

Social Networks

Joan de Martí and Yves Zenou

INTRODUCTION

A large body of research, first in sociology, then in physics, and more recently in economics, has studied the importance of social networks in different activities. Social networks are indeed important in several facets of our lives. For example, the decision of an agent whether or not to buy a new product, attend a meeting, commit a crime, find a job is often influenced by the choices of their friends and acquaintances (be they social or professional). The emerging empirical evidence on these issues motivates the theoretical study of network effects. For example, job offers can be obtained from direct, and indirect, acquaintances through word-of-mouth communication. The spread of diseases, such as AIDS, also strongly depends on the geometry of social contacts. If the web of connections is dense, we can expect higher infection rates.

The study of social networks was indeed initiated by sociologists more than a century ago and has grown to be a central field of sociology over the past 50 years (see, for example, Wasserman and Faust, 1994). During that same period, a mathematical literature on the structure and properties of graphs has been developed and extensively studied (see Bollobás, 1998; Diestel, 2005).

A recent awakening of interest in social networks has occurred in the computer science and statistical physics literatures, mainly over the past five or six years (see Albert and Barabási, 2002; Newman, 2003, for an overview of these studies). While the importance of embeddedness of economic activity in social settings has been fundamental to sociologists (and to some extent to applied mathematicians) for some time, it was largely ignored by economists until the last decade. Indeed, studies of networks with economic perspectives and using game-theoretic modeling techniques have only emerged over the last decade (see Goyal, 2007; Jackson, 2007, 2008; Vega-Redondo, 2007).

A network is an abstract object that models these social interactions. In particular, a network is formed by *nodes* (or *vertices*) that represent the actors involved, and *edges* (or *links*) that express the linkage among these nodes. Networks provide a simplified geometrical representation of a complex magma of social relationships. However, if social interactions represent a first-order driving force for the problem under consideration, a detailed study of the characteristics of the network should reveal some relevant features of social structure that induce the resulting outcomes. In the job market example, it is of paramount importance to know the

geometric characteristics of the network that induce job market outcomes and how different individuals end up in different situations due to their asymmetric positions in the network of personal relationships. Isolated individuals, or individuals with low-quality links with the rest of the community, have weaker positions in the network and are therefore more prone to be, and to stay, unemployed for a long period of time since they do not obtain (valuable) job information from their contacts.

Different fields have different approaches obviously. Most sociologists would explain (most) networks as an *unintended* outcome of other kinds of activities that individuals engage in. Individuals grow up in certain neighborhoods, they attend certain schools, they take jobs at certain workplaces, etc. and as a by-product of this, they get friends and acquaintances that become nodes in their networks. Their choice of friends and acquaintances is based rarely on their instrumental usefulness. On the contrary, economists would explain networks as an *intended* outcome stemming from strategic interactions. As Coleman (1988) puts it:

There are two broad intellectual streams in the description and explanations of social action. One, characteristic of the work of most sociologists, sees the actor as socialized and action as governed by social norms, rules, and obligations. The principal virtues of this intellectual stream lie in its ability to describe action in social context and to explain the way action is shaped, constrained, and redirected by the social context. The other intellectual stream, characteristic of the work of most economists, sees the actor as having goals independently arrived at, as acting independently, as a whole self-interested. Its principal virtue lies in having a principle of action, that of maximizing utility. This principle of action, together with a single empirical generalization (declining marginal utility) has generated the extensive growth of neoclassical economic theory, as well as the growth of political philosophy of several varieties: utilitarianism, contractarianism, and natural rights.

The aim of this chapter is to survey the literature on social networks putting together the economics, sociological, and physics/applied mathematics approaches, showing

their similarities and differences. In the next section, we will first present some measures of characteristics of social networks, mainly introduced by sociologists, which will be useful later on. Then we will expose and develop the two main ways of modeling network formation. While the physics/applied mathematics approach is capable of reproducing most observed networks, it does not explain why they emerge. On the contrary, the economics approach is very precise in explaining why networks emerge but does a poor job in matching real-world networks. The fourth section will analyze behaviors on networks, which take as given networks and study the impact of their structure on individuals' outcomes. We will use the game-theoretical framework, characteristic of economics, and will compare the results obtained to those in sociology. The last section will then conclude.

SOCIAL NETWORK ANALYSIS

Many techniques and concepts have been developed continuously over the years during the last century and there is right now a powerful machinery available under the corpus of social network analysis. It is not the aim of this survey to treat with full generality all relevant contributions in this field.¹ In particular, in this section, we are going to focus on a small subset of concepts, a set of *centrality measures*, that have been introduced in the literature to capture in a numerical form the prominence of actors inside a network.

Centrality measures aim at ranking individuals in terms of their relevance due to their position in the network. Social network analysis also introduces other characteristics of social structure such as the concept of *structural equivalence* (Lorrain and White, 1971) and *block modeling* techniques (White et al., 1976). The concept of structural equivalence tries to uncover similar roles and social positions shared by different actors in a network. The hypothesis behind structural

equivalence is that actors who share a similar position inside a network are going to end up with similar outcomes. On the other hand, block modeling techniques aim at disentangling different roles in a network when considering social network data.

For convenience, in this section, we are going to interpret links as communication lines. This unifies the treatment and interpretation but has no technical implication on the concepts defined.

There are several ways of constructing numerical statistics that can give a measure of the relevance of an individual embedded in a complex web of social relationships. The simplest and a most natural way is simply to count the number of connections an agent has. This measure is called *degree centrality*. Under degree centrality, agents who have a higher degree enjoy a better position inside the network. This measure is associated to the idea that “the more connections one has, the better it is.” For example, more connected individuals can have access to more information sources that can translate into better socioeconomic outcomes.

While degree distribution seems an appealing concept, it is only considering direct benefits derived from connections, abstracting from the potential contribution of indirect benefits derived from indirect connections. In that respect, degree centrality can be quite misleading in some situations.

Granovetter’s (1973, 1983) seminal work extends this concept by distinguishing between strong and weak ties.² This taxonomy of social interactions determines different functions of different kinds of connections. Granovetter’s thesis is that the strength of a tie among two different actors is proportional to the level of overlapping of their local social capital. Indeed, in a close network, everyone knows each other, information is shared and so potential sources of information are quickly shaken down, the network quickly becomes redundant in terms of access to new information. In contrast, Granovetter stresses the strength of weak ties involving a secondary ring of acquaintances

that have contacts with networks outside ego’s network and therefore offer new sources of information on job opportunities. As Granovetter (1973) claims “whatever is to be diffused can reach a larger number of people, and traverse greater social distance (...), when passed through weak ties rather than strong ties” (p. 1366). In substance, weak ties foster widespread diffusion, while strong ties breed local sharing.

If as Granovetter suggests, increasing the number of informational sources can have some impact on outcomes, then direct links can be of different quality depending on whether they are providing access to new secondary sources of information or not. In the same way aggregating all kind of connections irrespective of their strength, as degree centrality does, can be quite misleading as a measure of influence and power inside the network. To understand this, let us look at *clustering* coefficients. An agent has a high level of clustering if most of the friends of his/her friends are also his/her friends. If this is the case, the local ego-centered network of an individual is weaker than that of another agent with the same number of connections but with a smaller clustering. This effect is, precisely, the redundancy of information associated to strong ties stressed by Granovetter.

It is in fact possible to incorporate indirect connections in a centrality measure in several ways. When considering direct and indirect benefits, it is natural to assume that the benefits of any connection dilute with the distance between agents involved in it. It is indeed not the same to receive information immediately from a direct connection (i.e., a close friend) and indirectly from individuals located several links away. The quality of information decays with distance. As a result, it seems quite natural to construct the following centrality measure. Give to any individual a particular numerical value for each of his/her direct connections. Then, give a smaller value to any connection at distance two and an even smaller value to any connection at distance three, etc.³ When adding up all these values, we end up with a new numerical

value that is now capturing both direct and indirect connections of any order. This is the idea behind Katz's (1953) and Bonacich's (1987) network centrality measure.⁴ This is not at all an ad hoc construction. It arises naturally in the dynamic settings of social influence (Friedkin, 1991) or, as we will see below, as the equilibrium outcome of games with network complementarities (Ballester et al., 2006).

The degree and Katz–Bonacich centrality measures capture the power and influence of individuals as recipients of information. Although this can be the natural way to introduce power in a number of situations, there are other settings in which the power of an individual is induced by his/her position as an intermediate in the communication process. When an agent is relevant for most communication processes inside the network, he/she can exert an important role since he/she can deter or even prevent information transmission. Betweenness centrality, introduced by Freeman (1977, 1979), measures precisely this source of influence inside the network. Essentially, betweenness centrality calculates the relative number of indirect connections (or paths) in which the actor in consideration is involved in with respect to the total number of paths in the network. This strongly relates to the concept of *structural holes* due to Burt (1992). Holes in social structure emerge in situations in which there is a lack of connections between different subgroups. Burt claims that agents who bridge these structural holes can extract a disproportionate benefit compared to other individuals who share a similar social position but who are not linking holes. Betweenness centrality provides a mathematical way to characterize agents that are bridging structural holes. To illustrate the concepts presented in this section, consider the network shown in Figure 16.1.

In this network, the agent represented by a white circle is less relevant when considering degree centrality. He/she has a degree equal to two, while the rest of agents have a degree at least equal to 3. Instead, we can ensure that this same agent is the more central one

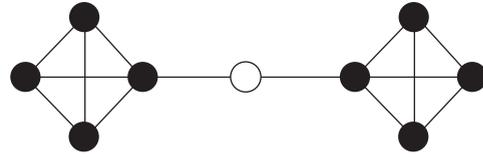


Figure 16.1 Bridge network and centrality measures

when considering betweenness centrality. He/she is in a bottleneck position inside the network. All communication from one side to the other side of the network necessarily has to pass through him/her, and, if controlling information is important, he/she can extract large positive benefits. He/she is bridging a structural hole. If we now consider the Bonacich centrality measure, this individual can either be the less or the more central agent inside the network depending on the value of the discount factor. For small discount factors (i.e., indirect links give less benefits), the agent is the less central one while for high levels of discount (i.e., direct links are weighted less), this agent is the most central.

Hence, the three centrality measures rank agents differently. The adequacy of one measure over the other very much depends on the context and interpretation given to the network.

NETWORK FORMATION

One of the main goals of the analysis of social networks is to shed some light on the mechanisms explaining how and why networks form. If social networks are relevant, we need to understand how networks emerge and the forces determining their shape.

One possible reason why a link is formed is *pure chance*. Two individuals randomly meet and create a link between them, which can represent friendship or a stable working relationship. A set of different models have arisen based on this assumption. They are called *random models of network formation*.

Another possible reason for the formation of a link is strategic interactions. Individuals carefully decide with whom to interact and this decision entails some consent by both parts in a given relationship. *Strategic network formation models* are, precisely, grounded on this premise. We now discuss with some detail both domains of research.

Random models of network formation

The simplest useful model of a random network (and one of the oldest) is the *Bernoulli random graph*, often just called the *random graph* for short (Bollobás, 2001; Erdős and Rényi, 1959, 1960; Solomonoff and Rapoport, 1951). In this model, a certain number of vertices (or nodes) are taken and edges (or links) are created between them with independent probability p for each vertex pair. When p is small, there are only a few edges in the network and most vertices exist in isolation or in small groups of connected vertices. Conversely, for large p , almost every possible edge is present between the possible vertex pairs and all or almost all of the vertices join together in a single large connected group called a *giant component*.

Understanding when a network is going to be fully connected is important, for example, to characterize the possible dynamics of infection rates in a population. There are, in fact, several common characteristics shared by most social networks. While the random graph model illuminates our understanding about when a giant component emerges, it cannot correctly mimic other critical aspects of real-world networks. The recent physics/applied mathematics literature have proposed different models that solve this weaknesses of the Erdős–Rényi framework.

One of the main characteristics most networks have is that their *degree distribution* (i.e., the number of links each node has) is scale-free,⁵ which in a mathematical form means that it follows a *power law*.⁶ A power law implies that the probability that an agent

is intensively connected to other agents in the society is non-negligible and strictly above zero. In particular, if a network is scale-free, we should expect that a small number of agents in the network show a high number of connections compared to the majority of agents who are involved in a small number of connections.

We can find empirical evidence of this pattern when analyzing, for example, the pattern of adult human sexual contacts. Liljeros et al. (2001) studied the 1996 Swedish survey of sexual behaviour and provided neat evidence that the distribution of sexual partners follows a power law distribution, hence, validating the scale-free hypothesis. Figure 16.2 taken from Liljeros et al. (2001, Figure 2), plots the cumulative distribution of sexual connectivity.

More precisely, Figure 16.2a shows that, for male and female respondents, the cumulative distribution of the number of sexual partners over a period of 12 months before the survey while Figure 16.2b displays the same statistics for a much longer period, namely for the respondent's life up to the time of the survey. It is easy to see that, in both figures, the data closely follow a straight line in a double-logarithmic plot, which is consistent with power-law dependence and the fact that the degree distribution has fat tails. This means, in particular, that most of the people in the sample had very few sexual partners while a small group had a very high number of sexual partners over their lifetime (up to 600 for the most “active” ones!). Other networks such as the internet, the movie actor collaboration, or linkage due to scientific co-authorship (Albert and Barabási, 2002; Goyal et al., 2006) also display the same properties of scale-free networks and power laws.

It is interesting to compare the results of Liljeros et al. (2001) to those of Bearman et al. (2004). The latter is a thorough analysis of the structure of an adolescent romantic and sexual network in a population of over 800 adolescents residing in a mid-sized town in the midwestern United States.⁷

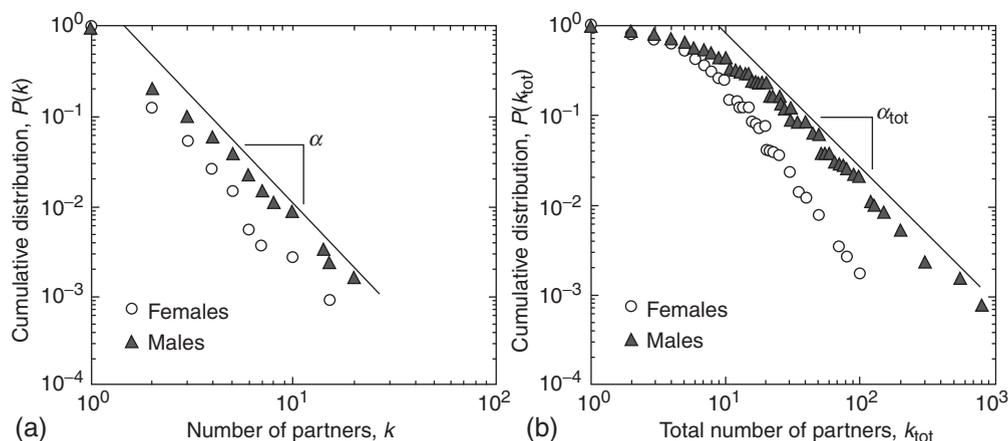


Figure 16.2 Scale-free distribution of the number of sexual partners for females and males. **a**, Distribution of number of partners, k , in the previous 12 months. Note the larger average number of partners for male respondents: this difference may be due to 'measurement bias' – social expectations may lead males to inflate their reported number of sexual partners. Note that the distributions are both linear, indicating scale-free power-law behaviour. Moreover, the two curves are roughly parallel, indicating similar scaling exponents. For females, $\alpha = 2.54 \pm 0.2$ in the range $k > 4$, and for males, $\alpha = 2.31 \pm 0.2$ in the range $k > 5$. **b**, Distribution of the total number of partners k_{tot} over respondents' entire lifetimes. For females, $\alpha_{tot} = 2.1 \pm 0.3$ in the range $k_{tot} > 20$, and for males, $\alpha_{tot} = 1.6 \pm 0.3$ in the range $20 < k_{tot} < 400$. Estimates for females and males agree within statistical uncertainty. Network of sexual partners in Sweden.

The authors analyzed reports of relationships of this adolescent network that occurred over a period of 18 months between 1993 and 1995. Interestingly, they found that the geometry of this network of adolescent sexual contacts was quite different from that of Liljeros et al. Most of its young members were included in a single giant component, which means that either directly or indirectly they were all connected to the majority of the students in the high-school. This, in particular, implies that the structure of this network looks very much like what the Erdős–Rényi model predicts: the degree distribution of each adolescent follows a gaussian (normal) distribution rather than a power law. These striking differences lead to distinct policy implications when trying to combat, for example, the spread of diseases related to sexual contacts. For adults, as suggested by Liljeros et al. (2001), it might be more productive to focus all efforts into targeting individuals in the upper tail of the degree distribution, that is,

the small number of people that have a high number of sexual contacts. This targeting would ensure that those nodes from which spread and contagion can be more harmful are controlled for or at least stimulated to take preventive measures. However, because of the results of Bearman et al. (2004), for adolescents, given that they are all essentially on equal footing within the network, a policy based on widespread education on healthy habits may be much more efficient.

The Erdős–Rényi model cannot rationalize scale-free degree distributions. In particular, it predicts that the probability that an individual has a large number of connections goes to zero as this number of connections increases. Barabási and Albert (1999) propose a model which is able to reproduce the scale-free nature of degree distributions and explain the fat-tailed degree distributions observed in some networks (like the sexual contact network described above). In the Barabási–Albert model, the population is

increasing over time and agents show *preferential attachment*, meaning that new agents in the network are more likely to connect to agents that are already well connected. The authors show that the preferential attachment mechanism naturally induces the emergence of a power-law degree distribution. Although, other alternative mechanisms have been considered in both the theoretical and empirical literatures, preferential attachment has been frequently invoked as a probable source of scale-free properties in self organization of social, and other types of, networks.

Unfortunately, while the preferential-attachment model does a good job in trying to fit the empirical evidence on the scale-free nature of connectivity, it does not satisfy another widespread characteristic of many real-world networks: the *small-world* property.⁸ A possible interpretation of this property is that agents in a network tend to show similar levels of connectivity and a high level of *clustering* (i.e., if John is friends with Patricia and if Patricia is friends with Patrick, then a high level of clustering means that John and Patrick are, most probably, also friends), and furthermore that the diameter of the network is quite small.

Watts and Strogatz (1998) provide a model that incorporates elements of both social structure and randomness to obtain networks that have the small-world properties. In this model, the “social structure” is represented by a uniform one-dimensional lattice (the circle), where each node is connected to

its k nearest neighbors on the lattice, and “randomness” is characterized by a tunable parameter p , which specified the probability that a link in the lattice would be *randomly rewired* (see Figure 16.3 below). To be more precise, start with a strongly *regular* network in which nodes are located in a circular network and are also connected to the neighbors of their neighbors (in the panel at the extreme left of the figure, each node is connected to its two direct neighbors and to its two neighbors of length 2). Then, take randomly some of the existing links, just a small number of them, and rewire them across the network. This rewiring process leads to a new geometrical arrangement where the clustering level is still high, due to the initial topology of connections, and where the value of the diameter is reducing due to the rewiring process which randomly connects distant agents in the initial network. Figure 16.3, taken from Watts and Strogatz’s work, provides a graphical intuition of the consequences of the rewiring mechanism.

The process starts with the network at the left of the picture. If no kind of rewiring and randomness are introduced, then the same regular network is sustained. When the level of rewiring/randomness starts to increase, we evolve to a situation in which only a small number of links are rearranged, leading to a characteristic small-world pattern, represented at the center of the picture. If there is too much rewiring/randomness (p close to 1), then, as can be seen from the right panel of

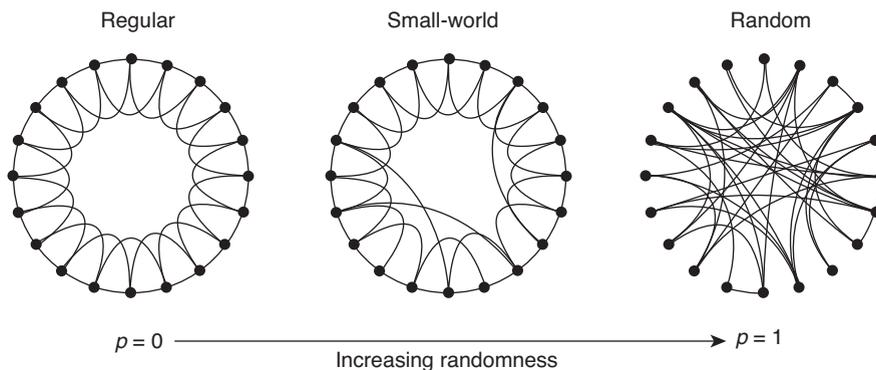


Figure 16.3 Networks with and without small world properties

the figure, the process turns back to a purely random process à la Erdős–Rényi. However, because of the mechanism that governs this network formation process, Watts and Strogatz (1998) are not able to obtain the scale-free degree distribution property.

A recent paper by Jackson and Rogers (2007) overcomes this last problem by providing a hybrid model that has most of the properties shared by real-world networks, that is, low diameter, high clustering, small-world and scale-free properties. In this paper, Jackson and Rogers (2007) combine the pure random model of Erdős–Rényi with the preferential attachment model of Barabási and Albert (1999). More precisely, they develop a dynamic model of network formation where nodes find other nodes with which they form links with in two ways: some are found uniformly at random (as in the Erdős–Rényi model) while others are found by searching locally through the current structure of the network, for example, meeting friends of friends (i.e., preferential attachment).

The authors then use their model to calibrate it with data from widely varied applications, in particular, the network of romantic relationships among high-school students exposed above (Bearman et al., 2004). They also find that the high-school romance networks are almost uniformly random. Interestingly, clustering is absent since the high-school romance network is predominantly a heterosexual (bipartite) network.

Strategic network formation

We would like to expose now the economic approach of network formation where links are not formed at random or in a rather mechanical way (as in the preferential attachment model), but by a process for which agents form links that maximize their own well-being. Indeed, the main advantages of random graph models are that: (i) they generate large networks with well identified properties; (ii) they mimic real networks (at least in some characteristics); and (iii) they

tie a specific property to a specific process. However, as pointed out by Jackson (2007), they are unable to answer the *why* behind network formation. For example, in Watts and Strogatz (1998), the “social structure” is represented by a uniform one-dimensional lattice. This is an assumption and we do not know why and under which conditions this network structure prevails. Similarly, Barabási and Albert (1999) never justify why agents behave according to the preferential attachment rule. On the contrary, by focusing on the optimal behavior of agents in making links, we can understand why certain network structures emerge. The economics approach also has its limits since it is unable to derive “equilibrium” networks that have the properties of most real-world networks and, in particular, cannot tell which degree distribution should emerge.

The strategic component in network formation models relies on the fact that the utility function⁹ of an agent depends on the activity undertaken by the rest of the agents with whom he/she is linked to. Hence, agents have to choose with whom they would like to form a link. At the same time other agents are considering the same kind of decision. Altogether, this specifies a game.

Game theory provides a set of rigorous and powerful tools to analyze such a situation (see, for example, Myerson, 1991). Strategic network formation analysis borrows from this literature and has also contributed to its recent enlargement. The Nash equilibrium is the key concept in game theory.¹⁰

Myerson (1977, 1991) provides an early formulation of a network formation game. The structure of the game is simple: agents have to decide about their potential partners and their strategies consist in naming those with whom they want to form a link with. For a link to be formed, it has to be that two individuals name each other, that is, there needs to be *mutual consent* in link creation. The Nash equilibrium concept can then be used to find out which strategy profiles are stable and, hence, which networks are the possible outcomes of the game.

The main problem of using the Nash equilibrium concept is that it exacerbates the coordination problems that arise when all agents are simultaneously deciding about friendship ties. In particular, the empty network (i.e., nobody forms a link) is always a Nash equilibrium of this game since no deviation is profitable. Indeed, deviating means here to name someone as a friend but this will not generate a link because of mutual consent. In general, with the Myerson game, there are many equilibria, which reduces the attractiveness of this approach.

A first solution to this problem has been proposed by Jackson and Wolinsky (1996). In their seminal paper, they introduce an alternative solution concept for network formation games, namely *pairwise stability*. To be pairwise stable, a network has to satisfy two conditions. First, no agent has incentives to sever any of the existing links, and, second, no pair of agents has incentives to create a non-existing link among them. Pairwise stability is a relatively weak solution concept that captures the essence of mutual consent. It still has limited prediction power, but introduces in network formation models the tendency to form mutually beneficial links and severing links that are counterproductive for at least one of the two parts of the relationship.

A second alternative to the Myerson game is to consider *directed* instead of *undirected* networks,¹¹ which solves the coordination problem of mutual consent. Bala and Goyal (2000) extend the Myerson model described above by only considering directed networks, that is, individuals form links *unilaterally* without requiring the consent of the other party to create a link.

Jackson and Wolinsky (1996) also introduce the so-called *connections model*, which uses a simple linear utility function aiming at capturing the usefulness of *direct* and *indirect* connections to gain access to a number of different resources (such as, for example, information on job openings) and the costs of establishing these connections. In particular, any connection, direct or indirect, is valuable, although more distant connections contribute less to an agent's well-being. Only direct links are, however, costly. Agents have to consider the costs and benefits of creating new links in a decentralized manner and, as a result, the set of possible outcomes can be characterized. Several network typologies can be pairwise stable and, quite naturally, the set of possible outcomes depends on the costs and benefits of forming a link.

One of the main virtues of using pairwise stability as a solution concept for network formation games is that it is relatively easy to check when a network is pairwise stable or not. For example, consider the star network in Figure 16.4, with a center (agent 1) connected to three peripheral agents (agents 2, 3, and 4).

In the connections model, connectivity rewards are computed through a discount factor $0 < \delta < 1$. A direct connection provides a reward equals to δ ; a friend of a friend gives δ^2 ; a link with an indirect friend of length k gives δ^k ; and so forth. Only the cost of forming direct friendships is considered and is denoted by $c > 0$. For the sake of exposition, let us consider the case in which $\delta = 0.5$ and $c = 0.6$. We can then see that the star network depicted above is not pairwise

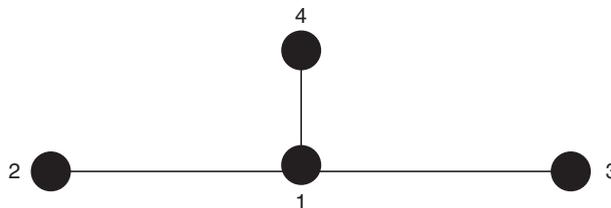


Figure 16.4 A star network with four agents

stable because the central agent (agent 1) would improve his/her utility by cutting any of his/her existing links with agents 2, 3, and 4. Indeed, the reward of any of such connections is 0.5 while the cost is 0.6. This generates a disutility of -0.1 . Therefore, for this particular parametric setup, given that the center has incentives to sever some of his/her existing links in the network, this star network is not pairwise stable.

Besides this pure positive analysis, a utilitarian perspective is introduced to conduct a *normative* analysis, trying to evaluate which networks are socially optimal. The tractability of the “connection model” offers a first reasonable benchmark to gain insight on this issue. Jackson and Wolinsky (1996) provide a characterization of the set of all efficient networks.¹² It is shown that the set of efficient networks usually does not equal, or only partially include, the set of pairwise stable networks.¹³

For example, coming back to the star network described above, it can be proved that this network is socially efficient when $\delta = 0.5$ and $c = 0.6$, while we have seen that it was not pairwise stable. Indeed, the central agent 1 is paying for three different connections, while the peripheral agents (2, 3, and 4) only pay for one, but they all benefit from indirect connections due to their direct link to the center (peripheral agents are all at a distance 2 from each other). It is the positive spreading of externalities from indirect neighbors that makes the star network efficient. However, this network is not pairwise stable because the cost the central agent is paying is too large compared to the benefit of maintaining links with peripheral agents. In that case, decentralized behavior cannot induce a socially optimal geometry of connections.

This analysis reveals that, in general, individual and social incentives are not compatible, and that decentralized behavior usually results in an inefficient arrangement of social ties.

This last result is particularly remarkable. This tension between individual incentives and social welfare leads quite naturally to

the following question: is there any way this tension can be weakened? In other words, can we find welfare-enhancing policies that ensure that the resulting network of relationships is optimal, or closer to an optimum, for the society as a whole? For example, it is clear that identity issues shape friendship formation during childhood and adolescence (see, for example, Akerlof and Kranton, 2000; De Martí and Zenou, 2009). The assignment of students to classrooms can be then considered as a potential policy tool that impacts on the outcome of friendship formation mechanisms in school and schooling outcomes. Strategic network formation models can thus shed some light on these kinds of debates, in particular, in situations where individual decisions are driven by both private interests and social forces that can be incorporated into a utility function.

BEHAVIOR ON NETWORKS

In this section, we would like to take the network as given and study how its structure impacts on outcomes and individual decisions.

The theory of “games on networks” aims at determining the outcome of a game with n players or agents (that can be individuals, firms, regions, countries, etc.) who are all linked (directly or indirectly) together in a network and are choosing an action (that could be search effort, job market decisions, engagement in criminal activities, investment in R&D, etc.) that maximizes their utility. What is interesting in these games is that agents are utility maximizers that take into account the interdependencies generated by the social network structure. Social capital is incorporated in the utility function of each actor.

Most of this rather small and new literature has focused on *static* models (only what happens today prevails so that there is no time dimension) with *perfect information*, that is, all agents perfectly know the structure

of the network and the actions taken, not only by their direct friends, but by all agents in the network. Even so, both premises can be relaxed to accommodate more realistic assumptions. For example, Galeotti et al. (2009) develop a general model of games played on networks, in which players have *private* and *incomplete information* about the network.¹⁴ In this paper, agents do not know the whole network but are informed only of their own degree (i.e., number of links). Moreover, they assume that agents have beliefs about the rest of the network that depend on their own degree, which are summarized by a probability distribution over the degrees of their neighbors.

A Networked Model of Peer Effects

Consider an individual connected by a network of peer influences. In this network each agent reaps complementarities¹⁵ from all his/her direct network peers. Ballester et al. (2006) computed the Nash equilibrium of this peer-effect game when agents choose their peer effort simultaneously. In their setup, restricted to linear-quadratic utility functions, they established that the peer-effect game described above has a unique Nash equilibrium¹⁶ where each agent strategy is proportional to his/her Katz–Bonacich centrality measure, already introduced above. This is the main theoretical result of Ballester et al. (2006), a good example of how the economic and sociological literature on networks can fruitfully borrow from each other. At equilibrium, individual decisions emanate from all the existing network chains of direct and indirect contacts stemming from each node, which is a feature characteristic of Katz–Bonacich centrality. Consider for example the network in Figure 16.5.

This network results from the overlap between two different dyads (dyad of agents 1 and 2 and dyad of agents 1 and 3) with a common partner, agent 1. Agent 2 reaps *direct* complementarities from agent 1 in one dyad whom, in turn, reaps *direct* complementarities from agents 2 and 3 in both dyads. Thus, through the interaction with the central agent, peripheral agents end up reaping complementarities *indirectly* from each other. For this reason, the equilibrium decisions in each dyad cannot be analyzed independently of each other. Rather, each dyad exerts a strategic externality on the other one, and the equilibrium effort level of each agent reflects this externality.

What kinds of predictions can we learn from this model? Let us think of crime. It is well established that crime is, to some extent, a group phenomenon, and the source of crime and delinquency is located in the intimate social networks of individuals (see, for example, Sutherland, 1947; Sarnecki, 2001; Warr, 2002). Indeed, criminals often have friends who have themselves committed several offences, and social ties among criminals are seen as a means whereby individuals exert an influence over one another to commit crimes. In fact, not only friends but also the structure of social networks matters in explaining individuals' own criminal behavior.¹⁷

The model presented above predicts that a key determinant of one's criminal activities is the position in the network, that is, the more central (in terms of the Katz–Bonacich centrality index) a person is in a network, the higher his/her level of criminal activity is. In other words, it gives a causal relationship between network's location and criminal activities.^{18,19}

But, in fact, the model can predict even more. It can propose new policies aimed

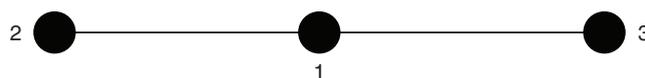


Figure 16.5 A start network with three agents

at reducing crime. Indeed, in the standard economic literature on crime, where criminals decide to commit crime based on a cost–benefit analysis (see, for instance, the literature surveys by Garoupa, 1997, and Polinsky and Shavell, 2000), the standard policy tool to reduce aggregate crime relies on the deterrence effects of punishment, that is, the planner should increase uniformly punishment costs.

The model above, though, associates a distribution of delinquency efforts to the network connecting them. In particular, the variance of delinquency efforts reflects the variance of network centralities. In this case, a targeted policy that discriminates among delinquents depending on their relative network location, and removes a few suitably selected targets from this network, alters the whole distribution of delinquency efforts, not just shifting it. In practice, the planner may want to identify optimal network targets to concentrate (scarce) investigatory resources on some particular individuals, or to isolate them from the rest of the group, either through leniency programs, social assistance programs, or incarceration.

This type of discriminating policy has been studied in detail by Ballester et al. (2006, 2010). To characterize the network optimal targets, they propose a new measure of network centrality, the so-called *optimal inter-centrality measure*, which does not exist in the social network literature. This measure solves the planner’s problem that consists in finding and getting rid of the key player, that is, the criminal who, once removed, leads to

the highest aggregate crime reduction. They show that the key player is, precisely, the individual with the highest optimal inter-centrality in the network.

The ranking of criminals according to their individual optimal inter-centrality measures, relevant for the selection of the optimal network target, needs not always to coincide with the ranking induced by individual equilibrium-Katz–Bonacich centralities. In other words, *the key player is not necessarily the most active criminal*. Indeed, removing a criminal from a network has both a direct and an indirect effect. First, fewer criminals contribute to the aggregate crime level. This is the direct effect. Second, the network topology is modified, and the remaining criminals adopt different crime efforts. This is the indirect effect. The key player is the one with the highest overall effect.

To illustrate this result, consider the network of 11 criminals in Figure 16.6. We can distinguish three different types of equivalent criminals in this network, which are shown in Table 16.1.

We can distinguish three different types of equivalent criminals in this network, which are the following:

We can identify the key player in this network of criminals.²⁰ If the choice of the key player were solely governed by the *direct* effect of criminal removal on aggregate crime, type-2 criminals would be the natural candidates. Indeed, these are the ones with the highest number of direct connections. But the choice of the key player needs also to take into account the *indirect* effect on

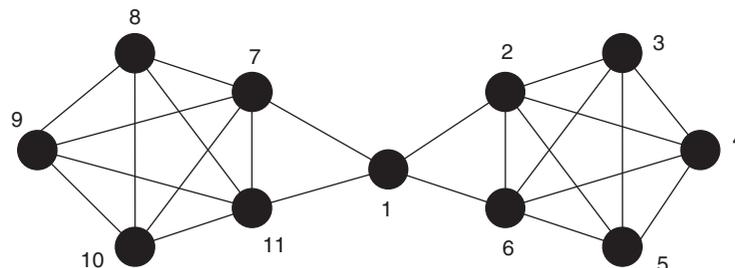


Figure 16.6 A bridge network with eleven criminals

Table 16.1 Types of criminals in the bridge network

Type	Criminals
1	1
2	2, 6, 7, 11
3	3, 4, 5, 8, 9, 10

aggregate crime reduction induced by the network restructuring that follows the removal of one criminal from the original network. Because of his/her communities' bridging role, criminal 1 is also a possible candidate for the preferred policy target.

It can be showed that type-2 criminals always display the highest Katz–Bonacich centrality measure.²¹ These criminals have the highest number of direct connections. Besides, they are directly connected to the bridge criminal 1, which gives them access to a very wide and diversified span of indirect connections. Altogether, they are the most central criminals and thus those who commit most crimes. However, the most active criminals are not the key players. Indeed, because indirect effects matter a lot, eliminating criminal 1 has the highest joint direct and indirect effect on aggregate crime reduction. Criminals counter the higher deterrence they face by spreading their know-how further away in the network and establishing synergies with criminals located in distant parts of the social setting. In that case, the optimal targeted policy is the one that maximally disrupts the crime network, thus harming the most its know-how transferring ability.²²

Public Good Provision and Innovation Processes

Efforts can also be strategic substitutes²³ when we focus on a network model of public goods.²⁴ Examples of public goods abound. When a person plants a garden, his/her neighbors benefit. When a jurisdiction institutes a pollution abatement program, the benefits also accrue to nearby communities.

When individuals innovate – for example, experiment with new technology or generate new information – the results are often non-excludable in certain dimensions and thus others may benefit from the innovation. For example, in agriculture, one farmer's experience with a new crop can benefit other farmers, and the physical and social geography of the countryside can influence experimentation and learning. In industry, it has been long posited that research findings spill over to other firms.

The questions that these types of models seek to answer are the following: How does the social structure affect the level and pattern of public good provision? Do people exert effort themselves or rely on others? These are similar questions to the previous model in fourth section but the answers will be different because of the strategic substitutability assumption.

Bramoullé and Kranton (2007) develop a model along these lines. People who are linked directly or indirectly by a network of relationships desire a public good which is costly to produce. Individuals decide how much effort they want to exert to contribute to this good. The authors characterize the Nash equilibria of this game. Their analysis yields three main insights. First, networks can lead to *specialization*. Indeed, in any network there is a Nash equilibrium where some individuals contribute to the public good and others completely free ride (this was not possible in the case of strategic complementarities). In many networks, this extreme situation is the only equilibrium outcome. In all networks, such patterns are the only stable outcomes. Hence, agents' positions in a network can determine whether or not they contribute to the public good. Second, specialization can have welfare benefits. This outcome arises when contributors are linked, collectively, to many people in society. Finally, new links can reduce overall welfare. A new link increases access to the new information/public good, but also reduces individual incentives to contribute. Hence, overall welfare can be higher when there are *holes* in a network.

To illustrate some of these results, let us take an example. Consider the two *complete* networks with four individuals shown in Figure 16.7 (where the numbers are the effort levels provided by each individual).

Bramoullé and Kranton (2007) show that for complete networks, in any Nash equilibrium, the total aggregate effort is 1, and it can be split in any way among the agents. For example, as in the left panel, efforts could be equally distributed, so that each agent exerts $1/4$, or, as in the right panel, one agent could be a specialist and provide an effort of 1 whereas all other individuals free ride by providing no effort. Consider then on the contrary star networks as shown in Figure 16.8.

In that case, only specialized profiles are equilibria. There are just two Nash equilibria: either the center is a specialist, or the three agents at the periphery are specialists. What do we learn from these simple examples? We learn that the structure of the network is the main determinants of the equilibria; denser networks can lead to less overall experimentation, and effort sharing is not always possible.

There is a growing field in economics studying innovation and diffusion of information, and there are many studies that suggest that social structures affect experimentation and spread of information. The present analysis suggests that individuals who have active social neighbors should have high benefits but exert little effort. Also, individuals who have prominent social positions are expected to bear less of the effort costs and instead rely more on others' efforts. This model indicates the importance of strategic and network efforts. For instance, in their study of the adoption of pineapple for export in Ghana, Conley and Udry (2001,

2010) find that, for a given farmer, (i) he/she is more likely to change his/her fertilizer use after his/her information neighbors who use similar amounts of fertilizer achieve lower than expected profits; and (ii) he/she increases (decreases) his/her use of fertilizer after his/her information neighbors achieve unexpectedly high profits when using more (less) fertilizer than he/she did.

Collective Action, Coordination, and Social Structure

Another important issue of games on networks is *collective action*. A group of people face a collective action problem in that an individual wants to participate only if joined by enough other individuals. We can therefore define a threshold for each individual as the number of people necessary for him/her to participate. However, a group's behavior can strongly depend on its internal configuration. Hence, when considering collective action problems, it is of paramount importance to understand how the arrangement of intra-group social ties is going to shape the decisions of their members.

Granovetter (1978) and Schelling (1978) provide dynamic models to analyze such issues, avoiding a clear linkage of outcomes to the social network structure in society. In particular, a snowball effect might be generated only if there are initially enough activists who are willing to participate, independently of others decisions. Once these activists start the process, other agents will join as far as the number of people who have already joined is higher than their thresholds.

Chwe (1999, 2000) provides an alternative, static, formulation of the rising of such



Figure 16.7 Two complete networks in a public-good model

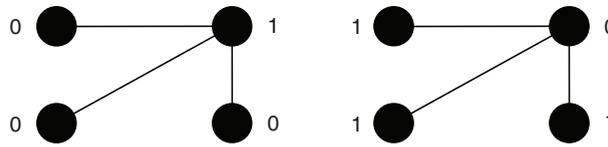


Figure 16.8 Two star networks in a public-good model

processes rooted at the societal network geometry. He studies the internal structure of a group and analyzes how network configuration can lead to massive participation or to a failure in coordinating individual activity. Arguably, it does not contradict Granovetter’s and Schelling’s models but enriches the understanding of how coordination might be necessary to join the first *snowflakes*, which will then generate a snowball effect in collective action.

More precisely, in Chwe’s model, agents in the network have to decide whether or not to participate in a common activity. They are only willing to participate in the activity if a minimum number of individuals also participate. There is a coordination problem since this information is not known in advance. Each person only knows the thresholds of his/her neighbors (i.e., direct links) in the social network (“local knowledge”). Individuals can communicate to their acquaintances the value of this minimum threshold on activists necessary for him/her to join the activity. The network, thus, shapes individual beliefs. To understand the mechanisms of the model, consider the two networks shown in Figure 16.9a and the situation in which any agent is willing to join the activity if and only if he/

she knows for sure that two more individuals are going to joint it too.

In the first network, the circle, all agents share a symmetric situation. Agent 1 knows his/her own threshold and the threshold of his/her neighbours in the network, that is, individuals 2 and 4. Since all of them know the threshold of two other members of the network, it is immediate that all four agents are going to join the activity. Indeed, the three of them know that each one needs two other individuals to decide to join. Since this adds up to three, this suffices to confirm their engagement. Common knowledge emerges due to the configuration of the network geometry.

In the second network, the kite, the first three agents share a similar position inside the network while agent 4 is isolated with only one connection to agent 2. As in the circle, agents 1, 2, and 3 are going to join the activity because their respective thresholds are common knowledge to all of them. However, agent 4 is not going to join them. He/she knows that agent 2 needs two other agents to engage into the activity, but he/she does not know much about agents 1 and 3. Since agent 4 can not exclude the fact that agents 1 and 3 need four or more agents to join, he/she is not sure that either 1 or 3 are going to engage in the activity. Thus, agent 4 decides not to participate.

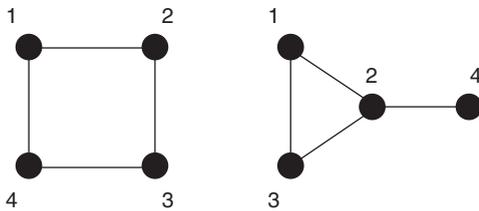


Figure 16.9 Collection action outcomes with different network structures

Observe that in both networks, the circle and the kite, there are four agents and four links. The different arrangement of these links, which differentiates both situations, are responsible for the fact that agent 4 has a different behavior in the two networks. Of course, more links are always going to increase the possibility that common knowledge emerges spontaneously.

Job market outcomes

So far, we have focused on static outcomes of agents who were connected by a network of relationships. Let us examine what happens in the long run when the network is assumed to be fixed and the same over time. We present the model by Calvó-Armengol and Jackson (2004) who discuss their framework in terms of the labor market, a natural application of dynamic processes since workers' employment status are not static but change over time.²⁵

The authors study the evolution over time of the employment statuses of n workers connected by a network of relationships. They are able to determine how the network of relationships influences these outcomes. To be more precise, the network starts with some agents being employed and others not. Next, information about job openings is obtained. In particular, any given individual hears about a job opening with some probability. If the individual is unemployed, then he/she will take the job. However, if the individual is already employed, then he/she will pass the information along to a *direct* unemployed friend. Furthermore, when an employed worker hears about a job but all his/her friends (i.e., direct links) are already employed, then the job is lost for the period. It is also assumed that all unemployed neighbors are treated on an equal footing, meaning that the employed worker who has the job information does not favor any of his/her direct neighbors.

A network thus summarizes the links of all agents, where a direct link indicates that two individuals know each other, and share their knowledge about job information.

It is shown that, in steady-state, there is a *positive correlation in employment status between two path-connected workers*. This

result is not trivial since, in the short run, the correlation is negative. Indeed, in a static model, if an employed worker is directly linked to two unemployed workers, then if he/she is aware of a job, he/she will share this job information with his/her two unemployed friends. These two individuals, who are path-connected (path of length two), are thus in competition with each other and only one of them (randomly chosen) will obtain the job and be employed while the other will remain unemployed. So their employment status will be negatively correlated (see Boorman, 1975; Calvó-Armengol, 2004; Calvó-Armengol and Zenou, 2005).

To get the main intuition on why this negative correlation result does not hold in a dynamic labor-market model, we focus on the case with three workers with a star-shaped network as shown in Figure 16.10.

In this figure, a black node represents an employed worker (individual 1), while unemployed workers (2 and 3) are represented by white nodes. In fact, even though in the short run, individuals 2 and 3 are “competitors” for the job information that is first heard by individual 1, this is not true in the long-run since individual 2 can benefit from individual 3's presence. Indeed, individual 3's presence helps improve individual 1's employment status. Moreover, when individual 3 is employed, individual 2 is more likely to hear about any job that individual 1 hears about. These aspects of the problem counter the local (conditional) negative correlation and help induce a positive correlation between the employment status of individuals 2 and 3.

To illustrate the fact that not only the network is of importance but also its structure, let us consider a network with four workers, that is, $n = 4$. Table 16.2 depicts the value of unemployment probabilities of worker 2,

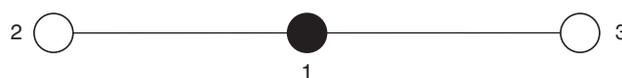


Figure 16.10 A star network with employed and unemployed workers

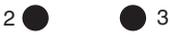
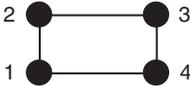
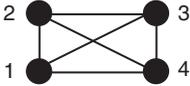
and the correlations between workers 1 and 2, and between workers 1 and 3, in the long-run steady state. These results are calculated using numerical simulations repeated for a sufficiently long period of time as shown in Table 16.2.

First, when there is no social network so that no information is exchanged between workers, the unemployment rate of each agent is just equal to its steady-state value. Thus, the probability of being unemployed for each worker is 13.2 percent, given that they cannot rely on other workers to get information about jobs and the only chance they can have of obtaining a job is by direct methods. Imagine now that only one link is added in this network so that workers 1 and 2 are directly linked to each other. Steady-state unemployment decreases substantially for workers 1 and 2, from 13.2 percent to 8.3 percent. When more links are added, the unemployment rate for each worker decreases even more from 13.2 percent when there are no links to 5 percent when the social network is complete. This table also shows the positive correlation between the employment statuses of different workers already mentioned before.

This model provides a rationale, for example, of why ethnic minorities, who tend to have friends who are of the same ethnicity, have difficulties in finding a job. Since employment statuses between direct and indirect friends of the same network are correlated, then, like a disease, unemployment will spread among all individuals belonging to this network. Indeed, if at the beginning some individuals do not obtain a job (say because they are discriminated against), then they cannot help their friends to obtain a job, who themselves cannot help their own friends to obtain a job, etc.

This model is consistent with what research in sociology has found. First, its implications are similar to those from the *homophily* literature, that is, that links between similar people typically occur at a much higher rate than among dissimilar people (see McPherson et al., 2001, who provide an extensive review of this literature in sociology).²⁶ Furthermore, the result above that employment statuses between path-connected friends (i.e., weak ties) are correlated is consistent with Grannovetter’s claim that weak ties are extremely important to provide information about jobs.²⁷

Table 16.2 Workers’ outcomes with different network structures

	Probability of being employed for individual 1	Correlations in employment statuses between 1 and 2	Correlations in employment statuses between 1 and 3
	0.132	–	–
	0.083	0.041	–
	0.063	0.025	0.019
	0.050	0.025	0.025

The model goes even further in that it can give interesting policy predictions. As we have seen, *ex ante* identical individuals connected through a network can end up with very different outcomes. This *ex post* heterogeneity is mainly due to differences in geometry across local ego-centred networks. As a result, public policy needs to be tailored to the explicit role of the network. In particular, the planner can alter the allocation of network externalities across individuals and implement optimal redistributive schemes by suitably manipulating the network locally. Interventions to improve and sustain the employment status of a given agent also improve the outlook for the social acquaintances of the agents targeted by the intervention – a contagion effect in reverse. In a networked society, education subsidies and other labor market regulation policies display local increasing returns for the social topology. Targeted interventions magnify the initial effect of the policy, and the more so, the more the subsidies and programs are circumscribed to tightly clustered agents in a network. For this reason, targeting is more efficient than spreading resources more broadly, and the more concentrated the interventions, the higher the efficiency gains. This is similar to the targeted policies on crime (i.e., the key player policy) highlighted earlier.

CONCLUDING REMARKS

As we have seen, there are varied ways of approaching social networks. Our aim here has not been to cover exhaustively all possible contributions to this literature. Instead, we have tried to single out some of the representative works in each of the fields interested in understanding the shape and role of social networks. By doing so, we could pinpoint the main virtues of each approach and analyze how they can benefit from each other.

Some of the research presented in this chapter already merge concepts and techniques from different fields. For example,

economists and physicists alike are increasingly being aware that sociologists have provided over the years a rich set of tools analyzing the structure of social networks, which can be adopted in their modelling choices. Economists have contributed to the literature by enriching the behavior of actors inside the network. Utilitarianism has proved effective in incorporating individual incentives and social preferences that are, arguably, at the core of social network formation when agents act in a decentralized manner. Finally, physicists and applied mathematicians have tackled issue of complexity of webs with a large number of nodes. Their mathematical tools provide a systematic way of understanding the pattern of giant components in networks and give a neat picture of some of the relevant characteristics of these networks for applied works in other disparate areas such as, for example, medicine or criminology.²⁸

We hope and expect that an increasing effort will be devoted in the coming years to build stronger connections between the different approaches. This can only enrich our view and understanding of social networks. It is only through this knowledge that we will be able to adopt more realistic policies and recommendations for practitioners.

NOTES

1 See Wasserman and Faust (1996) for an excellent and exhaustive survey of social network analysis.

2 Roughly speaking weak ties are indirect links such as acquaintances while strong ties are direct links such as close friends or family.

3 A simple way to construct these values is through a discount factor $0 < \delta < 1$. If a direct connection has value α , any indirect connection of order two then has a value equals to $\delta\alpha < \alpha$, an indirect connection of order three has a value of $\delta^2\alpha < \delta\alpha$, and so on.

4 These centrality measures are closely related to the so-called eigenvector centrality measures (see Bonacich and Lloyd, 2001) that lie at the core of ranking methods (see, for example, Palacios-Huerta and Volij, 2004), like Google's PageRank algorithm, or on recently developed segregation measures

based on social interactions (Echenique and Fryer, 2007).

5 Of course, this is not a property of universal compliance in all kind of networks. There is empirical evidence against it (see, for example, the work of Amaral et al., 2000, on three different examples of social networks) and also some literature that partially corroborates this rule, with some inconsistency in the tails of the distribution (see Newman, 2001, for an example of a network that satisfies a kind of perturbed power law distribution when considering scientific collaborations).

6 A *power law* is any polynomial relationship that exhibits the property of scale invariance.

7 Data are drawn from the National Longitudinal Study of Adolescent Health (Add Health).

8 A network that displays small-world properties has the characteristic that most nodes can be reached from every other by a small number of steps. More precisely, a network displays a small-world property if the mean *geodesic distance* between vertex pairs (the geodesic distance between two vertices is the length of any shortest path between them) is small compared with the size of the network as a whole (often measured by its diameter, which is the length of the largest geodesic distance between any pair of vertices in the network). In a famous experiment conducted in the 1960s, the psychologist Stanley Milgram (1967) asked participants (located in the United States) to get a message to a specified target person elsewhere in the country by passing it from one acquaintance to another, stepwise through the population. Milgram's remarkable finding that the typical message passed through just six people on its journey between (roughly) randomly chosen. The research was groundbreaking in that it suggested that human society is a small-world type network characterized by short path lengths. The experiments are often associated with the phrase "six degrees of separation," although Milgram did not use this term himself.

9 A *utility function* is a mathematical expression that assigns real numbers (utils) to all possible choices.

10 Indeed, in game theory, the *Nash equilibrium* (named after John F. Nash) is a solution concept of a game involving two or more players, in which no player has anything to gain by changing only his or her own strategy (i.e., by changing unilaterally). If each player has chosen a strategy and no player can benefit by changing his or her strategy while the other players keep theirs unchanged, then the current set of strategy choices constitute a Nash equilibrium.

11 A network is *undirected* if links are reciprocal, that is if i is linked to j , then it has to be that j is linked to i . On the contrary, a network is *directed* if links are not reciprocal. In the latter case, a link has two distinct ends: a head (the end with an arrow) and a tail.

Each end is counted separately. The sum of head endpoints count toward the *indegree* and the sum of tail endpoints count toward the *outdegree*. Friendships are usually modelled using undirected networks while, for the world wide web (i.e., a link is when someone clicks on a web page), directed networks are more appropriate.

12 Even if a utilitarian approach is considered valid, this is not necessarily the adequate one. We skip this debate here, and refer the interested reader to Moulin (2003).

13 See also Bala and Goyal (2000) for an analysis of the connections model in directed networks.

14 There are other studies of games with incomplete network knowledge in economics (see, in particular, Galeotti and Vega-Redondo, 2005; Sundararajan, 2006; Jackson and Yariv, 2007).

15 The efforts of two or more agents are called *strategic complements* if they mutually reinforce one another. For example, in a crime context, the effort decisions are strategic complements if an increase in the crime effort of one individual increases the marginal utility of the others, because, in that case, the others will have an incentive to provide more crime effort too.

16 Under the condition that network complementarities are smaller than the inverse of the largest eigenvalue of the graph (i.e., network).

17 See Glaeser et al. (1996) who model criminal activity through imitation in stylized social networks, and Calvó-Armengol and Zenou (2004) for an economic model of externalities in crime and social networks. For recent empirical studies on crime and social interactions, see, in particular, Patacchini and Zenou (2008, 2012), and Bayer et al. (2009).

18 The main obstacle to empirically identifying peer and social network effects is what Manski (1993) termed the *reflection problem*: absent further assumptions on the nature of peer influence, it is impossible to distinguish the causal impact of peers' behavior on an individual's behavior from the causal impact of peers' background characteristics on an individual's behavior. Manski refers to this type of causal influence as an *exogenous social effect*. It is also difficult to distinguish between peers' behavior and individual's behavior, independent of any other individual, peer, or contextual factors. This is referred to as an *endogenous social effect*. Consider, for example, a child's decision to initiate drug use. Is it because his/her next-door neighbor decided to initiate drug use (endogenous social effect)? Or is it due to the fact that some peer background characteristic, such as a substance-abusing parent, caused both children to adopt the same behavior (exogenous social effect)? The distinction between these explanations is important for policy purposes. When true contagion effects operate, intervening to alter one child's behavior may have the net effect of changing several children's behavior. In other cases,

such as when children initiate substance use because adults in their neighborhood provide opportunities to do so, these so-called *multiplier effects* do not exist (Brock and Durlauf, 2001). Correctly distinguishing endogenous from exogenous social effects is necessary for any effort to gauge the true net impact or social benefit of any behavioral intervention. Bramoullé et al. (2009) show how the simple existence of some asymmetry in a peer network can be a sufficient condition to ensure that endogenous and exogenous peer effects can be effectively identified. A nice literature survey on the empirics of social interactions can be found in Durlauf and Ioannides (2010).

19 Haynie (2001) tests the impact of the Bonacich centrality on delinquent behaviors using the AddHealth data (these data were already used by Bearman et al., 2004). Fixing arbitrarily the discount factor to 0.1 for all networks, she finds a positive and significant relationship. Using a more structural approach by estimating rather than fixing the discount factor (so that different networks have different discount factors), Calvó-Armengol et al. (2005, 2009) test the same relationship. They show that, after controlling for observable individual characteristics and unobservable network specific factors, a standard deviation increase in the Katz–Bonacich centrality increases the level of individual delinquency by 44 percent of one standard deviation.

20 Observe that, from a macro-structural perspective, type-1 and type-3 criminals are identical: they all have four direct links, while type-2 criminals have five direct links each. From a micro-structural perspective, though, criminal 1 plays a critical role by bridging together two closed-knit (fully intraconnected) communities of five criminals each. By removing criminal 1, the network is maximally disrupted as these two communities become totally disconnected, while by removing any of the type-2 criminals, the resulting network has the lowest aggregate number of network links.

21 See Ballester et al. (2006) for details.

22 A policy targeting a *group* of criminals rather than one criminal can also be implemented. See Ballester et al. (2010) for details.

23 Contrary to complementarities, effort decisions are substitutes if an increase in one individual's effort decreases the marginal utility of the others, giving them an incentive to provide less effort.

24 In economics, a *public good* is a good that is non-rival and non-excludable. This means, respectively, that consumption of the good by one individual does not reduce availability of the good for consumption by others; and that no one can be effectively excluded from using the good.

25 Sociologists and labor economists have produced a broad empirical literature on labor market networks. In fact, the pervasiveness of social networks and their relative effectiveness varies with the

social group considered. For instance, Holzer (1987, 1988) documents that among 16–23 year old workers who reported job acceptance, 66 percent use informal search channels (30 percent direct application without referral and 36 percent friends/relatives), while only 11 percent use state agencies and 10 percent newspapers. See also Corcoran et al. (1980) and Granovetter (1995). More recently, Topa (2001) argues that the observed spatial distribution of unemployment in Chicago is consistent with a model of local interactions and information spillovers, and may thus be generated by an agent's reliance in informal methods of job search such as networks of personal contacts. Similarly, Bayer et al. (2008) document that people who live close to each other, defined as being in the same census block, tend to work together, that is, in the same census block. For a survey of the literature on social interactions and the labor market, see Ioannides and Loury (2004), and for other theoretical models, see Montgomery (1991, 1992, 1994), Calvó-Armengol (2004), Calvó-Armengol and Zenou (2005), Calvó-Armengol et al. (2007), Calvó-Armengol and Jackson (2007).

26 See also the recent paper by Currarini et al. (2009).

27 Interestingly, Patacchini and Zenou (2008) also find that weak ties matter in transmitting information about crime in adolescent networks.

28 A recent paper by König et al. (2009) combine the applied-mathematic and economic tools to tackle the issue of dynamic network formation with strategic interactions. Indeed, they combine the network game introduced by Ballester et al. (2006) where the Nash equilibrium action of each agent is proportional to his/her Bonacich centrality (see fourth section) with an endogenous network formation process where links are formed and severed on the basis of agents' centrality, reflecting strategic interactions between agents. A remarkable feature of their dynamic network formation process is that, at each period of time, the network is a *nested split graph*, a graph that has very nice mathematical properties. They show that there exists a unique stationary network (which is a nested split graph) whose topological properties completely match features exhibited by real-world networks (i.e., small world properties, that is, high clustering and low diameter, and power law distribution).

REFERENCES

- Akerlof, G.A. and R. Kranton (2002) 'Identity and schooling: some lessons for the economics of education', *Journal of Economic Literature*, 40: 1167–1201.

- Albert, R. and A.-L. Barabási (2002) 'Statistical mechanics of complex networks', *Reviews of Modern Physics*, 74: 47–97.
- Amaral, L.A.N., A. Scala, M. Barthélémy and H.E. Stanley (2000) 'Classes of small world networks', *Proceedings of the National Academy of Sciences of the USA*, 97: 11149–11152.
- Bala, V. and S. Goyal (2000) 'A non-cooperative model of network formation', *Econometrica*, 68: 1181–1229.
- Ballester, C., A. Calvó-Armengol and Y. Zenou (2006) 'Who's who in networks. Wanted: the key player', *Econometrica*, 74: 1403–1417.
- Ballester, C., A. Calvó-Armengol and Y. Zenou (2010) 'Delinquent networks', *Journal of the European Economic Association*, 8: 34–61.
- Barabási, A.-L. and R. Albert (1999) 'Emergence of scaling in random networks', *Science*, 286: 509–512.
- Bayer, P., S.L. Ross and G. Topa (2008) 'Place of work and place of residence: Informal hiring networks and labor market outcomes', *Journal of Political Economy* 116: 1150–1196.
- Bayer, P., R. Hjalmarsen and D. Pozen (2009) 'Building criminal capital behind bars: Peer effects in juvenile corrections', *Quarterly Journal of Economics*, 124: 105–147.
- Bearman, P.S., J. Moody and K. Stovel (2004) 'Chains of affection: The structure of adolescent romantic and sexual networks', *American Journal of Sociology*, 110: 44–91.
- Bollobás, B. (2001) *Random Graphs*. 2nd edition. New York: Academic Press.
- Bonacich, P. (1987) 'Power and centrality: a family of measures', *American Journal of Sociology*, 92: 1170–1182.
- Bonacich, P. and P. Lloyd (2001) 'Eigenvector-like measures of centrality for asymmetric relations', *Social Networks*, 23: 191–201.
- Boorman, S. (1975) 'A combinatorial optimization model for transmission of job information through contact networks', *Bell Journal of Economics*, 6: 216–249.
- Bramoullé, Y. and R.E. Kranton (2007) 'Public goods in networks', *Journal of Economic Theory*, 135: 478–494.
- Bramoullé, Y., H. Djebbari and B. Fortin (2009) 'Identification of peer effects through social networks', *Journal of Econometrics*, 150: 41–55.
- Brock, W. and S.E. Durlauf (2001) 'Discrete choice models with social interactions', *Review of Economic Studies*, 68: 235–260.
- Burt, R.S. (1992) *Structural Holes*. Cambridge, MA: Harvard University Press.
- Calvó-Armengol, A. (2004) 'Job contact networks', *Journal of Economic Theory*, 115: 191–206.
- Calvó-Armengol, A.E. and M.O. Jackson (2004) 'The effects of social networks on employment and inequality', *American Economic Review*, 94: 426–454.
- Calvó-Armengol, A.E. and Y. Zenou (2004) 'Social networks and crime decisions. The role of social structure in facilitating delinquent', *International Economic Review*, 45: 939–958.
- Calvó-Armengol, A.E. and Y. Zenou (2005) 'Job matching, social network and word-of-mouth communication', *Journal of Urban Economics*, 57: 500–522.
- Calvó-Armengol, A.E. and M.O. Jackson (2007) 'Social networks in labor markets: Wage and employment dynamics and inequality', *Journal of Economic Theory*, 132: 27–46.
- Calvó-Armengol, A., E. Patacchini and Y. Zenou (2005) 'Peer effects and social networks in education and crime', CEPR Discussion Paper No. 5244.
- Calvó-Armengol, A.E., Verdier, T. and Y. Zenou (2007) 'Strong ties and weak ties in employment and crime', *Journal of Public Economics*, 91: 203–233.
- Calvó-Armengol, A., E. Patacchini and Y. Zenou (2009) 'Peer effects and social networks in education', *Review of Economic Studies*, 76: 1239–1267.
- Chwe, M. (1999) 'Structure and strategy in collective action', *American Journal of Sociology*, 105: 128–156.
- Chwe, M. (2000) 'Communication and coordination in social networks', *Review of Economic Studies*, 67: 1–16.
- Coleman, J.S. (1988) 'Social capital in the creation of human capital', *American Journal of Sociology*, 94: S95–S120.
- 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: 668–673.
- Conley, T.G. and C.R. Udry (2010) 'Learning about a new technology: Pineapple in Ghana', *American Economic Review*, 100: 35–69.
- Corcoran, M., L. Datcher and G.J. Duncan (1980) 'Most workers find jobs through word of mouth', *Monthly Labor Review*, 103: 33–35.
- Currarini, S., M.O. Jackson and P. Pin (2009) 'An economic model of friendship: Homophily, minorities, and segregation', *Econometrica*, 77: 1003–1045.
- De Martí, J. and Y. Zenou (2009) 'Ethnic identity and social distance in friendship formation', CEPR Discussion Paper No. 7566.
- Diestel, R. (2005) *Graph Theory*. Heidelberg: Springer-Verlag.
- Durlauf, S.N. and Y.M. Ioannides (2010) 'Social interactions', *Annual Review of Economics* 2: 451–478.

- Echenique, F. and R. Fryer (2007) 'A measure of segregation based on social interactions', *Quarterly Journal of Economics*, 122: 441–485.
- Erdős, P. and A. Rényi (1959) 'On random graphs', *Publicationes Mathematicae*, 6: 290–297.
- Erdős, P. and A. Rényi (1960) 'On the evolution of random graphs', *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5: 17–61.
- Freeman, L.C. (1977) 'A set of measures of centrality based upon betweenness', *Sociometry*, 40: 35–41.
- Freeman, L.C. (1979) 'Centrality in social networks: Conceptual clarification', *Social Networks*, 1: 215–239.
- Friedkin, N. (1991) 'Theoretical foundations for centrality measures', *American Journal of Sociology*, 96: 1478–1504.
- Galeotti, A. and F. Vega-Redondo (2005) 'Local externalities and complex networks,' Unpublished manuscript. University of Essex.
- Galeotti, A., S. Goyal, M.O. Jackson, F. Vega-Redondo, and L. Yariv (2009) 'Network games', *Review of Economic Studies*, 77:218–244.
- Garoupa, N. (1997) 'The theory of optimal law enforcement', *Journal of Economic Surveys*, 11: 267–95.
- Glaeser, E., B. Sacerdote and J. Scheinkman (1996) 'Crime and social interactions', *Quarterly Journal of Economics*, 111: 507–548.
- Goyal, S. (2007) *Connections: An Introduction to the Economics of Networks*. Princeton: Princeton University Press.
- Goyal, S., M.J. van der Leij, and J.L. Moraga-Gonzalez (2006) 'Economics: An emerging small world', *Journal of Political Economy*, 114: 403–412.
- Granovetter, M.S. (1973) 'The strength of weak ties', *American Journal of Sociology*, 78: 1360–1380.
- Granovetter, M.S. (1978) 'Threshold models of collective behavior', *American Journal of Sociology*, 83: 1420–1443.
- Granovetter, M.S. (1983) 'The strength of weak ties: A network theory revisited', *Sociological Theory*, 1: 201–233.
- Granovetter, M.S. (1995) *Getting a Job: A Study of Contacts and Careers*. 2nd edition. Chicago: Chicago University Press.
- Haynie, D.L. (2001) 'Delinquent peers revisited: Does network structure matter?', *American Journal of Sociology*, 106: 1013–1057.
- Holzer, H.J. (1987) 'Informal job search and black youth unemployment', *American Economic Review*, 77: 446–452.
- Holzer, H.J. (1988) 'Search method used by unemployed youth', *Journal of Labor Economics*, 6: 1–20.
- Ioannides, Y.M. and D.L. Loury (2004) 'Job information networks, neighborhood effects, and inequality', *Journal of Economic Literature*, 42: 1056–1093.
- Jackson, M.O. (2007) 'The economics of social networks', in R. Blundell, W.K. Newey, and T. Persson (eds) *Advances in Economics and Econometrics. Theory and Applications*, Ninth World Congress, Vol. I. Cambridge: Cambridge University Press. pp. 1–56.
- Jackson, M.O. (2008) *Social and Economic Networks*. Princeton: Princeton University Press.
- Jackson, M.O. and B.W. Rogers (2007) 'Meeting strangers and friends of friends: How random are socially networks?', *American Economic Review*, 97: 890–915.
- Jackson, M.O. and A. Wolinsky (1996) 'A strategic model of social and economic networks', *Journal of Economic Theory*, 71: 44–74.
- Jackson, M.O. and L. Yariv (2007) 'Diffusion of behavior and equilibrium properties in network games', *American Economic Review (Papers and Proceedings)*, 97: 92–98.
- Katz, L. (1953) 'A new status index derived from sociometric analysis', *Psychometrika*, 18: 39–43.
- König, M., C. Tessone and Y. Zenou (2009) 'A dynamic model of network formation with strategic interactions', CEPR Discussion Paper No. 7521.
- Liljeros, F., C. Edling, L.A.N. Amaral, H.E. Stanley and Y. Åberg (2001) 'The web of human sexual contacts', *Nature*, 411: 908–909.
- Lorrain, F. and H.C. White (1971) 'Structural equivalence of individuals in social networks', *Journal of Mathematical Sociology*, 1: 49–80.
- Manski, C.F. (1993) 'Identification of endogenous effects: The reflection problem', *Review of Economic Studies*, 60: 531–542.
- McPherson, M., L. Smith-Lovin and J.M. Cook (2001) 'Birds of a feather: Homophily in social networks', *Annual Review of Sociology*, 27: 415–444.
- Milgram, S. (1967) 'The small world problem', *Psychology Today*, 2: 60–7.
- Montgomery, J.D. (1991) 'Social networks and labor-market outcomes: Toward and economic analysis', *American Economic Review*, 81: 1408–1418.
- Montgomery, J.D. (1992) 'Job search and network composition: Implications of the strength-of-weak-ties hypothesis', *American Sociological Review*, 57: 586–596.
- Montgomery, J.D. (1994) 'Weak ties, employment, and inequality: An equilibrium analysis', *American Journal of Sociology*, 99: 1212–1236.
- Moulin, H.J. (2003) *Fair Division and Collective Welfare*. Cambridge, MA: MIT Press.

- Myerson, R. (1977) 'Graphs and cooperation in games', *Mathematics of Operations Research*, 2: 225–229.
- Myerson, R. (1991) *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press.
- Newman, M.E.J. (2001) 'The structure of scientific collaboration networks', *Proceedings of the National Academy of Sciences of the USA*, 98: 404–409.
- Newman, M.E.J. (2003) 'The structure and function of complex networks', *SIAM Review*, 45: 167–256.
- Palacios-Huerta, I. and O. Volij (2004) 'The measurement of intellectual influence', *Econometrica*, 72: 963–977.
- Patacchini, E. and Y. Zenou (2008) 'The strength of weak ties in crime', *European Economic Review*, 52: 209–236.
- Patacchini, E. and Y. Zenou (2012) 'Juvenile delinquency and conformism', *Journal of Law, Economics, and Organization*, forthcoming.
- Polinsky, M.A. and S. Shavell (2000) 'The economic theory of public enforcement of law', *Journal of Economic Literature*, 38: 45–76.
- Sarnecki, J. (2001) *Delinquent Networks: Youth Co-Offending in Stockholm*. Cambridge: Cambridge University Press.
- Schelling, T.S. (1978) *Micromotives and Macrobehavior*. New York: Norton.
- Solomonoff, R. and A. Rapoport (1951) 'Connectivity of random nets', *Bulletin of Mathematical Biophysics*, 13: 107–117.
- Sundararajan, A. (2006) 'Local network effects and network structure,' Unpublished manuscript. New York University.
- Sutherland, E.H. (1947) *Principles of Criminology*. 4th edition. Chicago: J.B. Lippincott.
- Topa, G. (2001) 'Social interactions, local spillovers and unemployment', *Review of Economic Studies*, 68: 261–295.
- Vega-Redondo, F. (2007) *Complex Social Networks*. Econometric Society Monograph Series. Cambridge: Cambridge University Press.
- Warr, M. (2002) *Companions in Crime: The Social Aspects of Criminal Conduct*. Cambridge: Cambridge University Press.
- Wasserman, S. and K. Faust (1994) *Social Network Analysis. Methods and Applications*. Cambridge: Cambridge University Press.
- Watts, D.J. and S.H. Strogatz (1998) 'Collective dynamics of "small-world" networks', *Nature*, 393: 409–410.
- White, H.C., S. Boorman and R. Breiger (1976) 'Social structure from multiple networks: I. Blockmodels of roles and positions', *American Journal of Sociology*, 81: 730–780.