

Diffusion, Strategic Interaction, and Social Structure

Matthew O. Jackson* and Leeat Yariv†

April, 2008

1 Introduction

How we act, as well as how we are acted upon, are to a large extent influenced by our relatives, friends and acquaintances. This is true of which profession we decide to pursue, whether or not we adopt a new technology, as well as whether or not we catch the flu. In this chapter we provide an overview of research that examines how social structure impacts economic decision making and the diffusion of innovations, behaviors and information.

We begin with a brief overview of some of the stylized facts on the role of social structure on diffusion in different realms. This is a rich area of study that includes a vast set of case studies, as well as some important regularities. With that empirical perspective, we then discuss insights from the epidemiology and random graph literatures that help shed light on the spread of infections throughout a society. Contagion of this form can be thought of as a basic, but important, form of social interaction, where the social structure largely determines patterns of diffusion. This literature presents a rich understanding of questions such as: “How densely connected does a society have to be in order to have an infection reach a nontrivial fraction of its members?”, “How does this depend on the infectiousness of the disease?”, “How does it depend on the particulars of the social network in place?”,

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

†Division of the Humanities and Social Sciences, Caltech, Pasadena, CA 91125.
<http://www.hss.caltech.edu/~lyariv/index.htm> e-mail: lyariv@hss.caltech.edu

“Who is most likely to become infected?”, and “How widespread is an infection likely to be?”, among others. The results on this subject apply beyond infectious diseases, and touch upon issues ranging from the spread of information to the proliferation of ideas.

While such epidemiological models provide a useful look at some types of diffusion, there are many economically relevant applications where a different modeling approach is needed, and, in particular, where the interaction between individuals requires a game theoretic analysis. In fact, though disease and the transmission of certain ideas and bits of information can be modeled through mechanical or purely probabilistic sorts of diffusion processes, there are other important situations where individuals take decisions and care about how their social neighbors or peers behave. This applies to decisions of which products to buy, which technology to adopt, whether or not to become educated, whether to learn a language, how to vote, and so forth. Such interactions involve equilibrium considerations and often have multiple potential outcomes. For example, an agent might care about the proportion of neighbors adopting a given action, or might require some threshold of stimulus before becoming convinced to take an action, or may want to take an action that is different from that of his or her neighbors (e.g., free-riding on their information gathering if they do gather information, but gathering information him or herself if neighbors do not). Here we provide an overview of how the recent literature has modeled such interactions, and how it has been able to meld social structure with predictions of behavior.

2 Empirical Background: Social Networks and Diffusion

There is a large body of work that identifies the effects of social interactions on a wide range of applications spanning fields: epidemiology, marketing, labor markets, political science, and agriculture are only a few.

While some of the empirical tools for the analysis of social interaction effects have been described in Chapter ?, and many of their implementations for research on housing decisions,

labor markets, addictions, and more have been discussed in Chapters ??, we now describe the type of empirical work that ties directly to the models that are discussed in the current chapter. In particular, we discuss several examples of studies that illustrate how social structure impacts outcomes and behaviors.

The relevant studies are broadly divided into two classes. First, there are cross-sectional studies that concentrate on a snapshot of time and look for correlations between social interaction patterns and observable behaviors. This class ties one to one to the analysis below of strategic games played by a network of agents. While it can be very useful in identifying correlations, it is important to keep in mind that identifying causation is more complicated without the right exogenous variation or structural underpinnings. Second, there are longitudinal studies that take advantage of the inherent dynamics in diffusion. Such studies have generated a number of interesting observations and are more suggestive of some of the insights the theoretical literature on diffusion has generated.

The empirical work on these topics is vast and we provide here only a narrow look of the work that is representative of the *type* of studies that have been pursued and relate to the focus of this chapter.

2.1 The Effects of Networks in a Cross Section

Studies that are based on observations at one point of time most often compare the frequency of a certain behavior or outcome across individuals that are connected as opposed to ones that are not. For example, Glaeser, Sacerdote, and Scheinkman (1996) showed that the structure of social interactions can help explain the cross-city variance in crime rates in the U.S.; Bearman, Moody, and Stovel (2004) examined the network of romantic connections in high-school, and its link to phenomena such as the spread of sexually transmitted diseases (see the next subsection for a discussion of the spread of epidemics). Such studies provide important evidence for the correlation of behaviors with connected individuals, and in the case of disease (given the nature of contagion) provide some direct evidence of diffusion.

With regards to labor markets, there is a rich set of studies showing the importance of

social connections for diffusing information about job openings, dating from Rees (1966) and Rees and Schultz (1970). Influential studies by Granovetter (1973, 1985, 1995) show that even casual or infrequent acquaintances (“weak ties”) could be quite important in diffusing information. Those studies were based on interviews that directly ask subjects how they obtained information about their current jobs. Other studies, based on outcomes, such as Topa (2001), Conley and Topa (2002), and Bayer, Ross, and Topa (2006), identify local correlations in employment status within neighborhoods in Chicago (and consider neighborhoods that go beyond the geographic but also include proximity in other socioeconomic dimensions), examining the extent to which local interactions are important in employment. Bandiera, Barankay, and Rasul (2006) create a bridge between network formation (namely, the creation of friendships amongst fruit pickers) and the effectiveness of different labor contracts.¹ This rich empirical literature documents the important role of social connections in transmitting information about jobs, and also differentiates between different types of social contacts and shows that even weak ties can be important ones.

There is further (and earlier) research that examines the different roles of individuals in diffusion. Important work by Katz and Lazarsfeld (1955) (building on earlier work of Lazarsfeld, Berelson and Gaudet (1944), Merton (1948), and others), identifies the roles of “opinion leaders” in the formation of various beliefs and opinions. Individuals are heterogeneous (at least in behaviors), and some specialize in becoming well-informed on certain subjects, and then information and opinions diffuse to other less-informed individuals via conversations with these opinion leaders. Lazarsfeld, Berelson and Gaudet (1944) study voting decisions in an Ohio town in the 1940 U.S. presidential campaign, and document the presence and role of such opinion leaders. Katz and Lazarsfeld (1955) interviewed women in Decatur, Illinois, and asked about a number of things such as their views on household goods, fashion, movies, and local public affairs. When women showed a change in opinion in follow-up interviews, Katz and Lazarsfeld traced influences that led to the change in opinion, again finding evidence of the presence of opinion leaders.

¹See Chapter ? for an extended discussion of network effects on employment.

Diffusion of new products is understandably a topic of much research. Rogers (1995) discusses numerous studies illustrating the impacts of social interactions on the diffusion of new products, and discussing various factors that impact which products succeed and which products fail. For example, related to the idea of opinion leaders, Feick and Price (1987) surveyed 1531 households and provided evidence that consumers recognize and make use of particular individuals in their social network termed “market mavens,” those who have a high propensity to provide marketplace and shopping information. Whether or not products reach such mavens can influence the success of a product, independently of the product’s quality. Tucker (2007) uses micro-data on the adoption and use of a new video-messaging technology in an investment bank consisting of 2118 employees. Tucker notes the effects of the underlying network in that employees follow the actions of those who either have formal power, or informal influence (which is, to some extent, endogenous to a social network).

In the political context, there are several studies focusing on the social sources of information electors select, as well as on the selective misperception of social information they are exposed to. A prime example of such a collection of studies is Huckfeldt and Sprague (1995) who concentrated on the social structure in South Bend, Indiana, during the 1984 elections. They illustrated the likelihood of socially connected individuals to hold similar political affiliations. In fact, the phenomenon of individuals connecting to individuals who are similar to them is observed across a wide array of attributes and is termed by sociologists *homophily* (for a review, see McPherson, Smith-Lovin, and Cook, 2001, as well as the discussion of homophily in Chapter ?).

While cross sectional studies are tremendously interesting in that they suggest dimensions on which social interactions may have an impact, they pose many empirical challenges. Most notably, correlations between behaviors and outcomes of individuals and their peers may be driven by common unobservables and therefore be spurious. Given the strong homophily patterns in many social interactions, individuals who associate with each other often have common unobserved traits, which could lead them to similar behaviors. This makes empirical analysis of the social impact on diffusion of behaviors based on cross sectional data especially

challenging. Laboratory experiments are consequently quite useful in eliciting the effects of real-world networks on fully controlled strategic interactions, and are being increasingly utilized.

For instance, Leider, Mobius, Rosenblat, and Do (2007) elicited the friendship network among undergraduates at a U.S. college and illustrated how altruism varies as a function of the social proximity. In a similar setup, Goeree, McConnell, Mitchell, Tromp, and Yariv (2007) elicited the friendship network in an all-girls school in Pasadena, CA, together with girls' characteristics and later ran dictator games with recipients who varied in social distance. They identified a "1/d Law of Giving," in that the percentage given to a friend was inversely related to her social distance in the network.²

The types of conclusions that have been reached from these cross sectional studies can be roughly summarized as follows. First, in a wide variety of settings, associated individuals tend to have correlated actions and opinions. This does not necessarily embody diffusion or causation, but as discussed in the longitudinal section below, we will see that there is significant evidence of social influence in diffusion patterns as well. Second, individuals tend to associate with others who are similar to each other, in terms of beliefs and opinions. This has an impact on the structure of social interactions, and can impact diffusion. It also represents an empirical quandary (that can partly be sorted out with longitudinal data) of the extent to which social structure influences opinions and behavior versus the reverse. Third, individuals fill different roles in a society, with some acting as "opinion leaders," and being key conduits of information and potential catalysts for diffusion.

2.2 The Effects of Networks over Time

Longitudinal data can be especially important in diffusion studies, as they provide information on how opinions and behaviors move through a society over time. They also help sort out issues of causation as well as supply specific information about the extent to which

²See Chapter ? for related experimental studies in the context of social learning. For a look at a few network experiments that are not based on a real-world social structure, see Kosfeld (2003).

behaviors and opinions are adopted dynamically, and by whom.

As mentioned, longitudinal data are especially important in going beyond the documentation of correlation between social connections and behaviors, and illustrating that social links are really the conduits for information and diffusion. For example, Conley and Udry (2005) show that pineapple growers in Ghana tend to follow those farmers who succeed in their use of fertilizers. Through careful examination of local ties, and the timing of different actions, they can trace the influence of the outcome of one farmer's crop on subsequent behavior of other farmers.

More generally, diffusion of new technologies is extremely important when looking at transitions in agriculture. Seminal studies by Ryan and Gross (1943) and Griliches (1957) examined the effects of social connections on the adoption of a new behavior, specifically the adoption of hybrid corn in the U.S. Looking at aggregate adoption rates in different states, these authors illustrated that the diffusion of hybrid corn followed an S-shape curve over time: starting out slowly and accelerating, and then ultimately decelerating.³ Foster and Rosenzweig (1995) collected a household-level panel data from a representative sample of rural Indian households having to do with the adoption and profitability of high-yielding seed varieties (associated with the Green Revolution). They identified significant learning-by-doing, where some of the learning was through neighbors' experience. In fact, the observation that adoption rates of new technologies, products, or behaviors exhibit S-shaped curves can be traced at least to Tarde (1903), who discussed of the importance of imitation in adoption. Such patterns are found across many applications (see Mahajan and Peterson, 1985 and Rogers, 1995).

Understanding diffusion is particularly important for epidemiology and medicine. On the one hand, it is important to understand how different types of diseases spread in a population. On the other (and very different) hand, it is crucial to gain insights into how new medicine gets adopted. Colizza, Barrat, Barthelemy, and Vespignani (2006, 2007) tracked the spread of severe acute respiratory syndrome (SARS) across the world combining census data with

³See Young (2006) for a complementary analysis to that of Griliches (1957).

data on almost all air transit during the years 2002-2003. They illustrated the importance of structures of long-range transit networks for the spread of an epidemic. Christakis and Fowler (2007) studied the social network of 12,067 individuals in the U.S. assessed repeatedly from 1971 to 2003 as part of the Framingham Heart Study. Concentrating on body-mass index, Christakis and Fowler found that a person's chances of becoming obese increased by 57% if he or she had a friend who became obese, by 40% if he or she had a sibling who became obese, and by 37% if they had a spouse who became obese. Since the study controls for various selection effects, it suggests that obesity spreads very much like an epidemic. In particular, the underlying social structure appears to play a critical role.

Coleman, Katz, and Menzel (1966) is one of the first studies to document the role of social networks in diffusion processes. The study looked at the adoption of a new drug (tetracycline) by doctors and highlighted two observations. First, as with hybrid corn, adoption rates followed an S-shape curve over time. Second, adoption rates depended on the density of social interactions. Doctors with more contacts (measured according to the trust placed in them by other doctors) adopted at higher rates and earlier in time.⁴

With respect to how individuals befriend others, we mentioned before that the idea of homophily, that people tend to connect to others similar to them, can be plagued with endogeneity issues when we discuss issues of similarity in behaviors rather than some exogenous traits. Namely, it is often hard to disentangle which of two processes is at the root of observed similarity in behavior between connected agents: a process of selection (assortative pairing), in which similarity precedes association, or a process of socialization, in which association leads to similarity. One way to overcome this issue is to consider attributes that are not malleable.⁵ Another approach is to track connections over time. Kandel (1978) provides an example of such a use of longitudinal sociometric data. She concentrated on adolescent

⁴As a caveat, Van den Bulte and Lilien (2001) add controls having to do with marketing exposure of the doctors in the study and show that the social effects may be mitigated. Nonetheless, further studies such as Nair, Manchanda, and Bhatia (2006) have again found evidence of such effects after more carefully controlling for the marketing and other characteristics in a more recent and much larger data set.

⁵For instance, Goeree, McConnell, Mitchell, Tromp, and Yariv (2007) observe homophily on height.

friendship pairs and examined the levels of homophily on four attributes (frequency of current marijuana use, level of educational aspirations, political orientation, and participation in minor delinquency) at various stages of friendship formation and dissolution. She noted that observed homophily in friendship dyads resulted from a significant combination of both types of processes, so that individuals emulated their friends, but also tended to drop friendships with those more different from themselves and add new friendships to those more similar to themselves.

In summary, let us mention a few of the important conclusions obtained from longitudinal studies of diffusion. First, not only are behaviors across socially connected individuals correlated, but individuals do in fact influence each other. While this may sound straightforward, it takes careful control to ensure that it is not unobserved correlated traits or influences that lead to similar actions by connected individuals. Second, in various settings, more socially connected individuals adopt new behaviors and products earlier and at higher rates. Third, diffusion exhibits specific patterns over time, and specifically there are many settings where an “S”-shaped pattern emerges, with adoption starting slowly, then accelerating, and eventually asymptoting.

3 Models of Diffusion and Strategic Interaction Absent Network Structure

Let us now discuss various models of diffusion. We start with some of the early models that do not account for the underlying network architecture per-se. These models incorporate the empirical observations regarding social influence through the particular dynamics assumed, or preferences posited, and generate a match to the *aggregate* empirical observations regarding diffusion over time of products, diseases, or behavior. For example, the so-called *S*-shaped adoption curves.

3.1 Marketing

One of the earliest and still widely used models of diffusion is the Bass (1969) Model. This is a parsimonious model, which makes predictions of the percentage of potential adopters of a product or behavior who will have adopted by a given time. The current rate of change of adoption depends on the current level and two critical parameters. These two parameters are linked to the rate at which people innovate or adopt on their own, and the rate at which they imitate or adopt because others have, thereby putting into (theoretical) force the empirical observation regarding peers' influence.

If we let $F(t)$ be the percentage of agents who have adopted by time t , and m be the fraction of agents in the population who are potential adopters, a discrete time version of the Bass model is characterized by the difference equation

$$G(t) = G(t - 1) + p(m - G(t - 1)) + q(m - G(t - 1)) \frac{G(t - 1)}{m},$$

where p is a rate of innovation and q is a rate of imitation. To glean some intuition, note that the expression $p(m - G(t - 1))$ represents the fraction of people who have not yet adopted and might potentially do so times the rate of spontaneous adoption. In the expression $q(m - G(t - 1)) \frac{G(t-1)}{m}$, the rate of imitation is multiplied by two factors. The first factor, $(m - G(t - 1))$, is the fraction of people who have not yet adopted and may still do so. The second expression, $\frac{G(t-1)}{m}$, is the relative fraction of potential adopters who are around to imitate. If we set m equal to 1, and look at a continuous time version of the above difference equation, we get

$$g(t) = (p + qG(t))(1 - G(t)), \tag{1}$$

where g is the rate of diffusion (time rate of change of G). Solving this when $p > 0$ and setting the initial set of adopters at 0, $G(0) = 0$, leads to the following expression:

$$G(t) = \frac{1 - e^{-(p+q)t}}{1 + \frac{q}{p}e^{-(p+q)t}}.$$

This is a fairly flexible formula that is good at fitting time series data of innovations. By estimating p and q from existing data, one can also make forecasts of future diffusion. It has




Figure 1: *S*-shape Adoption

been used extensively in marketing and in the analysis of diffusion (e.g., Rogers (1995)), and has spawned many extensions and variations.⁶

Generally, $G(t)$, is *S*-shaped (see Figure 1), matching one of the main insights of the longitudinal empirical studies on diffusion discussed above.

The Bass model provides a clear intuition for why adoption curves would be *S*-shaped. Indeed, when the adoption process begins, imitation plays a minor role (relative to innovation) since not many agents have adopted yet and so the volume of adopters grows slowly. As the number of adopters increases, the process starts to accelerate as now innovators are joined by imitators. The process eventually starts to slow down, partly because there are fewer agents left to adopt (the term $1 - G(t)$ in (1) eventually becomes small). Thus, we see a process which starts out slowly, then accelerates, and then eventually slows and asymptotes.

⁶For some recent models, see Leskovec, Adamic, and Huberman (2007) and Young (2006).

3.2 Collective Action, Fashion, and Fads

The Bass model described above is mechanical in that adopters and imitators are randomly determined, rather than choosing actions strategically. The empirical observation that individuals influence each other through social contact can be derived through agents' preferences, rather than through some exogenously specified dynamics.

Diffusion in a strategic context was first studied without a specific structure to interactions. Broadly speaking, there were two approaches in this early literature. In the first, all agents are connected to one another (that is, they form a complete network). Effectively, this corresponds to a standard multi-agent game in which payoffs to each player depend on the entire profile of actions played in the population. The second approach has been to look at interactions in which agents are matched to partners in a random fashion.

Diffusion on Complete Networks. Granovetter (1978) considered a model in which N agents are all connected to one another and each agent chooses one of two actions: 0 or 1. Associated with each agent i is a number n_i of agents such that action 1 leads to a higher payoff if at least n_i other agents choose the action 1, and action 0 leads to a higher payoff otherwise (so that the game has strategic complementarities). For instance, suppose that the utility of agent i when the profile of actions is $(x_1, \dots, x_N) \in \{0, 1\}^N$ is given by:

$$u_i(x_1, \dots, x_N) = \left[\frac{1}{N-1} \left(\sum_{j \neq i} x_j \right) - c_i \right] \mathbf{1}(x_i = 1), \quad (2)$$

where c_i is randomly drawn from a distribution F over $[0, 1]$. c_i stands for a cost agent i experiences upon choosing the action 1 (e.g., a one time switching cost from one technology to the other, the potential time costs of political revolt, etc.). Thus, the utility of agent i is normalized to 0 when choosing the action 0. When choosing the action 1, agent i experiences a benefit proportional to the fraction of other agents choosing the action 1 and a cost of c_i .

Granovetter considered a dynamic in which at each stage, each agent best responds to the previous period's distribution of actions. Then, if at period t there was a fraction x^t of agents choosing the action 1, then at period $t+1$ all agents i , who have chosen $x_i^t \in \{0, 1\}$ in period t ,

with a cost lower than $\frac{Nx^t - x_i^t}{N-1}$ would choose the action 1. For a large population, $\frac{Nx^t - x_i^t}{N-1} \simeq x^t$ and $x^{t+1} \simeq F(x^t)$. The equation $x^* = F(x^*)$ then corresponds to the (approximate) equilibria of the system.

The shape of the distribution F determines which equilibria are *tipping points*: equilibria such that only a slight addition to the fraction of agents choosing the action 1 shifts the population, under the best response dynamics, to the next higher equilibrium level of adoption (we return to a discussion of tipping and stable points when we consider a more general model of strategic interactions on networks below.)

Note that while in the Bass model the diffusion path was determined by $G(t)$, the fraction of adopters as a function of time, here it is easier to work with $F(x)$, corresponding to the fraction of adopters as a function of the previous period's fraction x . Nonetheless, S -shaped adoption can be identified in Granovetter's model and traced to the attributes of the cost distribution F . In this model diffusion occurs when tipping points are passed. If F is S -shaped in between a tipping point and the next higher equilibrium, mapping adoption across time will match Figure 1 above.

Fashions and Random Matching. Pesendorfer (1995) considered a model in which individuals are randomly matched and new fashions serve as signaling instruments for the creation of matches. He identifies particular matching technologies that generate fashion cycles. Pesendorfer describes the spread of a new fashion as well as its decay over time. In Pesendorfer's model, the price of the design falls as it spreads across the population. Once sufficiently many consumers own the design it is profitable to create a new design and thereby render the old design obsolete. In particular, demand for any new innovation eventually levels off as in the above two models.

4 Models of Diffusion and Strategic Interaction Embedded in a Social Setting

Let us now turn to models that explicitly consider social structure in examining diffusion patterns. We start with models that stem mostly from the epidemiology literature and account for the underlying social network, but are mechanical in terms of the way that disease spreads from one individual to another (much like the Bass model described above). We then proceed to models where players make choices that depend on their neighbors' actions as embedded in a social network; for instance, only adopting an action if a certain proportion of neighbors do (as in Granovetter's setup), or possibly not adopting an action if enough neighbors do so.

4.1 A Unified Setting

Many models of diffusion and strategic interaction on networks have the following common elements.

There is a finite set of agents $N = \{1, \dots, n\}$.

Agents are connected by a (possibly directed) network $g \in \{0, 1\}^{n \times n}$. We let $N_i(g) \equiv \{j : g_{ij} = 1\}$ be the neighbors of i . The degree of a node i is the number of her neighbors, $d_i \equiv |N_i(g)|$.

When links are determined through some random process, it is often useful to summarize the process by the resulting distribution of degrees P , where $P(d)$ denotes the probability a random individual has a degree of d .^{7 8}

Each agent $i \in N$ takes an action $x_i \in X_i$. In order to unify and simplify the description of various models, we focus on binary actions, so that $X_i = \{0, 1\}$. Actions can be metaphors for becoming "infected" or not, buying a new product or not, choosing one of two activities, etc.

⁷Such a description is not complete, in that it does not specify the potential correlations between degrees of different individuals on the network. We will return to this point.

⁸In principle, one would want to calibrate degree distributions with actual data. The literature on network formation, see Chapter ??, suggests many insights on plausible degree distributions $P(d)$.

4.2 Epidemiology Models

4.2.1 Random Graph Models

Technically, some basic insights about the extent to which behavior or an infection can spread in a society build from random graph theory. Random graph theory provides a tractable base for understanding things like the structure and size of the *components* of a network, maximally connected subnetworks.⁹

Before presenting some results, let us talk through some of the ideas in the context of what is known as the Reed-Frost model.¹⁰ Consider, for example, the spread of a disease. Initially, some individuals in the society get infected through mutations of a germ or other exogenous sources. Consequently, some of these individuals' neighbors are infected through contact, while others are not. This depends on how virulent the disease is, among other things. In this application, it makes sense (at least as a starting point) to assume that becoming infected or avoiding infection is not a choice; i.e., contagion here is non-strategic. In the simplest model, there is a probability $\pi \geq 0$ that a given individual is immune (e.g., through vaccination). If an individual is not immune, it is assumed that he or she is sure to catch the disease if one of his or her neighbors ends up with the disease. In this case, in order to estimate the volume of those ultimately infected, we proceed in two steps, depicted in Figure 2. First, we delete a fraction π of the nodes that will never be infected (these correspond to the dotted nodes in the Figure). Then, we note that the remaining components of the network that contain the originally infected individuals comprise the full extent of the infection. In particular, if we can characterize what the components of the network look like after removing some portion of the nodes, we have an idea of the extent of the infection. In Figure 2 we start with one large connected component (circumvented by a dotted line) and two small connected components. After removing the immune agents, there is still a large connected component (though smaller than before), and four small components.

⁹Formally, these are the subnetworks containing maximal sets $C \subseteq N$ of nodes such that for any $i, j \in C$, there exists $i_1, \dots, i_k \in C$ such that $g_{ii_1} = g_{i_1 i_2} = \dots = g_{i_{k-1} i_k} = g_{i_k j} = 1$.

¹⁰See Jackson (2008) for a more detailed discussion of this and related models.




Figure 2: Network Components and Immune Agents

A starting point for the formal analysis is on a network which is formed by randomly placing a link between two nodes with a probability of p , independently of any other link. This is known as a “Poisson random network” (and also sometimes referred to as an “Erdős-Renyi random graph”, see Chapter ??). So the analysis boils down to considering a network on $(1 - \pi)n$ nodes with an independent link probability of p , and then measuring the size of the component containing a randomly chosen node, proxying for the probability that a random agent gets infected.

Clearly, with a fixed set of nodes, and a positive probability p that lies strictly between 0 and 1, every conceivable network is possible. Given this, many of the characterization results in the literature are limiting results that focus on the probability of certain types of networks as the size of the population grows. For example, if we think of letting n grow, and then examine p as a function of n , we can ask for which p 's a nonvanishing fraction of nodes will become infected with a probability bounded above (thus, not tending to) 0. So, let us consider a sequence of societies indexed by n and corresponding probabilities of links $p(n)$.

Erdos and Renyi (1959, 1960) proved a series of results that characterize some basic properties of such random graphs. In particular,¹¹

- The threshold for the existence of a giant component, a component that contains a fraction of the population bounded away from 0, is $1/n$. That is, if $p(n)$ over $1/n$ tends to infinity, the the probability of having a giant component tends to 1, while if $p(n)$ over $1/n$ tends to 0, then the probability of having such a giant component tends to 0.
- The threshold for the network to be connected (so that every two nodes are path connected) is $\log(n)/n$.

The logic for the first threshold is easy to explain, although the proofs of these results are more involved. To see the threshold for emergence of a giant component, consider the question of following a link out of a given node, and asking whether or not one would expect

¹¹See Chapter 4 in Jackson (2008) for a fuller discussion and proofs of these results.

to be able to continue to find a link to another node from that one. If the expected degree is much smaller than 1, then following the few links from any given node is likely to lead to dead-ends. In contrast, when the expected degree is much higher than 1, then from any given node, one expects to be able to reach more nodes and so the component should expand outward.

Adjusting for the factor π of the number of immune nodes does not affect the above thresholds, since it is simply a constant.

Between these two thresholds, there is only one giant component, so that the next largest component is of a size that is a vanishing fraction of the giant component. This is intuitively clear, as to have two large components requires many links within each component but no links between the two components, which is an unlikely event. In that sense, the image that emerges from Figure 2 of one large connected component is quite general.

These results then tell us that in a random network, if average degree is quite low (smaller than 1), then any initial infection is likely to die out. In contrast, if average degree is quite high (larger than $\log(n)$), then any initial infection is likely to spread to infect all of the susceptible individuals, i.e., a fraction of $1 - \pi$ of the population. In the intermediate range, there is a probability that the infection will die out and also a probability that it will infect a nontrivial, but limited, portion of the susceptible population. There, it can be shown that for large n , the fraction of nodes in the giant component of susceptible nodes is approximated by the nonzero q which solves

$$q = 1 - e^{-q(1-\pi)np}. \quad (3)$$

Then, q is an approximation of the probability of the infection spreading to a non-trivial fraction of nodes, and also of the percentage of susceptible nodes that would be infected.¹²

This provides a rough idea of the type of results that can be derived from random graph theory. There is much more that is known, as one can work with other models of random graphs (other than ones where each link has an identical probability), richer

¹²Again, see Chapter 4 in Jackson (2008) for more detail.

models of probabilistic infection between nodes, as well as derive more information about the potential distribution of infected individuals. It should also be emphasized that while the discussion here is in terms of “infection,” the applications clearly extend to many other things like transmission of ideas and information. Fuller treatment of behaviors where individual decisions depend in more complicated ways on neighbors’ decisions are treated in Section 4.3.

4.2.2 Diffusion with Recovery

The above analysis of diffusion presumes that once infected, a node eventually infects all of its susceptible neighbors. This misses important aspects of many applications. In terms of diseases, infected nodes can either recover and stop transmitting a disease, or die and completely disappear from the network. Similarly, if we think about behaviors, it might be that the likelihood that a node is still actively transmitting a bit of information to its neighbors decreases over time.

Ultimately, we will switch to models that allow for rather general strategic impact of peer behavior (a generalization of the approach taken by Granovetter). The epidemiology literature, however, takes steps forward in that direction by considering two alternative models that keep track of the state nodes are in. A terminology for those states are: *susceptible*, where a node is not currently infected or transmitting a disease but can catch it; *infected*, where a node has a disease and can transmit it to its neighbors, and *removed* (or *recovered*) where a node has been infected but is no longer able to transmit the disease and cannot be re-infected.

The first of the leading models is the “SIR” model (dating to Kermack and McKendrick, 1927), where nodes are initially susceptible but can catch the disease from infected neighbors. Once infected, a node continues to infect neighbors until it is randomly removed from the system. This captures some childhood diseases, such as the chicken pox, where one can only be infected once.

The other mode is the “SIS” model (see Bailey, 1975), where once infected, nodes can

randomly recover, but then they are susceptible again. This corresponds well with an assortment of bacterial infections, viruses, and flus, where one transitions back and forth between health and illness.

The analysis of the SIR model is a variant of the component-size analysis discussed above. The idea is that there is a random chance that an “infected” node infects a given “susceptible” neighbor before becoming “removed.” This can be shown to be equivalent to an analysis where one examines component structures where instead of removing *nodes* randomly, one removes *links* randomly from the network. Basically, this results in variations on the above sorts of calculations, where there are adjusted thresholds for infection depending on the relative rates of how quickly infected nodes can infect their neighbors compared to how quickly they are removed.

In contrast, the SIS model involves a different sort of analysis. The canonical version of that model is best viewed as one with a random matching process rather than a social network. In particular, suppose that a node i in each period will have interactions with d_i other individuals from the population. Recall our notation of $P(d)$ describing the proportion of the population that has degree d (so d interactions per period). The matches are determined at random, in such a way that if i is matched with j , then the probability that j has degree $d > 0$ is given by¹³

$$\tilde{P}(d) = \frac{P(d)d}{\langle d \rangle}, \quad (4)$$

where $\langle \cdot \rangle$ represents the expectation with respect to P . This reflects the fact that an agent is proportionally more likely to be matched with other individuals who have lots of matches. To justify this formally, one needs an infinite population.¹⁴

Individuals who have high degrees will have more interactions per period and will generally be more likely to be infected at any given time. An important calculation then pertains to the chance that a given meeting will be with an infected individual. If the infection rate of degree d individuals is $\rho(d)$, then the probability that any given meeting will be with an

¹³Consider only individuals who have degree $d > 0$, as others do not participate in the society.

¹⁴See the appendix of Currarini, Jackson and Pin (2007) for some details along this line.

infected individual is θ , where

$$\theta = \sum_d \tilde{P}(d)\rho(d) = \frac{\sum P(d)\rho(d)d}{\langle d \rangle}. \quad (5)$$

The chance of meeting an infected individual in a given encounter thus differs from the average infection rate in the population, which is just $\rho = \sum P(d)\rho(d)$, because θ is weighted by the rate at which individuals meet each other.

The standard version of contagion that is studied is one where the probability of an agent of degree d becoming infected is

$$\nu\theta d, \quad (6)$$

where $\nu \in (0, 1)$ is a rate of transmission of infection in a given period, and is small enough so that this probability is less than one. If ν is very small, then this is an approximation of getting infected under d interactions with each having an (independent) probability θ of being infected and then conditionally (independently) having a probability ν of getting infected through contact with a given infected individual. Assume further that in any given period an infected individual recovers and becomes susceptible with a probability $\delta \in (0, 1)$.

If such a system operates on a finite population, then eventually all agents will become susceptible and that would end the infection. If there is a small probability of a new mutation and infection in any given period, then the system will be ergodic and always have some probability of future infection.

To get a feeling for the long run outcomes in large societies, the literature has examined a steady state (i.e., a situation in which the system essentially remains constant) of a process which effectively can be thought of as operating on an infinite (continuum) population. Formally, a steady-state is defined by having $\rho(d)$ be constant over time for each d . Working with an approximation at the limit (termed a “mean-field” approximation that in this case can be justified with a continuum of agents, but with quite a bit of technical detail), a condition for steady state can be derived to be

$$0 = (1 - \rho(d))\nu\theta d - \rho(d)\delta \quad (7)$$

for each d . $(1 - \rho(d))\nu\theta d$ is the rate at which agents of degree d who were susceptible become infected and $\rho(d)\delta$ is the rate at which infected individuals of degree d recover. Letting $\lambda = \frac{\nu}{\delta}$, it follows that

$$\rho(d) = \frac{\lambda\theta d}{\lambda\theta d + 1}. \quad (8)$$

Solving (5) and (8) simultaneously leads to a characterization of the steady-state θ :

$$\theta = \sum_d \frac{P(d)\lambda\theta d^2}{\langle d \rangle (\lambda\theta d + 1)}. \quad (9)$$

This system always has a solution, and therefore a steady-state, where $\theta = 0$ so there is no infection. It can also have other solutions where θ is positive (but always below 1 if λ is finite). Unless P takes very specific forms, it can be difficult to solve for this analytically.

Special cases have been analyzed, such as the case of a power distribution (e.g., see Pastor-Satorras and Vespignani, 2000, 2001) where $P(d) = 2d^{-3}$. In that case, there is always a positive steady-state infection rate. More generally, Lopez-Pintado (2007) addresses the question of when it is that there will be a positive steady-state infection rate. To get some intuition for her results, let

$$H(\theta) = \sum \frac{P(d)d}{\langle d \rangle} \left(\frac{\lambda d\theta}{\lambda d\theta + 1} \right) = \sum \tilde{P}(d) \left(\frac{\lambda d\theta}{\lambda d\theta + 1} \right), \quad (10)$$

so that the equation $\theta = H(\theta)$ corresponds to steady states of the system. We can now expand on the type of analysis Granovetter (1978) performed as $H(\theta)$ accounts for network attributes. However, while the fixed point equation identifying Granovetter's stable points allowed for rather arbitrary diffusion patterns (depending on the cost distribution F), here the governing H has very specific attributes.

In particular, suppose we examine what infection rate would result if we start at a rate of θ and then run the system on an infinite population for one period. Noting that $H(0) = 0$, it is clear that 0 is always a fixed point and thus a steady-state. Since $H(1) < 1$, and H is increasing and strictly concave in θ , there can be at most one fixed point besides 0. For there to be another fixed point (steady-state) above $\theta = 0$, it must be that $H'(0)$ is above 1, or else given the strict concavity it would be that $H(\theta) < \theta$ for all positive θ . Moreover, in cases

where $H'(0) > 1$, a small perturbation away from a 0 infection rate will lead to increased infection. In the terminology we have introduced above, 0 would be a *tipping point*. Since

$$H'(0) = \lambda \frac{\langle d^2 \rangle}{\langle d \rangle}, \quad (11)$$

we have a simple way of checking whether we expect a positive steady-state infection or a 0 steady-state infection. This simply boils down to a comparison of the relative infection rate λ and $\frac{\langle d \rangle}{\langle d^2 \rangle}$ so that there is a positive infection rate if and only if

$$\lambda > \frac{\langle d \rangle}{\langle d^2 \rangle}. \quad (12)$$

Higher infection rates lead to the possibility of positive infection, as do degree distributions with high variances (relative to mean). The idea behind having a high variance is that there will be some “hub nodes” with high degree, who can foster contagion.

Going back to our empirical insights, note that while this analysis captures the observations that highly linked individuals are more likely to get infected and experience speedier diffusion, the aggregate behavior does not exhibit the *S*-shape that is common in many real-world diffusion processes. From the concavity of the function H , it follows that diffusion is at first quick and later slows down (matching only part of the *S*-shape).

Beyond the extant empirical studies, this analysis provides some intuitions behind what is needed for an infection to be possible, but it does not provide an idea of how extensive the infection rate will be and how that depends on network structure. While this does not boil down to as simple a comparison as (12), there is still much that can be deduced using (9), as shown by Jackson and Rogers (2007). While one cannot always directly solve

$$\theta = \sum_d \frac{P(d)\lambda\theta d^2}{\langle d \rangle (\lambda\theta d + 1)},$$

the right hand side is an increasing and convex function of d (when pulling out the $P(d)$), and it can be ordered when comparing different degree distributions in the sense of stochastic dominance (we will return to these sorts of comparisons in some of the models we discuss below). The interesting conclusion regarding steady-state infection rates is that they depend

on network structure in ways that are very different at low levels of the infection rate λ compared to high levels.

4.3 Graphical Games

While the above models provide some ideas about how social structure impacts diffusion, they are limited to settings where, roughly speaking, the probability that a given individual adopts a behavior is simply proportional to the infection rate of neighbors. Especially when it comes to things like opinions or adoption of a technology or buying a product, an individual’s decision can depend in much more complicated ways on the behavior of his or her neighbors. Such interaction naturally calls on game theory as a tool for modeling these richer interactions.

We start with static models of interactions on networks that allow for a rather general impact of peers’ actions on one’s own optimal choices, as suggested by much of the empirical literature on the topic.

The first model that explicitly examines games played on a network is the model of “graphical games” as introduced by Kearns, Littman, and Singh (2001), and analyzed by Kakade, Kearns, Langford, and Ortiz (2003), among others. In analogy to some of the epidemiological models discussed above in which an agent’s probability of catching a disease depended on the extent to which the disease had spread in their immediate neighborhood, the underlying premise in the graphical games models is that an agent’s *payoffs* depend on their own actions and the actions of their direct neighbors, as determined by the network of connections.¹⁵

Formally, the payoff structure underlying a graphical game is given as follows. The payoff to each player i when the profile of actions is $x = (x_1, \dots, x_n)$ is

$$u_i(x_i, x_{N_i(g)}),$$

where $x_{N_i(g)}$ is the profile of actions taken by the neighbors of i in the network g .

¹⁵There are also models of equilibria in social interactions, where players care about the play of certain other groups of players. See Glaeser and Scheinkman (2000) for an overview.

Most of the empirical applications discussed earlier entailed agents responding to neighbors' actions in roughly one of two ways. In some contexts (such as adopting a new product or new agricultural grain, joining a criminal network, or joining the workforce) it appears that when neighbors choose a particular action, one best responds with a similar action and payoffs exhibit strategic complementarities. In other contexts (such as experimentation on a new drug, or contribution to a public good) it seems that one's best response is negatively correlated with neighbors action choices and there is strategic substitutability. The graphical games environment allow for the analysis of both types of setups, as the following example (taken from Galeotti, Goyal, Jackson, Vega-Redondo, and Yariv, 2007) illustrates.

Example 1 (Payoffs Depend on the Sum of Actions) Player i 's payoff function when he or she chooses x_i and her k neighbors choose the profile (x_1, \dots, x_k) is:

$$u_i(x_i, x_1, \dots, x_k) = f\left(x_i + \lambda \sum_{j=1}^k x_j\right) - c(x_i), \quad (13)$$

where $f(\cdot)$ is non-decreasing and $c(\cdot)$ is a "cost" function associated with own effort (more general, but much in the spirit of (2)). The parameter $\lambda \in \mathbb{R}$ determines the nature of the externality across players' actions. The shape and sign of λf determine the effects of neighbors' action choices on one own's optimal choice. In particular, the example yields (strict) strategic substitutes (complements) if, assuming differentiability, $\lambda f''$ is negative (positive).

There are several papers that analyze graphical games for particular choices of f and λ . To mention a few examples, the case where f is concave, $\lambda = 1$, and $c(\cdot)$ is increasing and linear corresponds to the case of information sharing as a local public good studied by Bramoullé and Kranton (2007), where actions are strategic substitutes. In contrast, if $\lambda = 1$, but f is convex (with $c'' > f'' > 0$), we obtain a model with strategic complements, as proposed by Goyal and Moraga-Gonzalez (2001) to study collaboration among local monopolies. In fact, the formulation in (13) is general enough to accommodate a numerous further examples in the literature such as human capital

investment (Calvo-Armengol and Jackson (2008)), crime networks (Ballester, Calvo-Armengol, and Zenou (2006)), some coordination problems (Ellison (1993)), and the onset of social unrest (Chwe (2000)).

The computer science literature has focused predominantly on the question of when an efficient (polynomial-time) algorithm can be provided to *compute* Nash equilibria of graphical games.

The economics literature has concentrated on characterizing equilibrium outcomes for particular applications (see example above), or describing general comparative statics, with respect to location within a given network and with respect to the network structure itself, matching to some of the empirical insights.

Information players hold regarding the underlying network (namely, whether they are fully informed of the entire set of connections in the population, or only of connections in some local neighborhood) ends up playing a crucial role in the scope of predictions network interactions generate. Importantly, graphical games, in which agents have complete information about the networks in place, suffer from inherent multiplicity problems, as clearly illustrated by one of the manifestations of (13) analyzed by Bramoulle and Kranton (2007), and illustrated in the following example:

Example 2 (Multiplicity – Complete Information) Suppose that in (13), we take $f(x) \equiv$

x , $c(x) \equiv c$, where $0 < c < 1$, and $\lambda = 1$. This game, often labeled the *best-shot game*, may be viewed as a game of local public-good provision. Each agent would choose the action 1 (say, experiment with a new grain) if they were alone (or no one else experimented), but would prefer one of their neighbors to experience the cost c that the action 1 entails (when experimentation is observed publicly).

Note that, since $c < 1$, in any Nash equilibrium, for any player i with k neighbors, it must be the case that one of the agents in the neighborhood chooses the action 1. That is, if the chosen profile is (x_1, \dots, x_k) , then $x_i + \sum_{j=1}^k x_j \geq 1$. In fact, there is a very rich




Figure 3: Multiplicity of Equilibria with Complete Information

set of equilibria in this game. To see this, consider a star network and note that there exist two equilibria, one in which the center chooses 0 and the spokes choose 1, and a second equilibrium in which the spoke players choose 0 while the center chooses 1. Figure 3 illustrates these two equilibria. In the first, depicted in the left panel of the Figure, the center earns more than the spoke players while in the second equilibrium (in the right panel) it is the other way round.

The above example illustrates that even in the simplest network structures equilibrium multiplicity may arise and the relation between network architecture, equilibrium actions, and payoffs exhibits no systematic pattern.

4.4 Network Games

While the complete information regarding the structure of the social network imposed in graphical game models may be very sensible when the relevant network of agents is small, in large groups of agents (such as a country's electorate, the entire set of corn growers in

the 50's, the world-wide web, or academic economists), it is very likely that individuals have noisy perceptions of their network's architecture. Technically, as the discussion above stressed, complete information poses many challenges in models that aspire to predict outcomes because of the wide occurrence of equilibrium multiplicity.

A *network game* is a modification of a graphical game in which agents can have private and incomplete information regarding the realized social network at place. We describe here the setup corresponding to that analyzed by Galeotti, Goyal, Jackson, Vega-Redondo, and Yariv (2007) and Jackson and Yariv (2005, 2007), restricting attention to binary action games.¹⁶

Uncertainty is operationalized by assuming the network is determined according to some random process yielding our distribution over agents' degrees, $P(d)$, which is common knowledge. Each player i has d_i interactions, but does not know how many interactions each neighbor has. Thus, each player know their local neighborhood (the number of direct neighbors they have), but only the distribution of links in the remaining population.

Consider now the following utility specification, a generalization of (2). Agent i has a cost of choosing 1, denoted c_i . Costs are randomly and independently distributed across the society, according to a distribution F^c . Normalize the utility from the action 0 to 0 and let the benefit of agent i from action 1 be denoted by $v(d_i, x)$, where d_i is i 's degree and she expects each of her neighbors to independently choose the action 1 with probability x . Agent i 's added payoff from adopting behavior 1 over sticking to the action 0 is then $v(d_i, x) - c_i$.

This captures how the number of neighbors that i has, as well as their propensity to choose the action 1, affects the benefits from adopting 1. In particular, i prefers to choose the action 1 if

$$c_i \leq v(d_i, x). \tag{14}$$

This is a simple cost-benefit analysis generalizing Granovetter (1978)'s setup in that benefits can now depend on one's own degree (so that the underlying network is accounted

¹⁶There are also other variations, such as by Galeotti and Vega-Redondo (2005) and Sundararajan (2006), who study specific contexts, compatible with particular utility specifications.

for). Let $F(d, x) \equiv F^c(v(d, x))$. In words, $F(d, x)$ is the probability that a random agent of degree d chooses the action 1 when anticipating that each neighbor will choose 1 with an independent probability x .

Note that $v(d, x)$ can correspond to very different utilities – ranging from step functions (much like Granovetter, 1978) to situations in which agents’ payoffs depend on the expected number of neighbors adopting, dx .

Existence of symmetric Bayesian equilibria follows standard arguments. In cases where v is non-decreasing in x for each d , it is a direct consequence of Tarski’s Fixed Point Theorem. In fact, in this case, there exists an equilibrium in pure strategies. In other cases, provided v is continuous in x for each d , a fixed point can still be found by appealing to standard theorems (e.g., Kakutani) and admitting mixed strategies.¹⁷

Atomic Costs. Suppose first that F^c is atomic, and that all individuals experience the same cost $c > 0$ of choosing the action 1 (much like in Example 2 above). In that case, as long as $v(d, x)$ is monotonic in d (increasing or decreasing), equilibria are characterized by a threshold. Indeed, suppose $v(d, x)$ is increasing in d , then any equilibrium is characterized by a threshold d^* such that all agents of degree $d < d^*$ choose the action 0 and all agents of degree $d > d^*$ choose the action 1 (and agents of degree d^* may mix between the actions). In particular, notice that the type of multiplicity that appeared in Example 2 no longer occurs (provided degree distributions are not trivial). It is now possible to look at comparative statics of equilibrium behavior and outcomes using stochastic dominance arguments on the network itself. For ease of exposition, we illustrate this in the case of non-atomic costs (see Galeotti, Goyal, Jackson, Vega-Redondo, and Yariv, 2007, for the general analysis).

Non-atomic Costs. Consider the case in which F^c has no atoms. A simple equation is sufficient to characterize equilibria. Let x be the probability that a randomly chosen neighbor chooses the action 1. Then $F(d, x)$ is the probability that a random (best responding)

¹⁷In such a case, the best response correspondence (allowing mixed strategies) for any (d_i, c_i) as dependent on x is upper hemi-continuous and convex-valued. Taking expectations with respect to d_i and c_i , we also have a set of population best responses as dependent on x that is upper hemi-continuous and convex valued.

neighbor of degree d chooses the action 1. We can now proceed in a way reminiscent of the analysis of the SIS model. Recall that $\tilde{P}(d)$ denoted the probability that a random neighbor is of degree d (see equation (4)). It must be that

$$x = \phi(x) \equiv \sum_d \tilde{P}(d) F(d, x). \quad (15)$$

Again, a fixed point equation captures much of what occurs in the game. In fact, equation (15) characterizes equilibria in the sense that any symmetric equilibrium results in an x which satisfies the equation, and any x that satisfies the equation corresponds to an equilibrium where type (d_i, c_i) chooses 1 if and only if inequality (14) holds. Given that equilibria can be described by their corresponding x , we often refer to some value of x as being an “equilibrium.”

Consider any symmetric equilibrium generating a probability of x for a random neighbor to choose action 1. If v is non-decreasing in d , then the expected payoff of a degree $d + 1$ agent is $v(d + 1, x) \geq v(d, x)$ and so $F^c(v(d + 1, x)) \geq F^c(v(d, x))$ and agents with higher degrees choose 1 with weakly higher probabilities. Since an agent of degree $d + 1$ can imitate the decisions of an agent of degree d and gain at least as high a payoff. Thus, if v is non-decreasing (or, in much the same way, non-increasing) in d for each x , then any symmetric equilibrium entails agents with higher degrees choosing action 1 with weakly higher (lower) probability. Furthermore, agents of higher degree have higher (lower) expected payoffs.

Thinking back to our empirical insights regarding the higher propensities of adoption by more connected individuals, these observations suggest that affecting v (namely, increasing v) may prove beneficial in enlarging the volume of agents taking the action 1.

Much as in the analysis of the epidemiological models, the multiplicity of equilibria is determined by the properties of ϕ , which, in turn, correspond to properties of \tilde{P} and F . For instance,

- if $F(d, 0) > 0$ for some d in the support of P , and F is concave in x for each d , then there exists at most one fixed point, and

- if $F(d, 0) = 0$ for all d and F is strictly concave or strictly convex in x for each d , then there are at most two equilibria - one at 0, and possibly an additional one, depending on the slope of $\phi(x)$ at $x = 0$.¹⁸

In general, as long as the graph of $\phi(x)$ crosses the 45 degree line only once, there is a unique equilibrium (see Figure 4 below).¹⁹

The set of equilibria generated in such network games is divided into stable and unstable ones (those we have already termed in Section 3.2 as *tipping points*). The simple characterization given by (15) allows for a variety of comparative statics on fundamentals pertaining to either type of equilibria. In what follows, we show how these comparative statics tie directly to a simple strategic diffusion process. Indeed, it turns out there is a very useful technical link between the static and dynamic analysis of strategic interactions on networks.

4.5 Adding Dynamics – Diffusion and Equilibria of Network Games

One of the earliest contributions to the study of diffusion of strategic behavior allowing for general network architectures was done by Morris (2000). Morris (2000) considered coordination games played on networks. His analysis pertained to identifying social structures conducive to *contagion*, where a small fraction of the population choosing one action leads to that action spreading across the *entire* population. The main insight from Morris (2000) is that maximal contagion occurs when local interaction is sufficiently uniform and there is low neighborhood growth, i.e., the number of players of distance d does not grow exponentially in d .

The analysis of Morris (2000) can be thought of as identifying conditions on the underlying network guarantying uniqueness of a stable equilibrium consisting of all agents picking

¹⁸as is standard, the slope needs to be greater than 1 for there to be an additional equilibrium in the case of strict concavity and less than 1 in the case of strict convexity.

¹⁹Morris and Shin (2002, 2003) consider uncertainty on payoffs rather than on the underlying network. In coordination games, they show the class of payoff shocks that lead to a unique equilibrium. Heterogeneity in degrees combined with uncertainty plays a similar role in restricting the set of equilibria. In a sense, the analysis described here is a generalization in that it allows studying the impact of changes in a variety of fundamentals on the set of stable and unstable equilibria, regardless of multiplicity, in a rather rich environment. Moreover, the equilibrium structure can be tied to the network of underlying social interactions.

the same action.

In order to identify the full set of stable of equilibria using the above formalization, consider a diffusion process governed by best responses in discrete time (following Jackson and Yariv, 2005, 2007). At time $t = 0$, a fraction x^0 of the population is exogenously and randomly assigned the action 1, and the rest of the population is assigned the action 0. At each time $t > 0$, each agent, *including the agents assigned to action 1 at the outset*, best responds to the distribution of agents choosing the action 1 in period $t - 1$, accounting for the number of neighbors they have and presuming that their neighbors will be a random draw from the population.

Let x_d^t denote the fraction of those agents with degree d who have adopted behavior 1 at time t , and let x^t denote the link-weighted fraction of agents who have adopted the behavior at time t . That is, using the distribution of neighbors' degrees $\tilde{P}(d)$,

$$x^t = \sum_d \tilde{P}(d)x_d^t.$$

Then, as deduced before from equation (14), at each date t ,

$$x_d^t = F(d, x^{t-1}).$$

and therefore

$$x^t = \sum_d \tilde{P}(d)F(d, x^{t-1}) = \phi(x^{t-1}).$$

As we have discussed, any rest point of the system corresponds to a static Bayesian equilibrium of the system.

If payoffs exhibit complementarities, then convergence of behavior from any starting point is monotone, either upwards or downwards. In particular, once an agent switches behaviors, the agent will not want to switch back at a later date.²⁰ Thus, although these best responses are myopic, any eventual changes in behavior are equivalently forward-looking.

²⁰If actions are strategic substitutes, convergence may not be guaranteed for all starting points. However, whenever convergence is achieved, the rest point is an equilibrium, and the analysis can therefore be useful for such games as well.




Figure 4: The Effects of Shifting $\phi(x)$ Pointwise

Figure 4 depicts a mapping ϕ governing the dynamics. Equilibria, and resting points of the diffusion process, correspond to intersections of ϕ with the 45 degree line.

The figure allows an immediate distinction between two classes of equilibria that we discussed informally up to now. Formally, an equilibrium x is *stable* if there exists $\varepsilon' > 0$ such that $\phi(x - \varepsilon) > x - \varepsilon$ and $\phi(x + \varepsilon) < x + \varepsilon$ for all $\varepsilon' > \varepsilon > 0$. An equilibrium x is *unstable* or a *tipping point* if there exists $\varepsilon' > 0$ such that $\phi(x - \varepsilon) < x - \varepsilon$ and $\phi(x + \varepsilon) > x + \varepsilon$ for all $\varepsilon' > \varepsilon > 0$. In the figure, the equilibrium to the left is a tipping point, while the equilibrium to the right is stable.

The composition of the equilibria set hinges on the shape of the function ϕ . Furthermore, note that a pointwise shift of ϕ (as in the figure, to a new function $\tilde{\phi}$) shifts tipping points to the left and all stable points to the right, loosely speaking making adoption more likely. This simple insight allows for a variety of comparative statics.

For instance, consider an increase in the cost of adoption, manifested as a First Order Stochastic Dominance (FOSD) shift of the cost distribution F^c to \bar{F}^c . It follows immediately that:

$$\bar{\phi}(x) = \sum_d \tilde{P}(d) \bar{F}^c(v(d, x)) \leq \sum_d \tilde{P}(d) F^c(v(d, x)) = \phi(x)$$

and the increase in costs corresponds to an increase of the tipping points and decrease of the stable equilibria (one by one). Intuitively, increasing the barrier to choosing the action 1 leads to a higher fraction of existing adopters necessary to get the action 1 to spread even more.

This formulation also allows for an analysis that goes beyond graphical games regarding the social network itself, using stochastic dominance arguments (following Jackson and Rogers, 2007, and Jackson and Yariv, 2005, 2007). For instance, consider an increase in the number of expected neighbors each agent has. That is, suppose \tilde{P}' FOSD \tilde{P} and, for illustration, assume that $F(d, x)$ is non-decreasing in d for all x . Then, by the definition of FOSD,

$$\phi'(x) = \sum_d \tilde{P}'(d) F(d, x) \geq \sum_d \tilde{P}(d) F(d, x) = \phi(x),$$

and under P' tipping points are lower and stable equilibria are higher.

Similar analysis allows for comparative statics regarding the distribution of links, by simply looking at Mean Preserving Spreads (MPS) of the underlying degree distribution.²¹

Going back to the dynamic path of adoption, we can generalize the insights derived from the Granovetter (1978) model. Namely, whether adoption paths track an S -shaped curve now depends on the shape of ϕ , and thereby on the shape of both the cost distribution F and agents' utilities. Indeed, S -shaped adoption could be observed if ϕ is S -shaped in between a tipping point and the next higher equilibrium.

²¹In fact, Jackson and Yariv (2007) illustrate that if $F(d, x)$ is non-decreasing and convex, then power, Poisson, and regular degree distributions with identical means generate corresponding values of ϕ^{power} , $\phi^{Poisson}$, and $\phi^{regular}$ such that

$$\phi^{power}(x) \geq \phi^{Poisson}(x) \geq \phi^{regular}(x)$$

for all x , thereby implying a clear ranking of the tipping points and stable equilibria corresponding to each type of network.

5 Closing Notes

There is now a growing body of theoretical work studying the potential impacts of interactions that occur on a network of connections. This work builds on the empirical observations of peer influence and generates a rich set of individual and aggregate predictions. The two main insights that have been shown consistently in real-world data pertain to the higher propensities of contagion (of a disease, an action, or behavior) in more highly connected individuals, as well as the aggregate *S*-shape of diffusion curves. The theoretical analysis opens the door to many other positive results, e.g., those regarding comparative statics *across* networks, payoffs, and cost distributions (when different actions vary in costs). Future experimental and field data will hopefully complement some of these theoretical insights.

The principle shortcoming of the analyses described in this chapter is that the foundation for modeling the underlying network is a random graph. Nonetheless, the empirical literature illustrates again and again the presence of homophily, people's tendency to associate with other individuals who are similar to themselves. In terms of infections or diffusion of behaviors, that means that people who are more likely to, say, be immune may be more likely to be connected to each other, and people who are more likely to be susceptible to infection are more likely to be connected to each other.²² Substantially enriching the interaction structure in that direction would be crucial for deriving more accurate diffusion predictions. This is an active area of current study (e.g., see Currarini, Jackson, and Pin, 2007, Peski, 2007, and Baccara and Yariv, 2008).

Ultimately, the formation of a network and the strategic interactions that occur amongst individuals is a two-way street. Indeed, one could imagine that particular structures will be more likely to be observed when they are welfare enhancing. Contemplating endogenous formation of networks, allowing for impacts of the strategic interactions to be engaged in, is an interesting direction for future work, both empirical and theoretical.

²²The mechanism through which this occurs can be rooted in more fundamental attributes such as risk aversion. Risk averse individuals may connect to one another and be more prone to protect themselves against diseases by, e.g., getting immunized.

References

- [1] Baccara, M. and Yariv, L. (2008), “Preference for Similarity and Information Collection in Groups,” mimeo.
- [2] Bailey, N.T.J. (1975), *The Mathematical Theory of Infectious Diseases and Its Applications*, London Griffin.
- [3] Ballester, C., Calvó-Armengol, A., and Zenou, Y. (2006), “Who’s who in networks. Wanted: The Key Player,” *Econometrica*, 74, 5, 1403-1417.
- [4] Bandiera, O., Barankay, I. and Rasul, I. (2006), “How to Make Friends and Influence Them: Evidence on the Formation of Social Networks and Peer Effects,” mimeo.
- [5] Bass, F. (1969), “A new product growth model for consumer durables,” *Management Science*, 15(5), 215-227
- [6] Bayer, P., S. Ross, and Topa, G. (2006), “Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes,” mimeo.
- [7] Bearman, P. S., Moody, J., and Stovel, K. (2004), “Chains of Affection: The Structure of Adolescent Romantic and Sexual Networks,” *American Journal of Sociology*, 110(1), 44-91.
- [8] Bollobás, B. (1980): “A probabilistic proof of an asymptotic formula for the number of labelled regular graphs”, *European Journal of Combinatorics* **1**, 311-16.
- [9] Bramoullé, Y. and Kranton, R. (2007), “Strategic Experimentation in Networks,” *Journal of Economic Theory*, forthcoming.
- [10] Calvó-Armengol, A. and Jackson, M. O. (2008), “Like Father, Like Son: Labor Market Networks and Social Mobility,” *mimeo*, Stanford and Universitat Autònoma de Barcelona.

- [11] Christakis, N. A. and Fowler, J. H. (2007), “The Spread of Obesity in a Large Social Network over 32 Years,” *The New England Journal of Medicine*, 357(4), 370-379.
- [12] Chwe, M. S-Y. (2000), “Communication and Coordination in Social Networks,” *The Review of Economic Studies*, 67, 1-16.
- [13] Coleman, J.S., Katz, E., and Menzel, H. (1966), *Medical Innovation: A Diffusion Study*, Bobbs-Merrill: Indianapolis.
- [14] Colizza, V., Barrat, A., Barthélemy, M., and Vespignani, A. (2006), “The role of the airline transportation network in the prediction and predictability of global epidemics,” *Proceedings of the National Academy of Sciences USA*, 103, 2015-2020.
- [15] Colizza, V., Barrat, A., Barthelemy, M., and Vespignani, A. (2007), “Predictability and epidemic pathways in global outbreaks of infectious diseases: the SARS case study,” *BMC Medicine*, 5, 34.
- [16] Conley, T. G. and Topa, G. (2002), “Socio-Economic Distance and Spatial Patterns in Unemployment,” *Journal of Applied Econometrics*, 17(4), 303-27.
- [17] Conley, T. G. and Udry, C. (2005), “Learning About a New Technology: Pineapple in Ghana,” *mimeo*.
- [18] Currarini, S., Jackson, M. O., and Pin, P. (2007), “An Economic Model of Friendship: Homophily, Minorities, and Segregation,” *mimeo*.
- [19] Ellison, G. (1993), “Learning, Local Interaction, and Coordination,” *Econometrica*, **61**, 1047-1071.
- [20] Erdős, P. and Rényi, A. (1959), “On random graphs,” *Publicationes Mathematicae*, 6, 290–297.
- [21] Erdős, P. and Rényi, A. (1960), “On the evolution of random graphs,” *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5, 17–61.

- [22] Feick, L. F. and Price, L. L. (1987), “The Market Maven: A Diffuser of Marketplace Information,” *Journal of Marketing*, **51(1)**, 83-97.
- [23] Foster, A. D. and Rosenzweig, M. R. (1995), “Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture,” *Journal of Political Economy*, **103(6)**, 1176-1209.
- [24] Galeotti, A., Goyal, S., Jackson, M. O., Vega-Redondo, F., and Yariv, L. (2007), “Network Games,” *mimeo*.
- [25] Galeotti, A. and Vega-Redondo, F. (2005), “Strategic Analysis in Complex Networks with Local Externalities,” *HSS-Caltech Working Paper 1224*.
- [26] Glaeser, E. L. and Scheinkman, J. A. (2000), “Non-market Interactions,” NBER Working Paper Number 8053.
- [27] Glaeser, E. L., Sacerdote, B., and Scheinkman, J. A. (1996), “Crime and Social Interactions,” *The Quarterly Journal of Economics*, 111(2), 507-548.
- [28] Goeree, J. K., McConnell, M. A., Mitchell, T. , Tromp, T., and Yariv, L. (2007), “Linking and Giving among Teenage Girls,” *mimeo*, Caltech.
- [29] Goyal, S. and Moraga-Gonzalez, J. L. (2001), “R&D Networks,” *Rand Journal of Economics*, **32(4)**, 686-707.
- [30] Granovetter, M. (1973), “The Strength of Weak Ties,” *American Journal of Sociology*, 78(6), 1360-1380.
- [31] Granovetter, M. (1978), “Threshold Models of Collective Behavior,” *American Journal of Sociology*, 83(6), 1420-1443.
- [32] Granovetter, M. (1985), “Economic action and social structure: the problems of embeddedness,” *American Journal of Sociology*, 91(3), 481-510.

- [33] Granovetter, M. (1995), *Getting a Job: A Study in Contacts and Careers*, University of Chicago Press.
- [34] Griliches, Z. (1957), “Hybrid corn: an exploration of the economics of technological change,” *Econometrica*, 25, 501-522.
- [35] Huckfeldt, R. R. and Sprague, J. (1995), *Citizens, Politics and Social Communication*, Cambridge Studies in Public Opinion and Political Psychology.
- [36] Jackson, M.O. (2008), *Social and Economic Networks*, Princeton University Press: NJ.
- [37] Jackson, M. O. and Rogers, B. W. (2007), “Relating Network Structure to Diffusion Properties through Stochastic Dominance,” *The B.E. Press Journal of Theoretical Economics*, Vol. 7, Iss. 1(Advances), Art. 6, pp. 1-13.
- [38] Jackson, M. O. and Yariv, L. (2005), “Diffusion on Social Networks,” *Economie Publique*, 16(1), 69-82.
- [39] Jackson, M. O. and Yariv, L. (2007), “Diffusion of Behavior and Equilibrium Properties in Network Games,” *The American Economic Review (Papers and Proceedings)*, 97(2), 92-98.
- [40] Kakade, S., Kearns, M., Langford, J., and Ortiz, L. (2003), “Correlated Equilibria in Graphical Games,” ACM Conference on Electronic Commerce, New York.
- [41] Kandel, D. B. (1978), “Homophily, Selection, and Socialization in Adolescent Friendships,” *The American Journal of Sociology*, 84(2), 427-436.
- [42] Katz, E. and Lazarsfeld, P. F. (1955), *Personal influence: The part played by people in the flow of mass communication*, Glencoe, IL: Free Press.
- [43] Kearns, M., Littman, M., and Singh, S. (2001), Graphical Models for Game Theory, in Jack S. Breese, Daphne Koller (eds.), *Proceedings of the 17th Conference on Uncer-*

- tainty in Artificial Intelligence*, 253-260, San Francisco: Morgan Kaufmann University of Washington, Seattle, Washington, USA, August 2-5, 2001.
- [44] Kermack, W. O. and McKendrick, A. G. (1927), “A Contribution to the Mathematical Theory of Epidemics,” *Proceedings of the Royal Society of London*, A 115, 700-721, 1927.
- [45] Kosfeld, M. (2003), “Network Experiments,” *mimeo*.
- [46] Lazarsfeld, P. E., Berelson, B., & Gaudet, G. (1944), *The people’s choice: How the voter makes up his mind in a presidential campaign*, New York: Duell, Sloan and Pearce.
- [47] Leider, S., Mobius, M. M., Rosenblat, T., and Do, Q-A. (2007), “How Much is A Friend Worth? Directed Altruism and Enforced Reciprocity in Social Networks,” *mimeo*.
- [48] Leskovec, J., Adamic, L.A., and Huberman, B.A. (2007) “The dynamics of viral marketing,” *ACM Transactions on the Web (TWEB) archive*, 1:1, Article No. 5.
- [49] Lopez-Pintado, D. (2004) “Contagion in Complex Networks,” *mimeo*.
- [50] Mahajan, V. and Peterson, R. A. (1985), *Models for Innovation Diffusion (Quantitative Applications in the Social Sciences)*, Sage Publications, Inc.
- [51] McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001), “Birds of a Feather: Homophily in Social Networks,” *Annual Review of Sociology*, 27, 415-444.
- [52] Merton, R. (1948), “Patterns of Influence: a study of interpersonal influence and of communications behavior in a local community,” In Lazarsfeld, Paul and Stanton, Frank (eds.), *Communication Research*, NewYork: Harper and Brothers. pp180-215.
- [53] Morris, S. (2000), “Contagion,” *The Review of Economic Studies*, 67(1), 57-78.
- [54] Morris, S. and Shin, H. (2002), “Heterogeneity and Uniqueness in Interaction Games,” *mimeo*.

- [55] Morris, S. and Shin, H. (2003), “Global Games: Theory and Applications,” in *Advances in Economics and Econometrics* (Proceedings of the Eighth World Congress of the Econometric Society), edited by M. Dewatripont, L. Hansen and S. Turnovsky. Cambridge: Cambridge University Press, 56-114.
- [56] Nair, H., Manchanda, P., and Bhatia, T. (2006), “Asymmetric Effects in Physician Prescription Behavior: The Role of Opinion Leaders,” mimeo.
- [57] Pastor-Satorras, R. and Vespignani, A. (2000) “Epidemic Spreading in Scale-Free Networks,” *Physical Review Letters*, 86:14, 3200-3203.
- [58] Pastor-Satorras, R. and Vespignani, A. (2001) “Epidemic dynamics and endemic states in complex networks,” *Physical Review E*, 63, 066117.
- [59] Pesendorfer, W. (1995), “Design Innovation and Fashion Cycles,” *The American Economic Review*, 85(4), 771-792.
- [60] Peski, M. (2007), “Complementarities, Group Formation, and Preferences for Similarity,” mimeo.
- [61] Rees, A. and Schultz, G. P. (1970), *Workers and Wages in an Urban Labor Market*, University of Chicago Press.
- [62] Rogers, E. M. (1995), *Diffusion of Innovations*, Free Press.
- [63] Ryan, B. and Gross, N. C. (1943) “The diffusion of hybrid seed corn in two Iowa communities,” *Rural Sociology*, 8, 15–24.
- [64] Schelling, T. C. (1978), *Micromotives and Macrobehavior*, W. W. Norton.
- [65] Solomonoff, R. and Rapoport, A. (1951), “Connectivity of Random Nets,” *Bulletin of Mathematical Biophysics*, **13**, 107-117.
- [66] Sundararajan, A. (2005), “Local network effects and network structure,” mimeo.

- [67] Tarde, G. (1903), *Les Lois de L'Imitation: Etude Sociologique*, Elibron Classics, translated to English in *The Laws of Imitation*, H. Holt and company.
- [68] Topa, G. (2001), "Social Interactions, Local Spillovers and Unemployment," *Review of Economic Studies*, 68, 2, 261-295.
- [69] Tucker, C. (2007), "Formal and Informal Influence in the Adoption of Technologies with Network Effects," mimeo.
- [70] Van den Bulte, C. and Lilien, G. L. (2001), "Medical Innovation Revisited: Social Contagion versus Marketing Effort," *American Journal of Sociology*, 106 (5), 1409-1435.
- [71] Young, P. (2006), "The Spread of Innovations by Social Learning," mimeo.