

An Overview of Social Networks and Economic Applications*

Matthew O. Jackson[†]

Final Version: June 25, 2010
Written for the *Handbook of Social Economics*[‡]

Abstract

In this chapter I provide an overview of research on social networks and their role in shaping behavior and economic outcomes. I include discussion of empirical and theoretical analyses of the role of social networks in markets and exchange, learning and diffusion, and network games. I also include some background on social network characteristics and measurements, models of network formation, models for the statistical analysis of social networks, as well as community detection.

Keywords: Social Networks, Network Games, Graphical Games, Games on Networks, Diffusion, Learning, Network Formation, Community Detection,

JEL Classification Codes: D85, C72, L14, Z13

1 Introduction

The people with whom we interact on a regular basis, and even some with whom we interact only sporadically, influence our beliefs, decisions and behaviors. Examples of the effects of social networks on economic activity are abundant and pervasive, including roles in transmitting information about jobs, new products, technologies, and political opinions.

*Financial support from the NSF under grant SES-0647867 is gratefully acknowledged. Thank you to Matthew Elliott, Carlos Lever, and Yves Zenou for comments on earlier drafts.

[†]Department of Economics, Stanford University, Stanford CA 94305, USA.
<http://www.stanford.edu/~jacksonm> e-mail: jacksonm@stanford.edu

[‡]Edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, forthcoming, Elsevier Press.

They also serve as channels for informal insurance and risk sharing, and network structure influences patterns of decisions regarding education, career, hobbies, criminal activity, and even participation in micro-finance. Beyond the role of “social” networks in determining various economic behaviors, there are also many business and political interactions that are networked. Networks of relationships among various firms and political organizations affect research and development, patent activity, trade patterns, and political alliances.¹ Given the many roles of networks in economic activity, they have become increasingly studied by economists.

There are two important aspects of the study of networks from an economist’s perspective. The first is understanding how network structure influences economic activity. Some examples of questions examined in the literature provides an idea of why this is an important issue:

- How does the role of social networks in disseminating job information affect wages and employment?
- How do terms of trade in networked markets depend on the network structure and compare to centralized markets?
- How are education and other human capital decisions influenced by social network structure?
- Will the networks that are formed be the efficient ones in terms of their implications for economic activity?
- ...

The second important aspects of the study of networks from an economist’s perspective is that economic tools are very useful in analyzing both network formation and network influence, and these tools are quite complementary to those from the many other disciplines that also study social networks. That is, even beyond the eventual implications for economic activity and welfare, what can we say about how people will self-organize and why certain patterns will emerge? For example, why do people tend to associate with other people who

¹References to studies on some of these subjects are provided below. Related chapters on labor markets (Topa (this volume)), networks in developing countries (Munshi (this volume)), risk sharing (Fafchamps (this volume)), diffusion and social structure (Jackson and Yariv (this volume)), learning (Goyal (this volume)) provide additional references, and Jackson (2008) provides a more detailed look at some applications and an extensive bibliography.

are similar to them along a number of dimensions and what will this imply for behavior? Why is the social distance between people (in terms of the shortest path in the social network in which they are embedded) so small even in very large societies?

In this chapter I provide background on networks, both in terms of how they influence social and economic activity as well as how they can be modeled and analyzed. In doing so, I cover areas branching out from what an economist might usually be exposed to because the study of networks is so naturally interdisciplinary and multidisciplinary. Part of this is due to the fact that networked interactions have wide-ranging applications, from purely social, to economic, to political, and even to biological. This is also partly due to the diverse set of tools that are useful in analyzing networks, including: anthropological case studies, sociological survey methods, mathematical analyses of random graphs, techniques from statistical physics as well as computer science for analyzing complex interactive systems, and models of strategic interaction from economics. Given the broad scope of network analysis, the chapter can only provide an introduction to and glimpse of the research in this area, discussing a few examples of the empirical and theoretical literature and giving a feel for the importance of the subject.²

2 Social Networks and Networked Markets

I start with a discussion of some of the various settings in which networked interactions are prominent and important.

2.1 Some Background on Networked Markets

Although we often idealize markets as centralized, many goods and services are contracted upon through networks of bilateral relationships. A manufacturer might have specific relationships with a few suppliers of raw materials, that also supply other firms and possibly even the manufacturer's competitors. Firms subcontract and out-source some of their business. The terms of trade, prices and products that emerge in such a world can depend on who is connected to whom. To understand this, it is useful to note the ingredients that comprise an idealized market: large numbers of well-informed agents in semi-anonymous settings, goods that are observable and of verifiable qualities, and contracts that are fully specified and costlessly enforced. In reality, even very large markets involve frictions that lead them to be at least partly decentralized. For example, in many labor markets jobs come with various

²See Jackson (2008) for a much more detailed look at the topics covered here as well as many that are not discussed here.

idiosyncrasies (including location, skills required, work environment, compensation, etc.), as do workers in terms of their backgrounds. Such heterogeneity in the market means that information can be critical in properly matching workers to positions.³ Social networks fill such a role, both in communicating information to workers about the specifics of various job opportunities, as well as communicating information to firms about the potential fit of various workers; mitigating substantial search frictions. This role of enhancing a matching is not specific to labor markets, as the same can be said of a broad range of markets, as most goods and services involve some heterogeneity both in terms of what is supplied and what is demanded. Also, this is not the only role that networks play in labor and other markets. Repeated interactions with specific partners help mitigate a number of problems related to moral hazard and adverse selection, and thus long-term economic relationships with known partners can dominate shorter-term anonymous transactions.

In this section I begin with a discussion of a few of the many studies that illustrate these roles of networks in market settings, and then I go on to discuss the more general roles of social networks in diffusing information and in shaping behavior.

2.1.1 Social Networks in Labor Markets

The role of social networks in labor markets deserves our attention for at least two reasons: first, because of the central role networks play in disseminating information about job openings they place a critical role in determining whether labor markets function efficiently; and second, because network structure ends up having implications for things like human capital investment as well as inequality. As such, this application is a great example of why economists should care about networks and why networks deserve careful study and should be incorporated into our modeling of markets. As this topic is also discussed in Topa (This volume) and Munshi (This volume), I will keep this discussion rather brief.⁴

The fact that social networks are an important conduit of information about and access to jobs is evident to anyone who has ever looked for employment in almost any profession. The role of networks in labor markets has been extensively documented, with early studies including that of Myers and Shultz (1951) who interviewed textile workers in a New England mill town. Myers and Shultz observed that 62 percent of the interviewees had found out about and applied to their first job through a social contact, while only 23 percent had applied directly on their own and 15 percent had found their job through an agency, ads,

³For example, see Pissarides (2000).

⁴For more extensive background, the reader is also referred to the survey article by Ioannides and Datcher-Loury (2004), and Chapter 10 in Jackson (2008).

or other means. This sort of pattern is not specific to the textile industry, but is typical. There are some variations in the role of networks as a source of job information as we make comparisons across professions, locations, ethnicities of workers, and other attributes; but networks play a substantial role in essentially all of the labor markets that have been studied, regardless of the skill level, location, or population of workers (e.g., see Rees and Shultz (1970), Montgomery (1991), Pellizzari (2009), Corcoran, Datcher, and Duncan (1980), and Bentolila et al (2009)).⁵

The fact that information about jobs is passed through a social network becomes interesting because of its implications for wage and employment dynamics and patterns. As a starting point, models of social networks in labor markets imply that peoples' wages and employment will be related to that of their friends and acquaintances. The basic driver for this is that when unemployed workers will generally obtain more job information if their social contacts are employed than if their social contacts are unemployed. This gives rise to robust and strong forms of correlation in the wages and employment of linked individuals in a network, as studied in some detail by Calvó-Armengol and Jackson (2004, 2007). Although the theoretical implications can be cleanly established, testing for such correlations is not so easy since there is an absence of longitudinal data sets that include detailed social network observations together with employment and wage data. As such, much of the empirical work testing for social influence on wages and employment has tried to proxy for the social network by using some other observable such as the proximity of individuals, their ethnicity, or other attributes. For example, Bayer, Ross and Topa (2005) make use of census data to demonstrate higher correlation in employment among people living in the same city block compared to correlations among those living on different city blocks (but still close enough to avoid a geographical employment effect), and controlling for workers' characteristics. Thus, an individual's employment outcome is not simply based on (observed) characteristics or the employment of the general geographic area, but also the state of employment of his or her peers at very close proximity. This is not conclusive evidence for social network effects, as one cannot completely rule out some other unobserved characteristic that accounts for employment outcomes and also is related to peoples' tendency to live on the same city block, but it provides evidence that goes beyond previous studies. This challenge of alternative explanations of unobserved characteristics that are correlated with social interaction is endemic to peer effect studies, and in particular has pushed studies of social effects in labor

⁵The resilience of networks as job contact information conduits is the subject of an analysis by Casella and Hanaki (2008), and there is even evidence that the use of social networks as information sources is positively related to unemployment, as found by Galeotti and Merlino (2008) looking at U.K. data.

markets to be increasingly detailed and careful in their design with the intent of ruling out other potential explanations for observed patterns.⁶

In addition to the correlation of wage and employment outcomes across agents, the fact that jobs are accessed through social networks has implications for the time series of employment. In particular, as Calvó-Armengol and Jackson (2004) point out, duration dependence in unemployment arises in a networked model of employment. Duration dependence (e.g., see Schweitzer and Smith (1974) and Heckman and Borjas (1980)) refers to the stylized fact that the longer a given individual is unemployed, the higher the chance that the individual will still be unemployed in the next period, even after controlling for various characteristics of the individual. There are various reasons for duration dependence given in the labor economics literature, including characteristics about the worker that potential employers can observe but the researcher cannot (a form of “unobserved heterogeneity”), so that conditioning on a longer spell of unemployment it is more likely that the worker is unattractive to employers. It can also be that workers lose skills from being out of work, or that longer spells of unemployment are correlated with unobserved local labor market conditions (see Lynch (1989) for more discussion and references). The role of social networks as a conduit of information about jobs adds a new facet to understanding duration dependence. Duration dependence arises in a job-contact network context, because longer spells of unemployment are more likely to occur to an agent who has fewer (employed) friends, all else held equal. Thus, the longer we see an agent being unemployed, the lower the conditional probability that the agent has many employed friends, and so the lower the probability that the agent will have access to job information in the near future.

Beyond the implications for employment and wages, studies of job contact networks also have led to other observations and implications. Perhaps the best-known example of this is Granovetter’s (1973) observation of the “strength of weak ties.” Granovetter interviewed people in Amherst Massachusetts and asked not only how they found out about their jobs, but also how frequently they interacted with the people through whom they heard about their jobs. He called relationships “strong” if the two people interacted at least twice per week on average, “medium” if the pair interacted less than twice per week but more than once per year, and “weak” if they interacted less than once per year. Based on 54 interviewees who found their most recent job through a social contact, Granovetter found that 16.7 percent had found their job through a strong tie, 55.7 percent through a medium tie, and 27.6 percent through a weak tie. While people might tend to have many more medium and weak ties

⁶For some other clever recent approaches to isolating social effects, see Lachever (2007), Munshi (2003), and Beaman (2007), as well as the discussion in Topa (This volume).

than strong ones, it is still significant that people with whom an individual interacted so infrequently could play such an important and instrumental role. Granovetter (1973) goes on to discuss that the importance of weak ties can be in part traced to the fact that they often connect an individual to parts of a social network to which that individual would otherwise be quite far from. Thus, weak ties play a sort of bridging role, and can provide access to information that an individual might not find through other means. Much of the impact of Granovetter’s work has come from the wide application and evaluation of the “strength of weak ties” idea in other contexts, and the important observation that the strength of ties have consequences and so it can be useful to keep track of tie strength.

Additional implications of networked labor markets beyond their direct impact on employment and wages concern things like workers’ decisions to drop out of the labor force, and decisions invest in education and other forms of human capital. As Calvó-Armengol and Jackson (2004, 2007, 2009), point out, the fact that social networks are important in conveying job information results in complementarities in investment decisions between friends and acquaintances. As more of an agent’s friends are in the labor market, an agent has a better chance of hearing about jobs which increases that agent’s payoff from remaining in the labor force and may also increase that agent’s returns from education. Modeling these implications of networked labor markets is discussed in more detail below in Section 5.5.

2.1.2 Other Networked Markets

Beyond labor markets, there are many other market settings where networks of relationships play an important role. An interesting and illustrative example is that of the garment industry. Uzzi’s (1996) study of the New York garment industry uncovered several aspects of the role of networks. In that industry there are a set of firms that manufacture garments and others that are contracting to buy particular garments. A typical relationship involves a contractor coming to a manufacturer with a given design and contracting to buy a given number of garments produced to the design specification. Some such arrangements are fairly straightforward, whereas others are more complex and involve some idiosyncracies in production potentially requiring special investments, uncertainties, or other things that might lead to less than perfect contracting. Uzzi’s interviews suggest that some relationships function well as one-time “arm’s length” or “market” interactions, while others are “special” or “close” and involve repeated interaction, trust, fine-grained information transfer, and joint problem solving. Uzzi (1996) also explores the extent to which having close relationships versus market relationships is related with a firms survival. In particular he examines data

from 1991 from the International Ladies Garment Workers' Union which cover most of the active firms in New York in that year. Uzzi then documents that 125 of the 479 contractors with complete records in the data set failed to survive during this year. He regresses whether or not a firm survives on a variable that keeps track of the extent to which a firm is involved in close and repeated relationships, as well as a series of other background variables (geographic location, age of the firm, size of the firm, network centrality measures, and other neighborhood variables). He finds a positive and significant relationship between the survival variable and the variable measuring how concentrated a firm's contracts are. A firm that has completely dispersed contracts (so no repeat business) is almost twice as likely to exit the market in the year as a firm that is completely monogamous and contracts with only one other firm.

Clearly, there are many potential explanations for the correlation between a firm's survival and how many other firms it contracts with, and so we cannot deduce the causal relationship from these data.⁷ Nevertheless, whatever the causal chain, there is still a statistically significant negative relationship between the degree of a firm in terms of how many other firms it contracts with and the firm's chance of survival. Thus, network structure is playing some role, either affecting the survival rate, or being an outcome of the survival prospects of a firm, or relating to some other unobserved characteristic that affects the survival rate. Some of the models discussed below provide specific hypotheses about how firm characteristics relate to the network structure, as well as the terms of trades. These models, together with further empirical and experimental studies, should help us to better understand the role of networks of relationships in various markets.

2.2 Social Networks in Learning and Diffusion

Another important role of social networks is in influencing learning as well as the diffusion of technology, opinions and behaviors. This is discussed in more detail in Goyal (This volume) and Jackson and Yariv (This volume), and I will mention a few of the primary points here.

Given the obvious importance of the diffusion of a technology, product, disease, or opinion, diffusion has been extensively studied from a variety of perspectives. There are early studies such as Ryan and Gross's (1943) research on the diffusion of hybrid corn seed among farmers, and Hagerstrand's (1970) study of the diffusion of the telephone, as well as a range

⁷The role of networks of relationships between firms and their survival is still not extensively modeled. Some work on this appears in Allen and Babus (2007), but it is still an under-developed area of study, especially given the potential role of network structure in financial contagions.

of other case studies (see Rogers and Rogers (2003)). In terms of some of the interesting issues, on a most basic level there is a simple question of establishing that a given individual's behavior is influenced by that of his or her social contacts. There are many different challenges in such an analysis. If we observe that an individual's behavior is correlated with that of his or her friends and acquaintances can we conclude that they influenced it? There is an ever-looming question that we may not have observed all of the characteristics that influence behavior. That is, if we see that two friends both buy a new product, one after the other, can we conclude that one's purchase had an influence on the other? The difficulty is that there can be many things that influence these individuals' decisions, such as their age, income, education, ethnicity, exposure to advertising, and so forth. To the extent that we can observe *all* of the relevant factors affecting behavior, we can test for peer influence then by seeing whether a peer's purchase of a product leads to an increase (or decrease) in an individual's propensity to purchase after accounting for all of those other factors. Of course, the difficulty is that we generally will not have observed all such factors. This is exacerbated by the fact that people's friendships tend to correlate with how similar they are, something termed "homophily" and discussed in more detail below. So, it could be that the individuals happen to be friends because they share some trait, and it is that trait which causes them both to buy the product. If we do not observe that trait we could mistakenly conclude that one's purchase influenced the other's.⁸

Overcoming this problem requires some clever collection of data or experimental analysis. A nice example of this using field data is a study of social learning by Conley and Udry (2001, 2004). They examine the use of fertilizer by pineapple farmers. In particular, they show that changes in the amount of fertilizer used by a given farmer are related to the success or failure of similar past changes in fertilizer use by other farmers. Having controlled experiments can substantially narrow down the range of explanations for observed peer correlations. For example, Hesselius, Johansson, and Nilsson (2009) examine absences in the workplace based on a randomized rule affecting about 3000 workplaces in Göteborg Sweden. Randomly assigned agents were allowed to have longer spells of absence from work (14 days) without having to produce a doctor's certificate than was the rule for the general population (8 days). This resulted not only in an increase in absences for the treated individuals (those allowed the extra time before producing a doctor's certificate), but also for non-treated individuals conditional on being in a workplace with many treated individuals. Interestingly, the affect of how many other treated individuals there were in the workplace did not significantly

⁸For an interesting recent study showing how strongly homophily can bias studies of such peer influence, see Aral, Muchnik, and Sundararajan (2009).

influence treated individuals' behavior. This allows them to distinguish between various ways in which the peer effects might work, ruling out things like enjoying time together and being more consistent with a fairness effect or related peer effect on preferences. This sort of study shows the power of (field) experiments in identifying peer effects.⁹

Of course, an alternative technique to working with controlled data is a structural approach where one works with a model that makes pointed predictions about patterns in the data depending on the mechanism at work. This can offer an improved understanding of the mechanisms at work in peer effects and diffusion, but depends on the plausibility of the model. For example, Banerjee, Chandrasekhar, Duflo and Jackson (2010) fit models of diffusion to patterns of microfinance participation in a set of villages in rural India. They take advantage of differences in predicted patterns of behavior as a function of the diffusion of basic information (awareness of microfinance) compared to peer effects (where agents are reluctant to participate unless they have a personal acquaintance who endorses microfinance). Based on such models, they find that patterns of microfinance are more in line with peer effects than basic information diffusion.

Another challenge faced in studying peer effects comes from dealing with problems where one does not directly observe the network of who interacts with whom, but instead proxies for this with aggregated behavior by a given individual's peers. This can lead to identification problems, such as the reflection problem pointed out by Manski (1993) and discussed in more detail in the chapters by Blume et al (This volume) and Durlauf (This volume).¹⁰

Beyond establishing the fact that individuals are influenced by their peers, we also wish to better understand how this depends on an individual's "position" in the network. An early study in this direction is by Coleman, Katz, and Menzel (1966) who examined the time at which different doctors first prescribed a new drug to one of their patients. Coleman, Katz, and Menzel first interviewed doctors to map out a social network in terms of whom they would turn to for advice and whom they had contact with socially. Coleman, Katz, and Menzel also kept track of the first date at which a doctor prescribed the new drug to a patient. They divided their sample of doctors into three groups: those named by at least three other doctors as an advisor or friend, those named by one or two other doctors as an advisor or friend, and those not named by any other doctors as an advisor or friend. They

⁹Taking advantage of randomized programs in the field has proven to be a useful strategy in identifying various peer effects, not only operating through preference interactions but also through the diffusion of information. For example, see Duflo and Saez (2003) and Bayer, Hjalmarsson, and Pozen (2009).

¹⁰See Bramoullé, Djebbari and Fortin (2009) for a discussion of how network information can overcome the reflection problem.

then tracked what fraction of the doctors in each of these three groups had prescribed the drug at least once. Essentially their finding was that the most highly connected group (named by three or more other doctors) had the greatest fraction who had prescribed the drug by each date, and the second most highly connected group was second, and the unconnected group lagged behind in the fraction that had prescribed the drug. Thus, according to one measure of how “connected” a doctor is, more connected doctors were significantly more likely to be prescribing this new drug at an earlier date. So, we see that position seems to matter; although it is not clear why it matters. Subsequent studies have worked to sort out why this might have occurred, and there one faces the same challenges as discussed above in interpreting these data (see Jackson (2008) for some discussion and references to follow ups). It is possible that there are other factors that lead doctors to begin prescribing a new drug, such as their exposure to drug companies’ marketing, or doctors’ attitudes towards change, and so forth; and these other factors could be spuriously correlated with the connectedness of the doctors. Again, by specifying the mechanism that one believes might be at work, one can then begin to distinguish between some of the potential explanations via more detailed observations, or controlled experiments in lab or field settings. Regardless of whether network position is directly causal, or only indirectly related via some other attributes, it is of interest to understand why network position ends up being related to activity.

The list of settings where peer effects, or network effects more generally, have been found to be important is a long and varied one. It includes a range of things from criminal behavior (Reiss (1980), Glaeser, Sacerdote and Scheinkman (1996), Kling, Ludwig and Katz (2005), Patacchini and Zenou (2008)), to education (e.g., Calvo-Armengol, Patacchini and Zenou (2008)), to risk-sharing and loan behavior (Fafchamps and Lund (2003), De Weerd (2004), Karlan, Mobius, Rosenblatt, Szeidl (2009)), to obesity (Christakis and Fowler (2008), Fowler and Christakis (2008), and Halladay and Kwak (2009)). (See Fafchamps (This volume), Ioannides (This volume), Jackson and Yariv (This volume), Munshi (This volume), Sacerdote (This volume), and Topa (This volume), for more examples and background on empirical evidence.)

3 The Structure of Social Networks

I now turn to discussing what is known about social networks in terms of their basic structure and how they can be usefully quantified. These issues are of interest from a pure social science perspective to those studying how humans self-organize, as well as a basic tool box for those wishing to further study the role of network structure in economic interactions.

3.1 Definitions and Graph Terminology

I begin with some definitions and terminology that will allow us to talk about network structure. Much of this terminology emerges from standard graph theory, with some variation in terms across disciplines.¹¹

A network is represented as a graph on a set N of *nodes*, with a finite number of members n . Nodes are also sometimes referred to as vertices, agents, or players.

A *graph* or *network* is a pair (N, \mathbf{g}) , where \mathbf{g} is an $n \times n$ adjacency matrix on the set of nodes, where g_{ij} indicates the relationship between nodes i and j . For most of the discussion here I focus on cases where $g_{ij} \in \{0, 1\}$ so that a relationship is either present ($g_{ij} = 1$) or absent ($g_{ij} = 0$), although weighted and/or directed (as well as dynamic!) cases are clearly of interest as well.

A graph is *undirected* if \mathbf{g} is required to be symmetric so that $g_{ij} = g_{ji}$, and is *directed* otherwise. Whether or not a network is directed or undirected depends on the application. In applications where mutual consent is required to maintain a relationship (friendships, alliances, partnerships, contracts, and so forth) it will often be most appropriate to represent these as an undirected graph, while there are other applications where unilateral relationships are possible (such as one author citing another or a web page linking to another).

It is generally useful to use the notation $ij \in \mathbf{g}$ to indicate that $g_{ij} = 1$ and $ij \notin \mathbf{g}$ to indicate that $g_{ij} = 0$, and one can represent a graph by the set of links that are present (so one could alternatively represent \mathbf{g} by its set of links). I alternatively view \mathbf{g} as a matrix or a set of links depending on which is more convenient, and thus abuse notation in what follows.

A relationship between two nodes i and j , represented by $ij \in \mathbf{g}$, is referred to as a link. Links are also referred to as edges or ties in various parts of the literature; and sometimes also directed links, directed edges, or arcs in the specific case of a directed network.

A *walk* in a network (N, \mathbf{g}) refers to a sequence of nodes, $i_1, i_2, i_3, \dots, i_{K-1}, i_K$ such that $i_k i_{k+1} \in \mathbf{g}$ for each k from 1 to K .¹² The *length* of the walk is the number of links in it, or $K - 1$.

A *path* in a network (N, \mathbf{g}) is a walk in (N, \mathbf{g}) , $i_1, i_2, i_3, \dots, i_{K-1}, i_K$, such that all the nodes are distinct.

¹¹See Chapter 2 in Jackson (2008) for additional background and references.

¹²Standard definitions of walks, paths, and cycles formally define them as subnetworks of the original network, in which case they are not simply sequences of nodes, but need to be specified as sets of nodes together with sets of links. The definitions here simplify notation, and for the purposes of this chapter the difference is inconsequential.

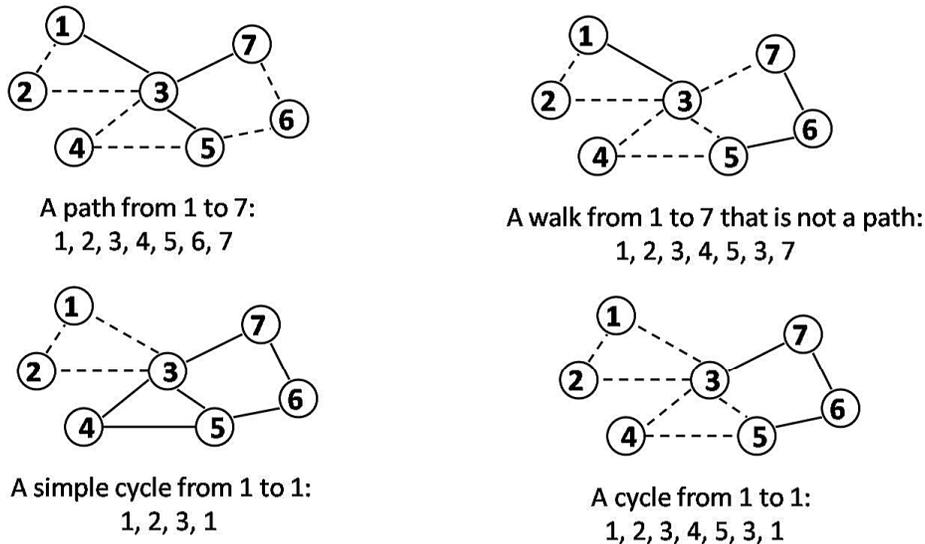


Figure 1: Paths, Walks, and Cycles

A *cycle* in a network (N, \mathbf{g}) is a walk in (N, \mathbf{g}) , $i_1, i_2, i_3, \dots, i_{K-1}, i_K$, such that $i_1 = i_K$.

A network (N, \mathbf{g}) is *connected* if there is a path in (N, \mathbf{g}) between every pair of nodes i and j .¹³

A *component* of a network (N, \mathbf{g}) is subnetwork (N', \mathbf{g}') (so $N' \subset N$ and $\mathbf{g}' \subset \mathbf{g}$) such that there is a path in \mathbf{g}' from every node $i \in N'$ to every other node $j \in N'$ ($j \neq i$), and such that every node $\ell \in N$ such that $\ell \notin N'$ has no link in \mathbf{g} to any node in N' . Thus, a component of a network is a maximal connected subgraph, so that the subgraph is connected and there is no way of expanding the set of nodes in the subgraph and still having it be connected.

The *distance* between two nodes in the same component of a network is the length of a shortest path (also known as a *geodesic*) between them.

¹³Each of these definitions has an analog for directed networks, simply viewing the pairs as directed links and then having the name directed walk, directed path, and directed cycle. In defining connectedness for a directed network one often uses a strong definition requiring a directed path from each node to every other node.

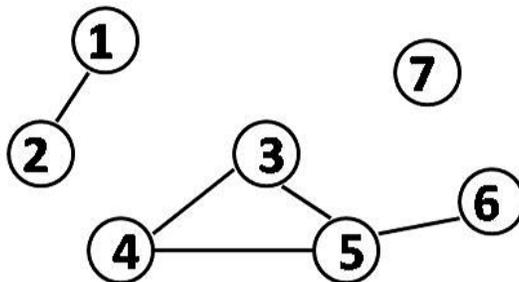


Figure 2: A Network with 3 Components.

The *neighbors* of a node i in a network (N, \mathbf{g}) are denoted¹⁴

$$N_i(\mathbf{g}) = \{j | ij \in \mathbf{g}\}$$

The *degree* of a node i in a network (N, \mathbf{g}) is the number of neighbors that i has in the network, so that $d_i(\mathbf{g}) = |N_i(\mathbf{g})|$.¹⁵

3.2 Degree Distributions

While the information contained in a full specification of all relationships, (N, \mathbf{g}) is sometimes very useful, it is generally too cumbersome when there are many nodes, and so descriptive statistics that capture facets of the network are used. For instance, knowing the average degree in the network $\sum_i d_i(\mathbf{g})/n$ gives some idea of the density of the connections in a network. However, often we need richer information and the distribution of the degrees of the nodes provides more substantial information about network structure.

The *degree distribution* of a network (N, \mathbf{g}) is the frequency distribution P of the degrees in the network. $P(d)$ indicates the fraction of nodes that have degree d .

Degree distributions vary across applications. One extreme distribution corresponds to a *regular network* such that all nodes have the same degree. A useful benchmark is a network where each link is formed at random with the same probability p and independently of all

¹⁴For the remaining definitions, I omit dependence on the set of nodes N , so for instance I write $N_i(\mathbf{g})$ rather than $N_i(N, \mathbf{g})$, as generally the set of nodes will be fixed and so only the set of connections will be varying.

¹⁵Unless otherwise stated, let us suppose that $g_{ii} = 0$, so that nodes are not linked to themselves.

other links in the network. In that case, the probability that a given node has degree d has a binomial distribution described by

$$\binom{n-1}{d} p^d (1-p)^{n-1-d}. \quad (1)$$

For large n and relatively small p , a standard approximation of a binomial distribution by a Poisson distribution applies and the probability that a node has d links is approximately

$$\frac{e^{-(n-1)p} ((n-1)p)^d}{d!}. \quad (2)$$

Such networks where all nodes are formed uniformly at random with the same probability have been studied extensively in random graph theory, including seminal papers by Erdős and Rényi (1959, 1960, 1961) and many others.¹⁶ They are often referred to as “Poisson random graphs,” due to the (approximate) degree distribution. They serve a useful benchmark and exhibit many properties that are common to many random graph models:¹⁷

- when p is very low (well below $1/n$) most nodes are completely isolated and only a few nodes are linked as pairs,
- as p increases (above $1/n$) a network begins to emerge in the sense that some nodes have more than one link and a large component (referred to as the *giant component*) begins to emerge and dominate the network,¹⁸ and cycles begin to occur,
- as p increases further (beyond $\log(n)$) the isolated nodes disappear and network begins to coalesce into a single connected component.

Another useful benchmark distribution is a power distribution such that

$$P(d) = cd^{-\gamma}$$

for some parameter γ and normalizing constant c , where the distribution is generally truncated at some upper bound. In settings where such degree distributions are prevalent it is often said that a *power law* is satisfied, and the distributions are referred to as being

¹⁶See the book by Bollobás (2001) for an overview of this literature.

¹⁷See Jackson (2008) for more detailed discussion and background. This class of networks is also referred to as Bernoulli random graphs, and even simply “G(n,p).”

¹⁸There will generally only be one large component as it is very unlikely to have two components that each have many nodes in them but with absolutely no links between the two components.

scale-free. The scale-free property refers to the fact that $P(d)/P(d') = P(kd)/P(kd')$ for any rescaling by a factor k . Such distributions have been found in a variety of settings, with prominent examples being the distributions of wealth noted by Pareto (1896) (for whom the related Pareto distribution is named), word usage, and city sizes (often referred to as Zipf's law - Zipf (1949)). An example of such a distribution in a network context was noted by Price (1965) who examined the network of citations among articles. Albert, Jeong and Barabasi (1999) and Huberman and Adamic (1999) found that a portion of the world wide web (examining links between web pages) fit a power distribution.

Power distributions have the nice feature that the frequency distribution can be rewritten as

$$\log(P(d)) = \log(c) - \gamma \log(d)$$

and so are linear when viewed on a log-log plot. An important feature of such a distribution is that it has “fat tails” relative to a Poisson distribution. Thus, the frequency of very high and very low degree nodes is greater than if links were formed uniformly at random, and correspondingly the frequency of nodes with degrees near the center of the distribution is lower than if links were formed uniformly at random. This distinction can lead the network to have very different properties, as very high degree nodes can serve as “hubs” and play prominent roles in different contexts as I discuss below.

There are many examples of networks whose degree distribution have fat tails, and so it sometimes said that a power law is satisfied by many networks (e.g., see Barabasi (2002)). Nevertheless, social networks exhibit a full spectrum of degree distributions across different applications, ranging from one extreme where the distribution of links is nearly as if they were formed uniformly at random (e.g., matched well by distributions of romances among high school students in the Add-Health data set) and another extreme where there the distribution is nearly scale-free (e.g., the www from Albert, Jeong, and Barabasi (1999) and Huberman and Adamic (1999)). Thus, although many networks have fatter tails than one would see uniformly at random, when statistically fitting degree distributions they can come out somewhere between the extremes of a scale-free and being formed uniformly at random, as discussed by Jackson and Rogers (2007).^{19,20}

¹⁹See also Pennock et al. (2002) for other examples.

²⁰There are also some interesting measurement biases to keep in mind here. If links are estimated by an interview process, degree can be underestimated either by some cap imposed by the interview or the memory of the interviewee; while if degree is estimated by some computer program that “crawls” from one web page to another then it can be biased towards finding larger degree nodes. There are techniques for limiting such biases, but measurement error is an endemic challenge in network estimation.

3.3 Average Distances and Small Worlds

How far apart are nodes on average in a network? Consider an individual who has 100 people with whom they are in regular contact. If each of them has 100 (different) people with whom they are in contact and so forth, then as a rough approximation the individual is at a distance of at most 2 from 10000 people, and at a distance of at most 3 from a million people, and 4 from 100 million people. With this sort of reasoning, there are on the order of $\bar{d}(\bar{d} - 1)^{k-1}$ nodes at a distance of k from a given node if each node has \bar{d} neighbors on average. With such expansion one reaches at least n nodes by moving out k steps where

$$\bar{d}^k = n$$

or

$$k = \log(n) / \log(\bar{d}).$$

This calculation is a rough one in at least two ways: first, it presumes that all people have the same number of friends and most applications exhibit substantial heterogeneity; and second, it does not account for cycles in the network in that there may be some overlap in the friends of friends. Still this calculation shows us why distances between nodes in social networks can be quite small relative to the number of nodes, since neighborhoods tend to expand exponentially as we radiate outwards from a given node. Perhaps surprisingly, $\log(n) / \log(\bar{d})$ is a very accurate estimate of the average distance between nodes and the diameter for a wide set of random networks, including the Poisson random graphs as originally shown by Erdős-Renyi (1959, 1960), networks with other sorts of degree distributions (e.g., see Chung and Lu (2002)), and even quite general ones where there are heterogeneous types of nodes and link formation is type dependent (see Jackson (2008b)).

Small average distances and diameter relative to the number of nodes has also been extensively explored empirically. One of the earliest and most famous experiments in the social network literature was conducted by Milgram (1967) and shed light on this phenomenon. Milgram had people in one part of the United States try to get a letter to a person in another part of the United States. The subjects were told limited information about the target, such as the target's name and some information about where the target lived (but not an address), and then were instructed to send the letter to someone who might be able to forward it to someone, who could forward it to someone, etc., with the intent of eventually finding someone who knew the target and could get it directly to the target. Roughly a quarter of the letters made it to their targets, with a median number of five links. This sort of result has also been many follow up studies on larger data sets, across countries, and with more

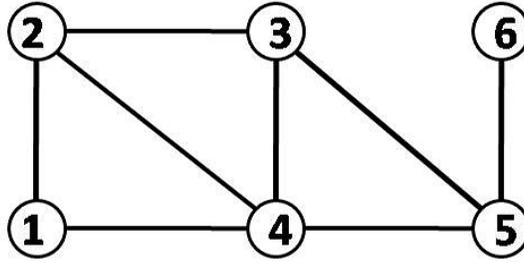


Figure 3.3.1. Clustering is 1 for nodes 1, is $2/3$ for nodes 2 and 3, is $1/2$ for node 4, $1/3$ for node 5, and 0 for node 6

detailed analyses of what strategies people used in selecting to whom they forwarded messages (e.g., see Watts (2004)). The small-world phenomenon applies not just to acquaintance networks, but also to things like links on web pages (e.g., see Adamic (1999)) and a variety of other networks (e.g., see Watts (1999, 2004)). The small average distances in networks has important implications for things like diffusion and contagion.

3.3.1 Clustering

An aspect of networks that can be important in social and economic settings is the extent to which relationships are transitive: that is, the extent to which if node i is linked to node j , and j is linked to k , then i is linked to k . The frequency with which such transitivity is present is referred to as *clustering* and is measured in various ways. For any given node, such as the node j above, we can measure the clustering relative to that node by measuring the fraction of all pairs of nodes that are both linked to j that are linked to each other. Averaging this measure of clustering across nodes then gives an idea of the extent to which such transitivity exists in a network.²¹

There are various reasons as to why clustering can be important in a network. For instance, it impacts how the extent to which connections reach out to new nodes and so can affect information transmission. It can also impact how able a group is to monitor

²¹That is, the clustering for node j is $\sum_{k \in N_j(g), i \in N_j(g), k \neq i} g_{ik} / \sum_{k \in N_j(g), i \in N_j(g), k \neq i} g_{jk} g_{ji}$. One can then average this across nodes. Alternatively, one can simply examine the overall fraction of pairs of adjacent links that are completed: $\sum_{i \neq j \neq k \neq i} g_{jk} g_{ji} g_{ik} / \sum_{i \neq j \neq k \neq i} g_{jk} g_{ji}$.

and enforce behaviors. For example, suppose that without any threat of punishment or loss of future access, an individual might have an incentive to cheat another individual in a transaction. If there is no clustering, and information only travels by word of mouth, then it might be that if an individual cheats another then he or she only faces retribution and punishment from that individual. If instead, there is substantial clustering, then the cheated individual might inform other people who are also involved in relationships with the cheating agent who can aid in retribution and punishment. The importance of clustering traces back to the pioneering social network research of Simmel (1908), and Coleman (1988) provides specific discussion of the role of clustering (or more general forms of “closure”) in enforcing social norms.²²

As Newman (2003) points out, there are a number of observed social networks that have much higher clustering than would be present in, for instance, a Poisson random network. For example, Newman discusses how networks of who has co-authored with whom in various research areas exhibit clustering rates ranging from around fifteen to fifty percent, while a Poisson random networks of similar size and density would have clustering close to zero.

Clustering can be traced to a variety of sources: it occurs quite naturally if friends meet new friends via their current friends (see Jackson and Rogers 2007). Institutional structures and geography can also affect who meets whom or who might benefit from interacting with whom.²³

Recently, Jackson, Rodriguez-Barraquer and Tan (2009) have proposed an alternative measure of closure within a network that they call *support*. A link ij in a network g is *supported* if there exists a node k who is a neighbor of both i and j , so that $ik \in g$ and $jk \in g$. One can then measure the fraction of links within a network that are supported.

Superficially, support and clustering seem to be similar measures, as they both involve triads. Nonetheless, they are quite different, and support can be quite higher than clustering. For example, consider a network with links $\{12, 13, 23, 14, 15, 45\}$. In this network all links are supported and so the support measure is 1. However, only 1/3 of the pairs of friends of agent 1 are friends with each other (e.g., 2 is a friend of 3, but not of 4 or 5). Thus the clustering in this network is much lower than the support measure.

Jackson, Rodriguez-Barraquer and Tan (2009) find that a theoretical model of favor exchange leads to specific predictions about support levels within a network, but does not

²²See Ali and Miller (2009) for a game-theoretic model of the role of clustering in enforcing cooperative behavior.

²³For more discussion of potential sources of clustering, see Watts (1999), Jackson and Rogers (2005), and Carayol and Roux (2003).

make predictions about clustering. Their model is based on examining the incentives for agents to perform costly favors for each other. In cases where the threat of losing one friend is not enough to induce an agent to perform a favor, having common friends who can provide incentives for agents to cooperate and perform costly favors via the threat of ostracism. Based on the insights from such a theory, Jackson, Rodriguez-Barraquer and Tan analyze data concerning exchange of favors as well as other relationships from 75 rural Indian villages. They find significantly higher levels of support than clustering, and that the fraction of links that are supported is higher when one examines relationships based on favor exchange than other sorts of more “hedonic” relationships.

3.3.2 Homophily

Beyond, the patterns of degrees, average distance, clustering, and other such measures that only concern network architecture, there are also patterns that relate to how links depend on other characteristics of nodes. For instance, if nodes are people, then they have identities that include things like their age, gender, ethnicity, profession, education level, as well as other behavioral attributes such as what their hobbies are, whether they smoke, their political attitudes, and so forth. When we keep track of these various characteristics of nodes, then we see further patterns in terms of which nodes are linked to which other ones. In particular, one of the most extensively studied and documented aspects of social networks structures is that nodes tend to be more frequently linked to other nodes that are similar to themselves in terms of their characteristics than to nodes that are less similar to themselves in characteristics. This is referred to as homophily, as originally named by Lazarsfeld and Merton (1954). McPherson, Smith-Lovin and Cook (2001) provide an overview of the many dimensions on which homophily has been observed and also discuss some of the potential reasons for it.

Homophily can impact behavior and welfare in a variety of ways.²⁴ For example, it can affect workers’ decisions of whether to drop out of the labor force. In particular, such decisions depend on the decisions of a worker’s friends and colleagues and are often complementary: the greater the drop-out rate of a worker’s social neighbors, the more attractive it becomes for the worker to drop out. When there is substantial homophily in a network, different groups can be quite isolated from each other and it might be that many individuals of one group drop out, while very few of another group drop out, even when the groups are otherwise similar (e.g., see Calvó-Armengol and Jackson (2004), Jackson (2007), as well as the discussion below). Homophily can similarly affect decisions of whether to invest in

²⁴See Jackson (2007) for more discussion of some of the effects mentioned here.

education (e.g., see Calvó-Armengol and Jackson (2009)). Homophily can also affect the speed of learning (e.g, see Golub and Jackson (2008)), as well as a variety of other network attributes (e.g., Currarini, Jackson, Pin (2007), Bramoullé and Rogers (2009)), and in field experiments has been found to affect things like how generous agents are towards each other (e.g., see Goeree et al (2008)).

A simple measure of homophily is as follows. Let us partition the set of nodes N into groups according to their characteristics, $N_1, \dots, N_i, \dots, N_m$, where all nodes in some group N_i have the same characteristics and nodes from different groups have different characteristics. Let n , and $n_1, \dots, n_i, \dots, n_m$ be the respective cardinalities. First let us examine how many links form in the network \mathbf{g} compared to how many could have formed, and denote this proportion by $p(\mathbf{g})$; so,

$$p(\mathbf{g}) = \frac{\sum_{j \in N} d_j(\mathbf{g})}{n(n-1)},$$

where $d_j(\mathbf{g})$ is node j 's degree. Next, let us do the same calculation but now seeing what proportion of links between nodes of the same types occur in the network \mathbf{g} compared to how many could have occurred and denote this by $p_s(\mathbf{g})$, where

$$p_s(\mathbf{g}) = \frac{\sum_{i=1, \dots, m} \sum_{j, k \in N_i} g_{jk}}{\sum_{i=1, \dots, m} n_i(n_i - 1)}.$$

Then let us define the homophily in the network to be

$$h(\mathbf{g}) = \frac{p_s(\mathbf{g})}{p(\mathbf{g})}$$

so that $h(\mathbf{g})$ is how relatively prevalent links among same-type nodes are compared to links in the network overall. If this measure turns out to be 1, then there is no bias in the link formation relative to these characteristics, at least on average. If the measure is above 1, then we observe what is generally referred to as *inbreeding homophily*, so that links are more likely to be formed within groups than across groups. It is also possible to have “out-breeding,” so that links across groups are relatively more likely than within groups, as would be the case in some sorts of trading networks or other bipartite networks. We can also examine similar measures group by group. That is, for a group $N_i \subset N$ we can define

$$p^i(\mathbf{g}) = \frac{\sum_{j \in N_i} d_j(\mathbf{g})}{n_i(n_i - 1)} \quad \text{and} \quad p_s^i(\mathbf{g}) = \frac{\sum_{j, k \in N_i} g_{jk}}{n_i(n_i - 1)}.$$

Then we define the homophily of group $N_i \subset N$ to be $h^i(\mathbf{g}) = \frac{p_s^i(\mathbf{g})}{p^i(\mathbf{g})}$.

As an illustration of homophily, Table 1 reports friendship patterns among high school students in the Adolescent Health data set that is based on interviews with over ninety thousand students at a representative sampling of U.S. high schools including urban and rural, private and public, large and small, religious and secular schools, from a variety of geographical regions, and including different ethnic and socio-economic mixes of students. Looking at 84 of the schools that had substantial network data, the following data summarizes the average of the homophily measure defined above, $h(g)$, across the 84 different high schools (so 84 networks or different g 's):²⁵

Table 1: Homophily in High School Friendships

	Groups Defined by:		
	Race	Sex	Grade
average homophily ($h(g)$) across schools	1.4	1.2	4.0
minimum homophily across schools	.99	1.0	1.5
maximum homophily across schools	2.7	1.5	5.6
standard deviation of homophily across schools	.43	.08	.90

From Table 1 we see that, on average across the schools, students are 4 times more likely to form friendships with a student in their own grade than with students overall, and are 1.4 more likely to form friendships with students of their own race than with students overall, and so forth. In these data and comparing across these characteristics, the strongest bias in relationships is by grade (that is, by year in school and so roughly by the students' age), with weaker biases by race and gender.^{26,27,28}

Homophily can occur for various reasons. For example, one would expect substantial age-based homophily in friendships of children due to the grouping of students into classes in

²⁵These data come from a joint project with Ben Golub, although we did not report them in the paper Golub and Jackson (2008). For more background on these data see Golub and Jackson (2008).

²⁶There are some measurement issues here. For example, subjects were asked to name up to five male friends and five female friends. So, a subject with fifteen male friends and five female friends would end up only naming five of each. Most subjects ended up below the caps, but there is some censoring of the data.

²⁷These data also contain information about intensity of relationships, including the number of various activities that pairs of individuals reported participating in together. If one looks only at relationships such that there are more than three interactions in a week, then the homophily along all dimensions becomes much more pronounced.

²⁸There are also various ways of normalizing homophily measures. For some discussion see Coleman (1958).

schools so that most of their contact is with other students of nearly the same age. Homophily can also be driven by preferences: students may prefer to associate with other students of the same age since their interests and maturity will be similar. The role of biases in contact as a source of homophily is discussed by Allport (1954), Blau (1977), Feld (1981), Rytina and Morgan (1982), and the role preferences for associating with individuals with similar traits, behaviors, or backgrounds is discussed by Cohen (1977), Kandel (1978), Knoke (1990), and Currarini, Jackson and Pin (2008, 2009, 2010) and Bramoullé and Rogers (2009). There are also other reasons that homophily might arise, including competition among groups (Giles and Evans (1986)), social norms and culture (Carley 1991)), institutional and organizational pressures (Meeker and Weiler (1970), Khmelkov and Hallinan (1999), Kubitschek and Hallinan (1998), Stearns (2004)). Empirical work based on models that allow for more than one source of homophily as in Currarini, Jackson and Pin (2009, 2010) and van der Leij and Buhai (2008) helps identify sources of homophily and can help explain various patterns of homophily.

Beyond some of the patterns discussed above, there are a variety of other patterns that can be important in various settings. For instance, we consider the diffusion of a disease through a network can be affected by whether or not high degree nodes are linked to other high degree nodes or to low degree nodes. Similarly, there are relationships between clustering and degree, degree and homophily, and a variety of other patterns in networks.²⁹

Understanding network structure is not simply interesting in its own right, but also because of its implications for decision-making and economic behavior. For example, as a society becomes more homophilous and groups become more segregated, the reaching of a consensus in word-of-mouth learning can slow significantly (e.g., see Golub and Jackson (2008)). I discuss some such implications below, but first begin with a discussion of how networks form and why they exhibit some of the features mentioned above.

4 Network Formation

Given the impact of network structure on various behaviors, it is important to understand how networks are formed and why they might have certain characteristics.

The literature on network formation has adopted three main approaches.

- One approach originates in random graph theory and is process-based. Such models begin with the classic work of Erdős and Rényi (1959, 1960, 1961) and provide an

²⁹See Jackson (2008) for more background.

understanding of how certain observed features of networks (such as fat-tailed degree distributions, high clustering, low diameter, and other properties) can be traced to processes governing how links form.

- The second approach is based on building statistical models for working with social network data. This approach develops models that are versatile in terms of estimation. That is, these are models that enable one to estimate which patterns and correlations among various features appear in social network data.
- The third approach is based economic fundamentals presuming that agents choose their relationships based on the payoffs that emerge as a function of the network. Such modeling incorporates game theoretic techniques and can help indicate why certain structures emerge.

Clearly, these approaches stem from very different perspectives and goals, and they offer different insights into social networks. Let me discuss each approach in turn.³⁰

4.1 Random Networks

Random network models originate in the random graph literature where the main focus is on how specific assumptions about the random emergence of links leads to various properties of network structure. These models analyze the outcomes of particular (stochastic) algorithms by which links are created.³¹ The following are some of the canonical models (see Newman 2003 and Jackson 2008 for more background).

4.1.1 Poisson Random Networks.

This is the most basic random graph model that was mentioned above, and was independently proposed by Solomonoff and Rapoport (1951), Gilbert (1959), and Erdős and Rényi (1959, 1960, 1961) who discovered some of the seminal results on the properties of the random

³⁰Parts of Section 4.1 on random networks were co-written with Leeat Yariv and originally appeared in Jackson and Yariv (This volume), but the section was a better fit here. Thank you to Leeat for her work on it and grace in moving it here.

³¹The focus is largely on identifying simple procedures generating certain classes of degree distributions. The literature has also tackled the converse question having to do with the feasibility of general degree distributions in a network. The configuration model and various relatives (see Bender and Canfield (1978) and Chung and Lu (2002)) are such that nodes are connected in a manner to realize a pre-specified degree distribution.

graphs. Given a finite set of nodes N , a link between nodes i and j is formed independently of all other pairs of nodes with a given probability p . The degree of any given node thus follows a binomial distribution, and as the number of nodes becomes large (provided p does not grow too quickly), this is well-approximated by a Poisson distribution.

The main insights from this literature are that there are specific thresholds in terms of the link probability, such that the networks have distinct properties above and below the thresholds. As mentioned before, there is a threshold at $p = 1/n$ (where each node expects to have a single neighbor) at which some cycles and a giant component emerge. That is, if p over $1/n$ goes to 0, then with a probability going to 1 there are no cycles and no component containing more than a vanishing set of nodes, while if p over $1/n$ goes to infinity, then with a probability going to 1 there are cycles and a (single) giant component containing a non-vanishing fraction of the networks. Such “large network” thresholds and properties are the primary basis for analysis of random networks.

4.1.2 The Small World model.

Watts and Strogatz (1998) noted that basic random network models failed to capture an important feature of many observed networks: the combination of relatively small diameters and high levels of clustering. Unless each node is connected to a nontrivial fraction of all other nodes (which is clearly not true of most large social networks), a Poisson random network will have vanishing clustering. So, in order to maintain the small diameters of a (connected) Poisson random network, but also to obtain high clustering Watts and Strogatz (1998) constructed a model in which nodes are initially connected according to a highly clustered lattice. For example, think of having nodes located on a circle, and each connected to neighbors that are of distance k steps or lower on the circle. This initial configuration has high clustering, but will not have a small diameter. Next, some fraction of links are severed and reconnected uniformly at random. That is, a given link is removed with probability π and then rewired to a node chosen uniformly at random from those to which it is not already connected. The probability π is a measure of the randomness of the network. For $\pi = 0$, the network is a lattice, while for $\pi = 1$ the network is effectively a Poisson random network. The small world model is so named since even for small re-wiring probabilities π , the distance between any two nodes on the network is significantly smaller than in the original lattice and similar to the average distance of a Poisson random network, but while keeping the high clustering of the lattice. As Watts and Strogatz (1998) show, there is a fairly wide range of parameters for which the network maintains the dual properties of high clustering of the

lattice and relatively low average distance between nodes of a Poisson random network.

4.1.3 Preferential Attachment

While the Watts and Strogatz (1998) model exhibits high clustering and small average distances, it does not match some other characteristics of observed social networks. In particular, the degree distribution can be quite unlike that of most observed networks as it looks like a mix between a regular network (where all nodes have the same degree) and a Poisson random network, and certainly does not exhibit the fatter tailed distributions seen in some applications. In order to generate degree distributions with such fat tails, like a power distribution, one needs other sorts of formation processes.

An early version of such a process is suggested by Price (1973) in the context of citation networks, and Barabási and Albert (1999) show how a simple model can be applied quite generally to result in networks where the degree distribution is scale-free; that is, satisfies a power law so that the frequency of degree d is proportional to $d^{-\gamma}$ for a parameter γ . The two essential features of the model³² are that (i) nodes enter over time, and so we can index them by their date of birth $i = 1, 2, \dots, n$, and (ii) each newborn node forms a given number of links, say m , to the existing nodes in a manner that Barabási and Albert (1999) refer to as “preferential attachment.” The idea is that the newborn node i chooses which m of the existing nodes $1, \dots, i - 1$ to link to with probabilities proportional to the number of links that those nodes already have. For example, if node $j < i$ has 10 links and node $k < i$ has 5 links, then node i is twice as likely to attach a given link to node j as to node k . This sort of process exhibits a “rich-get-richer” pattern so that nodes that have high degree grow in degree over time more rapidly than nodes with low degree, leading to the fat-tails in the distribution. Such networks end up with a sort of “hub-and-spoke” pattern to them, with some nodes with very high degrees that act like “hubs” and help connect the large number of small degree nodes to the rest of the network. These networks can have even smaller average distances than in Poisson random networks with similar average degree.

4.1.4 Richer Sequential Link Formation Models.

While the preferential attachment model provides insight into what might generate fat-tailed degree distributions, it exhibits negligible clustering in large networks and hence also fails to match many observed networks when considering the broader set of characteristics that

³²The study of power distributions has a rich history, and these features of the model reflect those found to generate power laws in other settings. See Mitzenmacher (2004) for an overview.

they exhibit. Moreover, as pointed out by Pennock et al (2002), many observed networks have degree distributions that lie somewhere between that of a Poisson random network and one formed by preferential attachment.

Models developed by Vasquez (2003) and Jackson and Rogers (2007) span between uniformly random link formation and something like preferential attachment. The key to these models is that they have some combination of links formed uniformly at random and others that are based on existing network structure. For instance, by first finding some nodes uniformly at random, and then finding others by meeting some of those nodes friends, the friends of friends that are met will tend to be those who have many friends. That is, if we locate nodes by meeting friends of friends, then a node with twice as many friends as another node is twice as likely to be found via such a process. Thus, one ends up with a sort of preferential attachment because nodes are found via the existing network structure. As the ratio of how many links are formed through uniformly random meetings, and how many are formed by searching through the existing network, these models span a set of degree distributions, with extremes of a scale-free distribution and a growing version of the Erdős-Renyi uniformly random world. Moreover, as Jackson and Rogers (2007) point out, some versions of these models also have naturally high clustering, as well as correlations in degrees among neighbors, decreasing clustering with degree, and other features matching observed networks. For example, high clustering emerges naturally since some links are formed via meeting friends of friends, and such links naturally result in a triad and so result in clustering. Such models can be fit to observed networks to estimate the extent to which links are formed uniformly at random versus through meetings determined by the existing network, as shown by Jackson and Rogers (2007). In fitting such models to data, it becomes clear that degree distributions are quite varied, with some friendship networks looking almost uniformly random and some other networks exhibiting formation based on existing network structure like a friends of friends meeting process.

4.2 Models for Statistical Analysis

Although some of the random-graph-style models listed above can be used in empirical analyses of social networks, they are stark in terms of the characteristics they incorporate. They are useful for understanding how social networks come to exhibit some features that they do, but the models need to be enriched in order to examine things like homophily or how various node characteristics, and specific local network patterns, influence network formation. Such an emphasis has led to the development of another class of models for network analysis

are that I will refer to as “statistical models,” since they were developed specifically for the empirical analysis of social networks. As social networks are naturally complicated, models that look for regularities, patterns, and uncover various formation properties, can be vital to understanding social networks. This is a large topic on its own, and so here I offer overviews of a couple of the most prevalent classes of statistical models: exponential random graph models and community detection models.

4.2.1 Exponential Random Graph Models and \mathbf{p}^* models

The starting point of exponential random graph models, called “ERGMs” for short, is to express the probability, $\Pr(\mathbf{g})$, that a given network \mathbf{g} arises as a function of a set of K different statistics of the network $\{s_k(\mathbf{g})\}_1^K$. For instance, the statistics could include the number of links in the network, the number of triads (completely connected triples of nodes) in the network, and so forth. The purpose of doing this is to test for various correlation patterns. For example, are networks with certain patterns, e.g., clustering or other forms of closure, more likely to appear than networks without such patterns? While such models are not necessarily well-suited for identifying causal relationships, they can be quite useful for identifying certain patterns in networks.

A standard formulation of this class of models one where

$$\Pr(\mathbf{g}) = \frac{\exp(\sum_k \beta_k s_k(\mathbf{g}))}{c(\boldsymbol{\beta})}. \quad (3)$$

So, the probability that a given network is formed depends on certain patterns that it exhibits, which are specified as the statistics s_k ’s on the right hand side.

For example, the special case of Erdős-Rényi random graphs with a probability of a link of p is expressed by setting $K = 1$ and letting $s_1(\mathbf{g})$ be the number of links in the network. In that case,

$$\Pr(\mathbf{g}) = p^{s_1(\mathbf{g})} (1 - p)^{n(n-1)/2 - s_1(\mathbf{g})}.$$

If we let $\beta_1 = \log(p/(1 - p))$ and $c(\boldsymbol{\beta}) = (1 - p)^{-n(n-1)/2}$, then this is expressed as in the form of (3).

More generally, the point of allowing for a range of statistics in (3) is to investigate other patterns that might be present in the network. In the Erdős-Rényi random graph setting the links are independent and so any clustering that occurs is simply uniformly at random and simply occurs based on the link probability. As already discussed, clustering in many observed social networks is significantly higher than would occur uniformly at random. Allowing for richer statistics to govern network formation probability, allows one to incorporate

a range of dependencies into a network. For example, if we want to see whether clustering is statistically significant, possibly in the presence of other attributes of a network, then we can include a statistic s_k which counts the number of triads in the network. Other statistics that are often included are the number of various types of “stars”, where there is some node connected to some given number of other nodes. For instance if we are examining a network of marriages, then we should not see any stars, whereas in networks that are nearly scale-free we would expect to see some very large stars. Some of the seminal work on this subject, by Frank and Strauss (1986), built a model that included triangles and various stars.

There are several practical difficulties in estimating an ERGM. The first of these depends on the model’s specification. In looking at the model specification in (3) we see that the network appears on both the left and implicitly on the right hand sides of the expression. Since this is a nonlinear specification, it is generally not possible to reduce this to a simple expression. Moreover, when estimating an ERGM we are nominally working with a single observation as we generally see only one realized network, \mathbf{g} . Thus in order for this to make statistical sense, implicitly there must be much more information that we take advantage of than just one observation, and in particular the formulation in terms of statistics generally includes information about local parts of the network (links, triads, local star formations, etc.) that can lead to many implicit observations in a single network. For example, in the case of an Erdős-Rényi random graph, we can think of the observation taking place at the link level, and so we have $n(n - 1)/2$ independent observations if there are n nodes. This allows for a very accurate estimation of the link probability. The extent to which we are at one extreme with the observed network as a single datum, or the other extreme with each link being an independent observation, or somewhere in between, depends on the specification of the ERGM and how rich and interdependent the statistics $s_k(\mathbf{g})$ are. If the statistics involve large parts of the network, or correlations between various portions of the network, then we cut down on the number of independent observations that are in the single network. The richer the model becomes in terms of interdependencies, the fewer implicit observations there are from a given network on which to estimate the model.

Beyond limitations of observations, a central difficulty in estimating the coefficients in (3), and one that is often most limiting in practice, comes from the fact that the normalizing coefficient $c(\boldsymbol{\beta})$ is effectively impossible to compute for large networks. In particular, for the right hand side of (3) to be a probability it must be that when we sum the expression in (3)

across networks g that it sums to 1. Thus, the normalizing coefficient must satisfy

$$c(\boldsymbol{\beta}) = \sum_{\mathbf{g}} \exp \left(\sum_k \beta_k s_k(\mathbf{g}) \right).$$

The summation on the right hand side is over all possible networks, of which there are $2^{n(n-1)/2}$ in the undirected case and $2^{n(n-1)}$ in the directed case. With only 30 nodes, this is already more than 2^{435} networks in the undirected case, which is more than the estimated number of atoms in the universe (on the order of 2^{270})!. Thus, unless there is some intuitive way to deduce the normalizing coefficient, one is forced to avoid the use of the normalizing coefficient in estimating the coefficients in (3). Thus in order to estimate a model of the form (3), we have to get around calculating the normalizing coefficient.

To get work around the normalizing coefficient, it is useful to work at the link level and to think about the probability that a given link ij takes on a certain value conditional on the rest of the network:

$$\Pr(g_{ij} = 1 | \mathbf{g}_{-ij}) = \frac{\exp(\sum_k \beta_k s_k(g_{ij} = 1, \mathbf{g}_{-ij}))}{\exp(\sum_k \beta_k s_k(g_{ij} = 1, \mathbf{g}_{-ij})) + \exp(\sum_k \beta_k s_k(g_{ij} = 0, \mathbf{g}_{-ij}))}.$$

We no longer have the normalizing parameter in this equation. We can similarly deduce the odds ratio

$$\frac{\Pr(g_{ij} = 1 | \mathbf{g}_{-ij})}{\Pr(g_{ij} = 0 | \mathbf{g}_{-ij})} = \frac{\exp(\sum_k \beta_k s_k(g_{ij} = 1, \mathbf{g}_{-ij}))}{\exp(\sum_k \beta_k s_k(g_{ij} = 0, \mathbf{g}_{-ij}))}.$$

Thus, the log odds ratio is

$$\log \left(\frac{\Pr(g_{ij} = 1 | \mathbf{g}_{-ij})}{\Pr(g_{ij} = 0 | \mathbf{g}_{-ij})} \right) = \sum_k \beta_k [s_k(g_{ij} = 1, \mathbf{g}_{-ij}) - s_k(g_{ij} = 0, \mathbf{g}_{-ij})].$$

or

$$\log \left(\frac{\Pr(g_{ij} = 1 | \mathbf{g}_{-ij})}{\Pr(g_{ij} = 0 | \mathbf{g}_{-ij})} \right) = \sum_k \beta_k \Delta s_k(g_{ij}, \mathbf{g}_{-ij}),$$

where $s_k(g_{ij}, \mathbf{g}_{-ij}) = s_k(g_{ij} = 1, \mathbf{g}_{-ij}) - s_k(g_{ij} = 0, \mathbf{g}_{-ij})$. (4.2.1) looks almost like a standard logit calculation used in a logistic regression. If the links were independent of each other, then this would be a standard logistic regression. Indeed, one technique for estimating the coefficients in (3) is a “pseudolikelihood” technique, where one effectively ignores the interdependencies in (4.2.1) and works with a formulation of the form:

$$\log \left(\frac{\Pr(g_{ij} = 1)}{\Pr(g_{ij} = 0)} \right) = \sum_k \beta_k \Delta s_k(g_{ij}, \mathbf{g}_{-ij}),$$

where we have eliminated the conditional distributions on the left hand side. This can then be maximized by mimicking standard techniques of maximum likelihood. Unfortunately, this can lead to (very) inaccurate estimates since it is ignoring the interdependencies that were the real purpose for exploring such a model in the first place, and conditions for it to be a reasonable technique are not well understood (e.g., see van Duijn, Giles, and Handcock (2009)).

What has led to the recent surge in the use of ERGMs are advances in Monte Carlo simulation have provided techniques for estimating such models, and in particular MCMC (Markov chain Monte Carlo) techniques which are becoming increasingly manageable and in some cases much more accurate than the pseudolikelihood techniques that ignore dependencies. A key technique along this line was introduced by Snijders (2002) and in a different variation by Handcock (2003). Let me describe the basic ideas.

The method relies on generating a distribution of different g 's that emerges for any given specification of β s. Then one can search over the set of β s (more on this below) to find one that leads to the highest likelihood of getting a network that looks similar to the observed g . So, how do we generate a distribution of different g 's for a given specification of β ? We can do this by fixing a starting network g^0 . Then we randomly pick a link to change, ij . Then, based on (4.2.1), one can randomly put the link ij in or out with with the appropriate probability given the profile of parameters β and given the g_{-ij}^0 . This leads to a new network g^1 . Now, let us iteratively do this, cycling through the different links ij (possibly randomly). This results in a Markov chain over the resulting networks, and over time the probability that we visit any given network approaches that of its steady-state distribution. Provided the model is such that only a small number of networks get visited very frequently (the critical condition for this technique to work well), then this will converge reasonably well in a limited time and so for each given β we get an estimation of the relative likelihood of different networks. Then we search across β 's by various techniques (e.g., Metropolis-Hastings algorithm or Gibbs Sampling, etc., see Snijders et al. (2006)) to find a specification that leads to the highest likelihood of the observed network or something similar to it.

These techniques may or may not end up overcoming the difficulties in estimating the ERGM model. Even with an MCMC method, we are still only sampling relatively few networks relative to the huge number possible (recall that there are 2^{435} networks on just 30 nodes...). As such, the Markov chain might not converge to the steady state distribution, or even close to it, in a reasonable time. This problem is particularly acute if networks that are quite different from each other can have similar likelihoods and/or there are local basins of attraction among networks (e.g., the distribution over networks is multi-modal)

which can happen quite easily when there are complementarities or interdependencies among links. With very large data sets or challenging specifications of an ERGM this can be almost hopeless. There are many researchers working on improving techniques, but there are still many cases without convergence in reasonable time. This can make the output from such analyses difficult to interpret, and some of the significance tests that have been developed should be interpreted with the appropriate caution since some of them are based on asymptotic properties that may not be well-satisfied in practice.

In spite of the challenges that accompany their estimation, the use of ERGM and related models is rapidly increasing since they help researchers detect a multitude of network patterns and can be tailored to look for very specific sorts of effects. Moreover, in addition to statistics about network structure, we can also include various observed attributes of agents such as socio-economic and demographic data as independent variables affecting the likelihood of links.

4.2.2 Fitting Random and Strategic Models

Quite complementary to statistical models such as the ERGMs, there are important reasons for also using “structural estimation” methods of network analysis based on more foundational models. The point is a familiar one: structural models can help disentangle causal relationships into which other models might offer little insight. As an example, consider homophily. Working with the adolescent Health data set discussed above various researchers, including Moody (1991) and Goodreau, Kitts, Morris (2009), have fit statistical network models and found that races have different propensities to form friendships with each other. While such analyses show that propensities to form friendships depend on whether the students are of the same race, such models cannot identify whether such homophily is due to preferences for same-race friendships, or instead due to differences in how frequently students of different races meet each other, or some combination of both. Is the lower rate of inter-racial friendships due to the segregation of students through the classes and activities that they participate in, so that students rarely interact with students of other races, or do students prefer to form friendships with others of their own race? Currarini, Jackson and Pin (2009, 2010) develop models of network formation that allow for biases in both preferences and meeting probabilities. Through a characterization of equilibrium conditions, Currarini, Jackson and Pin show that preference biases can be identified from patterns in the number of friendships formed based on the mix of races in a school, and biases in meeting probabilities can be identified via patterns in the homophily as a function of racial composition

of a school. They find that both preference and meeting biases are significant and that the extent and structure of these biases differ significantly across races.

Of course, any model, statistical or structural, necessarily omits some relationships and thus can lead to incorrect conclusions. This makes it important to combine the use of statistical models, which help uncover patterns and critical correlations among variables, with structural models that can help identify the factors underlying the patterns.

4.2.3 Community Detection

Another type of statistical modeling of network analysis relates to “community structure” and detection. The basic idea is that a society has natural underlying “communities” that the researcher wants to discover from examining social network data. This can be seen as a special case of detecting some latent structure that generates or influences observed behaviors or data; an idea that has a long history in anthropology and sociology (e.g., see Lévi-Strauss (1958) and Lazarsfeld and Henry (1968)). As an illustration, given a network of scientific collaborations one might wish to identify the natural communities or disciplines that influence the relationships, as those might not coincide with standard disciplinary boundaries and organizations. Or, one might have data on a labor market and wish to investigate whether there are biases in hiring. Such techniques can also be useful in reducing very large networks to smaller networks between communities. The literature evolved to include a variety of approaches to detecting or identifying the communities that underly a network. I will not survey that literature here (see Newman (2004) and chapter 13 in Jackson (2008) for more background), but let me provide a quick summary of the landscape.

Let us think of a community structure as a partition of the set of nodes, with each element of the partition referred to as a community. The idea of community detection is to uncover the community structure from an observed network. This has its roots in what is known as “block modeling” where one identifies blocks of nodes that are comparable or equivalent, as introduced by Lorrain and White (1971) and White, Boorman and Breiger (1976). It is rare to find nodes that are completely interchangeable in a network (so that their relationships with all other nodes is identical), and so strict definitions of equivalence are often too restrictive to be useful in identifying communities of nodes. Thus, one needs to loosen the approach to categorizing nodes as belonging to the same community or class of nodes. An early and popular method for doing this is based on an algorithm called CONCOR (for “convergence of iterated correlations”), as developed by Breiger, Boorman, and Arabie (1975). CONCOR is based on correlation patterns among the connections that nodes have. The idea is that

if two nodes are have similar connections, then they should belong to the same community. CONCOR thus loosens the idea of two nodes having identical relationships with other nodes to instead having a high level of correlation in their relationships. CONCOR iterates on this idea, studying the correlation patterns in the correlation patterns among nodes, with the idea that nodes in the same community should not only have similar connections, but also similar correlation patterns with other nodes, and correlation in correlation patterns, and so forth. There is a whole class of loosely related approaches that build up communities by grouping similar nodes together where similar is based on some measurement of the similarity of their connection patterns. Communities coalesce as nodes are groups of nodes are declared similar, and so one ends up with a hierarchy of communities depending on when one stops the process. A variety of such methods is known as “hierarchical clustering.”

Another branch of community detection methods originates from the computer science literature and amounts to repeatedly bisecting a network. This works by, for instance, minimizing the number of links between two comparably sized groups (see Newman (2004) for background on some of those algorithms). Here the idea is not so much based on similarity of nodes, but instead based on a notion that separate communities should have few links between them. A related approach in terms of a starting point that communities should be sets of nodes with few links between them is based on edge removal. For example, Girvan and Newman (2002) developed a popular algorithm that repeatedly deletes links from a network based on finding links that have very high measures of betweenness. The premise is that a high betweenness score means that a link must be joining two fairly disjoint groups, which could naturally be different collections of communities. As one iteratively removes edges, the network naturally fragments and the resulting component structure leads to a partition of the set of nodes which can be thought of as a community structure.

Each of the above-described methods faces a question of when to stop. One can either continue bisecting until each node is its own community, or one can keep putting nodes together in groups building up until all the nodes are in one community. Part of the difficulty with using these methods stems from the fact that there is no natural underlying notion of exactly what communities are or how the network was formed, and so the decision of when two nodes should be in the same or different communities is somewhat subjective.

A very different approach is based on a model of the role of community structures in generating a network. The idea is based on a random network formation model. As a simple example, consider a simple variation on a Poisson random network, where instead of all links forming with the same probability nodes within the same community are linked to each other with a higher probability than nodes in different communities. With such models, community

detection becomes a natural statistical exercise (e.g., see Holland, Laskey, and Leinhardt (1983), Snijders and Nowicki (1997)). This has been referred to as “a posteriori block modeling,” in the sociology literature.³³ Essentially, there is homophily, but the researcher does not directly observe the communities that underly the homophily and so seeks to recover it. A class of such community detection techniques based on maximum likelihood estimation have been axiomatized by Copic, Jackson, and Kirman (2009).

The basic challenge facing community detection techniques is similar to that facing the estimation of ERGMs: the potential number of community structures is a factorial function of the number of nodes and so exhaustive search for one that was most likely to generate the data is impossible when the network involves more than a handful of nodes. Thus, methods that are based on underlying models of community structure (e.g., the likelihood methods) have solid foundations but face difficulties in implementation, while in contrast methods that are defined by their algorithms can be quite tractable but can also be very ad hoc and difficult to interpret.

4.3 Strategic Network Formation

A completely different approach to modeling network formation originates in the economics literature and examines the consequences of agents’ choices of relationships. The basic premise is that agents choose relationships in order to maximize their well-being. These may be individuals choosing friendships that make them happy or otherwise benefit them, or firms choosing other firms with which to transact, or firms choosing which workers to hire, and so forth.

Externalities abound in network settings, as agents are generally impacted not only by their choice of friends, but also their friends’ choices of friends, and so forth. For example, in a co-authorship relationship an author is affected by how many other researchers with whom his or her co-author communicates. Those other relationships impact both the co-author’s experience and knowledge and also affect how busy the co-author might be, and thus such co-author of co-author relationships can have both positive and negative external effects. As another example, in terms of obtaining information and favors, an individual might prefer, all else held equal, to be friends with someone who has a larger number of contacts rather than

³³A more general form of this is referred to as “latent space estimation,” (e.g., see Hoff, Raftery, and Handcock (2002) and Hoff (2006)) where there may be a more specific spatial structure (where this may be a “social space”) that underlies the interconnection of nodes that one wishes to recover or detect based on social network data. An obvious case is where individuals belong to multiple groups at once, or have various attributes that affect their interconnectivity, rather than each residing in a single community.

a smaller number of contacts. A country that forms a military alliance with another country will care about which other military alliances are in place. It is easy to see that in many, if not most, networked settings an individual's decision of whom to maintain relationships with has both direct effects on those involved in the relationship and indirect effects on others in the network. Thus, understanding which relationships will tend to emerge when individuals react to such incentives is paramount to understanding which networks we expect to see and what the consequences will ultimately be for the society's welfare.

An early model that incorporated individual decision making in a network setting is due to Boorman (1975) who examined a labor market setting and the trade-offs between maintaining 'strong' and 'weak' ties. Boorman was interested in understanding the tradeoff that individuals faced in terms of maintaining a few strong ties versus many weak ties. Examples of individual choice of relationships also emerged in the cooperative game theory literature, including the formation of a graph in a context where the network of connections would impact the structure of the cooperative game and thus ultimately the payoffs of different agents, as studied by Aumann and Myerson (1988). The modeling of strategic formation in a general network setting originates in a paper by Jackson and Wolinsky (1996), who modeled payoffs to individuals as a function of the network, and then examined individual incentives to form networks. The literature on this subject is covered in more detail in Bloch and Dutta (This volume) (see also Jackson (2003, 2008)), and so here I just present an example that introduces some of the ideas and themes.

The main thrust of the early literature on this subject was to understand if and when individual incentives to form and maintain links would lead to socially efficient networks. The following example, from Jackson and Wolinsky (1996), illustrates some of the basic ideas and themes. It is one of the simplest possible cases where we begin to see network externalities emerge as it has just three agents, but it makes the issues very clear.

Each agent is a node and gets a payoff that depends on the network structure that emerges in the society. The payoff that an agent gets can depend on the full configuration of the network, and for instance agent 1's payoff might be affected by whether agents 2 and 3 are friends. A simple example of possible payoffs to nodes is pictured in Figure 4.3.

The arrows in Figure 4.3 indicate the incentives that agents have to add or delete links. An arrow points from one network to another network where some link is added whenever both of the agents involved in that link would weakly benefit from adding the link to the network with at least one of them benefiting strictly. An arrow points from one network to another network where some link is deleted if one of the two agents involved in the link would strictly benefit from deleting it. Sequences of such arrows leading from one network

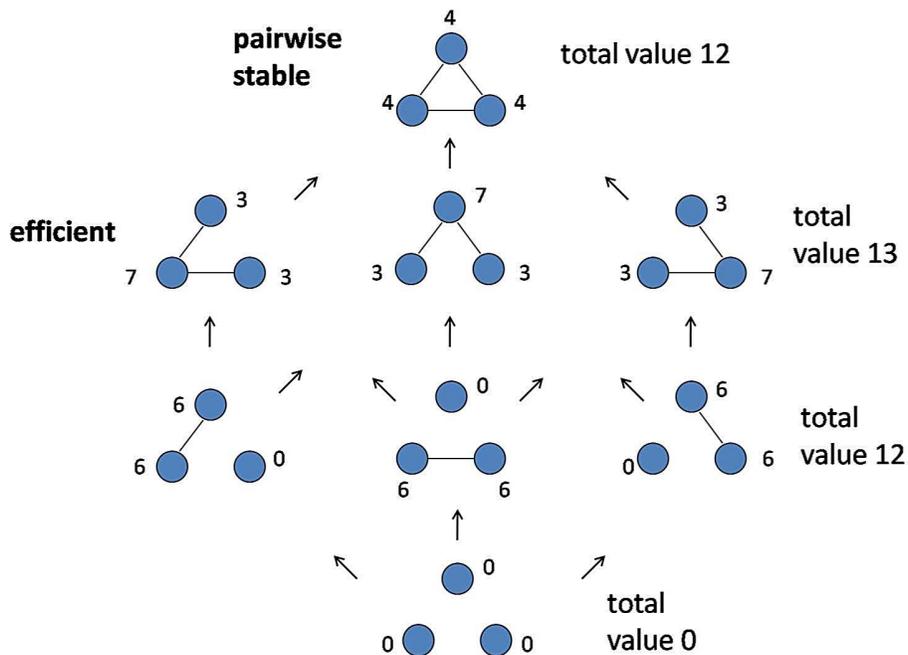


Figure 4.3. Payoffs to each node as a function of the network. Two-link networks result in the maximal overall payoffs. The arrows indicate changes networks that benefit both nodes associated with adding a link or at least one node who can delete a link. The unique pairwise stable network in this simple example is the complete network.

to another is what has been called an “improving path” by Jackson and Watts (2002), and any network with no arrows leaving it is called a pairwise stable network, as defined by Jackson and Wolinsky (1996).³⁴

If just two agents are connected, so that there is just one link in the network, then those two agents benefit equally from the relationship and each of their payoffs is 6. Adding a second link leads to an increase in payoff for the center node, but with a lower marginal payoff than the first link. The center node ends up with a payoff of 7 and the peripheral nodes get a payoff of 3 each. Here we see the incentives to form links: starting a one-link network, if a second link is added, then the agent who now has two links has seen an increase in his or her payoff (7 compared to 6), and the newly linked peripheral player gets a payoff

³⁴There are a variety of different solutions that can be used to model network formation, and many of them coincide in this example. For more background, see the chapter by Bloch and Dutta (This volume), and Chapters 6 and 11 in Jackson (2008). Bloch and Jackson (2006) provide comparisons of many of the solution concepts.

of 3, which is better than being isolated with a payoff of 0. There is a negative externality, however, from the addition of this second link. One of the agent's payoff goes down from 6 to 3 from the addition of the second link. The peripheral agents lose value from each other's presence. Such an effect could be present in a variety of settings, whenever agents compete with each other for the attention of the resources of an agent to whom they are both connected. Overall, however, the total payoff for the society is highest with two links: the benefit to the center node and second peripheral node outweigh the loss in payoff to the first peripheral node. This is not the end of the story however. The two peripheral agents now have an incentive to add the third possible link, as their payoffs each go up, from 3 to 4. Again, this has a negative externality, as even though each of their payoffs goes up, the center agent's payoff goes down from 7 to 4. In fact, adding the third link in the society reduces total payoff from 13 to 12. Effectively, the third link's cost outweighs its marginal contribution to the society. Nonetheless, the two peripheral agents end up gaining by adding the third link. This sort of effect can be present in many bargaining situations, where without the third link the center agent is in a strong bargaining position, and with the third link that agent's position is weakened and so the other agents have an incentive to add the third link to strengthen their bargaining position even though it is destructive in an overall sense for the society. This example exhibits negative externalities, in the sense that adding a link generally reduces the payoff of the third agent who is not involved in the link. In this example, the incentive is for agents to keep adding links and the the unique pairwise stable network is the complete network.

We can also consider a different variation on this simple example that has the same total payoff to society as a function of the network, but a different allocation of the payoffs to the agents in the two-link networks, as pictured in Figure 4.3. Here, in a two-link network the peripheral agents get a higher payoff than the center agent, who bears a higher total cost from maintaining two relationships rather than just one. In this case, we no longer see the incentive for the peripheral agents to add the third link, and so with this modified allocation of payoffs we no longer see an incentive to over-connect relative to what is efficient. However, in passing from a one-link to a two-link network, the center node bears most of the marginal cost of forming the second link and ends up with a lower payoff in the two-link network than in a one-link network. Thus, the center node would prefer to sever one of the links in a two-link network. In this variation of the example, the only pairwise stable networks are now the one-link networks, as we see via the arrows in Figure 4.3.

By changing the allocation of the payoffs in the two-link networks we saw a change from the pairwise stable networks being over-connected to being under-connected relative to that

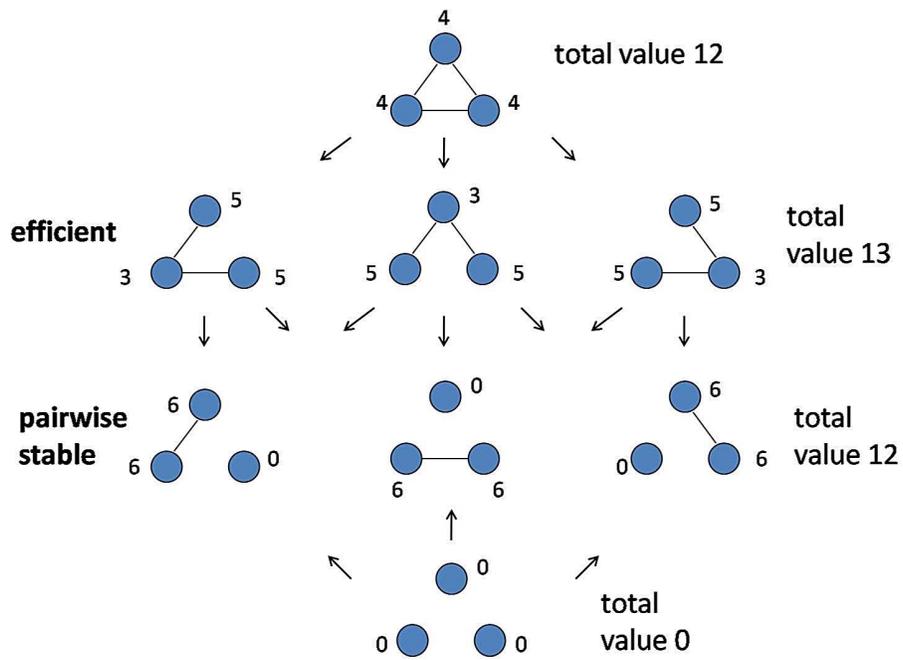


Figure 4.3. Payoffs to each node as a function of the network. Two-link networks result in the maximal overall payoffs. The arrows indicate changes in the network that benefit both nodes associated with adding a link, or one of the nodes involved in deleting a link.

The pairwise stable networks are the one-link networks.

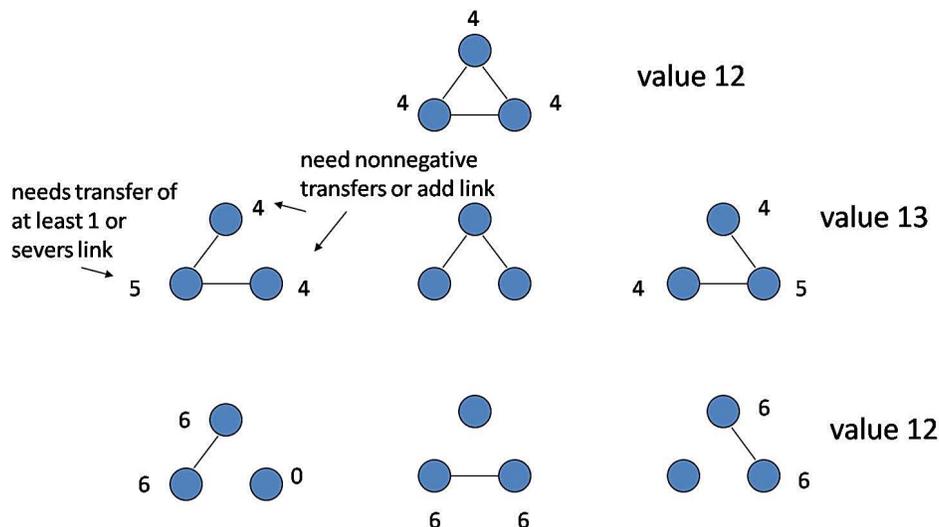


Figure 4.3. The impossibility of maintaining the total utility maximizing network as being pairwise stable, regardless of transfers between the peripheral nodes and center node.

which maximizes society’s total payoff. Is there any way in which we could allocate the payoff between the center agent and the two peripheral agents in the two-link networks so that they would be pairwise stable? It turns out that there is no way that this can be done without treating the peripheral players unequally!³⁵ This was shown in Jackson and Wolinsky (1996) and is seen in Figure 4.3. We need to give each of the peripheral agents at least a payoff of 4 in order to avoid over-connection, and we need to give the center agent a payoff of at least 6 to avoid under-connection. This sums to 14, which is greater than the total payoff of 13 which is available.

This example provides a glimpse of some of the issues that arise in strategic network formation. The tension between individual incentives and societal value is not new to economists, but there are new facets to it here. Generally, such inefficiencies arise in settings with some asymmetries of information or an inability to bargain. Here, even with full

³⁵One can solve the problem with unequal allocations, as discussed by Dutta and Mutuswami (1998). In this example, for instance, on a two link network give the center 6, and then give one peripheral agent 6 and the other 1. Such a network would be pairwise stable.

information and ability to reallocate payoffs up to some symmetry constraints,³⁶ we cannot reconcile individual incentives with economic efficiency.³⁷

This goes counter to what economists think of as a form of the “Coase theorem”: namely that with the ability to bargain or make transfers conditional on the setting, fully efficient outcomes should be obtained. The complexity of network settings, and in particular of the multilateral bargaining and multiple incentive constraints that need to be satisfied simultaneously, means that the simple logic of the Coase theorem in bilateral settings does not generalize to all multilateral settings. Whether or not transfers lead to efficient networks being stable depend on the nature of the externalities and the transfers available (e.g., see Jackson and Wolinsky (1996), Dutta and Mutuswami (2000), Currarini and Morelli (2000), Bloch and Jackson (2007)). It also depends on how agents behave when forming links, as it might be a dynamic process and they might be farsighted and anticipate each other’s reactions to their actions (e.g., see Mauleon and Vannetelbosch (2004) and Page, Wooders and Kamat (2005)), or they might be able to form links unilaterally (e.g., see Bala and Goyal (2000) and Dutta and Jackson (2000)). Broader surveys of this central and broad question of the tension between individual incentives and societal efficiency appear in Bloch and Dutta (this volume) and Jackson (2003, 2008).

Before moving on, let me make an important comment on strategic models of network formation. Such an approach to modeling is not only useful for examining tradeoffs between incentives and efficiency, but it also provides insight into some observed phenomena, like the small worlds referred to in Sections 3.3 and 4.1.2. For example, suppose that agents are located in some space, which might correspond to physical geography, but might also relate to their characteristics, such as age, profession, education, religion, etc., with closer agents being more similar than ones farther away. If the costs of links are low in terms of forming connections to nearby agents, then one would tend to see very dense connections on a local level, and thus high clustering. Links between agents that are farther apart would tend to be more costly to maintain and thus would be rarer given the higher cost. However, if there were no links that covered large distances, then there would be very substantial payoffs to

³⁶There is also a hidden constraint here in that I have not mentioned reallocating value on one-link networks. If we can give disconnected agents some of the payoffs from the network, then we can avoid the difficulty here. While that might be feasible in some cases, it is not in others. This corresponds to a component balance condition in Jackson and Wolinsky (1996).

³⁷There is also a question of which notion of efficiency is appropriate. Here we are considering maximizing the total payoff rather than using a notion of Pareto efficiency. However, the tension here generalizes to a version of Pareto efficiency that allows for reallocation of payoffs up to the symmetry and balance constraints, as argued in Jackson (2003).

such links, as they would provide closer access for an agent to many other agents that are far away. Thus, we would expect some long distance links to emerge, and at least a few such links should emerge precisely because they shorten the diameter of the network. In fact, as long as costs of long-distance links are not overwhelming, the diameter of the network is limited, since if we end up with too big a distance between some sets of nodes, then there would be substantial gains from at least one pair forming a link. Thus, a strategic formation model, with quite natural assumptions about costs and benefits can explain *why* we might expect small world network phenomena. This point appears in various forms in Johnson and Gilles (2000), Carayol and Roux (2003) and Jackson and Rogers (2005).

Thus strategic models of network formation are complementary to random-graph and statistical models of network formation, not only in their methods and approach, and the settings to which they might apply, but also in the types of insights that they provide into which network structures should emerge and why.

4.4 Mixing Random and Strategic Models

Jackson (2005) discusses the contrast and complementarities between random network and strategic models of network formation, and the need for some hybrid models. Serendipity plays a role in the relationships that people form, but so does choice. The relative roles of choice and chance can depend on the setting, and it can also be that the randomness that determines who meets whom is not uniform but depends on the context and the evolution of some process. It can also be that the choices that individuals make given their meeting opportunities end up being critical to determining the patterns of links that emerge. There are a few models of network formation that incorporate both randomness and choice, but still relatively few such models, especially given the insights that they can generate. Let me discuss a few of them.

One of the best-known such models is due to Schelling (1978), who considered a preference-based model of neighborhood formation taking into account agents' preferences to be close to other agents who are similar to themselves. Schelling devised a simple model illustrating how a city consisting of agents with different attributes (e.g., religion, race, age, etc.) who are initially integrated can “tip” into a highly segregated state just due to slight biases in preferences. This phenomenon turns out to be quite robust and provides important insights into homophily.

Schelling's basic model can be described as follows. Think of a city constructed as a checkerboard with 64 squares representing possible locations where an agent can reside.

There are two agent types: green and red. Assume the number of agents is smaller than 64 and that they are initially randomly distributed on the board. Suppose now that each agent is content if a fraction greater than $\alpha \in (0, 1)$ of his or her immediate neighbors is of his or her type. For instance, if $\alpha = \frac{1}{4}$, and all of the eight adjacent squares are occupied, then the player is content if at least two of the adjacent squares are occupied with agents of the same type. Schelling considered a simple dynamic in which at each stage one agent is randomly selected and can move if he or she is not content, and then moves to one of the nearest squares where she would be content.³⁸ Schelling inspected the effects of the initial constellation of parameters (distribution of players and taste parameters α 's). Even very integrated initial constellations can shift to very segregated ones with even slight preference biases towards being with own type: one individual moving can tip his or her former neighbors' neighborhoods past a threshold, which leads them to move, and leading to chain reactions.

As mentioned above, a model designed to directly investigate roles of choice and chance, Currarini, Jackson, and Pin (2009, 2010) consider a model in which individuals have types and preferences over the types of their friends. Friendships are formed through a random meeting process, but that process and the resulting friendships that form are influenced by individual decisions, and so there are two sources of bias that could potentially lead to homophily: bias in whom people wish to befriend, and biases in whom they meet.³⁹ They find that both of these biases are needed within the model in order to generate patterns that are consistent with two empirical facts that Currarini, Jackson and Pin identify in the data on high school friendships: as an ethnic group forms a larger percentage of a school it exhibits higher per-person average numbers of friends, and while almost all groups are homophilistic relative to the base-rate demographics, the most inbred groups are those which form middle-sized proportions of their school. They also find that these biases can differ significantly across races. Thus, at least in one setting there is some evidence suggesting that both choice and chance play important roles in determining the network that emerges, and they are responsible for different properties of the emerging network and understanding the roles of choice and chance can help explain differences in network structure across races.

Given that models that incorporate nontrivial randomness and heterogeneity along with

³⁸One can consider a variety of algorithms, with similar results. For instance, see Fagiolo, Valente and Vriend (2007).

³⁹Currarini, Jackson, and Pin (2009, 2010) do not model what underlies the preference to link with others with similar characteristics. For recent models leading to such a bias, see Peski (2007) and Baccara and Yariv (2008).

individual choice can be very hard to solve analytically, one can work with simulations. For example Carayol, Roux and Yildizoglu (2006, 2008) solve large versions of a connections model originating in Jackson and Wolinsky (1996) and are able to match some moments of various data. Such simulations provide a promising technique in fitting such models to data. The class of such models is also growing, both in terms of the development of more general statistical models (e.g., see Christakis, Imbens and Fowler (2010)) and the fitting of existing models to data (e.g., the papers of Carayol, Roux and Yildizoglu (2006, 2008) mentioned above and that of Comola (2009)).

5 Modeling the Impact of Networks

As discussed in the introduction, network structure is important because it impacts behavior and ultimately the welfare of a society. In understanding, modeling and measuring these sorts of effect, it is useful to distinguish between two sorts of situations. In one sort of situation, the impact on behavior is somewhat mechanical and not strategic. For example, in understanding the diffusion of a disease or an idea, or information about jobs, and so forth, network structure matters mainly as a conduit, and the transmission can be modeled probabilistically. In other situations, such as the trade of goods and services, the adoption of a technology, the provision of local public goods, and other decision making that is influenced by friends and acquaintances, network structure also matters but with the added features of strategic interactions between networked agents. In the first case where a network serves mainly as a conduit, much of the resulting behavior can be traced directly to network structure and attributes and some information about the process of diffusion or interaction. In the second case, the interaction between network structure and outcomes can be more complicated requiring some dynamic and/or equilibrium analysis. Let me discuss some of the issues in analyzing these sorts of effects and process, partly in the contexts of some more specific applications.

5.1 Diffusion

There are many situations where an idea, disease or behavior is transmitted from one person to another. Network structure is the primary determinant of whether diffusion occurs to a significant fraction of the society, how quickly diffusion occurs, what fraction ends up affected, and other related questions. This subject is discussed in Jackson and Yariv (This volume), and so I refer the interested reader there for more detail, and just outline some

basic points here. It is useful to start with the simplest case. Consider a situation where some nodes are initially “infected” with a disease, idea, or behavior. Suppose that then they spread this to each of their neighbors, and then those neighbors spread it to their neighbors, and so forth. In that case, it is clear that the extent of diffusion will depend on which nodes are initially affected and which components of the network they lie in. If the network is connected, then all nodes will eventually be “infected”. If not, then the extent of the diffusion will be determined by the component structure of the network. Thus, a direct description of the components is enough to understand the extent of diffusion in this simple case. Component structure of random graphs (and strategically formed graphs) is something that is well-studied and so predictions are easy to make for this simple case. This is discussed in more detail in Jackson and Yariv (This volume).

Of course, in many cases of interest, it might be that some nodes are immune to the infection, or would never choose to adopt the behavior, etc. It could also be that interaction among nodes is probabilistic or that transmission has some inherent randomness. Simple variations of diffusion incorporating variation in nodes immunity can be analyzed simply by deleting those nodes that would be immune to infection or face prohibitively high costs of adoption and considering the subnetwork that remains. Some situations where there is stochastic transmission can be incorporated via variations that allow for links to only be present with certain probabilities. It might also be that nodes can recover and become immune to becoming infected, or that nodes need repeated exposure in order to become infected. There are a variety of such models that have come out of studies of epidemiology as well as diffusion (again, see Jackson and Yariv (This volume) for background). These sorts of analyses show how network structure can directly translate into predictions about emergent behaviors in a society.

5.2 Learning

Word-of-mouth is an important means of communication and formation of opinions about subjects ranging from political elections to consumer products (see Katz and Lazarsfeld (1955) for seminal research on these subjects). Modeling the effects of network structure on how we learn and what opinions we hold involves complications relative to the pure diffusion setting above. The most obvious difference is that individuals no longer fall into simple categories such as “infected” or not, but instead might have beliefs or opinions that vary more continuously and are influenced in more complicated ways by interaction with their

neighbors. As Goyal (This volume) discusses this subject,⁴⁰ I simply focus on the major themes here.

How social learning works depends on a number of facets of the setting, including:

- is the learning observational, so that agents see others' decisions, or do agents directly communicate?
- does new information come in over time or just once?
- do agents repeatedly act or communicate?
- is the interaction simultaneous or in sequence?
- how do agents process information, via Bayesian updating or via some other process?
- do some agents have more precise information than others?

There are also a number of questions that we can answer about social learning:

- does the society reach a consensus and eventually hold similar beliefs and/or make similar decisions?
- how much influence does each agent have over societal outcomes and how does that depend on network position?
- do individuals end up accurately aggregating initially decentralized information?
- how quickly is information aggregated and how does that depend on network structure?

It is useful to begin by describing a benchmark model that comes from the seminal work on social learning of Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992). That model is observational and sequential: each agent sees one piece of information, makes a decision just once, and agents arrive in sequence with each agent getting to observe the choices (but not the information) of all of the previous agents. Agents choose between two actions, say A and B . Each agent $i \in \{1, 2, \dots\}$ sees a signal $s_i \in \{A, B\}$ that provides the agent with information about which action offers a higher payoff. Banerjee describes an example where agents are deciding between two restaurants, A and B . One restaurant has a better chef and all agents would like to go to the restaurant with the better chef, and

⁴⁰For additional background, see Sobel (2000) and Jackson (2008).

they each get an independent, equally accurate, but noisy signal about which restaurant has the better chef. Let us suppose that an agent's signal is correct in telling him or her which restaurant has the better chef with a probability $p > 1/2$, and suppose that without any signals the prior is that each restaurant is equally likely to have the better chef. Agents arrive in sequence and can see the other agents in each restaurant before making their choice. Suppose also, that in the case of indifference, an agent follows his or her own signal.⁴¹ Agents can deduce the first agent's information from observing her choice. Suppose, without loss of generality, that the choice is A . Now the second agent makes a choice. Since that agent follows his or her signal when indifferent, we can also deduce the second agent's choice from his choice. So, if the third agent observes that the first two choices were A, A , then she knows that there were two signals that indicated that A has the better chef. Regardless of the third agent's signal, she will choose A , since there are at least two signals in favor of A and at most one in favor of B . Thus, the third agent's choice will be A and will not provide any information about that agent's signal. That will be true from then on, so the society will then "herd" to the A restaurant. Agents will all observe that it has all of the agents, and correctly infer that it is more likely to be the better restaurant, based on what they can deduce from others' actions. If, instead the second agent saw a B signal, then that agent would have chosen B and so the first two choices would be A, B and would effectively cancel each other out. Those choices could effectively be ignored, and so it would be as if we started the process all over again. Almost surely, the society will eventually herd so that agents ignore their information and all end up going to the restaurant that ends up having at least two more agents than the other restaurant.

This benchmark case illustrates some of the potential outcomes of social learning. First, a consensus is eventually reached and agents end up all making the same decision. Second, and quite importantly, it is not necessarily the correct decision. It could be that restaurant B happens to have the better chef, but that the first two agents get signals saying that A has the better chef, and society herds on restaurant A even though it has the worse chef. Third, there can be some randomness in the process, so it might take some time before the herd forms, depending on the particular realization of the sequence of signals.

The conclusions in this canonical social learning example are sensitive to specifics of the setting. To begin with, suppose that agents got to observe previous agents' signals rather than their actions. In that case, by a law of large numbers, the probability that agents would be going to the restaurant with the better chef would converge to one over time. Even

⁴¹This does not affect the qualitative conclusions, but helps simplify the analysis.

without seeing signals, if periodically some agents do not observe anything and make choices just based on their own information, then simply observing such agents would eventually lead to accurate information about which restaurant had the better chef.⁴² Allowing some agents to have arbitrarily accurate signals can overcome herding, as they will have accurate enough information to lead them to go against a mistaken herd, and then this can be seen by subsequent agents providing they see the order in which actions are taken before them and are confident in the rationality of those agents. Allowing for heterogeneous preferences, and other idiosyncratic preferences that favor particular types of restaurants, can also change the results. These are the subjects of a set of papers such as Bikhchandani, Hirshleifer and Welch (1998), Smith and Sorensen (2000), Celen and Kariv (2004a,b), and Acemoglu et al (2008). For example, Acemoglu et al (2008) provide results on how social learning depends on how neighborhoods of which agents observe which other agents develop over time, and also on how accurate various agents' signals are. Although the results show that conclusions about social learning are somewhat case-specific, the results do have nice intuitions about things like the precision of information and who observes whom.

This canonical herding model has a sequentiality to it that means that social network structures do not play a prominent role. Social network structure can begin to play more of a role once agents either repeatedly take actions or repeatedly communicate with each other. A variation on the above model was investigated by Bala and Goyal (1998) and can be described as follows. Instead of taking actions just once, each agent takes an action at every date. Agents do not get signals about which is the better restaurant, but instead have experiences each time they eat, and those experiences are somewhat random but are correlated with the skill of the chef. If agents were simply acting alone, this would be a classic “two-armed-bandit” problem. For some number of periods agents would experiment and sample each of the restaurants, and eventually they would settle in to one of the restaurants.⁴³ However, agents are also connected to each other in a social network so that they observe their friends' choices and experiences at each date. Agents are boundedly rational so they do not infer anything from which choices their friends make,⁴⁴ but instead they just keep track of the

⁴²It is not even necessary that later agents know who these agents are, but just what fraction of the agents will make such choices, as then they can deduce from long enough histories whether an “incorrect” herd has occurred by noting whether there is a large enough set of agents who have gone against the herd.

⁴³This, of course, abstracts away from things like preferences for variety or non-stationary restaurant quality.

⁴⁴This poses a challenging Bayesian updating problem. If I see that a friend changes restaurants, that could tell me something about what she has learned from her friends, who could be people that I do not know. How should I weight that in my decisions? Bayes' rule provides an answer, but one that quickly

quality of all of the meals that they and their friends have experienced over time. A fairly intuitive and direct conclusion in such a setting is that the long run average outcome of all the agents will be the same. The idea is that if some agent is enjoying consistently better meals than one of his or her friends, then it must be that the agent is going to the better restaurant more frequently than his or her friend, but then that friend would come to observe this over time and so should change restaurants. Again, this does not imply that the agents all end up going to the better restaurant, but instead that there will be a consensus and in the long run all agents who lie in the same component of the network will eventually end up going to the same restaurant.⁴⁵ Another similarity to the canonical herding setting is that the conclusions depend on the specific assumptions (e.g., see Ellison and Fudenberg (1995), Gale and Kariv (2003), Rosenberg, Solan and Vielle (2007), Mueller-Frank (2009), Acemoglu, Dahleh, Lobel, Ozdaglar (2008), among others), but the basic idea that agents converge to some consensus action will carry through provided there is not too much heterogeneity in the payoffs across agents (so, for instance, agents are similar enough in what they view as a good or bad meal and restaurants treat agents similarly), and there is enough repeated viewing of neighbors' actions over time.

A challenge of Bayesian learning in social settings, both for the agents and the modeler, is that the updating becomes quite complex very quickly.⁴⁶ Even after a few periods, one is faced with a rather complicated inference problem of deducing what an agent's action indicates about that agent's friends' information. Even with a handful of agents in the simplest settings, this quickly becomes intractable. Choi, Gale, and Kariv (2005, 2009) have done some laboratory experiments on simple variations on three agent networks and found that although the strategies that agents employ show qualitative features of those employed in an equilibrium setting with fully Bayesian rational agents, the strategies of individual agents can deviate substantially, especially when the computations involved become more complicated.

A leading alternative to Bayesian updating in network settings is a model that was partly

becomes intractable even in very simple networks.

⁴⁵If the restaurants happen to have exactly the same average quality, then it is possible that the agents frequent different restaurants but they will still enjoy the same quality of meals on average. But provided there is a difference the restaurants, then agents will converge to picking the same restaurant, almost surely.

⁴⁶This presumes some bounds on communication. If every agent can tell each neighbor exactly what they have seen in every period, and what they have heard from their neighbors about all of their information, and so forth, then there is no inference problem. Such communication is obviously burdensome, and it becomes especially complicated when information is not in the form of simple signals, but in terms of subjective perceptions of the world.

described in work by French (1956) and Harary (1959), and was more completely specified and developed by DeGroot (1974). It has been used and extended by Besag (1974), Krause (2000), Friedkin and Johnsen (1997), DeMarzo, Vayanos, and Zwiebel (2003), Lorenz (2005), and Golub and Jackson (2008, 2010), among many others. I will refer to it as the “DeGroot model.”

The DeGroot model is simple and tractable, and has a number of nice properties that make it a useful benchmark both in terms of positive and normative features. The setting is one where agents observe signals just once and then repeatedly communicate with each other and update their beliefs after every round of communication. The social network is described by a weighted and possibly directed “trust” matrix $\mathbf{T} \in [0, 1]^{n \times n}$.

The idea is that T_{ij} is the weight that person i places on person j ’s opinion. The matrix is stochastic, so that $\sum_j T_{ij} = 1$ for each i , so these are really relative weights.

A simple version of this model is where society is described by an undirected social network \mathbf{g} , with $g_{ij} = 1$ indicating that i and j are linked, and then setting⁴⁷

$$T_{ij} = \frac{g_{ij}}{d_i(\mathbf{g})}, \quad (4)$$

where recall that $d_i(g)$ is i ’s degree. Thus, i places equal weight on each of his or her friends. Of course, this is just an example, and more generally agents might place different weights on different friends based on frequencies of interaction, envisioned reliability, affinity, or other reasons.

In the DeGroot model, agents begin with some initial opinions described by a belief $b_i(0) \in [0, 1]$ and then update these over time. The updating rule is just

$$b_i(t) = \sum_j T_{ij} b_j(t-1)$$

which can be written as

$$\mathbf{b}(t) = \mathbf{T}\mathbf{b}(t-1)$$

or

$$\mathbf{b}(t) = \mathbf{T}^t \mathbf{b}(0).$$

We easily see why this model is so tractable, and provides a useful benchmark, since working with a matrix raised to a power allows us to draw on substantial mathematical structure

⁴⁷This presumes that the degree of i is not 0, and to fix ideas let us consider a case where we allow $g_{ii} = 1$ so that each agent pays attention to his or her own opinion.

and knowledge and so there is much that can be said about the process that is not so easily deduced about nonlinear processes, such as Bayesian updating.

There are many interpretations of this process. For example, one interpretation is that each agent is trying to estimate some unknown parameter μ . The initial signals (the $b_i(0)$ s) are independently distributed with mean μ . If these were normally distributed signals, then at least at a first step Bayesian updating would be such that each agent would take a weighted average of his or her neighbors' signals, where the weights would be related to the precision of various friends' signals. With equal precision, the weights would be exactly those in (4). The divergence from Bayesian behavior comes from the fact that agents do not adjust their updating rule over time to account for the network structure: some friends may be talking to more people over time than others, and so forth. Despite this boundedly rational behavior, there are still many situations where the society eventually reaches a consensus that correctly approximates the unknown μ , as shown by Golub and Jackson (2010). Whether or not the DeGroot process converges to accurate estimate (and hence the Bayesian estimate) depends on how well balanced the relative weights are that different groups of agents place on each other. If there is sufficient balance in the weights then an accurate consensus is reached over time, while if things are too imbalanced then some small subset of agents' information can dominate the eventual consensus. Another very different interpretation of this process is one of myopic best responses: instead of beliefs the b_i s represent some behavior and each agent wants to match the average behavior of his or her friends (e.g., as in some of the peer effects models like Manski (1993)).

There are several nice aspects of this process. Beyond the models tractability, it allows the network to enter in a nontrivial way. In the analysis of observational learning, the conclusions that a consensus was reached in the society did not really depend on network structure, and the nature of the consensus might depend on the network structure, but in ways that researchers have not been able to deduce. In contrast, the limiting behavior of (5.2) is very easily analyzed and depends on the network structure in interesting and intuitive ways.

To get a feeling for this, let us examine a simple example, as pictured in Figure 5.2. this corresponds to a network where each agent is connected to him or herself, and there are also links between agents 1 and 2, and between agents 2 and 3. Each agent places equal weight on each friend when updating.

We see how the DeGroot updating works for this case, when starting with an initial belief vector of $\mathbf{b}(0) = (0, 1, 0)$, so that agent 2 has an initial belief of 1 and the others have an initial belief of 0. In this case, agent 1's first period belief is the average of these beliefs, and so it

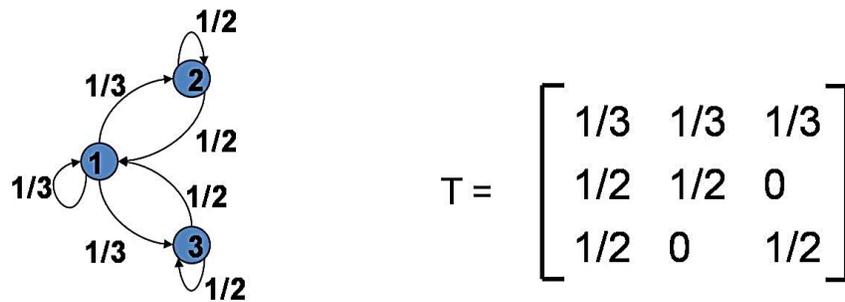


Figure 5.2. The DeGroot updating process for a case where agent 1 has links to all agents and agents 2 and 3 just link to agent 1 and themselves, and where agents equally weight all of their friends, as in (4).

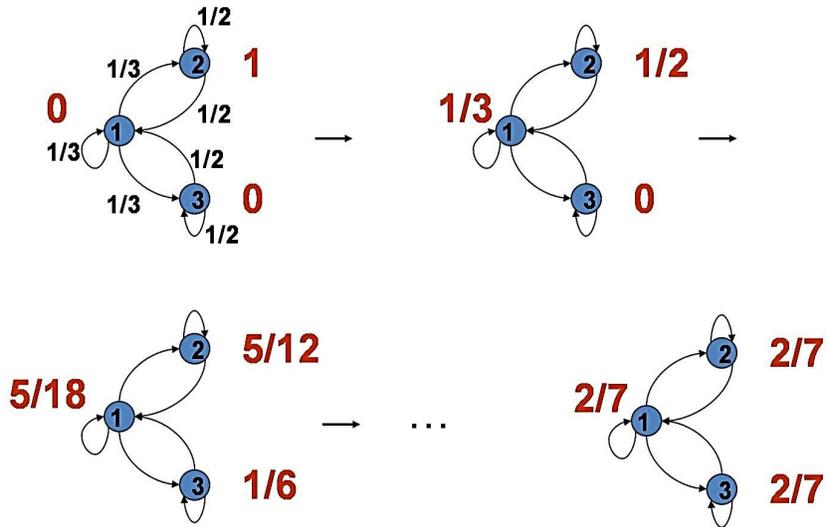


Figure 5.2. The DeGroot updating process over time for the system in Figure 5.2 when agent 2 starts with belief 1 and the other agents start with belief 0.

becomes $1/3$. Agent 2 averages beliefs of 1 and 0 and ends up at a new belief of $1/2$, while agent 3 averages two beliefs of 0 and so stays at 0. So the new beliefs are $\mathbf{b}(1) = (1/3, 1/2, 0)$. We now repeat this process and then beliefs become $\mathbf{b}(2) = (1/3, 1/2, 0)$, and iterating in this way the beliefs eventually converge to $\mathbf{b}(\infty) = (2/7, 2/7, 2/7)$, as pictured in Figure 5.2. This limit has a natural interpretation. Note that in this network agents 2 and 3 each have two links, while agent 1 has three links. There is a total of seven links, and each agent’s “influence” turns out to be proportional to the number of links that the agent has. Here agent 2 has $2/7$ of the links and that is agent 2’s influence. Agent 2 is the only agent with a positive initial belief, and the limit point is $2/7$ times agent 2’s initial belief, plus $2/7$ times agent 3’s initial belief and $3/7$ of agent 1’s initial belief.

This example illustrates a couple of features of the DeGroot process.

First, the agents’ beliefs converge to a consensus. This will be true of any strongly connected component such that there is a directed path from each agent to every other agent, and such that the component is aperiodic⁴⁸ (for which it is sufficient that at least

⁴⁸This requires that the least common divisor of the length of all the cycles in the component be one.

one agent place some weight on his or her own opinion). Thus, under very weak conditions, the society will reach a consensus in the DeGroot model. The intuition behind this result is straightforward: if we have not already reached a consensus then some agent who holds the highest belief at some point in time must be communicating with someone with a lower belief, and so that high-belief agent's belief will decrease, and similarly the agents with the lowest beliefs will have their beliefs move up over time. Since these are weighted averages of previous beliefs, they do not overshoot, and so the set of beliefs contracts over time (presuming the connectedness and aperiodicity conditions discussed above are satisfied).

Second, the influence that each agent has on the final consensus depends in very intuitive ways on the network structure. In the setting of (4) where agents place equal weights on each of their friends, the influence of an agent is proportional to his or her degree. An agent with twice as many friends as another agent has twice the influence on the eventual consensus. More generally, the influence will be related to the unit (left-hand) eigenvector of the trust matrix \mathbf{T} , that is, the unique vector \mathbf{s} such that $\mathbf{s}\mathbf{T} = \mathbf{s}$. This has the nice intuition that an agent's influence is related to the influence of those agents who trust him or her. Beyond this application, it also provides some foundation for eigenvector-based definitions of centrality that date to Katz (1953), and also is the reasoning behind things like Google's page rank system (e.g., see Langeville and Meyer (2006)).

Thus, we see that network structure plays an intuitive and tractable role in the DeGroot model of updating. This tractability enables a number of questions to be answered. To get a feeling for this, let us change the network in Figure 5.2 to include links between all of the agents, then we see that both the relative influences of the agents, the consensus, and the speed of convergence changes, as pictured in Figure 5.2.

Although very simple, this example shows that network structure affects both the consensus reached and how quickly it is reached. Understanding speed of convergence is important as a society may have limited iterations on communication, and so understanding the factors that affect rate of convergence can give an idea of the extent to which a consensus might emerge in the DeGroot model. As one might expect intuitively, the main thing that slows convergence is a split of the network into two or more groups that communicate more intensely within group than across groups. This is studied by Golub and Jackson (2008) who provide details as to when and to what extent homophily slows the convergence process. They show how the DeGroot updating process can be slowed by homophily due to the fact that convergence is dependent on the relative distribution of communication across groups, while there are other processes, such as those that simply depend on shortest paths in a network, which are unaffected by homophily.

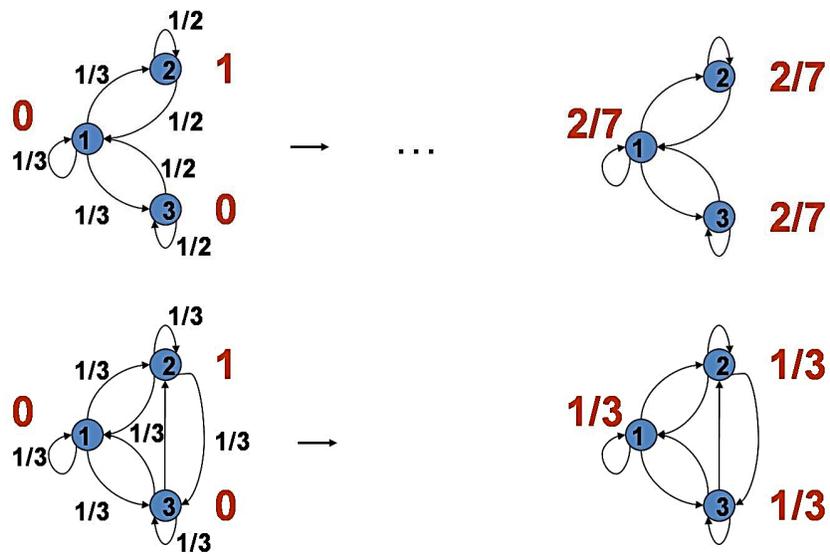


Figure 5.2. The DeGroot updating process for two different configurations: In the top setting agents 2 and 3 are not friends and the process is slower to converge to a consensus, in the bottom setting there is a complete network and consensus occurs in the first period.

There are many other questions that can be studied in the context of the DeGroot model and its variations. For example, Demarzo, Vayanos, and Zwiebel (2003) show how communication along many dimensions at once can reduce to convergence along a single dimension, providing insight into why complicated political landscapes often reduce to unidimensional discourses, where agents have the approximately the same relative positions in terms of their updated beliefs compared to the average belief across different dimensions after sufficient discourse time. The technical details of the proof involve working with the spectral decomposition of the trust matrix, but the intuition can be seen fairly easily from examples. As an illustration, suppose that society is divided into three groups, agents in group 1 talk evenly to all agents in groups 1 and 2, agents in group 3 talk evenly to all agents in groups 2 and 3, and agents in group 2 talk evenly with all agents. Agents in the middle group 2 will be at the average opinion, while agents in group 1 will be biased in a direction that under-weights group 3's average opinion, and group 3 agents will be biased in a direction that under-weights group 1's average opinion. Whatever the issue, groups 1 and 3 will be on opposite extremes and group 2 in the middle, and they will reach a consensus as groups 1 and 3 slowly move in towards the central group 2 average opinion.

The Bayesian and DeGroot analyses discussed above represent just some of the many in this growing area of analysis. Recent work has led to fuller understandings of what leads to a society to a consensus opinion or behavior and what the consensus will be, as well as how influential each agent in the society is. There are still many issues that remain open, including understanding issues of strategic transmission of information when agents have incentives to misrepresent or distort their information to try influence the eventual outcome (e.g., see Hagenbach and Koessler (2009), Lever (2009) and Acemoglu Ozdaglar, and ParandehGheibi (2009)), costly information acquisition in such contexts, endogenous formation of networks in the context of learning, and developing alternative models between the rational and DeGroot-style models (e.g., Jadbabaie, Sandroni, and Tahbaz-Salehi (2009)).

5.3 Bargaining and Trade through Networks

Another important application of network analysis is to understanding how terms of trade and transactions are affected by the network of relationships through which they take place. Moreover, once we understand how network structure affects trade we can then explore the incentives for agents to form trading networks and to see whether they lead to competitive or efficient outcomes.

In order to fix ideas, let us examine settings where there is some cost to establishing

communication between a potential buyer and seller. This cost might reflect many things. It could represent the opportunity cost of time spent learning about a product, adjustments for compatibility, or simply establishing and/or maintaining lines of communication. In a first stage, agents form a network of relationships and then in a second stage the agents can bargain and transact, but only with agents to whom they are linked.

The key to understanding the efficiency of trading networks comes through understanding the “externalities” in such settings and the splits of the gains from trade. Usually, we do not think of there being any externalities in the exchange of private goods. The consumption of a good by one agent does not affect another agent. However, when trading opportunities depend on which relationships are present in a society, choices of network ties have external effects. Thus, the (Pareto) efficiency that is the hallmark of the competitive exchange of private goods no longer applies when goods can only be traded through established relationships. The fact that inefficiencies arise when there is some limitation on the potential transactions that can occur is not new, as we see in the extensive literature on search. Moreover, such frictions serve as the basis for understanding various labor market imperfections as well as macroeconomic phenomena (see Rogerson, Shimer, and Wright (2005) for a recent survey). However, the way that inefficiencies arise and manifest themselves in network settings provide new insights into trading frictions, price dispersion, and bargaining power.

To get a feeling for some of the analysis in the literature, let us start with the simplest setting. Consider a benchmark case where each agent is either a “buyer” or a “seller”. When buyer i and seller j transact there is a total value to the transaction of v_{ij} . The value of that transaction can be split between them in any way. We can capture this via a price p so that the value of the transaction to the seller is p and to the buyer is $v_{ij} - p$. Buyers can only transact with sellers and vice versa, and each agent can participate in at most one transaction.

Let us start with a case where the transactions are homogeneous, so that the value of any transaction between a buyer and seller is 1 (so, all the v_{ij} s are 1) and where the cost is $1/2 > c_S > 0$ to the seller for each link, and $1/2 > c_B > 0$ to each buyer involved in a link. Clearly, the efficient network is to have pairs of linked buyers and sellers, so that no agent has more than one link and the number of links is equal to the minimum of the number of buyers and the number of sellers.⁴⁹ Extra links waste cost, and given that the total cost of

⁴⁹This takes an ex post perspective. From an ex ante point of view, if, for instance there are more buyers than sellers the buyers are not sure of which transaction might take place ex post, then one could imagine extra links being good from a buyer’s perspective. However, networks with extra links are Pareto dominated by randomly choosing which links get formed and then transacting at the expected value on the links that

a link is less than its value it makes sense to maximize the number of transactions taking place.

To see the basic issues that arise regarding efficiency, consider an example with just one seller and two buyers. The fact that the two-link network is inefficient does not prevent it from forming. Whether it forms when the buyers and sellers choose the links will depend on the the expected values that each of the agents gets as a function of the network structure. The key point is that we should expect the price that the seller gets (presuming the bargaining occurs after the network is formed and held fixed) to be higher when there are two links than when there is just one. If the bargaining takes place after the links have formed and the agents ignore the sunk costs of link formation, then with symmetry in the bargaining game, we would expect $1/2 - c_i$ to be the payoff to each of the agents in a single-link network. This is indeed the outcome of an alternating bargaining game like a variation on the Rubinstein-Stahl bargaining game considered in the network bargaining study of Corominas-Bosch (2004), if we examine the limit as the discount factor goes to 1, or we randomize in terms of who gets to make the first offer and examine the expected value of the bargaining. In the two-link network, we should expect that the seller will obtain a greater expected share of the transaction than in the one-link network, although exactly how much better the seller is will depend on the bargaining protocol. Let this share be $v \geq 1/2$, so that the payoffs are as pictured in Figure 5.3. The value of v determines which network we should expect to form.

If $1 - 2c_B > v > 1/2 + c_S$, then the unique pairwise stable network is the two-link network, as both the unlinked-buyer and seller gain from adding a second link to the network. If instead, $v > 1 - 2c_B$ or $1/2 + c_S > v$, then the one-link networks are the pairwise stable ones, as in the first case the buyers have a negative value from the two-link network and in the second case the seller is better off in a one-link network than a two-link network.

Which of these cases ensues depends on the specifics of the bargaining protocol. In an extreme case where the seller can get the buyers to bid against each other as in a sort of reverse Bertrand competition, the seller would extract all of the surplus when there are two links formed, and so $v = 1$, then only the efficient networks are pairwise stable, as the buyers would benefit from severing a link from the two-link network. This is the outcome under a core definition of trade (see Elliott (2008)) or else under the Corominas-Bosch (2004) alternating offer bargaining (with a limiting discount factor).

If in contrast, if the bargaining power in the two-link network is less extreme, but still

forms.

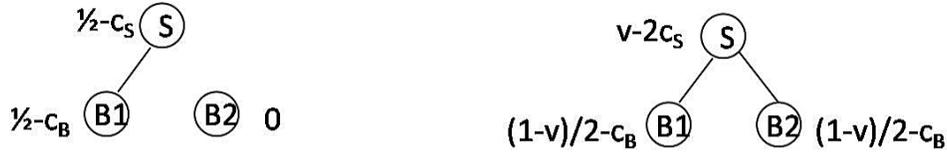


Figure 5.3. Payoffs to buyers and sellers as a function of the network.

favors the seller to some extent so that $1 - 2c_B > v > 1/2 + c_S$, then the two-link network will be the unique pairwise stable network. It is interesting to note that the seller would actually benefit from committing to lower his or her bargaining power if v exceeds $1 - 2c_B$. By committing to a bargaining procedure that leads to a v such that $1 - 2c_B > v > 1/2 + c_S$, the seller will obtain a greater surplus than if the bargaining procedure is more extreme.

Even in this very simple example, we see the role of the externalities and the potential for resulting inefficiencies. The buyers do not internalize the impact on each other of their decisions to form links. When we get to more complicated networks, how the allocation of utilities depends on the network structure depends on having a well-specified prediction for the bargaining outcome, and different bargaining protocols can lead to different conclusions regarding the efficiency of the stable networks.

Let me discuss a few of the ways in which such settings have been modeled. Corominas-Bosch (2004) examines settings such as those above where there is an identical value of 1 to each potential buyer-seller transaction. The prediction of the outcome of the bargaining depends on the network in a way that keeps track of relative balance of buyers and sellers in various subnetworks and is based on an alternating move bargaining game. A clever

algorithm identifies which buyers and sellers are evenly matched and which ones end up extracting full surplus. In that setting, if both buyers and sellers bear link costs (and without too much asymmetry in their costs) and the discount factor in the bargaining is high enough, then the only pairwise stable networks are the efficient ones, as buyers or sellers on the long side of the market do not get enough surplus to cover link costs.

Kranton and Minehart (2001) allow for heterogeneity in the realization of various buyers' valuations after links are formed and study settings where sellers' items are all identical and simultaneously auctioned off. They show that then the marginal expected value to a buyer from adding a link is exactly that link's social expected value. In that case, if only buyers pay costs (so that $c_S = 0$) then the network will be efficient, but if sellers pay costs it may not. Depending on the cost to sellers, the resulting network can be over- or under-connected relative to what would be total surplus maximizing. There are two effects: one is that sellers see changes in the relative competition in bidding as links are added and so have an incentive to add links, and the other is that the change in a seller's surplus can differ from the change in social value from adding a link, and so their incentives to add links can be distorted in either direction from the social incentive; compounded by the fact that buyers and seller are also internalizing only part of the cost of a link.

The model of Elliott (2008) allows for a full heterogeneity in each buyer-seller transaction value. This presents a challenge in predicting the relative splits of surplus as a function of the network or in working with any specific bargaining protocol. Elliott uses a core definition to sort this out, which provides a multiplicity of predictions as to the values realized to buyers and sellers as a function of the network, but the core has well-defined endpoints in terms of maximum and minimum outcomes for buyers and sellers. Through an illuminating algorithm, Elliott traces out how the determination of relative shares can be traced to chains of outside options that buyers and sellers have. Elliott also presents interesting results on the "price of anarchy," a term due to Papadimitriou (2001) and Roughgarden and Tardos (2002), which examines the extent to which the total surplus of the society can be dissipated when agents form the network. He shows that the full level of surplus can be dissipated by the inefficient network formation of the agents involved. He also traces the inefficiencies to over-investment (to improve bargaining position)⁵⁰ and under-investment where beneficial networks fail to form in cases where the agents do not see enough value from a transaction to cover their personal cost. Which form of inefficiency emerges depends on whether link costs are exogenous or can be negotiated.

⁵⁰See Jackson (2003, 2004) for more discussion of how externalities in determining allocations and how bargaining can lead to systematic over-connection of networks.

Kakade et al (2004) (as well as Kakade, Kearns and Ortiz 2004) examine a model of exchange on random graph-generated network structures. Their interest is in the extent to which there can exist price dispersion and how this depends on network structure. They examine a simple general equilibrium model, reminiscent of the Corominas-Bosch model described above, but where goods are fully divisible, so that a seller quotes a price so that she sells exactly her unit supply of the good, and buyers exchange their unit supply of money until they have exhausted it. Kakade et al (2004) find, roughly, that if the network is sufficiently symmetric in that buyers all have similar numbers of connections, and similarly for sellers, then there will be low levels of price dispersion, but as asymmetry increases so that some sellers have high degrees while others do not (or similarly for buyers), then substantial variation in prices can be observed across different parts of a network.

Beyond these models there are many other important questions to be investigated. For example, in the above analyses buyer-seller transactions occur just once and are exclusive, but one can also study settings where there are repeated transactions, or where there are multiple goods for sale. For example, Manea (2008) (see also Abreu and Manea (2008)) examines a model similar to the setting described above, except that any two agents can transact and once two agents transact, they are replaced by new agents. The bargaining game is such that each period a different link is recognized and then one of the agents is randomly selected to make an offer for trade to the other agent. This gives agents bargaining power in proportion to the number of links that they have, and although it does not directly correspond to any particular application, it provides a tractable model in which to examine the above issues of allocations as a function of network structure and the efficiency of network formation. One can also examine models of oligopoly on networks as in Nava (2008) to see when it is that competitive outcomes are reached and how that depends on network structure, or to see what incentives firms have to enter each other's markets as in Lever (2008), as well as to understand the role of middlemen in determining profits and achieving economic efficiency (e.g., Blume et al (2007)), or the role of collaboration among firms (Goyal and Joshi (2003)).

Beyond the models above, and the empirical studies referred to in Section 2.1, there are also a number of experimental studies regarding how surplus is split among agents in network situations who negotiate over potential transactions. Charness, Corominas-Bosch and Frechette (2005) provide some direct tests of the Corominas-Bosch model. Although the model does not fit exactly, they do find that the predictions of the model are accurate in terms of the directions of changes in surplus as a function of network changes. This comes on the heels of a fairly large literature on "exchange theory," which evolved from general forms of dyadic exchanges in Homans (1958, 1961), to more direct economic applications as

in Blau's (1964), and eventually included explicit consideration of social network structure as in Emerson (1972) and Cook and Emerson (1978) (see Cook and Whitmeyer (1992) for an overview). The exchange theory literature includes many experiments examining various network configurations and the exercise of bargaining power by agents as a function of their position in a network.

The growing catalog of studies on exchange through networks have shown that, beyond the basic point that network structure affects outcomes, full efficiency only arises in some specific circumstances. Looking across the studies one sees a theme that parallels one from the industrial organization literature: conclusions can be sensitive to details of the interaction. Nonetheless, there are some underlying regularities. Externalities are generally negative: agents have incentives to add relationships that might not be needed for efficient transactions but improve their bargaining power, and to the extent that agents do not fully see the value of potential transactions they may hurt others by not adding relationships that are needed to reach efficiency. Better connected agents (in precise network-defined senses) are relatively favored, and more asymmetric degree distributions can lead to greater inequality in outcomes. Looking forward, it seems that there is substantial promise in bringing network-based models of transactions to empirical studies of bargaining and trade, and measuring levels of inefficiencies. Given recent market failures, it also seems clear that a deeper understanding of interconnected liabilities and correlations in investments and financial contagion is needed.

5.4 Peer Interactions and Games on Networks

Beyond direct transactions through a network, there are many other contexts where agents make decisions that are influenced by the decisions of their friends and acquaintances. This includes whether we drop out of the labor force, what political opinions we hold, which music we listen to, whether or not we engage in criminal activity, and which products we buy, among a myriad of other behaviors. Once the payoff to the decision of one agent from a given choice depends on the actions of his or her neighbors, the decisions can be modeled as a game.

Simple variations of games on networks were studied in the computer science with respect to showing how hard it can be to compute Nash equilibria in n person games.⁵¹ In particular,

⁵¹Some specific examples of games on networks had been studied earlier, such as Ellison's (1995) and Young's (1998) studies of coordination games played between players in various lattice configurations (and see Jackson and Watts (2002b) and Goyal and Vega-Redondo (2005) for coordination games on more general

Kearns, Littman and Singh (2001) introduced a class of games, that they called “graphical games,” such that each agent chooses between two actions 0 and 1 and an agent’s payoff depends not only on his or her choice, but also on the decisions of his or her neighbors in a social network. The concern of that early literature was how hard it was to compute equilibria in cases where players’ payoffs as a function of their neighbors’ actions could be quite arbitrary. Despite the original paper’s results on the difficulties of computing equilibria in some graphical games with large numbers of players, the basic model is quite useful as a device for understanding peer effects and strategic peer interactions. In particular, in many situations with peer effects there is substantial structure in the way that payoffs behave, so that actions are strategic complements or substitutes. This makes equilibria much more manageable and interesting.

Some aspects of such games are discussed in other chapters by Goyal (This volume) and Jackson and Yariv (This volume),⁵² and so I just provide some illustrations of the central themes here.

It is useful to start with an example studied by Morris (2000) that provides some interesting insights. Consider a situation where agents are choosing two actions 0 and 1, which might be technologies, languages, fashions, etc. An agent can only choose one of the two and suppose that an agent prefers action 1 if and only if a fraction of at least q of his or her neighbors chooses action 1 and otherwise prefers action 0, where $1 > q > 0$.⁵³ For instance, it might only be worthwhile to adopt a new technology if a sufficient fraction of the agent’s neighbors adopt it. There are clearly multiple (strict) Nash equilibria to this game: if all agents choose action 1, then each agent will strictly prefer to take action 1; while if all agents choose action 0 then all agents strictly prefer to take action 0. When is it that equilibria exist where some agents take action 0 and others take action 1?⁵⁴

In Figure 5.4 we see an illustration of the multiplicity of equilibria in a setting where agents wish to match the majority of their neighbors’ actions (with a preference for action 0 if an agent’s neighbors are evenly split). In particular, it is possible to have equilibria where both actions are played simultaneously by different agents in the same component of a network, as we see in the two networks in the bottom part of Figure 5.4.

The survival of both actions at the same time in an equilibrium is not always possible. We see that in the top network in Figure 5.4 for a situation where agents are only willing to

and endogenous network structures).

⁵²See Chapter 9 in Jackson (2008) for more background and detail.

⁵³Letting q be an irrational number ensures that agents are never indifferent.

⁵⁴I restrict attention to pure strategy equilibria in this discussion.

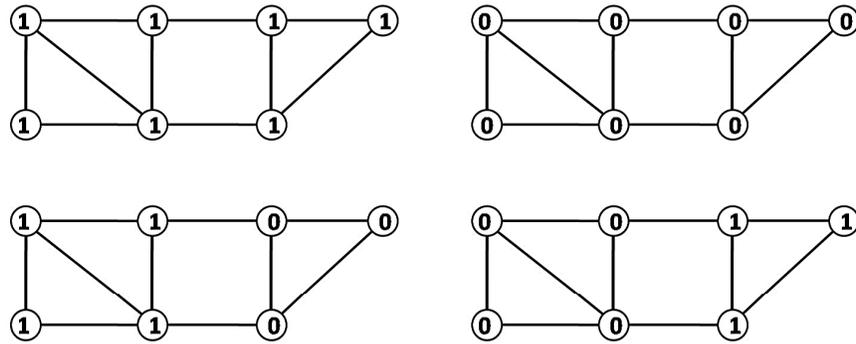


Figure 5.4. Equilibria when agents are willing to take action 1 if and only if more than half of their neighbors do.

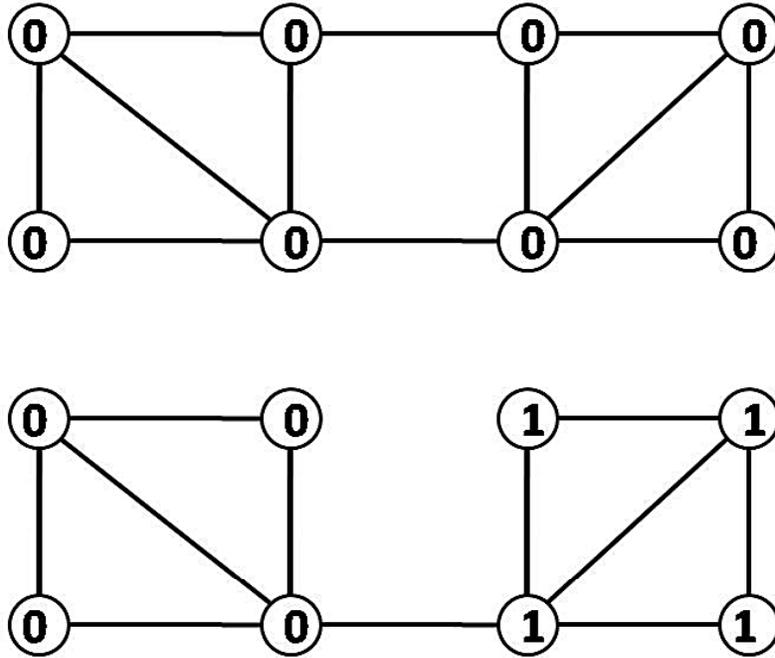


Figure 5.4. A game such that agents are willing to take action 1 if and only if more than seventy percent of their neighbors do. In the top network there does not exist an equilibrium where some agents play action 1 at the same time that other agents play action 0, while in the bottom network there exists such an equilibrium.

take action 1 provided more than seventy percent of their neighbors do. However, by severing one of the links we obtain a network, as pictured in the bottom of Figure 5.4, where there are two sufficiently isolated groups so that both actions can be sustained in equilibrium.

In general, the ability of a society to sustain equilibria where both actions are played by different segments of the society depends on the network structure and the preferences of the agents. As one can intuit from the above examples, one needs to find a splitting of the society such that each of the two groups of the society are sufficiently inward-looking. This is captured through an intuitive definition of “cohesiveness” introduced by Morris (2000). A group of agents is r -cohesive if each agent in the group has a fraction of at least r of his or her neighbors in the group. In order to have an equilibrium played with both actions played in a game where action 1 is preferred if and only if more than a fraction q of an agent’s neighbors take action 1, there must exist a partitioning of the agents into two groups: one group is more than q cohesive and ends up playing action 1, and the other group is at least

$1 - q$ cohesive and plays action 0. Morris also provides conditions under which if some small segment of a society is fixed to select action 1, that action will eventually spread to the entire society if agents (iteratively) best respond to their neighbors' actions.

There are variations on this sort of coordination game that have been studied in the statistical physics literature with the best known version being what is called the "voter model." That model dates to Clifford and Sudbury (1973) and was named by Holley and Liggett (1975). Early versions of the voter model were on lattices, but more recent studies have examined general network structures. The simplest version of the model is one where a node is randomly selected and then a neighbor is randomly selected and the node's state of 0 or 1 is changed to match the neighbor. Over time, if the number of nodes is finite and all nodes are path connected, the society will eventually reach a consensus and stay there, although the random time to reaching a consensus can depend on the network.⁵⁵ Many variations on the voter model have been considered, including ones where nodes match the majority of their neighbors. See Castellano, Fortunato, and Loreto (2009) for a survey.

These examples provide a flavor of the types of results that emerge from games on networks. In many situations there is an inherent multiplicity of equilibria and thus many ways in which behavior might evolve, and so the challenge is to get a handle on the set of possible outcomes. Despite the multiplicity, some settings still have some nice intuitions and results that can be derived. For example, in games with strategic complementarities, such as the examples above, where an agent's incentive to take a higher action increases as more of an agent's neighbors take higher actions, then the existence of pure strategy equilibria is guaranteed and some of the equilibria are quite easy to compute.⁵⁶ Moreover, in such settings a variety of dynamics naturally tend towards equilibrium behavior. There are also many applications where there is a nice relationship between network structure and

⁵⁵The idea behind the proof that a consensus will eventually be reached, for any network, is quite simple. With a finite number of nodes and a connected network, we can pick some node and then there is a positive probability that each one of its neighbors will be picked and matched to it before being matched to any other nodes, and then a positive probability that subsequently their neighbors will be picked and matched to those nodes before being matched to any other nodes, and so forth. Although the aggregate probability that some specific node's state will eventually overtake the whole network can be quite small, over time the needed sequence of matchings will eventually take place. So, with probability one the system will eventually reach a consensus.

⁵⁶The equilibria form a lattice, and the maximal and minimal equilibria are easily found. For example, start with all agents at the highest action. If some agents then strictly prefer to take a lower action, change their action. Iterating on this process, will eventually lead to a point where no agent wishes to lower his or her action. This is the maximal equilibrium in that it has higher actions for each agent than in any other equilibrium. Analogously, one can find a minimal equilibrium. For more details, see Jackson (2008).

equilibrium structure, such as the cohesiveness of different subgroups in the network and the sustainability of multiple actions in one equilibrium.

The analysis above pertains to a special case, and more generally agents might care about more than just the proportion of agents taking a given action. For example, it might be an absolute number of an agent's neighbors that matter, so that an agent might be willing to take action 1 provided at least some fixed number of the agent's neighbors take the action. For example, an agent might be willing to learn a new language if the agent has at least some number of friends who speak the language, or might be willing to take up a certain hobby if at least some number of friends partake in the hobby. In such cases, agents' decisions can depend on their degrees and the degrees of their neighbors, and so forth.

Beyond games of strategic complements, another important class of games on networks is that of strategic substitutes. In games of strategic substitutes, if the actions of an agent's neighbors increase, then the agent has a stronger preference for lower actions. This applies in many settings of local public good provision. For example, consider an example of buying a particular book. If an agent has a friend who buys the book, then the agent can free-ride and borrow the book and so does not need to buy it. This reversal in the direction of peer effects does not result in a simple variation on games of strategic complements: the analysis of games of strategic substitutes is quite different. In games of strategic complements, agents' preferences move together, while in games of strategic substitutes interactions and dynamics can be more complicated and existence and computation of equilibria can be significantly more challenging. Bramoullé and Kranton (2007b) provide an analysis of a class of such games, and show, among other things, that slight variations in network structure can lead to dramatic changes in equilibrium structure.

The multiplicity of equilibria, and the sensitivity to network structure, depends on the information structure. In the above discussions each agent is choosing an action that must be a best response given knowledge of his or her friends' actions. There are many applications where agents have less specific information when making a choice. In buying a new software program, an agent might not even know exactly with whom he or she might interact in the future, but might only have some idea of the number of interactions he or she is likely to have. The agent's decision may need to be based on some idea of the overall prevalence of the compatibility of the program with choices of other agents in the population. Galeotti et al. (2008) and Jackson and Yariv (2007) examine such settings and show that an incomplete information setting can actually simplify the analysis of games on networks. In particular, results can be derived showing how agents' actions vary with their degree. For instance, suppose that it is only worthwhile for an agent to adopt a new technology if he or she

expects to have at least k future neighbors adopt it. Then agents who expect to have more future interactions are more likely to exceed that threshold and thus have a stronger incentive to adopt the technology. This implies a monotonicity in decisions such that agents who expect to have more future interactions are those who will adopt and those who expect to have fewer interactions will not. Thus there is a threshold, such that the adopting agents are those who expect to have more than that number of interactions. This translates into a series of results of how equilibrium structure varies with the society's degree distribution, as changes in the distribution that increase the number of agents above the threshold lead to increased adoption, and changes that reduce the number of agents above the threshold lead to decreased adoption. Details are discussed in Jackson and Yariv (This volume).

As one adds more structure to the payoff structure to a game on a network, one can begin to obtain an explicit calculation of equilibrium. A nice example is a model by Ballester, Calvó-Armengol, and Zenou (2006) and is such that agents choose an intensity of an action from a continuum, and where there are local strategic complementarities and global substitution effects. In particular, Ballester, Calvó-Armengol, and Zenou work with a simple quadratic payoff structure that makes explicit equilibrium calculation quite easy, and yet insightful. The payoff structure captures an application of criminal behavior, where agents find increased benefits from engaging in crime as their friends' criminal activity increases, but there is also an overall competitive effect so that increased criminal activity on average in a society reduces the benefits to any individual's criminal behavior. Ballester, Calvó-Armengol, and Zenou show how equilibrium actions intuitively relate to a network centrality measure. Although the assumed functional form for payoffs is special, it encapsulates basic factors influencing choices and allows for a tractable analysis of how changes in network structure lead to changes in behavior.⁵⁷

While our knowledge of peer effects is growing, the complexities involved mean that there is still much to be learned. Useful tools in this area of research are laboratory and field experiments, where one can directly measure how agents' behaviors change as network structure is changed, or as a function of their position in a network. The use of experiments to study social networks has a long history including the seminal work of Milgram (1967) on small worlds and studies of exchange theory such as that of Cook and Emerson (1978). There is a growing literature using experiments to study strategic network formation (e.g., see Callendar and Plott (2005), Pantz and Zeigelmeyer (2003), Falk and Kosfeld (2003), Goeree, Riedl, and Ule (2003) and Charness and Jackson (2007)), learning (e.g., Choi, Gale,

⁵⁷For other examples of some of the advantages of working continuum models, see Rogers (2007) and Bloch and Dutta (2007).

and Kariv (2005, 2009) and Celen, Kariv and Schotter (2004), and as well as interaction on networks (see Kosfeld (2003) and Jackson and Yariv (this volume) for additional background).⁵⁸ In terms of studying how network structure influences behavior, Goeree et al (2007) and Leider et al (2009) show that giving behavior in dictator games is related to social distance, with agents being more generous to those who are nearby. Other examples include Kearns et al. (2009) who examine how a society's ability to reach a consensus in choosing an action depends on the network configuration and various agents' payoffs from their actions.

Another important aspect of strategic behavior in network settings that has been looked at but is far from being understood is the coevolution of networks and behavior. Much of the literature that I have discussed to this point focuses either primarily on the formation of a network or on the influence of a network on behavior. It is clear that there is feedback: people adjust their behaviors based on that of their friends and they choose their friends based on behaviors. Kandel (1978) provides interesting evidence suggesting that both effects are present in friendship networks, so that over time agents adjust actions to match that of their friends and are more likely to maintain and form new friendships with other individuals who act similarly to themselves. There has been some modeling of this in the context of the coevolution of behavior and networks in coordination games by Jackson and Watts (2002b), Goyal and Vega-Redondo (2005), and Ehrhardt, Marsili, and Vega-Redondo (2006), Fosco and Mengel (2009).⁵⁹ A message that emerges from those studies is that when networks co-evolve with actions, the actions that emerge can differ from what one sees with a fixed network. There is also a nascent literature on repeated games on networks, that provides insight into both how individuals behave and what sort of network structures are necessary to promote cooperation and pro-social behavior (e.g., see Jackson, Rodgriguez-Barraquer, and Tan (2010) and the references therein).

To get a feeling for this, consider the simple coordination game in Table 2.

Now let us consider this in the context of a social network. Each agent plays this game with each of his or her neighbors, and the agent just chooses one action, so that if the player chooses A then that is played against every one of the agent's neighbors. This game has two strict Nash equilibria: one where all players choose A, and another where all players choose

⁵⁸There is also a growing use of field experiments in combination with social network data, such as those by Duflo and Saez (2003), Karlan, Mobius, Rosenblat, and Szeidl (2009), Dupas (2010), Beaman and Magruder (2010), and Feigenberg, Field, and Pande (2010), Baccara, Imrohoroglu, Wilson, and Yariv (2009).

⁵⁹See also Jackson and Watts (2005) for an equilibrium analysis of a choice of partners together with a choice of behavior in a matching setting.

Table 2: A Coordination Game. The first entry in each cell is the payoff to Player 1 based on the combination of actions played.

		Player 2	
		A	B
Player 1	A	1, 1	-2, 0
	B	0, -2	0, 0

B.

Let us take a closer look at the incentives for various play in the context of a complete network, so that all agents play the game with every other agent. In this example playing B leads to a sure payoff of 0, whereas playing A can lead to 1, but could also lead to a payoff of -2. Playing B is the better response if an agent expects more than one third of the other players to play B, while A is the better response only if an agent expects less than one third of the other players to play B. Thus, if one starts with uniform uncertainty about what the other players are doing, B is the better response. In this sense every agent playing B is known as the “risk dominant” equilibrium. Refinements of equilibrium, such as stochastic stability, have found that if we add a bit of noise to the play of agents, and we examine a dynamic process where agents adjust their actions over time but with occasional errors so that the process never settles down completely, then play visits the “risk-dominant” equilibrium more often (e.g., see Kandori, Mailath and Rob (1993) and Young (1993)). The reasoning behind this is that if a society begins by all playing A, if one third of the population happened to tremble and switch to action B, then it would become a (myopic) best response for all players to play B. So errors made by one third of the population can lead away from the equilibrium where all play A. In contrast, if the society is playing equilibrium B, then it takes trembles by two thirds of the population to switch to A before it becomes a best response to play A. So, there is a precise sense in which it takes more perturbations to the system to transition from B to A than to move in the reverse direction, and so the equilibrium where all agents play B is more “robust,” in at least one particular sense. This is not great news for the society, since the equilibrium where all play A leads to a higher payoff for everybody involved.

However, as Jackson and Watts (2002b) point out, this conclusion is dependent upon the network structure. Consider the same game but played by a society arranged in a “star” network such that there is one central agent who plays with every other agent, and the

peripheral agents only play with the central agent. Then, it is much easier to transition between the two equilibria. Whatever action is chosen by the center of the star becomes the best response for all the other agents in the society. Here the society can transition back and forth between the two equilibria quite easily, simply by changing the action of the central agent, as the other agents' (myopic) best responses are to match whatever the central agent does.⁶⁰ This leads both of the equilibria to be stochastically stable.

This example illustrates that network structure can influence the play of a society, even in terms of selecting among (strict) Nash equilibria. This is true a fortiori when one endogenizes the network along with the play of the game. This echoes a point of Ely (2002), who studied the outcome of such games when agents could choose to change where they lived. All it takes is some agent to locate at an empty location and play A to have other agents want to move there, as they will then be playing equilibrium A rather than equilibrium B at their own location. When one endogenizes a network structure, agents have similar incentives: to form links with others who are playing action A and to sever links with those playing B. Exactly what emerges depends on a number of details including the cost structure to links, whether agents get a payoff that depends on the number of people they play with and not just the average play, how many links can be changed at once, and whether mutual consent is required to form a link (see Jackson and Watts (2002b) and Goyal and Vega-Redondo (2005) for more). There are also questions as to the relative rates at which actions change compared to network structure, and that can affect the overall convergence to equilibrium, as one sees in Ehrhardt, Marsili, and Vega-Redondo (2006) as well as Holme and Newman (2006).

5.5 Labor Markets

To see an illustration of the insights generated by analyzing games on networks, let us return to the role of social networks in labor markets. There are a variety of decisions that agents take, whether or not to drop out of the network, whether to become educated, whether to look for employment, and so forth, for which the payoff can be heavily influenced by the decisions and situations of an agent's friends. If an agent's friends have all dropped out of the labor force, then that makes it difficult for the agent to get information about

⁶⁰Intuitively, if the center of the star is forward looking, then he or she can actually consciously choose to steer the society towards one of the equilibria by choosing that action. Formalizing such behavior requires a careful definition of forward-looking behavior, and modeling it for all agents. For example, see and Mauleon and Vannetelbosch (2004) and Page, Wooders, and Kamat (2005) for such definitions.

job openings and so worsens that agent's future employment prospects and makes dropping out relatively more attractive. As Calvó-Armengol and Jackson (2004) point out, this is a (graphical) game where agents' decisions to be in the labor force or to drop out are strategic complements. The insights from games on networks are then very useful. If we begin with two groups of agents who have many inward connections and few cross group connections, and we start one group with many drop-outs and the other group with few drop outs, then we can see very different outcomes for the two groups, similar to what we see in the bottom of Figure 5.4. Agents react to their neighbors' decisions and some subgroups end up with high participation in the labor force while others end up with high drop-out rates. When coupled with historical patterns and initial conditions, this can help explain the significant differences in drop-out rates that are exhibited across races. In particular, as Jackson (2007) points out, given racial homophily patterns, so that individuals of a given race tend to be connected to others of the same race, very different drop out rates can emerge for different races, even after controlling for all individual characteristics. This can lead to persistent inequality across different social groups, and not just in drop out rates: the individuals of a race with a low participation rate who happen to be in the labor force, will then have less access to job information leading to more unemployment spells, worse matches with employers, and lower wages (see Arrow and Borzekowski (2004) for more on wage effects). Thus, such network-based complementarities lead to an explanation for some of the differences in labor force participation rates and employment outcomes that have been widely documented by Card and Krueger (1992), Chandra (2000), and others.

In addition to helping us to understand different behaviors across groups, these complementarities in decisions can also have implications for intergenerational correlations in behavior. For example, Calvó-Armengol and Jackson (2009) show that if a child's social network has overlap with that of his or her parents, that can lead the child's decisions regarding how much education to pursue to be correlated with that of the parent, even without any direct parental influence. This extends beyond education to any behavior that sees a network influence, and can have implications for general forms of social mobility. These examples show the usefulness of the theory of strategic interaction in networked settings.

6 Concluding Remarks

The study of social and economic networks has expanded rapidly in the past decades and naturally cuts across many disciplines. It is an exciting area not only because of the explosion of "social networking" that has emerged with the internet and other advances in commu-

nication, but because of the fundamental role that social networks play in shaping human activity. Social network analysis has already taught us a great deal and it holds tremendous potential for future application, especially in economics. Moreover, beyond the applications, and as I hope that this survey illustrates, the increasingly sophisticated tools emerging in a variety of fields promise to continue to improve our modeling and understanding of the patterns of human interaction.

7 Bibliography

Abreu, D., and M. Manea (2008) “Bargaining and efficiency in networks,” mimeo: Princeton University.

Acemoglu D, Dahleh M, Lobel I, Ozdaglar A. 2008. Bayesian learning in social networks., NBER Work. Pap. No. W14040

Acemoglu, D., A. Ozdaglar, and ParandehGheibi (2009) “Spread of (Mis)Information in Social Networks,” forthcoming: *Games and Economic Behavior*.

Adamic, L.A. (1999) The SmallWorldWeb, in *Proceedings of the ECDL. Lecture Notes in Computer Science 1696*, Berlin: Springer-Verlag.

Albert R, Jeong H, Barabási AL. 1999. Diameter of the World Wide Web. *Nature* 401:13031

Allen, F., and A. Babus (2007) Networks in Finance, in *Network-Based Strategies and Competencies*, P. Kleindorfer and J. Wind, eds., Philadelphia: Wharton School Publishing, forthcoming.

Allport, W.G., (1954) *The Nature of Prejudice*, Cambridge, MA: Addison-Wesley.

Aral, S., L. Muchnik, and A.Sundararajan (2009) “Distinguishing Influence Based Contagions from Homophily Driven Diffusion in Dynamic Networks,” *Proceedings of the National Academy of Sciences*.

Arrow, K.J., and R. Borzekowski (2004) Limited Network Connections and the Distribution of Wages, Federal Reserve Board Publication 41, Washington, D.C.

- Aumann R, Myerson R. 1988. Endogenous formation of links between players and coalitions: an application of the Shapley value. In *The Shapley Value*, ed. AE Roth, pp. 175 - 91. New York: Cambridge Univ. Press
- Baccara, M. and Yariv, L. (2009), "Similarity and Polarization in Groups," mimeo.
- Bala V, Goyal S. 1998. Learning from neighbors. *Rev. Econ. Stud.* 65:595621
- Bala V, Goyal S. 2000. A non-cooperative model of network formation. *Econometrica* 68:1181 - 30
- Bala V, Goyal S. 2001. Conformism and diversity under social learning. *Econ. Theory* 17:10120
- Ballester C, Calvó-Armengol A, Zenou Y. 2006. Whos who in networks: wanted the key player. *Econometrica* 74(5):140317
- Banerjee, A.V. (1992) A Simple Model of Herd Behavior, *Quarterly Journal of Economics* 107:797817.
- Banerjee, A.V., and D. Fudenberg (2004) Word-of-Mouth Learning, *Games and Economic Behavior* 46:122.
- Barabási A, Albert R. 1999. Emergence of scaling in random networks. *Science* 286:50912
- Barabási, A. (2002) *Linked*, Cambridge, Mass.: Perseus Publishing.
- Bayer, P.J., Hjalmarsson, R. and D. Pozen (2009), Building criminal capital behind bars: Peer effects in juvenile corrections, *Quarterly Journal of Economics* 124, 105-147.
- Bayer PJ, Ross SL, Topa G. 2005. Place of work and place of residence: informal hiring networks and labor market outcomes. *Econ. Growth Cent. Discuss. Pap.* 927. New Haven, CT: Yale Univ.
- Beaman L. 2007. Refugee Resettlement: The Role of Social Networks and Job Information Flows in the Labor Market. Northeast Univ. Dev. Consort. Conf., 2006. Sess. 20: Migration. mimeo: New Haven, CT: Yale Univ. , manuscript
- Bearman, P., J. Moody, and K. Stovel (2004) Chains of Affection: The Structure of Adolescent Romantic and Sexual Networks, Chicago, University of Chicago, manuscript

- Bender, E.A., and E.R. Canfield (1978) The Asymptotic Number of Labelled Graphs with Given Degree Sequences, *Journal of Combinatorial Theory A* 24:296-307.
- Bentolila, S., C. Michelacci and J. Suarez (2009), Social Contacts and Occupational Choice, *Economica*, forthcoming.
- Besag, J.E. (1974) Spatial Interaction and the Statistical Analysis of Lattice Systems (with Discussion), *Journal of the Royal Statistical Society B* 36:196-236.
- Bikhchandani, B., D. Hirshleifer, and I. Welch (1992) A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades, *Journal of Political Economy* 100(5):992-1026.
- Blau, P.M. (1964) *Exchange and Power in Social Life*, New York: Wiley
- Blau, P.M. (1977) *Inequality and Heterogeneity: A Primitive Theory of Social Structure*, New York: Free Press.
- Bloch F. 2004. Group and network formation in industrial organization. In *Group Formation in Economics; Networks, Clubs and Coalitions*, ed. G Demange, M Wooders, Chap. 11. Cambridge, UK: Cambridge Univ. Press
- Bloch, F. and B. Dutta (2009) ‘Communication networks with endogenous link strength,’ *Games and Economic Behavior*, Volume 66, Issue 1, 39 - 56.
- Bloch F, Genicot G, Ray D. 2008 “Informal Insurance in Social Networks,” *Journal of Economic Theory*, 143:1, 36-58.
- Bloch, F., and M.O. Jackson (2006) Definitions of Equilibrium in Network Formation Games, *International Journal of Game Theory* 34(3):305-318.
- Bloch F, Jackson MO. 2007. The formation of networks with transfers among players. *J. Econ. Theory* 133: 83 - 110.
- Blume, L. (1993) The Statistical Mechanics of Strategic Interaction, *Games and Economic Behavior* 5:387-424.
- Blume, L., D. Easley, J. Kleinberg, E. Tardos (2007) “Trading Networks with Price-Setting Agents,” *Games and Economic Behavior*, forthcoming.

- Bollobás, B. (2001) *Random Graphs*, second edition, Cambridge: Cambridge University Press.
- Bonacich P. 1972. Factoring and weighting approaches to status scores and clique identification. *J. Math. Sociol.* 2:113 - 20
- Bonacich P. 1987. Power and centrality: a family of measures. *Am. J. Sociol.* 92:117082
- Boorman S. 1975. A combinatorial optimization model for transmission of job information through contact networks. *Bell J. Econ.* 6:216 - 49.
- Borm, P., G. Owen, and S. Tijs (1992) On the Position Value for Communication Situations, *SIAM Journal on Discrete Mathematics* 5:305320.
- Bourdieu, P. (1986) *Forms of Capital*, in *Handbook of Theory and Research for the Sociology of Education*, J.G. Richardson, ed., Westport, Conn.: Greenwood Press.
- Bramoullé, Y., H. Djebbari and B. Fortin (2009), Identification of Peer Effects through Social Networks, *Journal of Econometrics* 150, 41-55.
- Bramoullé Y, Kranton R. (2007) Risk-sharing networks. *J. Econ. Behav. Organ.* 64(34): 275 - 94
- Bramoullé Y, Kranton R. (2007b) "Public Goods in Networks" *J. Econ. Theory* 135:1, 478 - 494
- Bramoullé, Y., and B. W. Rogers (2009) "Diversity and Popularity in Social Networks", Mimeo.
- Breiger, R.L., S.A. Boorman, and P. Arabie (1975) An Algorithm for Clustering Relational Data with Applications to Social Network Analysis and Comparison with Multidimensional Scaling, *Journal of Mathematical Psychology* 12:328383.
- Brock, W., and S.N. Durlauf (2001) *Interactions-Based Models*, in *Handbook of Econometrics*, 5, J. Heckman and E. Leamer, eds., Amsterdam: North-Holland.
- Burt R. 1992. *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard Univ. Press
- Calvó-Armengol A. 2004. Job contact networks. *J. Econ. Theory* 115: 191 - 206

- Calvó-Armengol A, Jackson MO. 2004. The effects of social networks on employment and inequality. *Am. Econ. Rev.* 94(3): 426 - 54
- Calvó-Armengol A, Jackson MO. 2007. Networks in labor markets: wage and employment dynamics and inequality. *J. Econ. Theory* 132(1):2746
- Calvó-Armengol A, Jackson MO. 2009. Like father, like son: labor market networks and social mobility. *Am. Econ. J. Microecon.* 1: 124 - 50
- Calvó-Armengol, A., E. Patacchini, and Y. Zenou (2009), "Peer effects and social networks in education," *Review of Economic Studies* 76, 1239-1267.
- Calvó-Armengol, A., and Y. Zenou (2004) Social Networks and Crime Decisions: The Role of Social Structure in Facilitating Delinquent Behavior, *International Economic Review* 45:935-954.
- Callander, S., and C. Plott (2005) Principles of Network Development and Evolution: An Experimental Study, *Study of Public Economics* 89:1469-1495.
- Carayol N, Roux P. 2003. Collective Innovation in a Model of Network Formation with Preferential Meet. mimeo: Univ. Louis Pasteur/Univ. Toulouse I. (Mimeo)
- Carayol N, Roux P, Yildizoglu M. 2006. In search of efficient network structures: the needle in the haystack. *Rev. Econ. Des.* 11(4): 434 - 70
- Carayol N, Roux P, Yildizoglu M. 2008. Inefficiencies in a model of spatial networks formation with positive externalities. *J. Econ. Behav. Organ.* 67(2):495-511.
- Card, D. and Krueger, A.B. (1992) School Quality and Black-White Relative Earnings: A Direct Assessment. *Quarterly Journal of Economics*, 7:1, 151-200.
- Castellano, C., S. Fortunato, and V. Loreto (2009) "Statistical physics of social dynamics" *Reviews of Modern Physics*, (eprint arXiv: 0710.3256).
- Casella, A., and J. Rauch (2002) Anonymous Market and Group Ties in International Trade, *Journal of International Economics* 58(1):194-7.
- Celen, B., S. Kariv, and A. Schotter (2004) Learning in Networks: An Experimental Study, New York and Berkeley, Columbia University, University of California at Berkeley, and New York University, mimeo.

- Chandra, A. (2000) Labor-Market Dropouts and the Racial Wage Gap: 1940-1990, *American Economic Review (Papers and Proceedings)*, 90:2, 333-338.
- Charness, G., and M.O. Jackson (2007) Group Play in Games and the Role of Consent in Network Formation, *Journal of Economic Theory* 136(1):417-445.
- Charness G, Corominas-Bosch M, Frechette GR. 2005. Bargaining on networks: an experiment. *J. Econ. Theory* 136:28 - 65
- Choi S, Gale D, Kariv S. 2005. Behavioral aspects of learning in social networks: an experimental study. In *Advances in Applied Microeconomics*, Vol. 13:, Behavioral and Experimental Economics, ed. J Morgan. Berkeley, CA: Bepress
- Choi S, Gale D, Kariv S. 2007. Social Learning in Networks: A Quantal Response Equilibrium Analysis of Experimental Data. Berkeley: Univ. Calif. Berkeley
- Christakis NA, Fowler JH. 2007. The spread of obesity in a large social network over 32 years. *N. Engl. J. Med.* 357:370-379
- Christakis NA, Fowler JH. 2008. The collective dynamics of smoking in a large social network. *N. Engl. J. Med.* 358(21):224-231
- Christakis NA, G.W. Imbens, and Fowler JH. (2010) "An Empirical Model for Strategic Network Formation" mimeo: Harvard University.
- Chung, F., and L. Lu (2002) The Average Distances in Random Graphs with Given Expected Degrees, *Proceedings of the National Academy of Sciences* 99:15879-15882.
- Clifford, P., and A. Sudbury, 1973, "A Model for Spatial Conflict," *Biometrika* 60:3, 581.
- Coase RH. 1960. The problem of social cost. *J. Law Econ.* 3:144
- Cohen, J., (1977) "Sources of peer group homogeneity", *Sociol. Educ.* 50:227-412
- Coleman, J., (1958) "Relational analysis: the study of social organizations with survey methods," *Human Organization*, 17, 28-36.
- Coleman, J.S. (1988) Social Capital in the Creation of Human Capital, *American Journal of Sociology* 94(Supplement: Organizations and Institutions: Sociological and Economic Approaches to the Analysis of Social Structure):S95-S120.

- Coleman, J.S. (1990) *Foundations of Social Theory*, Cambridge Mass.: Harvard University Press.
- Coleman, J.S., E. Katz, and H. Menzel (1966) *Medical Innovation: A Diffusion Study*, Indianapolis, Ind.: Bobbs-Merrill.
- Comola (2009) "The Network Structure of Informal Arrangements: Evidence from Rural Tanzania," mimeo: Paris School of Economics.
- Conley, T.G., and G. Topa (2001) Socio-Economic Distance and Spatial Patterns in Unemployment, *Journal of Applied Economics* 17(4):303327.
- Conley, T.G., and C.R. Udry (2001) Social Learning through Networks: The Adoption of New Agricultural Technologies in Ghana, *American Journal of Agricultural Economics* 83(3):668673.
- Conley, T.G., and C.R. Udry (2004) *Social Networks in Ghana*, New Haven, Conn., Economic Growth Center, Yale University, Discussion Paper 888.
- Conley, T.G., and C.R. Udry (2004) Learning About a New Technology: Pineapple in Ghana, Economic Growth Center Working Paper 817, New Haven, Conn., Yale University.
- Conley, T.G., and C.R. Udry (2004) The Adoption of New Agricultural Technologies in Ghana, *American Journal of Agricultural Economics* 83(3):668673.
- Cook, K.S., and R.M. Emerson (1978) Power, Equity and Commitment in Exchange Networks, *American Sociological Review* 43:721739.
- Cook, K.S., and J.M. Whitmeyer (1992) Two Approaches to Social Structure: Exchange Theory and Network Analysis, *Annual Review of Sociology* 18:109127.
- Cook, K.S., and R.M. Emerson (1978) Power, Equity and Commitment in Exchange Networks, *American Sociological Review* 43:721739.
- Copic, J., M.O. Jackson, and A. Kirman (2009) "Identifying Community Structures from Network Data," in the *B.E. Press Journal of Theoretical Economics*, Vol. 9, Iss. 1 (Contributions), Article 30.
- Corbae, D., and J. Duffy (2009) *Experiments with Network Economies*, Games and Economic Behavior.

- Corcoran M, Datcher L, Duncan G. 1980. Information and influence networks in labor markets. In *Five Thousand American Families*, Vol. 8, ed. G Duncan, J Morgan. Ann Arbor: Univ. Mich. 1 - 38
- Corominas-Bosch M. 2004. On two-sided network markets. *J. Econ. Theory* 115: 35 - 77
- Currarini, S., M.O. Jackson, and P. Pin (2007) "Long Run Integration in Social Networks," mimeo.
- Currarini S, Jackson M.O., Pin P. (2009) "An economic model of friendship: homophily, minorities and segregation," *Econometrica*, Vol. 77, No. 4, 1003-1045,
- Currarini, S., M.O. Jackson, and P. Pin (2010) "Identifying the roles of race-based choice and chance in high school friendship network formation," in the *Proceedings of the National Academy of Sciences*, 107(11): 4857 - 4861
- Currarini S, Morelli M. 2000. Network formation with sequential demands. *Rev. Econ. Des.* 5: 229 - 50
- DeGroot M.H. 1974. Reaching a consensus. *J. Am. Stat. Assoc.* 69: 118 - 21
- Demange, G. and M. Wooders. (2005) *Group Formation in Economics; Networks, Clubs and Coalitions*, G. Cambridge: Cambridge University Press.
- DeMarzo P., Vayanos D., Zwiebel J. 2003. Persuasion bias, social influence, and uni-dimensional opinions. *Q. J. Econ.* 1183: 909 - 68
- Deröian, F. (2003) Farsighted Strategies in the Formation of a Communication Network, *Economics Letters* 80:343349.
- De Weerd J. 2004. Risk-sharing and endogenous network formation. In *Insurance Against Poverty*, ed. S Dercon. Oxford: Oxford Univ. Press
- De Weerd J, Dercon S. 2006. Risk-sharing networks and insurance against illness. *J. Dev. Econ.* 81: 337 - 56
- Diestel, R. (2000) *Graph Theory*, Heidelberg: Springer-Verlag.

- Droste, E., R.P. Gilles, and C. Johnson (2000) Evolution of Conventions in Endogenous Social Networks, in Econometric Society World Congress 2000 Contributed Papers 0594, Seattle: Econometric Society
- Duflo, E. and E Saez 2003 'The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence From a Randomized Experiment,' Quarterly Journal of Economics, August 2003, Vol. 118, No. 3: 815 - 842
- Dutta, B., and M.O. Jackson (2000) The Stability and Efficiency of Directed Communication Networks, Review of Economic Design 5: 251272.
- Dutta, B., and M.O. Jackson (2003) Networks and Groups: Models of Strategic Formation, Heidelberg: Springer-Verlag.
- Dutta B, Ghosal S, Ray D. 2005. Farsighted network formation. J. Econ. Theory 122: 143 - 64
- Dutta B, Mutuswami S. 1997. Stable networks. J. Econ. Theory 76: 322 - 44
- Economides, N. (1996) The Economics of Networks, International Journal of Industrial Organization 16(4):673699
- Ehrhardt, G., M. Marsili, and F. Vega-Redondo (2006) "Diffusion and Growth in an Evolving Network," International Journal of Game Theory 34:383397.
- Elliott, M. (2009) "Inefficiencies in Trade Networks" mimeo: Stanford University.
- Ellison, G. (1993) Learning, Local Interaction, and Coordination, Econometrica 61:10471071.
- Ellison, G., and D. Fudenberg (1995) Word-of-Mouth Communication and Social Learning, Quarterly Journal of Economics 110:93126.
- Ely, Jeffrey C. (2002) "Local Conventions," Advances in Theoretical Economics: Vol. 2 : Iss. 1, Article 1.
- Emerson, R.M. (1962) Power-Dependence Relations, American Sociological Review 27(3):140.
- Emerson, RM (1967) Exchange Theory, Part I: A Psychological Basis for Social Exchange, and Exchange Theory, Part II: Exchange Relations and Networks, in Sociological Theories in Progress, J. Berger, M. Zelditch, Jr., and B. Anderson, eds., Boston: Houghton-Mifflin.

- Erdős P, Rényi A. 1959. On random graphs. *Publ. Math. Debrecen* 6: 290 - 97
- Erdős P, Rényi A. 1960. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.* 5: 17 - 61
- Erdős P, Rényi A. 1961. On the strength of connectedness of a random graph. *Acta Math. Acad. Sci. Hung.* 12: 261 - 67
- Fafchamps M, Lund S. 2003. Risk-sharing networks in rural Philippines. *J. Dev. Econ.* 71: 261 - 87
- Fagiolo, G., M. Valente, and N. Vriend (2007) "Segregation in Networks," *Journal of Economic Behavior and Organization*, 64: 316-336.
- Falk, A., and M. Kosfeld (2003) Its All About Connections: Evidence on Network Formation, Institute for the Study of Labor (IZA) Discussion Paper 777, Zurich IIEE Working Paper 146, Zurich, Switzerland.
- Feld, S.L., (1981) "The Focused Organization of Social Ties", *The American Journal of Sociology* 86 (5), 1015–1035
- Fong, E., and W.W. Isajiw (2000) "Determinants of friendship choices in multiethnic society," *Sociological Forum* 15 (2), 249–271
- Fosco, C. and F. Mengel (2009) "Cooperation through Imitation and Exclusion in Networks," mimeo Maastricht University.
- Fowler, J.H. and N.A. Christakis (2008) "Estimating Peer Effects on Health in Social Networks," *Journal of Health Economics* 27(5): 1400-1405.
- Frank, O., and D. Strauss (1986) Markov Graphs, *Journal of the American Statistical Association* 81:832842.
- Freeman, L.C. (2004) *The Development of Social Network Analysis: A Study in the Sociology of Science*, Vancouver: Empirical Press.
- French, J. (1956): A Formal Theory of Social Power, *Psychological Review*, 63: 181 - 194.
- Friedkin, N.E. and Johnsen, E.C. (1997) "Social Positions in Influence Networks," *Social Networks*, 19, 209-222.

- Furusawa T, Konishi H. 2005. Free trade networks. *Jpn. Econ. Rev.* 56: 144 - 64
- Gale D, Kariv S. 2003. Bayesian learning in social networks. *Games Econ. Behav.* 45(2): 329 - 46
- Galeotti A., Goyal S., Jackson M.O., Vega-Redondo F., Yariv L. (2009) "Network games," *Review of Economic Studies*, <http://www.jacksonm.edu/~jacksonm/networkgames.pdf>
- Garg, M. (2009) "Axiomatic Foundations of Centrality in Networks," mimeo: Stanford University.
- Giles, M.W., and A. Evans (1986) "The power approach to intergroup hostility," *Journal of Conflict Resolution* 30 (3): 469 - 486
- Gilles RP, Sarangi S. 2005. Stable networks and convex payoffs. In *Review of Economic Design*.
- Girvan, M., and M.E.J. Newman (2002) Community Structure in Social and Biological Networks, *Proceedings of the National Academy of Sciences* 99:78217826.
- Glaeser E, Sacerdote B, Scheinkman J. 1996. Crime and social interactions. *Q. J. Econ.* 111: 507 - 48
- Goeree, MA McConnell, T Mitchell, T Tromp, L Yariv (2007) 'Linking and giving among teenage girls,' mimeo: California Institute of Technology Working Paper.
- Goeree, J.K., A. Riedl, and A. Ule (2003) In Search of Stars: Network Formation among Heterogeneous Agents, Institute for the Study of Labor (IZA) Discussion Paper 1754, Bonn.
- Goeree, J.K., M.A. McConnell, T. Mitchell, T. Tromp, and L. Yariv (2007) Linking and Giving Among Teenage Girls, Pasadena, California Institute of Technology, mimeo.
- Golub B, Jackson MO. 2008. How Homophily affects the Speed of Learning and Diffusion in Networks. Stanford, CA: Stanford Univ, arXiv:0811.4013v2 [physics.soc-ph].
- Golub B, Jackson MO. 2010. "Naive learning and influence in social networks: convergence and wise crowds," *the American Economic Journal: Microeconomics*, 2(1): 11249, Feb. 2010.

- Goodreau, SM, JA. Kitts, and M Morris (2009) "Birds of a Feather, or Friend of a Friend? Using Exponential Random Graph Models to Investigate Adolescent Social Networks," *Demography*, 46(1): 103125.
- Goyal, S. (2008) *Connections: An Introduction to the Economics of Networks*, Princeton University Press.
- Goyal, S., and S. Joshi (2003) Networks of Collaboration in Oligopoly, Games and Economic Behavior 43:5785.
- Goyal, S., and F. Vega-Redondo (2005) "Learning, Network Formation and Coordination," *Games and Economic Behavior*, 50: 178207.
- Granovetter M. 1973. The strength of weak ties. *Am. J. Sociol.* 78: 1360 - 80
- Granovetter M. 1978. Threshold models of collective behavior. *Am. J. Sociol.* 83(6): 1420 - 43
- Granovetter M. 1985. Economic action and social structure: the problem of embeddedness. *Am. J. Sociol.* 91(3): 481 - 510
- Granovetter M. 1995. *Getting a Job: A Study of Contacts and Careers*. Chicago: Univ. Chicago Press. 2nd ed.
- Hagenbach, J. and F. Koessler (2009) "Strategic communication networks," Preprint, Université de Paris I - Sorbonne.
- Hagerstrand, T. (1970) "What about people in Regional Science?" *Journal Papers in Regional Science*, 24:1, 6-21.
- Halliday, T.J. and S. Kwak (2009) "Weight Gain in Adolescents and Their Peers," *Economics and Human Biology*.
- Heckman, J.J. and Borjas, G. (1980) Does Unemployment Cause Future Unemployment? Definitions, Questions and Answers from a Continuous Time Model of Heterogeneity and State Dependence. *Economica*, 47:187, 24783.
- Herings PJJ, Mauleon A, Vannetelbosch V. 2004. Rationalizability for social environments. *Games Econ. Behav.* 49: 135 - 56

- Hesselius, P., P. Johansson, and P. Nilsson (2009) "Sick of Your Colleagues Absence?" IZA DP No. 3960
- Hoff, P.D. (2006) Multiplicative Latent Factor Models for Description and Prediction of Social Networks, Seattle, University of Washington, mimeo.
- Hoff, P.D., A.E. Raftery, and M.S. Handcock (2002) Latent Space Approaches to Social Network Analysis, *Journal of the American Statistical Association* 97:1090-1098.
- Hojman D, Szeidl A. 2006. Endogenous networks, social games and evolution. *Games Econ. Behav.* 55:1: 112 - 30
- Holland, P.W., and S. Leinhardt (1977) A Dynamic Model for Social Networks, *Journal of Mathematical Sociology* 5:520.
- Holland, P.W., K.B. Laskey, and S. Leinhardt (1983) Stochastic Blockmodels: First Steps, *Social Networks* 5:109-137.
- Holley, R., and T. Liggett (1975) "Ergodic theorems for weakly interacting infinite systems and the voter model," *Ann. Probab.* 3, 643.
- Holme, P. and MEJ Newman (2006) "Nonequilibrium phase transition in the coevolution of networks and opinions," arXiv:Physics 060303 v3.
- Homans, C.G. (1958) Social Behavior as Exchange, *American Journal of Sociology* 62:597-606.
- Homans, C.G. (1961) *Social Behavior: Its Elementary Forms*, New York: Harcourt, Brace and World.
- Harary, F. (1959) 'Status and Contrastatus', *Sociometry* 22: 23 - 43.
- Holme, P. and Newman, M.E.J. (2006) "Nonequilibrium phase transition in the coevolution of networks and opinions," *PHYSICAL REVIEW E* 74, 056108
- Huberman, B.A. and L. Adamic (1999) "Growth dynamics of the World-Wide Web," *Nature*: VOL 401: 9 SEPTEMBER, page 131.
- Ioannides YM, Datcher-Lourey L. 2004. Job information networks, neighborhood effects and inequality. *J. Econ. Lit.* 42:4: 1056 - 93

- Jackson MO. 2003. The stability and efficiency of economic and social networks. In *Advances in Economic Design*, ed. S Koray, M Sertel. Heidelberg: Springer-Verlag. Reprinted in *Networks and Groups: Models of Strategic Formation*, ed. B Dutta, M.O Jackson. Heidelberg: Springer-Verlag, 99 - 140
- Jackson M.O. 2004. A survey of models of network formation: stability and efficiency. In *Group Formation in Economics; Networks, Clubs and Coalitions*, ed. G Demange, M Wooders. Cambridge, UK: Cambridge Univ. Press, 11 - 57.
- Jackson M.O. 2005a. Allocation rules for network games. *Games Econ. Behav.* 51: 128 - 54
- Jackson M.O. 2005b. The economics of social networks. Chapter 1 in *Proc. 9th World Congr. Econ. Soc.*, ed. R Blundell, W Newey, T Persson. Cambridge, UK: Cambridge Univ. Press
- Jackson M.O. 2007 *Social Structure, Segregation, and Economic Behavior*, Nancy Schwartz Memorial Lecture, given in April 2007 at Northwestern University, printed version: <http://www.stanford.edu/~jacksonm/schwartzlecture.pdf>
- Jackson M.O. (2008) *Social and Economic Networks*, Princeton, NJ: Princeton Univ. Press
- Jackson, M.O. (2008b) Average Distance, Diameter, and Clustering in Social Networks with Homophily, arXiv:0810.2603v1 [physics.soc-ph], in the *Proceedings of the Workshop in Internet and Network Economics (WINE 2008)*, *Lecture Notes in Computer Science*, Edited by C. Papadimitriou and S. Zhang, Springer Verlag, Berlin Heidelberg.
- Jackson, M.O., T. Rodriguez-Barraquer, and X. Tan (2009) "Social Capital and Social Quilts: Network Patterns of Favor Exchange," mimeo: Stanford University.
- Jackson M.O., Rogers B.W. 2005. The economics of small worlds. *J. Eur. Econ. Assoc. Pap. Proc.* 32(3): 617 - 27
- Jackson M.O., Rogers B.W. 2007. Meeting strangers and friends of friends: how random are social networks? *Am. Econ. Rev.* 97(3): 890 - 915
- Jackson M.O., Rogers B.W. 2007b. Relating network structure to diffusion properties through stochastic dominance. *B.E. Press J. Theor. Econ.* 7(1): 1 - 13

- Jackson M.O., van den Nouweland A. 2005. Strongly stable networks. *Games Econ. Behav.* 51: 420 - 44
- Jackson M.O., Watts A. 2002. The evolution of social and economic networks. *J. Econ. Theory* 106(2): 265 - 95
- Jackson M.O., Watts A. (2002b) On the Formation of Interaction Networks in Social Coordination Games, *Games and Economic Behavior* 41(2):265-291.
- Jackson M.O., Watts A. (2005) Social Games: Matching and the Play of Finitely Repeated Games, *Games and Economic Behavior*, forthcoming.
- Jackson M.O., Wolinsky A. (1996) "A strategic model of social and economic networks," *J. Econ. Theory* 71(1): 44 - 74.
- Jackson M.O., Yariv L. 2007. The diffusion of behavior and equilibrium structure properties on social networks. *Am. Econ. Rev. (Papers and Proceedings)* 97: 92 - 98
- Jackson, M.O. and L. Yariv (2010) "Diffusion, Strategic Interaction, and Social Structure," This volume.
- Jadbabaie, A., A. Sandroni, and A. Tahbaz-Salehi (2009) "Non-Bayesian Social Learning," mimeo: Northwestern University.
- Johari, R., S. Mannor, and J.N. Tsitsiklis (2006) A Contract-Based Model for Directed Network Formation, *Games and Economic Behavior* 56(2):201-224.
- Johnson, C., and R.P. Gilles (2000) Spatial Social Networks, *Review of Economic Design* 5:273-300.
- Kakade SM, Kearns M, Ortiz LE. 2004a. Graphical economics. *Proc. Annu. Conf. Learn. Theory* 23 3120: 17 - 32
- Kakade SM, Kearns M, Ortiz LE, Pemantle R, Suri S. 2004b. Economic Properties of Social Networks, *Proc. Neural Information Processing Syst., NIPS.*, Cambridge, MA: MIT Press
- Kandel, D.B. (1978) Homophily, Selection, and Socialization in Adolescent Friendships, *American Journal of Sociology* 14: 427 - 436.

- Kandori, M., G. Mailath, and R. Rob (1993) Learning, Mutation, and Long-Run Equilibria in Games, *Econometrica* 61:2956.
- Karlan, D., M. Mobius, T. Rosenblat, and A. Szeidl (2009) “Measuring trust in Peruvian shantytowns,” mimeo.
- Katz, E., and P.F. Lazarsfeld (1955) *Personal Influence: The Part Played by People in the Flow of Mass Communication*, New York: Free Press.
- Katz, L. (1953) A New Status Index Derived from Sociometric Analysis, *Psychometrika* 18:3943.
- Katz, M., and C. Shapiro (1994) Systems Competition and Networks Effects, *Journal of Economic Perspectives* 8:93115.
- Kearns, M.J., M. Littman, and S. Singh (2001) Graphical Models for Game Theory, in *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*, J.S. Breese and D. Koller, eds., San Francisco: Morgan Kaufmann.
- Kearns, M.J., S. Suri, and N. Montfort (2006) An Experimental Study of the Coloring Problem on Human Subject Networks, *Science* 313:824827.
- Kearns, M.J., S. Judd, J. Tan and J. Wortman (2009) “Behavioral Experiments on Biased Voting in Networks,” *PNAS*, 106:5, 13471352.
- Kets, W. (2008) *Networks and Learning in Game Theory*, dissertation, Tilburg University, Tilburg, The Netherlands.
- Khmelkov, V. T., and M. T. Hallinan (1999) “Organizational Effects on Race Relations in Schools”, *Journal of Social Issues*, 55 (4), 627-645
- Kirman AP. 1983. Communication in markets: a suggested approach. *Econ. Lett.* 12: 1-5
- Kirman AP, Oddou C, Weber S. 1986. Stochastic communication and coalition formation. *Econometrica* 54: 129-38
- Kling, J.R., Ludwig, J. and L.F. Katz (2005), Neighborhood effects on crime for female and male youth: evidence from a randomized housing voucher experiment, *Quarterly Journal of Economics* 120, 87 - 130

- Knoke, D., (1990) "Networks of Political Action: Towards Theory construction," *Social Forces*, 68, 1041-1063.
- Kochen, M. (1989) *The Small World*, Norwood, N.J.: Albex.
- Kogut, B. (2000) The Network as Knowledge: Generative Rules and the Emergence of Structure, *Strategic Management Journal* 21(3):405-425.
- Kosfeld, M. (2003) *Network Experiments*, Zurich IEER Working Paper 152, University of Zurich.
- Kranton R, Minehart D. 2001. A theory of buyer-seller networks. *Am. Econ. Rev.* 91(3): 485-508
- Krause, U. (2000) A Discrete Nonlinear and Nonautonomous Model of Consensus Formation, in *Communications in Difference Equations*, S. Elaydi, G. Ladas, J. Popenda, and J. Rakowski, eds., Amsterdam: Gordon and Breach.
- Warren N. Kubitschek and Maureen T. Hallinan (1998) "Tracking and Students' Friendships," *Social Psychology Quarterly*, Vol. 61, No. 1, 1 - 15.
- Lamberson, PJ (2009) "Social Learning in Social Networks," mimeo: MIT
- Langville, A.N. and Meyer, C.D. (2006), *Google's PageRank and Beyond: The Science of Search Engine Rankings*, Princeton University Press, Princeton, N.J.
- Laschever R. 2007. The doughboys network: social interactions and labor market outcomes of World War I veterans. *Work. Pap.*, Univ. Ill., Champaign-Urbana
- Lazarsfeld, P.F. and N.W. (1968) *Latent structure analysis* Houghton, Mifflin.
- Lazarsfeld, P.F., and R.K. Merton (1954) Friendship as a Social Process: A Substantive and Methodological Analysis, in *Freedom and Control in Modern Society*, M. Berger, ed., New York: Van Nostrand.
- Leider, S., M.M. Mobius, T. Rosenblatt, and Q.-A. Do (2007) How Much Is a Friend Worth? Directed Altruism and Enforced Reciprocity in Social Networks, revision of NBER Working Paper 13135, Cambridge, Mass., National Bureau of Economics Research.
- Lévi-Strauss, C. (1958) *Anthropologie structurale*, Paris: Librairie Plon.

- Littman, M.L., M.J. Kearns, and S.P. Singh (2001) An Efficient, Exact Algorithm for Solving Tree-Structured Graphical Games, in *Advances in Neural Information Processing Systems*, Cambridge, Mass.: MIT Press.
- Lopez-Pintado D. 2008. Diffusion in complex social networks. *Games Econ. Behav.* 62(2): 573–90
- Lorenz, J. (2005) A Stabilization Theorem for Dynamics of Continuous Opinions, *Physica A* 355:217223.
- Lorrain, F., and H.C. White (1971) Structural Equivalence of Individuals in Social Networks, *Journal of Mathematical Sociology* 1:4980.
- Loury, G (1977). A Dynamic Theory of Racial Income Differences. Chapter 8 of *Women, Minorities, and Employment Discrimination*, Ed. P.A. Wallace and A. Le Mund. Lexington, Mass.: Lexington Books.
- Lynch, L.M. (1989) The Youth Labor Market in the Eighties: Determinants of Reemployment Probabilities for Young Men and Women, *Review of Economics and Statistics* 71:3754.
- Manea, M. (2008): “Bargaining on networks,” mimeo: Princeton University.
- Manski, C.F. (1993) Identification of Endogenous Social Effects: The Reflection Problem, *Review of Economic Studies* 60: 419–431.
- Marsden PV. 1987. Core discussion networks of Americans. *Am. Sociol. Rev.* 52: 122–31
- Marsden PV. 1988. Homogeneity in confiding relations. *Soc. Netw.* 10: 57–76
- Marsili, M., F. Vega-Redondo, and F. Slanina (2007) The Rise and Fall of a Networked Society: A Formal Model, *Proceedings of the National Academy of Sciences of the U.S.A.* 101:14391442.
- Mauleon A., and V.J. Vannetelbosch (2004) Farsightedness and Cautiousness in Coalition Formation, *Theory and Decision* 56:291324.
- McPherson M, Smith-Lovin L, Cook JM. 2001. Birds of a feather: homophily in social networks. *Annu. Rev. Sociol.* 27: 415–44

- Meeker, R. J., and D. M. Weiler (1970) A New School for the Cities.
- Milgram S. 1967. The small-world problem. *Psychol. Today* 2: 60–67
- Mitzenmacher, M. (2004) A Brief History of Generative Models for Power Law and Log-normal Distributions, manuscript available at <http://www.eecs.harvard.edu/~michaelm/ListByYear.h>
- Mobius, M.M., and T.S. Rosenblatt (2003) Experimental Evidence on Trading Favors in Networks, Cambridge, Mass., and Middletown, Conn., Harvard University and Wesleyan University, mimeo.
- Mobius, M.M., and A. Szeidl (2006) Trust and Social Collateral, Cambridge, Mass., and Berkeley, Harvard University and the University of California at Berkeley.
- Montgomery J. 1991. Social networks and labor market outcomes. *Am. Econ. Rev.* 81: 1408–18
- Moody, J. (2001) Race, School Integration, and Friendship Segregation in America, *American Journal of Sociology* 107(3):679–716.
- Morris, S. (2000) Contagion, *Review of Economic Studies* 67:577–8.
- Mueller-Frank, M. (2009) “A General Framework for Rational Learning in Social Networks,” mimeo: Northwestern University.
- Munshi K. 2003. Networks in the modern economy: Mexican migrants in the U.S. labor market. *Q. J. Econ.* 118(2): 549–97
- Mutuswami S, Winter E. 2002. Subscription mechanisms for network formation. *J. Econ. Theory* 106: 242–64
- Myers CA, Shultz GP. 1951. *The Dynamics of a Labor Market*. New York: Prentice-Hall
- Myerson R. 1977. Graphs and cooperation in games. *Math. Oper. Res.* 2: 225–29
- Nava, F. (2009) “Quantity Competition in Networked Markets,” mimeo: University of Chicago.
- Newman, M.E.J. (2002) The Spread of Epidemic Disease on Networks, *Physical Review E* 66:016128.

- Newman, M.E.J. (2003) *The Structure and Function of Complex Networks*, SIAM Review 45:167256.
- Newman, M.E.J. (2004) *Detecting Community Structure in Networks*, Physical Review E 69:066133.
- Newman M.E.J., AL Barabasi DJ. Watts (2006) *The Structure and Dynamics of Networks: (Princeton Studies in Complexity, Princeton University Press.*
- Page F, Wooders M, Kamat S. 2005. Networks and farsighted stability. *J. Econ. Theory* 120(2): 257–69.
- Pantz, K., and A. Ziegelmeyer (2003) *An Experimental Study of Network Formation*, Garching, Germany, Max Planck Institute, mimeo.
- Papadimitriou, C.H. (2001) *Algorithms, Games, and the Internet*, in *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing*, New York: ACM.
- Pareto, V. (1896) *Cours d'Economie Politique*, Geneva: Droz.
- Pastor-Satorras R, Vespignani A. 2000. Epidemic spreading in scale-free networks. *Phys. Rev. Lett.* 86: 3200–3
- Pastor-Satorras R, Vespignani A. 2001. Epidemic dynamics and endemic states in complex networks. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 63:066117
- Patacchini, E. and Y. Zenou (2008), *The strength of weak ties in crime*, *European Economic Review* 52, 209-236.
- Patacchini, E., and Y. Zenou (2006) “*Racial Identity and Education*,” SSRN Working Papers
- Pattison, P.E. (1993) *Algebraic Models for Social Networks*, Cambridge: Cambridge University Press.
- Pattison, P.E., and S. Wasserman (1999) *Logit Models and Logistic Regressions for Social Networks: II. Multivariate Relations*, *British Journal of Mathematical and Statistical Psychology* 52:169193.
- Pellizzari M. 2004. Do friends and relatives really help in getting a job? *Cent. Econ. Res. Discuss. Pap.* 623. London Sch. Econ.

- Pellizzari, M. (2009), Do friends and relatives really help in getting a good job?, *Industrial and Labor Relations Review*, forthcoming.
- Pennock DM, Flake GW, Lawrence S, Glover EJ, Giles CL. 2002. Winners dont take all: characterizing the competition for links on the Web. *Proc. Natl. Acad. Sci. USA* 99(8): 5207-11
- Peski, M. (2008), "Complementarities, Group Formation, and Preferences for Similarity," mimeo.
- Pissarides, C.A. (2000), *Equilibrium Unemployment Theory*, Second edition, Cambridge, MA: MIT Press.
- Price DJS. 1965. Networks of scientific papers. *Science* 149: 510-15
- Price DJS. 1976. A general theory of bibliometric and other cumulative advantage processes. *J. Am. Soc. Inf. Sci.* 27: 292-306
- Putnam, R. (2000), "Bowling Alone: The Collapse and Revival of American Community" (Simon and Schuster).
- Rauch JE (2007) *Missing Links: Formation and Decay of Economic Networks*, NY: Russell Sage Foundation Publications
- Rees AJ, Shultz GP. 1970. *Workers in an Urban Labor Market*. Chicago: Univ. Chicago Press
- Reiss AJ. 1980. Understanding changes in crime rates. In *Indicators of Crime and Criminal Justice: Quantitative Studies*. Washington, DC: Bur. Justice Stat.
- Reiss AJ. 1988. Co-offending and criminal careers. In *Crime and Justice: A Review of Research*, Vol. 10, ed. M Tonry. Chicago: Univ. Chicago Press
- Rogers BW. 2006. A strategic theory of interdependent status. PhD diss., Calif. Inst. Technol.
- Rogers, E.M., and E. Rogers (2003) *Diffusion of Innovations*, fifth edition, New York: Free Press.
- Rosenberg, D., E. Solan, and N. Vieille (2007) *Informational Externalities and Emergence of Consensus*, Paris, HEC School of Management, mimeo.

- Roughgarden, T., and E. Tardos (2002) How Bad Is Selfish Routing? *Journal of the ACM* 49(2):236259.
- Ryan, B., and N.C. Gross (1943) The Diffusion of Hybrid Seed Corn in Two Iowa Communities, *Rural Sociology* 8:1524.
- Rytina, S., and D. L. Morgan (1982) “The Arithmetic of Social Relations: The Interplay of Category and Network,” *The American Journal of Sociology*, 88 (1), 88–113
- Schelling, T.C. (1978) *Micromotives and macrobehavior*, New York. W. W. Norton.
- Schweitzer, S.O. and Smith, R.E. (1974) The Persistence of the Discouraged Worker Effect. *Industrial and Labor Relations Review*, 27:2, 24960.
- Skyrms, B., and R. Pemantle (2000) A Dynamic Model of Social Network Formation, *Proceedings of the National Academy of Sciences, U.S.A.* 97:93409346.
- Slate, J. R., and C. H. Jones (2005) “Effects of school size: A review of the literature with recommendations”, *Essays in Education*.
- Slikker M, van den Nouweland A. 2001. *Social and Economic Networks in Cooperative Game Theory*. Norwell, MA: Kluwer Acad. Publ.
- Smith, L. and P. Sorensen, (2000) “Pathological outcomes of observational learning,” *Econometrica* 68:2, 371-398.
- Snijders TAB, PE. Pattison, GL. Robins, MS. Handcock (2006), ‘New Specifications for Exponential Random Graph Models,’ *Sociological Methodology*, Vol. 36 : 99 - 153
- Snijders, T.A.B. (2002) Markov Chain Monte Carlo Estimation of Exponential Random Graph Models, *Journal of Social Structure* 3(2):240.
- Snijders, T.A.B., and K. Nowicki (1997) Estimation and Prediction for Stochastic Block Models for Graphs with Latent Block Structure, *Journal of Classification* 14:75100.
- Snijders, T.A.B., C.E.G. Steglich, and M. Schweinberger (2007) Modeling the Coevolution of Networks and Behavior, in *Longitudinal Models in the Behavioral and Related Sciences*, K. van Montfort, H. Oud, and A. Satorra, eds., New York: Routledge.
- Sobel, J. (2002) Can We Trust Social Capital? *Journal of Economic Literature* 40: 139154.

- Solomonoff, R., and A. Rapoport (1951) Connectivity of Random Nets, *Bulletin of Mathematical Biophysics* 13:107117.
- Stearns, E., (2004) “Interracial Friendliness and the Social Organization of Schools”, *Youth & Society*, 35 (4), 395–419
- Sundararajan A. 2005. Local network effects and network structure. New York Univ. Stern.
- Tesfatsion L. 1997. A trade network game with endogenous partner selection. In *Computational Approaches to Economic Problems*, ed. H Amman, B Rustem, AB Whinston, pp. 249–69. Norwell, MA: Kluwer Acad. Publ.
- Topa, G. (2001) Social Interactions, Local Spillovers and Unemployment, *Review of Economic Studies* 68:261296.
- Uzzi B. 1996. The sources and consequences of embeddedness for the economic performance of organizations: the network effect. *Am. Sociol. Rev.* 61: 674–98
- van der Leij MJ and IS Buhai (2008) “A Social Network Analysis of Occupational Segregation” *Fondazione Eni Enrico Mattei Working Paper* 192.
- Vázquez, A. (2003) Growing Network with Local Rules: Preferential Attachment, Clustering Hierarchy, and Degree Correlations, *Physical Review E* 67(5):056104.
- Vega-Redondo, F. (2007) *Complex Social Networks*, *Econometric Society Monographs*, Cambridge: Cambridge University Press.
- Wasserman, S., and K. Faust (1994) *Social Network Analysis: Methods and Applications*, Cambridge: Cambridge University Press.
- Wasserman, S., and P. Pattison (1996) Logit Models and Logistic Regressions for Social Networks: I. An Introduction to Markov Graphs and P*, *Psychometrika* 61:401425.
- Watts, A. (2001) A Dynamic Model of Network Formation, *Games and Economic Behavior* 34:331341.
- Watts DJ. 1999. *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton, NJ: Princeton Univ. Press
- Watts D.J. (2004) The New Science of Networks, *Annual Sociological Review* 30:243270.

Watts D.J., Strogatz S. 1998. Collective dynamics of. small-world.networks. Nature 393: 440 42.

White, H.C., S.A. Boorman, and R.L. Breiger (1976) Social Structure from Multiple Networks. I. Blockmodels of Roles and Positions, American Journal of Sociology 81(4):730780.

Young, H.P. (1993) The Evolution of Conventions, Econometrica 61:5784.

Young, H.P. (1998) Individual Strategy and Social Structure, Princeton, N.J.: Princeton University Press.

Zipf, G. (1949) Human Behavior and the Principle of Least Effort, Cambridge, Mass.: Addison-Wesley.