Towards a Model of the Link Economy

Robert Ackland*
Centre for Social Research
Research School of Social Sciences
The Australian National University
robert.ackland@anu.edu.au
http://acsr.anu.edu.au/staff/ackland

20th February, 2006
Preliminary and incomplete: Please do not quote without permission of the author

Abstract

The "link economy" is commonly used to describe activity on the World Wide Web (WWW) and refers to the value of linking to (and being linked to by) other websites. The distribution of links on the WWW clearly has real economic effects; commercial sites with higher indegrees will be more visible on the Web, and thus visited more frequently, because they will be ranked more highly by Google, for example. However, there has been a surprising lack of economic analysis of linking activity on the Web, with the majority of quantitative analysis using approaches from applied physics. Applied physics has provided the influential characterisation of the WWW as being scale-free, with a small number of sites receiving the lion's share of links. This has been explained using the concept of "preferential attachment" whereby new sites joining the network prefer to link to sites that already have high in-degree, thus leading to the "rich get richer" phenomenon. However, for researchers from the social and economic sciences, the preferential attachment argument lacks a strong behavioural foundation and does not directly incorporate the influence of human agency, for example business practice or politics, on the distribution of links. The aim of the present paper is to investigate how Web networks exhibiting the key empirical regularities that have been discovered in applied physics (such as power law distribution of links) can be modelled as the outcome of utility-maximising behaviour at the level of the individual. A main focus of the paper is a presentation of results from Jackson and Rogers (2004a, 2005) who provide new insights relevant to reconciling the "stylised empirical facts" of large-scale networks such as the Web with utility-maximising behaviour at the level of the individual.

keywords: network formation, WWW, power laws, scale-free networks, small worlds.

*This paper was originally presented at the Econophysics Colloquium, 14-18 November 2005, Canberra and was subsequently presented to the RSSS Economics Program seminar, 20th February, 2006. This research has been supported by Australian Research Council grant DP0452051.

1 Introduction

The increasing importance of the Internet and the WWW to our daily lives has led Lipnack and Stamps (1994) to suggest that we live in the "network age" and Castells (1996, p. 469) to write that networks "constitute the new social morphology of our societies". While social scientists (in particular, sociologists) have been studying networks for over half a century, the main focus of social science network research has been on small-scale networks of maybe tens or in extreme cases hundreds of nodes. Until very recently, the direct contribution of social science to empirical analysis of large-scale networks such as the Web had been almost negligible, with most contributions coming from applied physics and computer science.[1]

Applied physics has provided the influential characterisation of the WWW as being scale-free, with a small number of sites receiving the lion's share of links. This has been explained using the concept of "preferential attachment" whereby new sites joining a continuously growing network prefer to link to sites that already have high in-degree, thus leading to the "rich get richer" phenomenon (Barabási and Albert, 1999). However, for researchers from the social and economic sciences, the preferential attachment argument lacks adequate conceptualisation of the influence of human agency, for example business practice or politics, on the distribution of links.

While some social scientists have expressed unease with the adequacy of the preferential attachment argument for explaining why the Web has evolved into a scale-free network, there has been very little research expressing an alternative theoretical framework. This is particularly the case in the field of economics and it is quite surprising that while the term "link economy" one of the most pervasive descriptions of activity on the Web, there has actually been very little rigorous economic research into behaviour on the WWW.

In this paper, I present a brief review of the main empirical regularities or stylised facts of the Web that have been identified in applied physics and outline the influential preferential attachment model that has been used to explain the existence of power laws on the Web. I discuss some of the criticisms of the preferential attachment model that have been raised by researchers in both applied physics and the social sciences. With reference to recent research from Jackson and Rogers (2004a, 2005), I then investigate the potential for reconciling the stylised empirical facts of the Web with the predictions regarding network structure that arise from utility-based models of network formation.

2 Background

A *network* (or graph) is set of items called *vertices* (or nodes) with connections between them called *edges*. The World Wide Web is an example of an network where the vertices are web pages or web sites and the edges are hyperlinks between these entities. The WWW is an example of a *directed* graph because the edges (hyperlinks) point in only one direction. For example, a hyperlink on page i to page j indicates that the author of page i is aware of page j (and is referencing or citing it for a particular reason) but the author of page j may not even know of the existence of page i. The *degree* of a vertex

[1] Though these disciplines had admittedly built upon earlier contributions from the social sciences.

refers to the number of edges that connect to it (in directed graphs, we can distinguish both an in-degree and out-degree for each vertex, which are the numbers of incoming and outgoing edges, respectively). The *geodesic path* is the shortest path through the network from one vertex to another (note that there is often more than one geodesic path between two vertices). The *diameter* of a network is the length (in terms of the number of edges) of the longest geodesic path between any two vertices.

2.1 Empirical regularities of the Web

Applied physics has made a significant contribution to the development of methods for characterising the structure and behaviour of large-scale networked systems. With regards to the Web, applied physicists have provided several highly influential characterisations (or empirical stylised facts) regarding the structure of this important network, three of which are outlined in this section.[2]

2.1.1 The small-world effect

In the 1960s, the sociologist Stanley Milgram carried out experiments in which people were asked to send a letter to a friend or acquaintance who might know a designated target individual. The letters were passed on until the target individual was reached and the average number of steps was around six (leading to the famous phrase "six degrees of separation"). The *small-world effect* exists when most pairs of vertices in the network are connected by relatively short paths. Formally, define l to be the mean geodesic (i.e. shortest) distance between vertices in a directed network:

$$l = \frac{1}{n(n-1)} \sum_{i \neq j} d_{ij},$$

where d_{ij} is the geodesic distance from vertex i to vertex j.[3] A network is said to exhibit the small-world effect if the diameter of the network and l scale logarithmically or slower with network size (for a fixed mean degree).

2.1.2 Transitivity or clustering

In many networks it is found than if vertex i is connected to vertex j, and vertex j is connected to vertex k, then there is a heightened probability that vertex i will also be connected to vertex k (in social networks, this property is explained as "a friend of my friend is likely to also be my friend"). Transitivity is quantified with the clustering coefficient C:[4]

[2]The following review draws from Newman (2003).

[3]Note that the above definition of the mean geodesic distance is problematic in graphs that have more than one *component* (a component of a graph is a set of vertices that are connected to one another) since vertices in different components of a graph have infinite geodesic distance and the value of l then becomes infinite. To get around this, in graphs with more than one component l is calculated as the mean geodesic distance between all vertices that have a connecting path.

[4]This version of the clustering coefficient is defined for unweighted undirected graphs.

$$C = \frac{3 \times \text{number of triangles in the network}}{\text{number of connected triples of vertices}},$$

where a triangle is a set of three vertices with edges between each pair of vertices and a connected triple is a set of three vertices where each vertex can be reached from each other (directly or indirectly). It has been found that socially-generated networks such as the Web tend to have high clustering coefficients relative to what is found in random graphs with a similar number of edges and vertices.

2.1.3 Power laws in degree distributions

Define p_k as the fraction of vertices in the network that have degree k, so that $\sum_k p_k = 1$. A plot of p_k can be formed by making a histogram of the degrees of vertices, but for networks such as the Web degree data is generally presenting via the cumulative distribution function:

$$P_k = \sum_{k'=k}^{\infty} p_{k'}$$

which is equal to the probability that the degree is greater than or equal to k. A distribution follows a power law if $p_k \sim k^{-\alpha}$ for some constant exponent α. Power law distributions are also exhibited as power laws in the cumulative distributions, but with exponent $\alpha - 1$ rather than α:

$$P_k \sim \sum_{k'=k}^{\infty} k'^{-\alpha} \sim k^{-(\alpha-1)}.$$

Power laws are easily identified by straight-line plots of the cumulative distribution function on log-log scales and there is extensive evidence (Albert et al., 1999; Barabási et al., 2000) that the degree distribution on the Web follows a power law (at least in the upper tail).[5] An implication of a power law in the degree distribution on the Web is that a very small number of sites receive the vast majority of links.

2.2 Power laws and preferential attachment

Of the three network properties outlined above, it is the existence of power laws in degree distributions on the Web that has attracted the majority of research attention. A focus of research has been: why do scale-free networks develop?

The most prominent explanation for the existence of power laws on the Web has come from applied physics. The concept of *preferential attachment* (Barabási and Albert, 1999), posits that new sites joining a continuously growing network prefer to link to sites that already have high in-degree, thus leading to a "rich get richer" phenomenon.[6] With

[5] Networks with power-law degree distributions are sometimes referred to as scale-free networks.

[6] Without detracting from the achievements of Barabási and Albert (1999), Newman (2003, p.32) notes that preferential attachment is essentially a re-discovery of the concept of cumulative advantage, due to Price (1976), with the key difference that while Price (1976) studied directed graphs, Barabási and Albert (1999) focus on undirected graphs.

regards to linking on the Web, the intuition behind preferential attachment is that the probability of a person coming across a particular web page (and deciding to form a hyperlink to that page) will presumably increase with the number of other pages that hyperlink to it for two reasons. First, there are more pathways to a page that has many hyperlinks to it (and so it is more likely to be discovered). Second, major search engines such as Google all tend to rank more highly pages with higher in-degree.

Consider an undirected graph of n vertices. As before, p_k is the fraction of vertices in the network with degree k. In each period, new vertices are added to the graph and each new vertice has an initial degree of m where m is an integer and must have the value $m \geq 1$ (m is therefore the number of edges added to the network with each new vertex). The probability that a new edge attaches to an existing vertex of degree k is proportional to k:

$$\frac{kp_k}{\sum_k kp_k} = \frac{kp_k}{2m}.$$

Note that $\sum_k kp_k$ is equal to the mean degree of the network, which is $2m$ since there are m edges for each vertex added and (given the graph is undirected) each edge contributes two ends to the degrees of network vertices. It can be shown (see, for e.g., Newman, 2003, p.32) that in the limit of large k the degree distribution follows a power law $p_k \sim k^{-3}$ and hence the model produces a single fixed exponent $\alpha = 3$.

2.2.1 Reaction to the Barabási-Albert model from applied physics and computer science

The Barabási-Albert (BA) model has been hugely influential because it is an elegant and relatively simple framework for explaining how power laws can arise on large-scale networks. However, there are several aspects of the BA model that make it less appropriate for explaining the degree distribution of the WWW.[7] These criticisms of the BA model do not necessarily invalidate preferential attachment as the explanation for the existence of power-law degree distributions on the Web, but indicate that preferential attachment may not be the *sole* explanation for this phenomenon.

First, Krapvisky and Redner (2001) have shown that the BA model produces a correlation between the age of vertices and their degrees, with older vertices having higher mean degree. Adamic and Huberman (2000) have shown that there is no such correlation between between the age of a web page and its degree. Bianconi and Barabási (2001a,b) extended the basic BA model by incorporating the concept of "fitness" whereby the probability of acquiring new edges is also a function of a vertices intrinsic worth, and not just their current degree. This modification to the BA model results in vertex age and degree no longer being correlated.

A second problem with the BA model is that it is a model of an undirected network, while the WWW is a directed network. Newman (2003) argues that it is possible to in fact regard the BA model as a model of a directed network but where attachment is in proportion to the sum of in- and out-degrees of a vertex, however it is more realistic that attachment is in proportion to in-degree only. A related issue is that if the BA model is

[7]This section draws from Newman (2003, Section VII).

regarded as producing a directed graph, then it generates *acyclic* graphs whereas on the Web it is quite common for two pages to reference each other.

[NEED TO FINISH THIS SECTION]

2.2.2 Reaction to the Barabási-Albert model from the social sciences

While social scientists have investigated the existence of power laws on the Web, their focus has generally been on the political or social consequences of having an unequal distribution of links, rather than the behavioural processes that could be leading to the existence of power laws in degree distributions. Hindman et al. (2003), for example, analysed large collections of web pages related to political topics such as abortion and gun control and found that the degree distributions of these pages did broadly follow power laws. The authors argued that the existence of power laws in the networks of politically-oriented web content had serious implications for the visibility of different political messages or viewpoints because search engines such a Google generally rank more inlinked pages higher.

In the field of economics, there has been very little work addressing the question of why power laws exist on large networks such as the Web. At heart, there should be something very troubling for economists about the preferential attachment argument. First, if it is the case that nodes that have been in the network for longer are going to get most of the links, why would a rational agent even enter the network? This has been partly addressed by the concept of vertex fitness described above, but the behavioural foundations for this concept are not clear. Second, if a new agent is only concerned to link to the existing nodes with the highest in-degree, why would that agent randomly select across existing nodes (with probability of selection proportional to in-degree)? Surely, in such a model, a rational agent would link to the most inlinked node with probability 1.

Galitsky and Levene (2004) cited such concerns as motivation for their model where hyperlinks are viewed as investment instruments that are the subject of exchange. Kavassalis et al. (2004) draw from New Economic Geography (see, for e.g., Fujita et al., 1999), which attempts to explain why economic activity occurs where it does using the concept of increasing returns (or positive feedback loops) that reinforce economic success or aggravate economic loss, to provide an agent-based model where the agents are web sites and web users.

[NEED TO FINISH THIS SECTION]

3 An economic model of network formation on the WWW

Most economic research into networks has focused on issues of network stability and efficiency, mainly in the context of relatively small-sized networks. This contrasts with the network research in applied physics and computer science which has focused on mechanical or stochastic processes of network formation, and the study of the characteristics of networks formed under different processes. Two recent papers Jackson and Rogers (2004a, 2005) are important in that they attempt to cross the divide between economics and applied physics with regards to the modelling of networks. In this section, there is first a brief review of several economic models of network formation and then there is an overview of the insights provided by Jackson and Rogers (2004a, 2005).

3.1 Brief review of economic models of network formation

In this section, a brief review of several relevant utility models for social and economic networks is presented.[8]

Jackson and Wolinsky (1996) present a model of network formation where agents benefit from both their direct connections and also from indirect connections (that is, friends of a friend generate value). Their model incorporates both cooperative and noncooperative aspects of behaviour; links are non-directed and their formation requires the consent of both parties involved, but link severance can be done unilaterally. They investigate the relationship between the *efficiency* of a network (which in some applications is the aggregate of the utilities of individual nodes) and the *stability* of a network. With regards to the latter, they introduce the concept of a *pairwise stable network* - this is where no player stands to benefit from severing an existing link, and no two players would benefit from forming a new link. Comparative static analysis of effects of adding/deleting link at or near equilibrium is used to derive results about economic efficiency and stability of networks.

Jackson and Watts (2002) extend the Jackson-Wolinsky framework by modelling network formation as a dynamic process - agents form and sever links based on improvement the resulting network offers relative to current network. They define an *improving path* as a sequence of networks observed in a dynamic process as agents delete/add links, one at a time. They further assume myopic behaviour - agents do not forecast how decisions to form/sever link may influence the future evolution of network.

Johnson and Gilles (2000) introduce a "spatial cost topology" into the Jackson-Wolinsky model that may reflect differences between agents due to geography or personal characteristics. The more similar the agents (with regards to individual characteristics), the less costly it is to form links.

In Bala and Goyal (2000a), the costs of link formation are incurred only by the agent initiating the link and thus the network formation process is formulated as noncooperative game (resulting in so-called *Nash networks*). The authors assume that a link from agent i to agent j allows i to access the benefits that accrue to j via his own links (thus individual links generate externalities that depend on the level of *information decay* associated with indirect links). The benefits accruing to i from linking with j are assumed to be *non-rival* in that the benefit accruing to i from linking to j is not dependent on the number of other agents who may form links with j in the same time period. The authors study both one-way flows of benefits (where a link from i to j yields benefits solely to i) and two-way flows where benefits accrue to both agents. Bala and Goyal (2000b) extend the model of Bala and Goyal (2000a) by introducing the concept *imperfect reliability* in information networks, implement in terms of a positive probability that agent i will not benefit from forming a link to agent j.

[8]See Jackson (2004) for a complete review of these models.

3.2 Explaining small worlds - Jackson and Rogers (2004a)

Jackson and Rogers (2004a) extend the Jackson-Wolinsky model by incorporating a cost structure to forming links by assuming that agents are grouped on "islands" and costs of connection are relatively low within an island and relatively high across islands (the distance between agents could reflect geography or could reflect heterogeneity among agents in terms of social and political preferences, for example). The authors show that this cost structure, together with the benefits gained from indirect connections between vertices, generates networks exhibiting small-world characteristics.[9] High clustering results from the low cost of connecting to similar (or nearby) nodes. Low diameter and average path length results from the large benefit from attaching to dissimilar (or distant) nodes because these nodes provide substantial benefits from indirect access to other distant nodes. Only a small number of these distant links are formed (due to the highcd cost), but because of the high rate of connection at the local level, these distant links substantially decrease network diameter and average path length.

While Jackson and Rogers (2004a) are successful in developing an economic model of network formation that produces small-world networks, they acknowledge that their extended Jackson-Wolinksy model is not capable of producing a degree distribution that follows a power law. The authors conjecture that it might be necessary to introduce a random process governing how nodes meet to augment the economic and strategic aspect of node behaviour already incorporated in their model. A feature of the model of Jackson and Rogers (2004a) that further limits its applicability to the study of link formation on the WWW is the fact that the Jackson-Wolinsky model is a model of an undirected network, while the Web is a directed network.

3.3 Explaining power laws and small worlds - Jackson and Rogers (2005)

Jackson and Rogers (2005) provide a model of network formation that incorporates both random attachment and search behaviour and produces networks exhibiting both small-world characteristics and power-law degree distributions. Note that this model, unlike that of Jackson and Rogers (2004a), is not an extension of the Jackson-Wolinsky model, and importantly (for applications of the model to the Web), it is a model of a directed network. The following is a non-mathematical description of the model of Jackson and Rogers (2005).

New nodes enter into the network sequentially, and upon entering the network a new node will form links to existing nodes via two processes. First, it will find out about (or "meet") some existing nodes that are chosen completely at random (this is referred to as "random attachment"). Second, the new node will then meet some nodes that are neighbours of the nodes chosen in the first stage (this is referred to as "search"). There is then a probability that the new node and and any given node that it has met are compatible, and if so, then a link will be formed.

This model produces all three empirical characteristics of networks that were discussed

[9] Note that in contrast to the above review, Jackson and Rogers (2004a) state that a small world effect is the simultaneous existence of low network diameter and high clustering coefficient.

above. First, nodes with higher in-degree are more likely to be found through the local search process since there are more paths leading to them. Thus, the model incorporates a "preferential attachment" dimension to it and this leads to power law degree distributions in the upper tail of the distribution. Second, the local search process also leads to high clustering (a link may be formed to a "point of entry" to a particular neighbourhood, and then also to a node within that neighbourhood found by searching from the point of entry). Finally, the model produces networks with relatively small diameter due to the fact that many nodes will link to the same node via the process of local search, and at the same time these nodes will link randomly to other neighbourhoods.

While the model in Jackson and Rogers (2004a) is clearly derived from utility-maximising behaviour, the model of Jackson and Rogers (2005) is much more closely aligned to the stochastic processes that have featured in the applied physics and computer science literature. In particular, the search component of the model is essentially a form of preferential attachment, and as stated above, this is what generates the scale-free degree distributions. As noted by Jackson and Rogers (2005, p.7), their model is similar to that of Vázquez (2003) which also combines random meetings with local search of neighbourhoods, and also shares some characteristics with the model of Pennock et al. (2002).

What is apparent from Jackson and Rogers (2004a, 2005) is that in order to generate models exhibiting network characteristics that are observed in the real-world (such as small-world characteristics and power-law degree distributions), it seems to be necessary to move away from "pure" economic models of strategic behaviour (such as the Jackson-Wolinsky model) and incorporate features from the stochastic modelling from applied physics. However, this is not to say that economic theory does not have a unique and important role to play in understanding large-scale networks. One of the major contributions of Jackson and Rogers (2005) is that they provide results that explicitly relate variations in the network formation process (i.e. the role of random attachment versus search) to variations in network efficiency, as defined as the sum of the utilities of the individual nodes.

4 Conclusions

The starting point to this research was a desire to formulate an economic model of network formation that could generate some of the key empirical regularities that characterise large-scale networks such as the Web. While preferential attachment is a highly influential concept that has been used to explain the growth of scale-free networks, there isn't a strong behavioural foundation to this concept. The motivation behind this paper was therefore to produce an economic model where preferential attachment was recast in a utility-maximisation framework and to have this model produce networks that exhibited small-world characteristics and power-law degree distributions.

However, after an extensive review of the available literature (and several aborted attempts at model building), it appears that the above research agenda was quite ambitious. One of the co-authors of the influential Jackson-Wolinsky model, Matthew O. Jackson has recently extended his own model to incorporate costs of link formation and has found that while this model can produce a small-world effect it does not provide for a scale-free network and for this, incorporation of features from stochastic models of network formation is necessary: "...a variation of the model that marries a random process to how nodes meet with the economic and strategic concerns analyzed here would begin to account for the degree distribution and could result in features consistent with observables."(Jackson and Rogers, 2004a, p.16)

Jackson and Rogers (2005) present a model that incorporates both random attachment and local search and this model does produce networks that have characteristics that are consistent with empirical findings (i.e. small worlds and power laws). However, it is of note that the authors themselves emphasise (p. 5) that the nodes or agents in their model are "non-strategic" (and hence aren't exhibiting purely economic behaviour) and that "[while] this meeting process has many 'natural' characteristics, it is more in the tradition of the random graph literature in terms of a model largely based on a mechanical (stochastic) process" (Jackson and Rogers, 2005, p.5). Instead, the authors emphasise the contribution of economics as being at the stage of relating the stochastic processes back to the welfare of the individual nodes, and the implications for the efficiency characteristics of networks.[10]

In conclusion, the experience of writing and researching this paper has led the author to firmly agree with Jackson and Rogers (2004a) that economic models should be seen not as a substitute for more mechanical models of network formation, but rather as a complement. In particular, while mechanical models provide the randomness and heterogeneity to produce networks that exhibit characteristics close to real-world networks, economic models can provide more insights into why networks form as they do. A model that simultaneously provides real-world empirical features and insights into the motives of agents is most likely going to be a hybrid model such as that provided by Jackson and Rogers (2005), rather than a "pure" economic model of network formation.

References

Adamic, L. and Huberman, B. (2000). Power-law distribution of the world wide web. *Science*, 287:2115.

Albert, A., Jeong, H., and Barabasi, A.-L. (1999). Diameter of the World Wide Web. *Nature*, 401:130–131.

Bala, V. and Goyal, S. (2000a). A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229.

Bala, V. and Goyal, S. (2000b). A strategic analysis of network reliability. *Review of Economic Design*, 5:205–228.

Barabási, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286:509–512.

[10]It is of interest that in an earlier version of Jackson and Rogers (2005), dated May 26, 2004, the authors were much more prone to using "fighting words" to illustrate the potential for economics to further illuminate our understanding of large-scale networks: "...the [stochastic] processes for generating networks are attempts at answering the question of 'how' but not of 'why'. That is, they are not models of an actual formation process where actors are making some explicit, and say rational, decisions about how to connect the network. The are instead simple models of specific stochastic processes of wiring or rewiring a network that will exhibit some of the desired characteristics."(Jackson and Rogers, 2004b).

Barabási, A.-L., Albert, R., and Jeong, H. (2000). Scale-free characteristics of random networks: The topology of the world wide web. *Physica A*, 281:69–77.

Bianconi, G. and Barabási, A.-L. (2001a). Boseeinstein condensation in complex networks. *Phys. Rev. Lett.*, 86:56325635.

Bianconi, G. and Barabási, A.-L. (2001b). Competition and multiscaling in evolving networks. *Europhys. Lett.*, 54:436–442.

Castells, M. (1996). *The rise of the network society. The information age: Economy, society and culture Vol. I.* Blackwell, London.

Fujita, M., Krugman, P., and Venables, A. (1999). *The Spatial Economy: Cities, Regions and International Trade*. MIT Press, Cambridge, MA.

Galitsky, B. and Levene, M. (2004). On the economy of web links: Simulating the exchange process. *First Monday*, 9(1). Available at: `http://www.firstmonday.org/issues/issue9_1/galitsky/`.

Hindman, M., Tsioutsiouliklis, K., and Johnson, J. (2003). 'Googlearchy': How a few heavily-linked sites dominate politics on the Web. mimeograph, Princeton University, 2003.

Jackson, M. (2004). A survey of models of network formation : stability and efficiency. In Demange, G. and Wooders, M., editors, *Group formation in economics; Networks, Clubs and Coalitions*. Cambridge University Press, Cambridge, UK.

Jackson, M. and Watts, A. (2002). The evolution of social and economic networks. *Journal of Economic Theory*, 106:265–295.

Jackson, M. and Wolinsky, A. (1996). A strategic model of social and economic networks. *Journal of Economic Theory*, 71:440–74.

Jackson, M. O. and Rogers, B. W. (2004a). The economics of small worlds. mimeograph, Caltech.

Jackson, M. O. and Rogers, B. W. (2004b). Search and strategic formation of large networks: When and why do we see power laws and small worlds? mimeograph, Caltech.

Jackson, M. O. and Rogers, B. W. (2005). Search in the formation of large networks: How random are socially generated networks? mimeograph, Caltech.

Johnson, C. and Gilles, R. (2000). Spatial social networks. *Review of Economic Design*, 5:273–299.

Kavassalis, P., Lelis, S., Rafea, M., and Haridi, S. (2004). What makes a website popular? *Communications of the ACM*, 47(2):50–55.

Krapvisky, P. and Redner, S. (2001). Organization of growing networks. *Phys. Rev. E*, 63(066123).

Lipnack, J. and Stamps, J. (1994). *The Age of the Network*. Wiley, New York.

Newman, M. (2003). The structure and function of complex networks. arXiv:cond-mat/0303516 v1.

Pennock, D., Flake, G., Lawrence, S., Glover, E., and Giles, C. (2002). Winners don't take all: Characterizing the competition for links on the web. *PNAS*, 99(8):5207–5211.

Price, D. J. (1976). A general theory of bibliometric and other cumulative advantage processes. *J. Amer. Soc. Inform. Sci*, 27:292–306.

Vázquez, A. (2003). Growing network with local rules: Preferential attachment, clutering hierarchy, and degree correlations. *Physical Review E*, 67(5):056104.