MANCHESTER 1824

The University of Manchester

Ranking the Importance of Boards of Directors

Conyon, M. J. and Muldoon, M. R.

2006

MIMS EPrint: **2006.65**

Manchester Institute for Mathematical Sciences
School of Mathematics

The University of Manchester

Reports available from: http://eprints.maths.manchester.ac.uk/
And by contacting: The MIMS Secretary
School of Mathematics
The University of Manchester
Manchester, M13 9PL, UK

ISSN 1749-9097

Ranking the Importance of Boards of Directors

Martin J. Conyon[1]* and Mark R. Muldoon[2]

[1]The Wharton School, University of Pennsylvania, USA
[2]Department of Mathematics, University of Manchester, UK

August, 2004

Abstract

We measure the importance (centrality) of boards of directors using the PageRank algorithm from computational graph theory. PageRank is at the heart of the immensely successful Google web search engine and, we argue, can be naturally extended to social network settings. In this view, a board can be represented as part of an affiliation network or, in graph theoretic-terms, an undirected bipartite graph. But PageRank operates on directed graphs, so we develop a procedure to pass from an undirected bipartite graph to appropriately weighted, directed projections. Finally, we present the rankings of publicly traded US and UK firms using this method.

1 Introduction

Network researchers have long been interested in the structural properties of complex systems. A particularly salient stream of analysis has focused on the properties of vertices within networks.[1] A large number of so-called *centrality* measures have arisen which gauge the structural importance of a vertex relative to other vertices within a complex web of associations. Such centrality measures have been very useful in the field of social networks for understanding the roles played by different actors (Wasserman & Faust (1994)). In this paper we describe an alternative measure of centrality. It is the PageRank algorithm from computational graph theory originally described in Brin & Page (1998) and Brin, Page, Motwani & Winograd (1999). Our PageRank-based centrality measure has, as we will explain later in this paper, an interpretation in terms

*__Acknowledgments__: E-mail conyon@wharton.upenn.edu or M.Muldoon@umist.ac.uk with comments. We are indebted to Peter Cappelli, Lerong He, Simon Peck, Mike Useem and seminar particiapants at the Wharton School for comments and suggestions. We would like to thank the Leadeship Center of the Wharton School for financial support.

[1]In the Social Network literature a vertex is often referred to as an *actor*.

of a random walk—a random exploration of the social network. As such it is an addition to the circle of random-walk centrality measures recently proposed by Newman (2003*a*).

PageRank was designed for, and is at the heart of, the immensely popular Google search engine. The PageRank algorithm, as originally conceived, is a system designed to rank the overall importance of each page on the web—thus providing an index to order responses to a user's web search. The rank that Google uses can, in addition to its interpretation in terms of random walks, be interpreted as a measure of a web page's authority, overall importance, or influence: Higham & Taylor (2003) give a very clear exposition of both these interpretations, as well as an excellent introduction to the PageRank algorithm itself. Although the algorithm was originally designed to determine the importance of web pages, it can be naturally extended to other network settings, including social networks. In this paper our interest is in the social network of corporate governance.

The extension of PageRank to social networks, such as boards of directors, seems a fruitful avenue to pursue. A recurrent theme within the management and organization science literature has been to understand the power and influence of boards and directors as determined "structurally", that is, as it is determined by their positions as actors within a social network (for example see the review byPettigrew (1992)). The PageRank centrality measure seems to offer a fresh insight on this topic.

Previous research on power and influence in top management teams has focused on the interlocking directorship. (Mizruchi & Bunting (1981); Useem & Karabel (1986); Useem (1984); Pennings (1980)). A board interlock occurs when the boards of two separate organizations share a common director. The shared director, then, creates a link between the two boards. This link provides the channel by which organizations may potentially influence each other. Of course, many interesting issues arise when a link between two boards is created. For instance, how does one measure the directional flow of influence between two organizations? Also, how does one represent the fact that one organization might exert more influence than another? There is now a considerable literature that examines the formation of interlocks and evaluates their consequences. For instance, see the analysis and review by Mizruchi (1982) and Mizruchi (1996). To cite just a few research areas of interest: Hallock (1997) demonstrates that CEO pay is higher in interlocked firms after controlling for other economic determinants. Haunschild (1993) and Haunschild & Beckman (1998) shows that interlocks are important determinants of corporate acquisition activity. Davis & Greve (1997) show that interlocks are important in facilitating corporate governance changes. Overall, the extant literature points to the fact that corporate connectivity (interlocks) appear to matter significantly for corporate outcomes. One can easily recast the idea of a board "interlock" in terms of social network (or graph) theory[2] and interpret it as a *centrality* measure. We describe how in

[2] The central themes in social network analysis are comprehensively reviewed by Wasserman & Faust (1994). The theory of random graphs is given in Bollobás (2001). The application of

Section 2.

In this paper we propose measuring the importance (centrality) of the board of directors using a variety of related measures computed with the PageRank algorithm. The rank assigned to a board derives from two sorts of data: the first is a numeric measure intrinsic to the board and which one may choose to be the same for all boards, or, alternatively, which may be calibrated to reflect differences in firm attributes (such as market value, performance etc.). The second contribution is essentially a poll of the board's "neighbors" (those firms with which it is interlocked), weighted according to both the rank (i.e. esteem or influence) in which the neighboring board is held, as well as the relative "strength" of the connection between the boards. In essence, contributions from boards that are themselves held in high esteem count more towards one's own rank. A board's rank is defined recursively and, as we will show in Section 3, it can interpreted as a steady state consensus about which boards are deemed important.

As we mentioned above, PageRank is designed to operate on directed graphs (these and other graph-theoretic terms are defined carefully in Section 2), while the usual representation of the social network of corporate governance—as an affiliation network or bipartite graph—captures only symmetric, or non-directional relationships and so gives rise to undirected graphs. Accordingly, one of the challenges in the development of a measure one might term "BoardRank" is to find a way to incorporate extra information into the construction of the social network in such a way as to yield directed graphs. In Section 4 we propose such a method and, in Section 5, apply it to data on the corporate governance of firms in the United States and in the United Kingdom. Our procedure is not, however, limited to such data and should be useful to social network researchers who wish to develop PageRank-like centrality measures for arbitrary affiliation networks.

The data used in Section 5 are, for the US, a snap shot of the board memberships of approximately 1,700 publicly traded firms in 2003. The data for the UK, collected in 2002, are a snap shot of approximately 2,200 publicly traded firms. We document which firms receive the highest rank using the PageRank procedure. We also present a simple statistical model illustrating which firm-level factors (such as size, company performance etc.) help determine the PageRank of a firm.

The rest of this paper is organized as follows. In Section 2 we introduce some notions from the mathematical theory of graphs as applied to social networks and boards of directors. Then, in Section 3, we discuss the PageRank algorithm and touch on its application to social networks. In section 4 we describe a method to pass from unweighted bipartite graphs to weighted, directed projections. In section 5 we apply PageRank to board data from the United States & the United Kingdom and, finally, in section 6 we offer some concluding remarks.

graph theory to social networks, such as boards of directors, is comprehensively reviewed by (Newman 2003*b*).

2 Corporate Boards and Social Networks

In this section we define terms and briefly review some important features of graph theory as applied to social networks. A much more extensive discussion of graphs and their representation, manipulation and application to the social sciences appears in Wasserman & Faust (1994). Newman (2003*b*) gives an excellent review of recent developments in the field of complex systems. Newman, Watts & Strogatz (2002) discuss specifically the application of random graphs to social networks.

2.1 Basic terminology

A *network* (or *graph*) is a set of items termed *vertices* (or *nodes*) with connections between them called *edges*. In discussions about graphs representing social networks the nodes are sometimes also called *actors*. We will restrict our attention to networks derived from the world of corporate directorship and adopt the following conventions: our nodes will be of two types, either boards or the directors who sit on them. Edges will represent, among other things, membership of a board (in one sort of graph) or an interlock between boards (in another, related sort of graph).

The latter relationship, an interlock due to a shared director, is clearly a symmetric one: if board A is interlocked with board B, then B is automatically interlocked with A as well. This sort of reciprocal connection will be represented by an *undirected* edge. But we will also need the notion of a *directed* edge, which will represent a unidirectional connection. A typical example of such a relationship is "has influence on": one can easily imagine a setting in which firm A has influence on firm B—perhaps because A is a major shareholder in B—but this connection is not symmetric: there is no reason to imagine that B has influence on A. Both directed and undirected edges can also carry *weights*. In social networks these weights are usually a measure of the relative strength of the connection the edge represents and so a weight of zero is often taken to mean that the edge does not exist—that the corresponding connection is absent. We will use weights of this kind in PageRank algorithm (see Section 3), but we will also want consider another sort of weight—something one might term a "Boolean weight"—whose value is somewhat akin to that of a dummy variable in statistical modelling. These weights, whose role is discussed in Section 3.1, assume the values zero and one, but a Boolean weight of zero does not indicate the absence of the corresponding edge.

Finally, a graph whose every edge is directed is called a *directed graph*. If, in addition, all the edges have weights the graph is said to be a *weighted, directed graph*. Alternatively, a graph (such as the one pictured in Figure 1) in which all the edges are undirected and none of them have weights is an *unweighted, undirected graph* or, for short, an *undirected graph*.

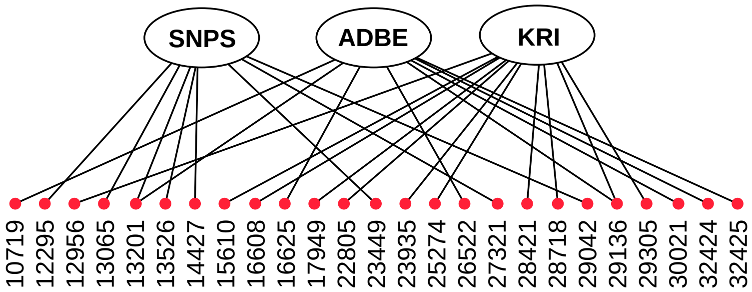

Figure 1: *A undirected, unweighted bipartite graph representing the boards of directors of Adobe Systems (ADBE) and Synopsis (SNPS), both software houses, and that of the Knight Ridder (KRI) chain of newspapers.*

2.1.1 Notation for edges and weights

Throughout this paper we will write $e_{j,k}$ to refer to a directed edge that connects vertex k (at the tail) to vertex j (at the tip). When there is a weight associated with the edge we will call it $w_{j,k}$. We will use similar notations, $E_{j,k}$ and $W_{j,k}$, for *undirected* edges. Of course, in this latter case $E_{1,2}$ is the same as $E_{2,1}$ and so $W_{1,2} = W_{2,1}$.

2.1.2 Paths, length and connectedness

A directed graph is said to contain a *path* between vertices a and b if it contains is a sequence of edges e_{a,k_1}, e_{k_1,k_2}, $\ldots$, $e_{k_l,b}$. That is, there is a path from a to b if one can get from one to the other by moving along the directed edges of the graph. A similar definition applies to undirected graphs. In either case, the *length* of the path is the total number of edges involved.

In an undirected graph two vertices are *connected* if there is a path between them and the *connected component* associated with a vertex is that part of the graph consisting of the vertex itself and all those others that can be reached by paths running along the edges of the graph. In a directed graph the notions of connectivity are slightly more complicated: node a is said to be *reachable* from node b if there is a path from b to a. But a has two, possibly distinct, connected components: those nodes reachable from a and those from which a can be reached.

2.1.3 Degree

The concept of *degree* will prove important. At it's simplest, in an undirected graph, a node's degree is the just number of edges connected to it.[3] By contrast, a node j appearing in a directed graph has both an *in degree* (the number of

[3]The notion of board "interlocks", discussed in Section 1, is simply the degree of the vertex. It is akin to the "degree centrality" measure introduced by Freeman (1978).

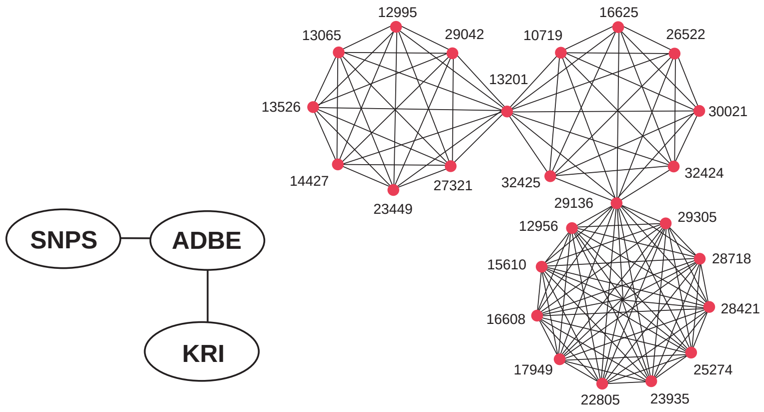

Figure 2: *The undirected graphs showing board interlocks (left) and the network of co-directorship as derived from the affiliation network illustrated in Figure 1.*

directed edges having j at their tips) and an *out degree* (the number of edges with j at their tails). These notions are generalized further, to accommodate weighted edges, in section 3 below.

2.2 Graphs of boards and directors

Data about boards of directors present an immediate problem: how should one draw a graph to represent it? The issue is that one could treat the board as the basic unit of analysis and form a graph whose vertices represent boards and whose edges represent interlocks (that is, shared directors). But alternatively, one could focus on the director and make a (generally much larger) graph whose vertices represent directors and whose edges represent shared board memberships. There is no obvious way to choose between these two representations and many authors simply analyze both. In Section 5 we will focus on the board-and-interlock graph, though our methods are equally applicable to the graph in which directors are vertices.

But the ambiguity about representation of the corporate world arises from the structure of the data: there really are two sorts of social entities here, the directors and the boards, and the network's edges represent membership of the former in the latter. The most natural representation of such a network, sometimes called an *affiliation* network, is a graph with two sorts of vertices—one each for boards and directors—that has edges connecting directors with the boards on which they sit. The result is an example of a *bipartite* graph: one whose vertices can be divided into two distinct sets and whose edges only make connections between the two sets. Figures 1 and 3 are examples. The board and director graphs mentioned above now appear as "projections" of the bipartite graph onto one of its two sets of vertices.

Figures 1 and 2 illustrate these issues for that part of social network of corporate governance connected to the board of Adobe Systems Inc., a software house. The former, Figure 1, shows the full bipartite graph, while Figure 2 shows the two projections. The graph appearing at the left of Figure 2 has the boards as its nodes and edges connecting interlocked boards while the graph at right has directors for nodes and includes an edge between two directors if they sit on the same board. Note that both of the projections are undirected graphs. But PageRank is, as we will see below, designed to operate on directed graphs. One of the main technical obstacles in adapting PageRank to the ranking of, for example, the boards, is to find a way to incorporate extra information into a bipartite graph such as the one in Figure 1 in such a way as to permit the construction a board projection that is a directed graph. This issue is touched on in the following section, then treated in detail in Section 4.

3 PageRank for Boards of Directors

The PageRank algorithm (Brin & Page (1998) and Brin et al. (1999)) assigns a numerical rank to each vertex in a directed graph. These ranks were originally intended as an aid to searching the World Wide Web and so have a natural interpretation in a graph whose vertices represent web pages and whose (directed) edges represent hyperlinks. PageRank then provides an assessment of the importance (or authority) of a Web page (node) that is largely independent of the page's content.

Here we describe a scheme that generalizes PageRank to the ranking of boards in the social network of corporate governance. The problem splits naturally into two pieces: deriving weighted, directed projections from the unweighted, undirected, bipartite graph that represents the affiliation network of boards of directors and (ii) computing PageRanks for the two projections.

3.1 Adding weights to a bipartite graph

The top panel of Figure 1 shows a small part of the world of the corporate governance in the US. The edges connecting boards to their directors in this graph are all the same: they are unweighted or, equivalently, all have the equal weight. More generally, one could assign a (non-negative) weight to each such edge. One might, for example, assign a Boolean weight of zero[4] to every edge that connects a non-executive director to the board on which he or she serves and a weight of one to those edges that represent the connections between executive directors and their boards. One can represent this graphically with something like Figure 3: in Section 4 below we will introduce a scheme that processes these weights (or any others one might choose to assign) and produces a pair of weighted, directed projections.

[4] As mentioned above, this is a somewhat non-standard use of the term "weight". In normal graph-theoretic usage an edge with zero weight is simply absent from in the graph, but here we wish to suppress the usual connection between an edge's weight and its existence. In the bipartite graph it will be possible for edges to exist, but have zero weight.

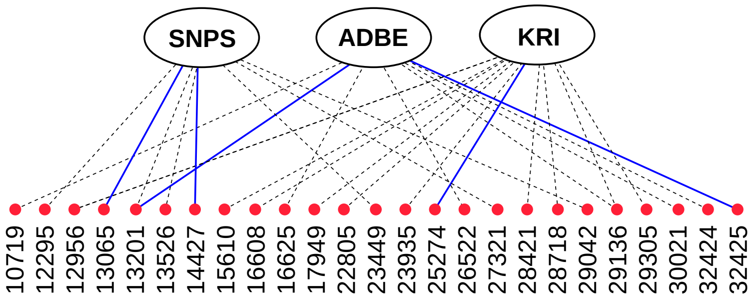

Figure 3: *Another representation of the bipartite graph in Figure 1, but now with edges weighted according to whether the director is an executive (weight is 1.0, plotted with a heavy blue edge) or a non-executive (weight is 1.0, plotted with a lighter dashed edge).*

3.2 PageRank

Here we offer a brief account of the calculations involved in generating PageRanks. For a longer discussion one might begin with the very clear and entertaining introduction by Higham and Taylor (Higham & Taylor 2003), then proceed to the original references, Brin et al. (1999) and Brin & Page (1998). A more general analysis of link-based algorithms, including PageRank, appears in Wang (2004).[5]

Suppose that one has a weighted, directed graph with N vertices and an adjacency matrix W whose entries are conventional, positive weights $w_{j,k} \geq 0$ satisfying

$$
\begin{aligned}
w_{j,k} > 0 &\qquad \text{if a (directed) edge connects vertex } k \text{ (tail) to vertex } j \text{ (tip);} \\
w_{j,k} = 0 &\qquad \text{if there is no edge between } j \text{ and } k.
\end{aligned}
$$

Then define o_k, the *weighted out-degree* of vertex k, to be

$$o_k = \sum_{j=1}^{N} w_{j,k}. \tag{1}$$

That is, o_k is the sum of the weights of all the edges that reach from k to some other vertex.

The PageRank algorithm uses these quantities, as well as an adjustable parameter $0 \leq d < 1$, to define a converging sequence of ranks r_j^n. Here the subscript j ranges over the vertices in the graph (that is $1 \leq j \leq N$) while the superscript n starts from zero and counts the number of times one has applied

[5]Other recent papers examining link-based algorithms and social networks include Ding, Zha, He, Husbands & Simon (2004) and Diligenti, Gori & Maggini (2004)

the PageRank update rule:

$$r_j^{n+1} = (1-d)s_j + d\sum_{k=1}^{N}\left(\frac{w_{j,k}}{o_k}\right)r_k^n \tag{2}$$

The quantities s_j are, for reasons that will become clear shortly, sometimes called *source strengths*: they should be positive and satisfy

$$N = \sum_{j=1}^{N} s_j. \tag{3}$$

The sum over k appearing in (2) is best thought of as a sum over (incoming) neighbors or, in graph-theoretic terms, *predecessors*, for only those edges that start at some vertex k and end at j will contribute to the PageRank of the j-th vertex (that is, only these edges will have weights $w_{j,k} \neq 0$).

Although the iterative rule (2) tells us how to generate r_j^{n+1} given r_j^n, it cannot tell us how to start the process. By convention one chooses $r_j^0 = 1\ \forall j$. That is, all the vertices start out with equal rank, then successive applications of (2) generate successive generations of ranks. It is not too hard to show that, provided $d < 1$, this iterative procedure will converge. That is, eventually it will be true that $r_j^{n+1} \approx r_j^n$. Indeed, something much stronger is true: given any small number $\delta << 1$, one can always choose an $n_\star$ sufficiently large that

$$\max_j ||r_j^{n_\star+t} - r_j^{n_\star}|| \leq \delta \qquad \forall t > 0.$$

In words, after $n_\star$ iterations *all* subsequent generations of ranks will be within δ of the $r_j^{n_\star}$. In practice, we choose some small tolerance δ and repeats (2) until the first n for which $\max_j ||r_j^{n+1} - r_j^n|| \leq \delta$.

3.2.1 Interpretation as a random walk

The intuitive idea behind PageRank is clearest in the algorithm's original context, the World Wide Web. Imagine a deeply indecisive individual who browses the web at random. He reads a web page, then chooses a new page at random by clicking, with equal probability, on any of the links on the page he has just finished (ignore, for the moment, the possibility that his page has no hyperlinks). If such a reader persisted in his efforts he would eventually visit a very large proportion of the web's pages, most of them many times over (also ignore the fact that pages are continually being added to the web or removed from it). If he kept a list of all the pages he ever visited and also kept track of how often he visited each one then, eventually, the ratio

$$\frac{\text{Number of visits to page } j}{\text{Total number of pages visited}} \tag{4}$$

would tend to a constant, converging in a manner reminiscent of the PageRanks.

Indeed, the ratio (4) would converge precisely to r_j/N where N is the number of pages in the web and is the PageRank r_j produced by applying (2) with $d = 1$. The parameter d is, in this view, a probability: it is the probability that our random reader, having finished his page, decides to proceed as described above. As an alternative, if $d < 1$, we could permit our reader to jump, with probability $(1 - d)$, to an arbitrary new page anywhere in the web, choosing the j-th page with probability (s_j/N). Here, N is the number of pages in the web and s_j is the source strength of the j-th page. If he pursues this mixed strategy—sometimes choosing a random link from the current page, sometimes jumping arbitrarily—then the ratio (4) will tend to r_j/N where r_j is the PageRank produced by repeated application of (2).

It is not hard to recast this view of PageRank to make it appear relevant to social networks. Suppose that instead of a random browser wandering idly through the web, we consider an item of news, information or gossip being relayed randomly along the connections of a social network. Each actor in the network, upon hearing the news, either passes it (with probability d) to a randomly chosen associate or (with probability $(1 - d)$), relays the news to some arbitrary third party (perhaps by posting a notice in some public place or writing a newspaper article). Although this is a drastically abstracted account of the propagation of information, it's not wholly implausible to imagine that it could capture some aspects of the diffusion of ideas as viewed in the large.

3.2.2 Interpretation as a weighted voting scheme

PageRank admits another interpretation suggestive of applications to social networks. If we consider a page's (or a board's) rank to be a measure of the esteem in which it is held, then the terms in the update rule (2) have natural interpretations. The source term $(1 - d)s_j$ represents a sort of natural, or intrinsic component of esteem: one may, democratically, set $s_j = 1$ for all boards or, if it seems more appropriate, assign some boards—perhaps those of particularly virtuous, innovative or profitable firms—a higher intrinsic esteem. The second term, the one involving the sum over k, is essentially a poll of the j-th node's neighbors, weighted according to both r_k^n, the esteem in which the neighboring k-th vertex is held, and the relative strength of the edge connecting vertex k to vertex j (this is the factor $w_{j,k}/o_k$, which compares the weight of the edge connecting k to j with the total strength of k's outgoing edges). That is, praise from the praiseworthy—contributions from nodes that are themselves highly esteemed—counts more.

In this view the iteration of the PageRank update corresponds to the gradual formation of a consensus about which nodes are most important. At the outset every node has the same rank, 1. Successive rounds of (2) redistribute esteem, treating a weighted, directed edge from k to j as a weighted vote of confidence by node k in favor of node j. Thus it is clear that the weights assigned to the edges have considerable influence on the final distribution of rank: in the next section we will describe a method for generating weights on the edges of the two projected networks from Boolean weights—the sort of Boolean weights

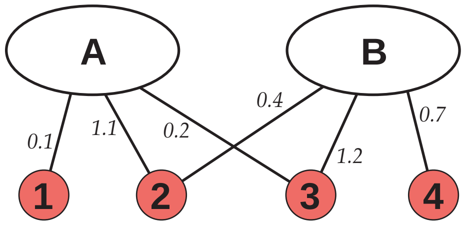

Figure 4: *A weighted, undirected bipartite graph in which directors 1–4 serve on boards A and B.*

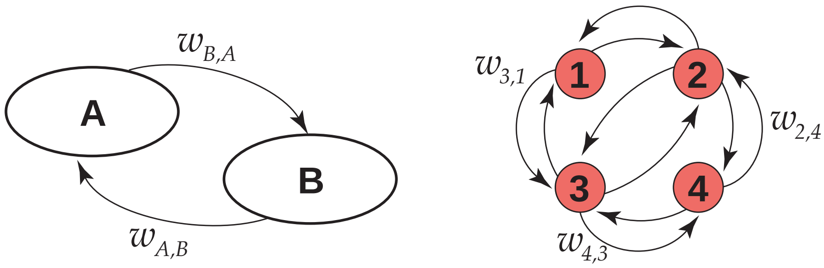

Figure 5: *Directed board (left) and director (right) projections derived from the bipartite graph of Figure 4. Although all the edges in the directed graph at right have associated weights, only a few are labelled explicitly.*

discussed in Section 3.1—on the edges of the bipartite graph.

4 Making weighted, directed projections

Consider Figure 4, which shows a small, weighted, undirected bipartite graph. Ignoring for the moment the question of how one assigns weights to the edges, this figure is an example of the sort of data from which one might hope to derive PageRank-like measures for the boards, directors, or both. In this section we develop a method to pass from graphs like that pictured in Figure 4 to the sort of weighted, directed projections appearing in Figure 5.

4.1 Two preliminary attempts and a formula

Here we build up gradually, by way of two intermediates, to our preferred formula for the weights of edges in the bipartite graph. The main observation is that edges in the projection arise from two-edge paths in the bipartite graph. So, for example, a pair of boards A and B are connected in the board projection if they share a director. But this is the same as saying that the bipartite graph contains a pair of undirected edges, say, $E_{A,k}$ and $E_{k,B}$, where the shared director has index k. It is natural to choose the weights for the two directed edges in the projection to be a linear combination of the weights in the bipartite graph

$$\begin{aligned} w_{A,B} &= \beta W_{B,k} + (1-\beta) W_{k,A} \\ w_{B,A} &= \beta W_{A,k} + (1-\beta) W_{k,B} \end{aligned} \tag{5}$$

where we have introduced a new parameter $0 \leq \beta \leq 1$ that controls the relative contribution of the two weights from the bipartite graph. This is our first preliminary attempt. Notice that the two lines above are really the same formula with the roles of A and B interchanged, so one need only state one of them. Note also that the sum of the weights is conserved. That is

$$\begin{aligned} w_{A,B} + w_{B,A} &= [\beta W_{B,k} + (1-\beta) W_{k,A}] + [\beta W_{A,k} + (1-\beta) W_{k,B}] \\ &= [\beta W_{A,k} + (1-\beta) W_{k,A}] + [\beta W_{B,k} + (1-\beta) W_{k,B}] \\ &= [\beta + (1-\beta)]\, W_{A,k} + [\beta + (1-\beta)]\, W_{B,k} \\ &= W_{A,k} + W_{B,k} \end{aligned}$$

where, in passing from the second line to the third, we have used the fact that, as the bipartite graph is undirected, $W_{A,k} = W_{k,A}$ and $W_{B,k} = W_{k,B}$.

Of course, as Figure 4 shows, two boards may share more than one director so one might generalize (5) to

$$w_{A,B} = \sum_{k \text{ shared}} \beta W_{B,k} + (1-\beta) W_{k,A} \tag{6}$$

where the sum runs over all shared directors: this is our second preliminary attempt. It retains the property that the sum $w_{A,B} + w_{B,A}$ is the same as the sum of all the weights on the edges that contribute to the formation of $e_{A,B}$ and $e_{B,A}$ in the projection.

Our second formulation, (6), is reasonably satisfactory, but in practice we prefer to make a slight modification:

$$w_{A,B} = \sum_{k \text{ shared}} \frac{\beta W_{B,k} + (1-\beta) W_{k,A}}{(W_{B,k} + W_{k,A})/2} \tag{7}$$

That is, we scale the contribution from each pair (in the bipartite graph) by the average of its weights. This scaling means that each shared director causes the sum $w_{A,B} + w_{B,A}$ to increase by 2.

This choice of scaling arises naturally in the analysis of Boolean-weighted bipartite graphs such as the one pictured in Figure 3. The aim of the rescaling is to permit a distinction between edges that don't exist in the bipartite graph (and so don't give rise to connections in the projections) and those that do exist (and so should contribute to connections in the projection), but have zero weight. Such graphs are discussed at greater length below, but we conclude this section with Table 1, which gives sets of weights for the edges in the projections pictured in Figure 5.

Director Projection

Weight	*Using (6)*	*Using (7)*
$w_{2,1}$	0.85	1.4167
$w_{3,1}$	0.175	1.1667
$w_{1,2}$	0.35	0.5833
$w_{3,2}$	1.425	1.9038
$w_{4,2}$	0.625	1.1364
$w_{1,3}$	0.125	0.8333
$w_{2,3}$	0.1475	2.0962
$w_{4,3}$	0.825	0.8684
$w_{2,4}$	0.475	0.8636
$w_{3,4}$	1.075	1.1316

Board Projection

Weight	*Using (6)*	*Using (7)*
$w_{A,B}$	1.375	1.8762
$w_{B,A}$	1.525	2.1238

Table 1: *Weights derived by applying (6) or (7) to the weighted bipartite graph shown in Figure 4. In both cases we used* $\beta = 0.25$.

4.2 Boolean weights: insiders and outsiders

Our original interest was to analyze weighted bipartite graphs such as the one illustrated in Figure 3. As we mentioned above, the weights here are somewhat unusual in that a weight of zero does not imply that the edge is absent. Rather, we imagine that the weights are Boolean variables: they reflect the answer to some "Yes"-"No" question such as "Is the director an executive director of the board to which she is connected?" One might refer to such graphs as *Boolean weighted bipartite graphs.*

Figure 6 shows three of the simplest possible such graphs: each contains two boards tied together by a single shared director. To illustrate the role of the parameter β let us compute the weights in the board projection using our preferred rule (7). The leftmost graph presents an immediate difficulty since:

$$\begin{aligned} w_{A,B} &= \sum_{k \text{ shared}} \frac{\beta W_{B,k} + (1-\beta) W_{k,A}}{(W_{B,k} + W_{k,A})/2} \\ &= \frac{\beta W_{B,1} + (1-\beta) W_{1,A}}{(W_{B,1} + W_{1,A})/2} \\ &= \frac{\beta \times 0 + (1-\beta) \times 0}{(0+0)/2} \\ &= 0/0. \end{aligned}$$

One way out of this problem is to adhere to the principle that each shared director should contribute 2 to the sum $w_{A,B} + w_{B,A}$: combining this with the observation that, in this problematic case, $W_{A,1} = W_{B,1}$ we'll adopt the convention that $w_{A,B} = w_{B,A} = 1$. This has a natural generalization to the case where the two boards share several directors.

The remaining cases are easier—the formula (7) yields a sensible result without any further contemplation—and the results are summarized in Table 2.

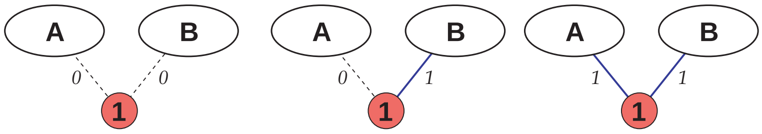

Figure 6: *Three small bipartite graphs with various sets of Boolean edge weights. The weights are shown numerically and, additionally, the graphs are colored with the conventions from Figure 3.*

Bipartite Graph		*Board Projection*	
$W_{A,1}$	$W_{B,1}$	$w_{A,B}$	$w_{B,A}$
0	0	1	1
0	1	2β	$2(1-\beta)$
1	1	1	1

Table 2: *Weights derived by applying (7) to the three weighted bipartite graphs shown in Figure 6.*

These results make qualitative sense in that if $W_{A,1} = W_{B,1}$, then $w_{A,B} = w_{B,A} = 1$. The more interesting case is when the director is an insider on only one of the two boards.

Suppose, for example, that the director is an outsider on board A, but an insider on board B. In this case the directed edges in the board projection receive different weights that depend on the parameter β. If $\beta \approx 0$ then $w_{A,B} \approx 0$ and $w_{B,A} \approx 2$, so the edge pointing from A to B is much more heavily weighted than the one running from B to A. In this case one might like to think of the heavier edge as indicating that board A is showing "esteem" for board B by recruiting one of B's executives. Alternatively, one might think of the heavily-weighted edge as indicating "influence": board A has a strong possibility of influencing board on B because one of A's directors is involved in the day-to-day running of B. When $\beta = 0.5$ both edges receive the same weight: $w_{A,B} = w_{B,A} = 1$ and, finally, when $\beta \approx 1$, then $w_{A,B} \approx 2$ and $w_{B,A} \approx 0$ and the imputations about esteem and influence run in the opposite direction. The relationships between β and the weights on the directed edges are summarized in Figure 7.

5 Ranking boards of directors

5.1 The data

To implement the PageRank algorithm for the social networks of boards of directors we use two distinct data sets. One is from the United States, kindly supplied by the *Corporate Library* and the other is data set is from the United Kingdom supplied by *Hemmington Scott* publishing. Both data sets contain an

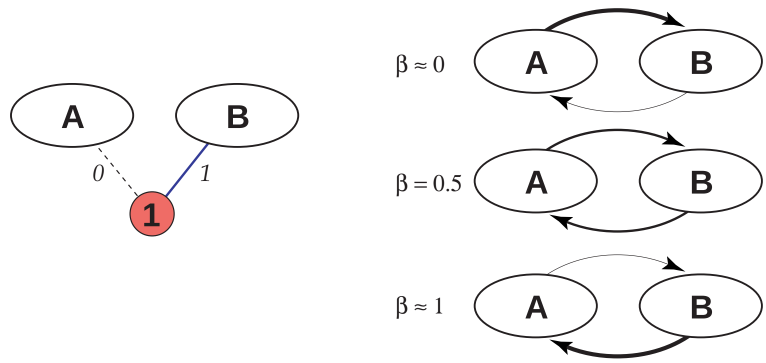

Figure 7: *At left, a shared director is an insider on board B, but an outsider on board A. The three projections at right show how the parameter β influences the weights of the directed edges arising from this interlock.*

expansive list of companies and the directors who serve on their boards.

The data for the United States are a snapshot (i.e. a cross section) of publicly traded US firms at February 2003. The data consist of 12,765 directors sitting on 1,731 boards. We report results for the largest connected component, which consists of 10,432 directors sitting on 1,452 boards. The United Kingdom (British) data are a snapshot of publicly traded UK firms at March 2002. These data describe 11,541 directors sitting on 2,236 boards. Once again we report results for the largest connected component, which here consists of 8,850 directors sitting on 1,732 boards.

5.2 Board ranks

The importance of each board (firm) is calculated using the PageRank update rule in Equation 2. We consider a Boolean weighted bipartite graph discussed in Section 3 where "yes" was the answer to "Is the director an executive director of the board to which she is connected?" We used a number proportional to the firm j's market capitalization as the source strength, s_j, in Equation (2). Of course, PageRanks for boards can be calculated for various values of the adjustable parameter, d: we set $d = 0.7$. [6] We then choose values of $0 < \beta < 1$, the relative contribution of the weights in the bipartite graph. Specifically, we examined $\beta = 0.5$, $\beta = 0.1$ and $\beta = 0.9$. In Tables 3 and 4 we report the 35 highest ranked board (firms) in the United States and the United Kingdom (where $d = 0.7$ and $\beta = 0.5$).

In Table 3 General Electric, Microsoft, Exxon Mobil, Pfizer and Wal-Mart

[6] The choice of $d = 0.7$ may seem arbitrary but choosing $d = 0.85$, a value typically mentioned in discussions of PageRank, does not qualitatively affect the rankings reported here. For example, in the US data the correlation coefficient between the board ranks calculated separately for $d = 0.7$ and $d = 0.85$, and using $\beta = 0.5$, is 0.97.

turn out to be the 5 highest ranked boards in the United States while, in Table 4, BP, GlaxoSmithKline, Vodafone Group, Lloyds TSB Group, and HSBC Holdings turn out to be the 5 top ranked boards in the United Kingdom. Our web sites (see Muldoon & Conyon (2004)) contain the complete list of board ranks for the firms in the largest connected component of our US and UK data sets. How are we to evaluate these results? The PageRank algorithm, applied to the board projection from a bipartite graph, gives the consensus (steady-state) solution as to which boards attract something that, in the informal motivation above, we termed "esteem". Of course, our choices of d, β and the source strengths are illustrative. We would encourage further investigation based on alternative d and β combinations and the selection of different sources strengths (e.g. profits or employees).

But the "esteem" we measure is in a certain narrow sense a structural feature of the corporate world: it depends simply on the list of companies, the membership of their boards and their market capitalizations. In this sense our measure of esteem contrasts to other measures which rely upon the judgments of people to say whether a board is highly regarded or not, or whether the board itself promotes itself as an ideal board. Because our centrality measure does not rely on self- or other-assessment, but computes a rank based mainly on board interlocks, it is in this respect harder to manipulate and less dependent on arbitrary judgments. This is a property inherited from Page and Brin's original PageRank, which is, by design, link-based rather than content-based.

Our rankings agree well with those reported in *Fortune* magazine's annual list of America's "Most Admired Companies". Each year *Fortune* asks a panel of executives, directors and security analysts to rank a firm according to eight criteria: innovation, employee talent, use of corporate assets, social responsibility, quality of management, financial soundness, long-term investment value, and quality of products and services. For the "top ten" survey, respondents are asked to select the ten companies they admire most in any industry. They chose from a list of corporations that ranked in the top 25% overall last year, plus any that finished in the top 20% of their category. Six of the *Fortune* top 10 firms appear in the list in Table 3. These are Wal-Mart, General Electric, Dell Inc., Microsoft Corp., Johnson and Johnson and IBM. Also, *Fortune* produces a global "World Most Admired Companies": all of the British companies that appear in the 2004 *Fortune* global top 50 also appear in Table 4.

Our rankings can also be compared to the "Governance Index" introduced by Gompers, Ishii & Metrick (2003). Using data from the *Investor Responsibility Research Center*, they identify 24 distinct corporate governance provisions. These include poison pills, director indemnification, golden parachutes, classified / staggered boards, anti-greenmail etc. For each firm they add one point to the index for every provision that restricts shareholder rights, or equivalently increases managerial power. This (inverse) measure of governance quality potentially ranges from zero to twenty-four for each firm. They identify IBM, Wal-Mart, PepsiCo, American International Group as firms with a low "G-Index". Again, these appear in Table 3.

Table 3: *America's highest ranked boards. The Rank is based on the implementation of the PageRank algorithm described in the text where* $d = 0.7$*, source strengths are market capitalization and the* β *values are as indicated in the table*

Company	Rank $\beta = 0.5$	Rank $\beta = 0.1$	Rank $\beta = 0.9$
General Electric Company	15.60	14.97	16.16
Microsoft Corporation	15.46	15.08	15.56
Exxon Mobil Corporation	14.50	14.77	13.99
Pfizer, Incorporated	13.98	13.89	13.92
Wal-Mart Stores, Incorporated	12.80	12.73	12.85
Citigroup, Incorporated	12.79	12.36	13.15
American International Group, Incorporated	10.58	10.23	9.94
Verizon Communications Incorporated	10.41	10.42	10.31
Johnson & Johnson	9.55	9.57	9.48
Coca-Cola Company (The)	9.14	8.81	9.47
Procter & Gamble Company (The)	9.09	9.51	8.64
International Business Machines Corporation	8.69	8.24	9.04
J.P. Morgan Chase & Co.	8.38	8.16	8.71
SBC Communications Incorporated	8.33	8.49	8.17
Bank of America Corporation	8.20	8.28	8.07
Merck & Co., Inc.	8.05	8.68	7.23
Cisco Systems, Incorporated	7.25	7.18	7.05
Fannie Mae	7.22	7.67	6.68
AOL-Time Warner, Incorporated	7.01	6.89	6.94
Viacom, Incorporated	6.65	6.29	6.84
Dell Computer Corporation	6.57	6.41	6.75
Wells Fargo & Company	6.57	6.60	6.49
ChevronTexaco Corporation	6.37	6.32	6.38
PepsiCo, Incorporated	6.28	6.20	6.35
Intel Corporation	6.26	6.38	6.13
Altria Group, Inc.	6.20	6.32	6.21
Eli Lilly & Company	5.96	6.41	5.26
Anheuser-Busch Companies	5.82	5.83	5.86
Home Depot, Inc. (The)	5.79	5.47	6.13
3M Company	5.66	5.30	5.99
Morgan Stanley	5.37	5.20	5.56
BellSouth Corporation	5.23	4.98	5.31
Amgen, Incorporated	5.20	5.47	4.92
Bristol-Myers Squibb Company	5.10	5.10	5.14
Allstate Corporation (The)	4.92	4.59	5.27

Table 4: *Britain's highest ranked boards. The Rank is based on the implementation of the PageRank algorithm described in the text where d = 0.7, source strengths are market capitalization and the β values are as indicated in the table*

Company	Rank $\beta = 0.5$	Rank $\beta = 0.1$	Rank $\beta = 0.9$
BP PLC	66.44	63.83	70.20
GlaxoSmithKline PLC	46.96	48.39	45.56
Vodafone Group PLC	36.16	36.49	35.19
Lloyds TSB Group PLC	33.51	33.58	30.97
HSBC Holdings PLC	33.24	34.61	31.85
Shell Transport and Trading Co PLC	25.04	26.18	24.22
AstraZeneca PLC	24.02	27.18	20.49
Royal Bank of Scotland Group (The) PLC	23.84	24.05	23.13
Rio Tinto PLC	21.38	20.12	19.67
Unilever	19.01	23.69	15.26
Diageo PLC	17.71	13.74	20.76
Barclays PLC	15.49	14.29	16.37
Reuters Group PLC	13.23	10.94	15.66
Anglo American PLC	13.16	13.82	12.61
Schroders PLC	12.72	12.77	12.41
HBOS PLC	12.57	12.14	12.53
BT Group PLC	12.53	10.64	14.67
Reckitt Benckiser PLC	12.38	12.95	10.47
Six Continents PLC	11.98	18.32	5.50
Prudential PLC	11.30	11.61	10.86
Standard Chartered PLC	10.60	7.74	11.17
Rolls-Royce PLC	10.31	14.46	7.84
Trinity Mirror PLC	10.14	11.25	9.27
BAA PLC	9.45	9.13	8.32
Johnson Matthey PLC	9.42	9.61	8.84
Boots Company (The) PLC	9.41	7.61	10.79
Legal & General Group PLC	9.18	8.92	9.38
British Airways PLC	9.07	8.52	9.19
Invensys PLC	8.87	9.50	8.30
Allied Domecq PLC	8.72	9.38	8.32
Cable and Wireless PLC	8.39	8.26	7.48
Marconi PLC	8.29	8.18	8.14
Close Brothers Group PLC	8.08	8.19	7.69
Smiths Group PLC	7.81	4.16	9.16
British Sky Broadcasting Group PLC	7.68	7.15	7.31

5.3 A simple "board rank" model

Having computed a rank for each board (vertex) in the social network, a natural question arises: What factors lead to, or are associated with, a high page rank? Here we only briefly investigate this question. To do so, we estimate a simple statistical model where the outcome variable is the rank of the j-th board (i.e. it's "Board Rank"). We identified a set of observable firm-level variables, described below, that might be thought to influence the rank of a board.

At this stage we can only estimate the model for the USA, as the necessary data for the UK were not available to us. We supplemented the US *Corporate Library* data with a secondary firm-level data set from Standard & Poors *Execucomp* database, which tabulates many potentially useful characteristics for each firm. We used the August 2004 release, which contains company information for fiscal year 2003, such as compensation, firm size, sector, etc. After combining the two data sets, we estimated the following simple linear model for the available data:

$$(\text{Board Rank})_j = \alpha + \gamma_1 x_{1j} + \epsilon_j \tag{8}$$

where x_1 is a matrix containing the following covariates:

(i) The degree of the vertex in the board projection. That is, the number of other boards with which a given board is interlocked.

(ii) The size of the firm, measured as the log of total sales.

(iii) The total compensation received by the firm's CEO. This is measured as the sum of salary, bonus, other payments and the Black-Scholes value of options granted during the fiscal year.

(iv) Firm performance measured as the five-year total return to shareholders (including reinvested dividends)

(v) The proportion of outsiders on the main board

(vi) The size of the board.

(vii) The "Governnace Index" defined by Gompers et al. (2003) and available through the *Investor Responsibility Research Center*.

(viii) A set of 64 separate industry dummy variables. These are defined at the 2-digit standard industrial classification level.

Finally, γ_1 is the parameter to be estimated and ϵ_j is a stochastic error term. The variance covariance is made stationary (i.e. robust to arbitrary heteroscedasticity) using the method of Huber (1964) and White (1980).

The results are contained in Table 5. Column (1) provides the means of the independent variables. Column (2) to (4) contain the results from the estimation. Each model in (2) through (4) is estimated under different assumptions about the weighting parameter β. The results indicate that firms with a higher

Table 5: *Board Rank model: Estimation of Equation 8. The model is estimated using United States data for varying levels of β as specified. Note, + is significant at the 10% level, $*$ significant at 5% and $**$ significant at 1%.*

	(1) Variable Mean	(2) $\beta = 0.5$ Influence	(3) $\beta = 0.1$ Influence	(4) $\beta = 0.9$ Influence
Boards degree	7.08	0.14**	0.14**	0.14**
		(0.01)	(0.01)	(0.01)
Log(Sales)	7.21	0.24**	0.24**	0.24**
		(0.05)	(0.05)	(0.05)
CEO pay	5.35m	0.04**	0.04**	0.04**
		(0.01)	(0.01)	(0.01)
Stock returns	7.01	0.01	0.01	0.02
		(0.02)	(0.02)	(0.02)
Proportion outsiders	0.66	-0.17	-0.23	-0.11
		(0.23)	(0.22)	(0.24)
Board size	9.97	0.04*	0.04*	0.04*
		(0.02)	(0.02)	(0.02)
Governance index	9.04	-0.08**	-0.08**	-0.08**
		(0.01)	(0.01)	(0.01)
Industry dummies		Yes	Yes	Yes
Observations		1263	1263	1263
R-squared		0.68	0.69	0.68

board degree centrality, greater sales, greater CEO compensation, and board size are likely to have higher recorded board ranks according to our method. The governance index (an inverse measure of quality) is correctly signed and significant. The stock returns and the proportion of outside directors variables are insignificant. We expected these indicators of corporate governance quality to contribute positively to board influence (rank). In separate regressions (not tabulated here) we found that re-estimating Equation 8, but excluding the board degree variable, resulted in positive and significant coefficient estimates for the proportion of outsiders. We would, therefore, encourage further modelling to build upon the preliminary results presented here.

6 Conclusions

In this paper we have used insights from complexity theory to revisit an important issue in management and organization research—the power and influence of boards of directors and top management teams. We have proposed a novel, essentially structural metric by which the authority, importance and influence of the board can be evaluated and we have argued that our measure, which is related to the PageRank algorithm, the system at the heart of the extremely popular Google search engine[7], is applicable to the social network of boards of directors.

Our contribution to social science research can be stated thus. First, we have reviewed some important features of mathematical graph theory which are germane to social networks. We began, by restating the idea that the affiliation network of the board of directors can be represented as a bipartite graph (namely two sets of vertices with edges running between unlike kinds). We illustrated in Figure 2 that the resulting projections from such a bipartite representation are undirected graphs.

We then introduced and explained the PageRank algorithm. It assigns a numerical value to each vertex in a directed graph according to the update rule given in Equation (2). We illustrated that the rank of a given vertex j depends recursively on an adjustable tuning parameter d, source strengths, s_j, and the sum over incoming neighbors for only those edges that start at some vertex k and end at j (only edges with weights $w_{j,k} \neq 0$ contribute to the rank of vertex j). We discussed two interpretations of PageRank first as a random walk and second as a weighted voting scheme. These interpretations have arisen in sister sciences, such as physics, applied mathematics and computing science, but their application to social sciences and management research is novel.

PageRank, as we discussed, is designed to operate on directed graphs. The simple bipartite representation in Figure 1 does not capture this directionality. This raises a technical obstacle in adapting PageRank to the ranking of social networks of boards of directors. One needs to find a way to incorporate extra information into a bipartite graph (like the one in Figure 1) in such a way as to generate a board projection that is a directed graph. In Section 4 we have

[7]Which one of the authors (MRM) uses as his home page.

proposed such a method. Our paper has therefore described, uniquely, a method for calculating the PageRank algorithm for weighted directed bipartite graphs.

Finally, we implemented our code and calculated PageRanks of the board projections for publicly traded firms in the United States and the United Kingdom. We documented which companies can be structurally classified as "esteemworthy" in the social network of corporate governance. In summary, we hope the procedure outlined in this paper is valuable to social network researchers investigating arbitrary affiliation networks since it permits the ready calculation of a PageRank centrality measure.

References

Bollobás, B. (2001), *Random Graphs*, second edn, Cambridge University Press, Cambridge.

Brin, Sergey & Lawrence Page (1998), 'The anatomy of a large-scale hypertextual web search engine', *Computer networks and ISDN systems* **30**, 107–117.

Brin, Sergey, Lawrence Page, Rajeev Motwani & Terry Winograd (1999), The PageRank citation ranking: Bringing order to the web, Technical Report SIDL-WP-1999-0120, Stanford Digital Libraries Project.

Davis, Gerald F. & Henrich R. Greve (1997), 'Corporate elite networks and governance changes in the 1980s', *American Journal of Sociology* **103**, 1–37.

Diligenti, Michelangelo, Marco Gori & Marco Maggini (2004), 'Neural computation, social networks, and topological spectra', *Theoretical Computer Science* **320**, 71–87.

Ding, Chris H.Q., Hongyuan Zha, Xiaofeng He, Parry Husbands & Horst D. Simon (2004), 'Link analysis: Hubs and authorities on the world wide web', *SIAM Review* **46**(2), 256–268.

Freeman, Linton C. (1978), 'Centrality in social networks: conceptual clarification', *Social Networks* **1**, 215–239.

Gompers, Paul, Joy Ishii & Andrew Metrick (2003), 'Corporate governance and equity prices', *Quarterly Journal of Economics* pp. 107–155.

Hallock, Kevin F. (1997), 'Reciprocally interlocking board of directors and executive compensation', *Journal of Financial and Quantitative Analysis* **32**, 331–344.

Haunschild, Pamela. R. (1993), 'Interorganizational imitation - the impact of interlocks on corporate acquisition activity', *Administrative Science Quarterly* **38**, 564–592.

Haunschild, Pamela. R. & Christine. M. Beckman (1998), 'When do interlocks matter?: Alternate sources of information and interlock influence', *Administrative Science Quarterly* **43**, 815–844.

Higham, Desmond J. & Alan Taylor (2003), 'The sleekest link algorithm', *Mathematics Today* **39**, 192–197.

Huber, Peter (1964), 'Robust estimation of a location parameter', *Annals of Mathematical Statistics* **35**, 73–101.

Mizruchi, Mark. S. (1982), *The American Corporate Network*, Sage, Beverly Hill, California.

Mizruchi, Mark. S. (1996), 'What do interlocks do? An analysis, critique, and assessment of research on interlocking directorates', *Annual Review of Sociology* **22**, 271–298.

Mizruchi, Mark S. & David Bunting (1981), 'Influence in corporate networks', *Administrative Science Quarterly* **26**, 475–489.

Muldoon, Mark R. & Martin J. Conyon (2004), 'Perl scripts for small worlds and page ranks', Muldoon and Conyon's Web sites. Please visit us at:.
URL: *http://www.ma.umist.ac.uk/mrm/SmallWorld/ and http://www-management.wharton.upenn.edu/conyon/*

Newman, Mark E. J. (2003*a*), 'A measure of betweenness centrality based on random walks', Submitted to Social Networks.
URL: *http://aps.arxiv.org/abs/cond-mat/0309045/*

Newman, Mark E. J. (2003*b*), 'The structure and function of complex networks', *SIAM Review* **45**, 167–256.

Newman, Mark E.J., Duncan J. Watts & Steven H. Strogatz (2002), 'Random graph models for social networks', *Proc. Nat. Acad. Sci. U.S.A.* **99**, 2566–2572.

Pennings, J. M. (1980), *Interlocking directorates*, Jossey-Bass, San Francisco, California.

Pettigrew, Andrew M. (1992), 'On studying managerial elites', *Strategic Management Journal* **13**, 163–182.

Useem, Michael (1984), *The Inner Circle*, Oxford University Press.

Useem, Michael & J. Karabel (1986), 'Pathways to top corporate management', *American Sociological Review* **51**, 184–200.

Wang, Ziyang (2004), 'Improved link-based algorithms for ranking web pages', *Lecture Notes in Computer Science* **30**, 291–302.

Wasserman, S. & K. Faust (1994), *Social Network Analysis: Methods and Applications*, Cambridge University Press, Cambridge.

White, Halbert (1980), 'A heteroscedasticity-consistent covariance matrix estimator and a direct test for heteroscedasticity', *Econometrica* **48**, 817–838.