Delta \mathbf{F}_x^H\}\}. \quad (4)$$

In addition, our prediction model can combine effects of historical financial statements of companies, such as the previous year's profit and the profit earned three years prior.

4.2 Prediction Model

After we have generated longitudinal network effects for each target node x , we integrate those values as features to learn and predict the future value of the company.

$$y_x^{t'} = f(\mathbf{F}_x^t, \beta) = \sum_k f_k \beta_k \quad (5)$$

Therein, $t' > t$, and f_k is the k -th effect from the historical network, and β_k is the importance of f_k . The prediction model is designed to learn the unknown parameter β from observed data, and it can use any up-to-date regression model. In this study, we use the support vector regression (SVR) model. We fit the predictive model to the observed dataset




Figure 2: Prediction of the mean profits of 20 Fortune companies.


of company value and effect variables. When an additional company's effect set is given, the model can predict the company value. Additionally, we can predict future values of a list of companies and understand the future trend of the industry.

4.3 Prediction Results


In this section, we evaluate the prediction results. The company values we try to learn and predict are obtained from the Fortune 500. The Fortune 500 list, published by Fortune magazine⁶, ranks top (gross revenue) American public corporations during 1955 to the current year. Therefore, we use longitudinal network effects generated during 1981–2009 to learn and predict the company's value in terms of profit. First, we learn and predict 20 Fortune companies' profits (and calculate their mean value). Then we calculate the mean profit with real profits (and also calculate the mean value). Second, we train the model for each company, i.e. *IBM* and *Intel*, and predict their profits. Finally, we compare different feature sets and parameters.

We select 20 large Fortune companies from different industries: *I.B.M.*, *Intel*, *Microsoft*, *GM*, *HP*, *Honda*, *Nissan*, *AT&T*, *Wal-Mart*, *Yahoo!*, *Nike*, *Dell*, *Starbucks*, *Chase*, *PepsiCo*, *Cisco*, *FedEx*, *Gap*, *AEP*, and *Sun*. Because these companies continually appeared on the Fortune list over several years, we have continuous records of both their company values and networks. We learn the profit model from every five years' networks, and predict the next year's profits; then we compare with the real value of the profit earned in that year. Figure 2 displays a plot of the mean value of the predicted profits of these 20 companies, as learned from past five years, and the mean value of the real profit earned in the year. We use *SVR* (with an *RBF* kernel) model to learn parameters and make predictions. We use the squared correlation coefficient (denoted as r^2) and mean squared error (MSE) to quantify the correlation and error between the predicted values and true values, respectively. It is apparent that the output of the predicted values can capture the profit trend of these companies over the years, where $r^2 = 0.440$ and $MSE = 0.437$. Only in 1995 was the prediction much lower than the real value, perhaps because these companies created profits but the intercompany relationships still suffered an impact from the previous years' networks.

We also predict profits for two individual companies, *IBM* and *Intel*, based on their embeddedness in longitudinal networks. We learn from the prior ten years' networks and predict next year's profit. Then we compare it with the real



(a) IBM profit prediction.



(b) Intel profit prediction.

Figure 3: Profit Prediction for *IBM* and *Intel*.

value of the profit. It is apparent that our prediction results can capture the trends of company profits moving with $r^2 = 0.888$ and $MSE = 0.108$ for *IBM*, and $r^2 = 0.243$ and $MSE = 0.251$ for *Intel*.

To evaluate the effectiveness of network features, we use different feature sets for predicting 20 companies' mean profits during several years, and take the average over years to compare the prediction performance by each feature set. The notations of "s", "t", "p", and "d" indicate that results are obtained only using current network structural features (i.e., $\mathbf{F}_x^t$), historical network structural features (i.e., $\mathbf{F}_x^H$), delta-changes in the network features (i.e., $\Delta\mathbf{F}_x^H$), and financial features only ($\mathbf{P}_x^H$), respectively. We also combine these features for prediction, such as "sp" implies combining current network features with financial features, and "stdp" signifies a combination of all features, both network and financial features. We use the *SVR* (with the *RBF* kernel) model to learn parameters from the past five years' networks and predict the next year's profit. As the results presented in Figure 4 reveal, for profit prediction, using the feature set "p" i.e., financial profile features only (e.g., last year's profit, revenue, etc.) has better performance ($r^2=0.383$, $MSE=0.287$) than that realized using "s", "t", and "d" features only. However, by combining structural and temporal features with financial features as in "sp", "tp", "dp", and "stdp", the prediction results will improve. Particularly the prediction performance realized using "stdp" ($r^2=0.512$, $MSE=0.363$), i.e., joint network and financial features, will outperform that realized using only the network features and only the financial features by 150% and 34%, respectively.

To tune the best parameters of historical *window size* and *delta size*, we compare their values that existed 1 and 3 years ago, and found that one is sufficient for profit prediction, which implies that using last year's network effects and the delta-change in them from last year to this year are better than those of the networks existing three or more years ago.

⁶<http://money.cnn.com/magazines/fortune/fortune500/2010/>.




Figure 4: Mean profit prediction for 20 Fortune companies using different feature sets. s, structural features; t, temporal features; d, delta-change in temporal features; p, financial profiles.

5 Conclusions

In this study, we explored a new analytical paradigm of using social networks of companies to predict a company's value. We developed an algorithm and system for inferring longitudinal intercompany networks from public news. We described *impact* relationships among companies and developed an extraction method based on the document-level and sentence-level co-occurrence and importance. Using the system, we can elucidate the evolution of company networks over the years with different structural characteristics. After constructing valued directed longitudinal company networks over several years, we defined and extracted network effects for each target company from the networks. We investigated network characteristics from local and global relationships, and combined historical structural effects as well as delta changes in structures to generate network effects. We applied an SVM regression model to learn and predict company values (i.e. profit).

Our prediction model can capture the trend of changes in the values of group companies or an individual company over years. Results showed that generic relational networks are good for predicting the company value, particularly the company profit. Profit prediction based on joint networks and financial analysis outperforms that based only on network effects by 150% and that based only on financial effects by 34%. By tuning the *window size* and the *delta size* of longitudinal network effects, we found that the last year's network embeddedness is good for predicting this year's company value. Networks make one-year slow impact on changes in a company's value. In this study, our networks are found to be suitable for predicting long-term changes in company value. Future studies will specifically examine detection of short-term changes in the company value based on intercompany networks and specifically examine real-time network-effect extraction and processing infrastructure.

References

- [Ben-Zvi, 2009] Tal Ben-Zvi. Network structure and centrality: A simulation experiment. In *8th IEEE/ACIS International Conference on Computer and Information Science (ICIS 2009)*, 2009.
- [Bernstein *et al.*, 2002] Abraham Bernstein, Scott Clearwater, Shawndra Hill, Claudia Perlich, and Foster Provost. Discovering knowledge from relational data extracted from business news. In *SIGKDD-2002 Workshop on Multi-Relational Data Mining*, 2002.
- [Doreian and Stokman, 1997] Patrick Doreian and Frans Stokman. *Evolution of Social Networks*. Amsterdam etc.: Gordon and Breach, 1997.
- [Feld, 1997] Scott L. Feld. Structural embeddedness and stability of interpersonal relations. *Social Networks*, 19:91–95, 1997.
- [Karamon *et al.*, 2008] Jun Karamon, Yutaka Matsuo, and Mitsuru Ishizuka. Generating useful network-based features for analyzing social networks. In *Twenty-Third Conference on Artificial Intelligence (AAAI-08)*, pages 1162–1168, 2008.
- [Katz and Proctor, 1959] Leo Katz and Charles H. Proctor. The concept of conformation of interpersonal relations in a group as a time-dependent stochastic process. *Psychometrika*, 24(4):317–327, 1959.
- [McCulloh and Carley, 2009] Ian A. McCulloh and Kathleen M. Carley. Longitudinal dynamic network analysis. Casos technical report, Center for the Computational Analysis of Social and Organizational Systems, 2009.
- [Snijders, 1997] Tom A.B. Snijders. Models for longitudinal network data. In Peter J. Carrington, John Scott, and Stanley Wasserman, editors, *Models and Methods in Social Network Analysis*, pages 148–161. Cambridge University Press, New York, 1997.
- [Tetlock *et al.*, 2008] Paul C. Tetlock, Maytal Saar-Tsechansky, and Sofus Macskassy. More than words: Quantifying language to measure firms' fundamentals. *Journal of Finance*, 63(3):1437–1467, 2008.
- [Uzzi, 1997] Brain Uzzi. Social structure and competition in interfirm networks: The paradox of embeddedness. *Administrative Science Quarterly*, 42:35–67, 1997.
- [Wang and Chan, 2007] Jar-Long Wang and Shu-Hui Chan. Stock market trading rule discovery using pattern recognition and technical analysis. *Expert Syst. Appl.*, 33:304–315, August 2007.
- [Wasserman and Faust, 1994] Stanley Wasserman and Katherine Faust. *Social network analysis. methods and applications*. Cambridge University Press, Cambridge, 1994.
- [Xiao and Dasgupta, 2009] Li Xiao and Subhashish Dasgupta. The effects of dynamic it capability and organizational culture on firm performance: An empirical study. In *Proc. ICIS2009*, 2009.