CHAPTER 9

Decisions, Behavior, and Games on Networks

Peers exert enormous influence on human behavior. It is easy to cite examples, ranging from which products we buy, whether we engage in criminal activities, how much education we pursue, to which profession we choose. There is a vast literature on the subject, including research by social psychologists, sociologists, researchers in education, and economists. We have already seen examples of studies that touch on these issues, and there is also a good bit known about statistical issues of identifying peer effects (see Chapter 13). The purpose of this chapter is to provide the foundations for understanding how the structure of social networks influences behavior. For example, if we change the network of social interactions, how does behavior change? This issue has a rich history in sociology and has more recently emerged in economics and computer science. Some aspects of this were touched on in Chapters 7 and 8 on diffusion and learning, but there are many situations in which social influences involve human decision making rather than pure contagion or updating. The focus in this chapter complements those earlier chapters by expanding the analysis to explicitly account for how individuals' decisions and strategies are influenced by those of their neighbors.

The main complication in the analysis of this chapter compared to the study of diffusion and learning is that behavior depends in more complicated ways on what neighbors are doing. For example, if an individual is choosing a piece of software or some other product and wants it to be compatible with a majority of neighbors, then this is a coordination game, and behavior can change abruptly, depending on how many neighbors are taking a certain action. It might also be that an individual only wants to buy a product or make an investment when his or her neighbors do not. These interactive considerations require game-theoretic reasoning, adapted and extended to a network setting.

The chapter begins by showing how a variation on the ideas encountered in the learning and diffusion analyses can be used to study behaviors. In particular, I start with a model in which people react to their neighbors in a way that can be captured probabilistically, predicting their actions as a function of the distribution of their neighbors' play. I then introduce game-theoretic settings, where richer behaviors are studied. There are two main types of situations considered that capture many of the applications of interest. One type deals with strategic complementarities, as in choosing compatible technologies or pursuing education, and players' incentives to take a given action increase as more neighbors take that action. This type leads to nice properties of equilibria, and we can deduce quite a bit about how players' strategies vary with their position in a network and how overall behavior in the society responds to network structural changes. The second type of strategic interaction considered has the opposite incentive structure: that of strategic substitutes. In it, players can "free ride" on the actions of their neighbors, such as in gathering information or providing certain services, and a player's incentive to take a given action decreases as more neighbors take that action. This dynamic leads to quite different conclusions about player behavior and equilibrium structure. I also discuss models that are designed to capture the dynamics of behavior.

9.1 ▪ Decisions and Social Interaction

Let me begin by discussing an approach to modeling interactive behavior that builds on the tools from Chapters 7 and 8 on diffusion and learning. To model decisions in the face of social interaction, one needs to characterize how a given individual behaves as a function of the actions of his or her neighbors. For instance, will the individual want to buy a new product if some particular subset of his or her neighbors are buying the product? Generally, the answer to such a question depends on a series of characteristics of the individual, the product, the alternative products, and on the set of neighbors buying the product.[1] As a useful starting point, we can think of this process as being stochastic, with a probability of the individual taking a given action depending on the actions chosen by the neighbors.

9.1.1 A Markov Chain

To fix ideas and provide a benchmark, consider a setting in which interaction is symmetric in the sense that any individual can be influenced by any (or every) other, and the particular identities of neighbors are not important, just the relative numbers of agents taking various actions. Each individual chooses one of two actions. This choice might be between smoking or not, adopting one technology or another, going to one park or another, voting or not, voting for one candidate or another, and so on.

The two actions are labeled as 0 and 1, and time evolves in discrete periods, $t \in \{1, 2, \ldots\}$. The state of the system is described by the number of individuals who are taking action 1, denoted by s_t, at the end of period t.

1. The implications of interdependencies in consumer choices have been studied in a variety of contexts, including the implications for how firms might price goods (e.g., see Katz and Shapiro [382]) and implications for increasing returns (e.g., see Arthur [21] and Romer [566]). For more background on consumers' behavior and network externalities, see the survey by Economides [216].

In some applications, if the state of the system is $s_t = s$ at time t, then there is a well-defined probability that the system will be in state $s_{t+1} = s'$ at time $t + 1$. That is, if we know how many people are taking action 1 at date t, then there is a well-defined distribution over the number of people taking action 1 at date $t + 1$. This might be deterministic or random, depending on what is assumed about how people react to others. For example, each person might look at the state at the end of time t and then choose the action that leads to the greatest benefit for himself, presuming that others will act as at time t; for instance if we pick one person at random from the society and ask him to update his choice. All individuals may respond, or each one may respond probabilistically. What is critical is that there are well-defined probabilities of being in each state tomorrow as a function of the state today. Let Π be the $n \times n$ matrix describing these *transition probabilities* with the entry in row s and column s' being

$$\Pr(s_{t+1} = s' \mid s_t = s).$$

This results in a (finite-state) Markov chain (recalling definitions from Section 4.5.8). If the Markov chain is irreducible and aperiodic, then it has a steady-state distribution described by the vector $\mu = (\mu_0, \ldots, \mu_n)$, where μ_s is the probability of state s, or that s players choose action 1. Thus when behavior can be described by a Markov chain, we have sharp predictions about behavior over the (very) long run. Let us now examine some applications of this reasoning.

9.1.2 Individual-by-Individual Updating

Consider a setting in which at the beginning of a new period, one individual is picked uniformly at random[2] and updates her action based on the current number of people in the society taking action 1 or 0.[3] How an individual updates her action depends on the state of the system. In particular, let p_s denote the probability that the individual chooses action 1 conditional on s out of the other $n - 1$ agents choosing action 1. This form of updating could be a form of best-reply behavior, in which an individual (myopically) chooses an action that gives the highest payoff given the current actions of other individuals. So, for instance, the individual decides on whether to go / to stay at the beach based on how many others are there. It could also be something similar to the contagion of a disease (see Section 7.2), but now the transition between infection and susceptible is a richer function of the state of the system.

The transition probabilities from one state to another are completely determined by the vector $p = (p_0, \ldots, p_{n-1})$. For example the probability of going from s to $s + 1$ is the probability that one of the $n - s$ players choosing 0 out of the n

2. The method is not important, as one can use any method of picking the individual as long as it is Markovian. It could weight different individuals differently and be dependent on the state.
3. Note that the interpretation of dates here is flexible. Effectively, time simply keeps track of the moments at which some agent makes a decision and need not correspond to any sort of calendar time. The arrival process of when decisions are made can be quite general, with the main feature here being that only one agent is updating at any given moment.

players in total is selected to update, and then that player has a probability of p_s of selecting action 1, and so the probability is $\frac{n-s}{n} p_s$. The full list of transition probabilities is

$$\begin{aligned}
\Pr\left(s_{t+1} = s+1 \mid s_t = s\right) &= \frac{n-s}{n} p_s, \text{ for } 0 \le s \le n-1, \\
\Pr\left(s_{t+1} = s-1 \mid s_t = s\right) &= \frac{s}{n}\left(1 - p_{s-1}\right), \text{ for } 1 \le s \le n, \\
\Pr\left(s_{t+1} = s \mid s_t = s\right) &= \frac{n-s}{n}\left(1 - p_s\right) + \frac{s}{n} p_{s-1}, \text{ for } 0 \le s \le n, \\
\Pr\left(s_{t+1} = s' \mid s_t = s\right) &= 0 \text{ if } s' \notin \{s-1, s, s+1\}.
\end{aligned}$$

This Markov chain has several nice properties. Provided that $1 > p_s > 0$ for each s, any state is eventually reachable from any other state, and so the Markov chain is irreducible. Moreover, there is a chance of staying in any state, and so the Markov chain is aperiodic. Thus it has a unique steady-state distribution over states.

In this setting, the Markov chain takes a particularly simple form as only one individual is changing actions at a time. At steady state, the probability of ending in state s is simply the probability of being in an adjacent state $s-1$ or $s+1$ and then getting one more or less individual to choose action 1, or else starting at s and staying there. Thus

$$\begin{aligned}
\mu_0 &= \mu_0(1 - p_0) + \mu_1\left(\frac{1}{n}\right)(1 - p_0), \\
\mu_s &= \mu_{s-1}\left(\frac{n-(s-1)}{n}\right) p_{s-1} + \mu_{s+1}\left(\frac{s+1}{n}\right)(1 - p_s) \\
&\quad + \mu_s\left(\frac{n-s}{n}\right)(1 - p_s) + \mu_s\left(\frac{s}{n}\right) p_{s-1}, \\
\mu_n &= \mu_{n-1}\left(\frac{1}{n}\right) p_{n-1} + \mu_n p_{n-1}.
\end{aligned}$$

Solving this system leads to

$$\frac{\mu_{s+1}}{\mu_s} = \left(\frac{n-s}{s+1}\right)\left(\frac{p_s}{1-p_s}\right) \tag{9.1}$$

for all $0 \le s \le n-1$; which, together with $\sum_s \mu_s = 1$, completely determines the solution.

Ants, Investment, and Imitation To see an application of such a Markov model of behavior, consider a special case due to Kirman [397] in which individuals imitate one another. In particular, Kirman's work was motivated by the observation that ants tend to herd on the food sources that they exploit, even when faced with equally useful sources. Moreover, ants also switch which source they exploit more intensively. Kirman discusses similar patterns of behavior in human investment and other forms of imitation. He models this behavior by considering a dynamic such that at each time interval an individual is selected uniformly at random from the population. With a probability ε the individual flips a coin to choose action 0 or 1, and with a probability $1-\varepsilon$ the individual selects another individual uniformly at random and mimics his or her action.

That model is a special case of the framework above, in which[4]

$$p_s = \frac{\varepsilon}{2} + (1-\varepsilon)\,\frac{s}{n-1}.$$

Then from (9.1) it follows that

$$\frac{\mu_{s+1}}{\mu_s} = \frac{n-s}{s+1}\left(\frac{(n-1)\varepsilon + 2s\,(1-\varepsilon)}{2(n-1)-(n-1)\varepsilon - 2s\,(1-\varepsilon)}\right). \tag{9.2}$$

Given the symmetry of this setting, the long-run probability (starting from a random draw of actions) of any individual choosing either action is $1/2$. If $\varepsilon = 0$, then this system is one of pure imitation, and eventually the system is absorbed into the state 0 or n and then stays there forever. If $\varepsilon = 1$ then the individuals simply flip coins to choose their actions irrespective of the rest of society. Then

$$\frac{\mu_{s+1}}{\mu_s} = \frac{n-s}{s+1}$$

and the system has a pure binomial distribution with a parameter of $1/2$, so that the probability of having exactly s individuals choosing action 1 at a given date in the long run is $\mu_s = \binom{n}{s}\left(\frac{1}{2}\right)^n$. In that case, the extreme states are less likely, and the most likely state is that half the society is choosing action 1 and half is choosing action 0.[5]

To obtain a uniform distribution across social states it must be that $\mu_{s+1} = \mu_s$ for each s. It is easily checked from (9.2) that this holds when $\varepsilon = 2/(n+1)$.[6] To model the behavior of the ants with tendencies to herd on action choice, the high and low states need to be more likely than the middle states. This requires that individuals pay more attention to others than in the situation just examined, in which the distribution across states is uniform. Thus the probability that an individual ignores society and flips a coin must be $\varepsilon < 2/(n+1)$. Then behavior is sensitive enough to the state of the society so that the more extreme states, in which all individuals follow one action or the other ($s = 0$ and $s = n$), are more heavily weighted. This distribution is pictured in Figure 9.1 for several different values of ε. Thus, to

---

4. Kirman also allows for the probability that an individual does not change his or her action, but that is equivalent in terms of long-run distributions, as it simply slows the system down and can be thought of as changing the length of the time period.

5. Note that each configuration of choices across individuals is equally likely, but there are many more configurations (keeping track of players' labels) in which half the society is choosing 1 and half 0 than there are configurations where all are choosing 0.

6. This is slightly different from Kirman's expression since ε here is the chance that an agent flips a coin, whereas it is the chance that an agent changes actions in Kirman's labeling.

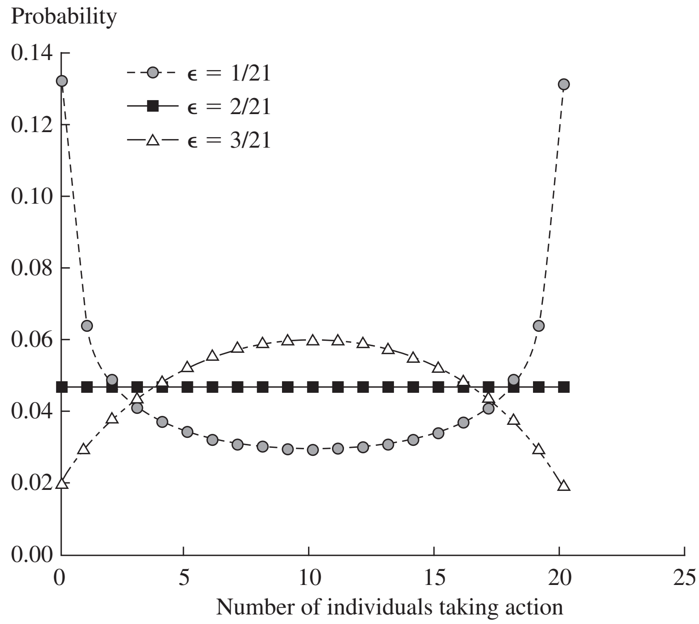

FIGURE 9.1 Kirman's [396] ant-imitation model for three levels of random behavior.

model herding, one needs to have a sufficiently high probability of imitating other agents so that the natural tendency toward even mixing is overturned.

Sensitivity to Societal Action and Herding As I now show, Kirman's result about the relationship between social sensitivity and herding can be formalized and is true for more general processes than the pure imitation process described above. To see this, consider a class of processes that treat actions 0 and 1 symmetrically. That is, suppose that $p_s = 1 - p_{n-1-s}$, so that the chance of choosing action 1 conditional on s out of the other $n-1$ agents choosing action 1 is the same as the chance of choosing action 0 conditional on s out of the other $n-1$ agents choosing action 0.

Say that a social system $p' = (p'_0, \ldots, p'_n)$ is *more socially sensitive* than another social system $p = (p_0, \ldots, p_n)$ if $p'_s \geq p_s$ when $s > (n-1)/2$ and $p'_s \leq p_s$ when $s < (n-1)/2$, with at least one strict inequality. Thus an individual in a society described by p' is more likely to choose the same action as the majority of the population than an individual in a society described by p. Note that this definition does not require that an individual want to match the majority of a society. It admits processes in which $p'_s < 1/2$ when $s > n/2$ so that an individual is actually choosing against the current of the society. The social sensitivity comparison is a relative comparison between two societies.

Proposition 9.1 *If two different societies, described by p and p', each treat actions 0 and 1 symmetrically and the process p' is more socially sensitive than p, then the steady-state distribution over numbers of agents taking action 1,*

$(\mu'_0, \ldots, \mu'_n)$ *corresponding to* p', *is a mean-preserving spread of the steady-state distribution* $(\mu_0, \ldots, \mu_n)$ *corresponding to* p. *In fact, there exists* $\overline{s} > n/2$ *such that* $\mu'_s > \mu_s$ *if and only if* $s > \overline{s}$ *or* $s < n - \overline{s}$.

The proof of Proposition 9.1 can be deduced from (9.1) and is left as Exercise 9.2. This result is not surprising, as it states that increasing the extent to which an individual's choice of action matches that of a majority of the society then increases the extent to which the society tends to extremes, in terms of spending more time in states in which higher concentrations of individuals choose the same action. Nevertheless, such a result provides insight as to the type of behavior that leads to herding (or lack thereof). This is a different sort of herding than the consensus formation we saw in our discussion of learning in Chapter 8,[7] as the society here oscillates between the actions over time, but with specific patterns of either herding to one action or the other, or else splitting among the two.[8]

Calvó-Armengol and Jackson [130] use such a model to study social mobility. Action 1 can be interpreted as pursuing higher education, and the n individuals as being the families in some community. When a family is randomly selected to make a new choice, it is interpreted as having a child in the family replacing the parent and making a choice. Social mobility patterns are determined by how likely it is that a child makes the same choice that his or her parent did. Provided p_s is nondecreasing in s, a parent's and child's decisions are positively correlated. The idea is that the child is most likely to choose action 1 when there are many others in the community who have chosen action 1, which is relatively more likely to happen if the parent also chose action 1. Thus parent and child actions are correlated even though there is no direct link between the two, only the fact that the surroundings that influence their decisions overlap. Thus this explanation of social mobility is complementary (no pun intended) to that of direct parent-child interaction.

For example, the probability that the parent and child both choose action 1 is $\sum_s \mu_s \frac{s}{n} p_{s-1}$. This comes from summing across states the probability of state s, μ_s, times the probability that the parent is a 1, $\frac{s}{n}$, times the probability that the child also chooses action 1, p_{s-1} (as there are $s-1$ others choosing action 1, given that the parent was choosing 1).

As a quick demonstration that such a model can easily generate outcomes consistent with observed patterns, Calvó-Armengol and Jackson [131] restrict attention to a special case for which there are just two parameters that govern the Markov chain. For this case $p_s = q$ for some $1 > q > 0$ when $s \geq \tau$, and $p_s = 1 - q$ when $s < \tau$, where $\tau \in \{0, \ldots, n-1\}$ is a threshold. Thus individuals choose action 1 with probability q if at least τ others have, and with probability $1 - q$ otherwise (see Exercise 9.3).

Table 9.1 shows a few representative observations of father-daughter education decisions from the Calvó-Armengol and Jackson [131] fitting exercise (see their

---

7. This is also distinct from the herding literature of Banerjee [36]; Bikhchandani, Hirshleifer, and Welch [65]; and others, who examine how situations with uncertainty and private information about the benefits of actions can lead individuals to follow a herd of others in choosing an action (see Exercise 8.6).

8. For example, if people wish to avoid congestion, so that they are likely to choose the action chosen by the minority, then the most likely state becomes that of $s = n/2$.

TABLE 9.1
Father/daughter education choices for selected European countries

	Data			Estimation	
AU	0	1		0	1
0	.903	.033	0	.902	.048
1	.054	.008	1	.048	.003
				$q = .95,\ \tau = 14$	
GR	0	1		0	1
0	.646	.192	0	.648	.157
1	.124	.038	1	.157	.038
				$q = .81,\ \tau = 15$	
UK	0	1		0	1
0	.246	.223	0	.230	.250
1	.254	.278	1	.250	.270
				$q = .52,\ \tau = 25$	

Note: AU, Austria; GR, Greece; UK, United Kingdom.

supplementary material). The data are the relative frequencies of observations of father and daughter education choices based on wave 5 of the European Community Household Panel data set. A value of 0 represents that the individual at most graduated from high school; while 1 represents pursuing some education beyond high school. The row choice of 0 or 1 is the father's choice, while the column is the daugher's.

The second column of the table displays the probabilities from the simple threshold model when the community size is 25 families, with the best fit q and τ (from a search on a grid of q to hundredths, and across each τ). The fitted values do not match exactly; for instance in the fitted matrices the probability of 0,1 is the same as 1,0, which is a function of the simplified threshold model. Allowing for richer choices of the p_s values would provide a better fit. This example shows that such simple models can lead to patterns that are quite close to observed ones, and that social interaction is a viable part of explaining parent-child correlations, among other things.

9.1.3 An Interaction Model with Network Structure

While the simple Markov model of social interactions discussed in the previous section provides insight into broad patterns of social behavior, it does not incorporate the micro-details of who interacts with whom. Such network relations can have a profound effect on the process. To incorporate networked interactions, we need a richer structure. Consider the following process, which allows us to incorporate a (possibly weighted and directed) network.

As before, individuals choose between two actions 0 and 1; now the social state needs to keep track of which agents are taking which actions. The social state is thus an n-dimensional vector $x(t)$, where $x_i(t)$ for $i \in \{1, \ldots, n\}$ is the action that agent i took at period t. Interaction is described by w, which is an $n \times n$-dimensional matrix, where entry $w_{ij} \in [0, 1]$ is a weight that describes the probability that individual i's choice in period $t+1$ is the action that j took in period t. This matrix is row stochastic.

In addition, let us allow for a probability $\varepsilon_i(1)$, that i chooses action 1 independently of the state of the system and a probability $\varepsilon_i(0)$, that i chooses action 1 independently of the state of the system. Letting $\Pr(x_i(t+1) = 1|x(t))$ denote the probability that $x_i(t+1) = 1$ given the state at time t is described by the vector $x(t)$, it follows that

$$\Pr\left(x_i(t+1) = 1|x(t)\right) = \varepsilon_i(1) + (1 - \varepsilon_i(1) - \varepsilon_i(0)) \sum_j w_{ij} x_j(t).$$

Allowing for $w_{ii} > 0$, we can encode the possibility that the agent does not update his or her action at all. The Kirman ants model considered in Section 9.1.2 is the special case for which $w_{ij} = 1/(n-1)$ for $j \neq i$ and $j \leq n$, and $\varepsilon_i(1) = \varepsilon_i(0) = \varepsilon/2$.

Much richer models than the ants model can now be encoded. To get a feel for the variables, consider the following example. With probability 1/4, individual i sticks with his or her previous action; with probability 1/4, i follows the action of agent $i-1$; with a probability 1/4, i follows the action of agent $i+1$, and with probability 1/4, i randomizes between 0 and 1. This strategy corresponds to $\varepsilon_i(1) = \varepsilon_i(0) = 1/8$ and then having equal (1/3) weights on i, $i+1$, and $i-1$ (mod n):

$$w = \begin{pmatrix} 1/3 & 1/3 & 0 & 0 & \ldots & 0 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 & \ldots & 0 & 0 & 0 \\ 0 & 1/3 & 1/3 & 1/3 & \ldots & 0 & 0 & 0 \\ & \vdots & & & \ddots & & \vdots & \\ 0 & 0 & 0 & 0 & \ldots & 1/3 & 1/3 & 1/3 \\ 1/3 & 0 & 0 & 0 & \ldots & 0 & 1/3 & 1/3 \end{pmatrix}.$$

This framework allows an individual's choice of actions to depend on arbitrary neighborhoods of others, placing varying weights on different agents. Correspondingly, the process becomes more cumbersome to deal with, as we now need to keep track of each agent's social state in each period, rather than just an aggregate statistic. Nonetheless, this is still a well-defined finite-state Markov chain, which allows us to deduce quite a bit about actions over time. First, in situations in which the system is irreducible and aperiodic, a steady-state distribution exists (see Section 4.5.8). Second, although the steady-state distribution is a potentially complex joint distribution on the full vector of agents' actions, it is easy to deduce the steady-state probability that any given agent takes action 1.

We determine this probability by using the following trick to encode the ε_i-based choices of the individuals. Expand the society to include two fictitious extra individuals, labeled $n+1$ and $n+2$. Agent $n+1$ always takes action 0, and agent $n+2$ always takes action 1. Now, we can encode an individual's actions entirely in terms of an $(n+2) \times (n+2)$ matrix W. This is done by setting $W_{i,n+1} = \varepsilon_i(0)$, $W_{i,n+2} = \varepsilon_i(1)$, and $W_{ij} = (1 - \varepsilon_i(0) - \varepsilon_i(1))w_{ij}$ for $j \le n$, and weights 1 on themselves for the extra individuals.

Then let $X(t) = (x_1(t), \ldots, x_n(t), 0, 1)$ be the $(n+2) \times 1$ vector representing the larger state space including the extra "constant" individuals' actions. Then

$$\Pr(X_i(t+1) = 1 | X(t)) = [WX(t)]_i \,.$$

More importantly, for any $t' > t$, it follows that

$$\Pr(X_i(t') = 1 | X(t)) = \left[W^{t'-t} X(t) \right]_i , \tag{9.3}$$

where $W^{t'-t}$ is the matrix W raised to the $t' - t$ power. In many cases, the vector of probabilities in (9.3) converges to a steady-state probability from any starting state. For example, if there is a directed path from each i to at least one of $n+1$ and $n+2$, then the limit is unique and well behaved.

Proposition 9.2 *If for each individual i either $\varepsilon_i(0) + \varepsilon_i(1) > 0$ or there is a directed path[9] from i to some j for whom $\varepsilon_j(0) + \varepsilon_j(1) > 0$, then there is a unique limiting probability that i chooses action 1, and this limiting probability is independent of the starting state. Moreover, the vector of limit probabilities is the unique (right-hand) unit eigenvector of W such that the last two entries ($n+1$ and $n+2$) are 0 and 1.*

Proof of Proposition 9.2. The convergence of W^t to a unique limit follows from Theorem 2 in Golub and Jackson [294]. Given that there is a directed path from each i to at least one of $n+1$ and $n+2$, and each of these places weight 1 on itself, the only minimal closed sets (minimal directed components with no directed links out) are the nodes $n+1$ and $n+2$ viewed as separate components. It then follows from Theorem 3 in Golub and Jackson [294] that $\lim_t W^t X(0)$ is of the form $W^\infty X(0)$, where

$$W^\infty = \begin{pmatrix} 0 & 0 & 0 \ldots & 0 & \gamma_1 & \pi_1 \\ 0 & 0 & 0 \ldots & 0 & \gamma_2 & \pi_2 \\ & \vdots & \ddots & & \vdots & \\ 0 & 0 & 0 \ldots & 0 & \gamma_i & \pi_i \\ & \vdots & \ddots & & \vdots & \\ 0 & 0 & 0 \ldots & 0 & \gamma_n & \pi_n \\ 0 & 0 & 0 \ldots & 0 & 1 & 0 \\ 0 & 0 & 0 \ldots & 0 & 0 & 1 \end{pmatrix}.$$

9. This path refers to a directed path in the directed graph on the original n individuals, where a link ij is present if $w_{ij} > 0$.

It also must be that $WW^{\infty} = W$. Thus $W\pi = \pi$, where $\pi = (\pi_1, \ldots, \pi_n, 0, 1)$, which implies that π is a right-hand unit eigenvector of w. ∎

To get a feel for Proposition 9.2, let us consider an application.

Example 9.1 (An Application: Probability of Action in Krackhardt's Advice Network) *We revisit Example 8.11, which concerned Krackhardt's [411] data on an advice network among managers in a small manufacturing firm on the West Coast of the United States. Suppose that the action is whether to meet to go to a bar after work. This happens repeatedly, and each day the managers find out who went the previous night. Consider the following action matrix based on Krackhardt's data. A manager with out-degree d chooses to go to the bar with probability $1/(d+2)$, not to go with probability $1/(d+2)$, and with the remaining probability of $d/(d+2)$ uniformly at random picks one of his or her neighbors and then goes to the bar if that neighbor did on the previous day. This strategy is used by all managers except the top-level ones (labeled 2, 7, 14, 18, and 21 in Example 8.11) who are biased toward going to the bar. The top-level managers have a similar rule except that they use weights $1/(d+1)$ and do not place any weight on action 0. We can then calculate the frequency with which each manager goes to the bar in the long run, as listed in Table 9.2. The probabilities were calculated via Matlab by finding the (right-hand) unit eigenvector of W.*

It is important to emphasize that Table 9.2 does not give us the joint distribution over people going to the bar. It is easy to see that there will be correlation and long periods of time when many people go to the bar, and then long periods of time where few people go. Moreover, even this myopic sort of behavior can lead individuals to coordinate. The technique of calculating the marginals, however, only allows us to see the individual probabilities and not the joint distribution. The joint distribution is over 2^{21} (more than 2 million) different states, and so it is a bit difficult to keep track of.

Beyond the difficulties in tracking the full joint distribution, there are two other limitations to this analysis, which motivates a game-theoretic analysis. The first is that the individuals are backward looking. That is, they look at what their neighbors did yesterday in deciding whether to go to the bar, rather than coordinating with their neighbors on whether they plan to go to the bar today. Second, agent i weights the actions of the others in a separable way. Alternatively, an individual might prefer to go when a larger group is going (or might, in contrast, want to avoid congestion and stay away on crowded nights). This sort of decision is precluded by the separability in the way that an individual i treats the actions of the others. In particular, it is that aspect of the framework above that allows us to solve for the probabilities of action. To see this explicitly, consider an example with $n = 3$ people. Suppose that individual 1 has $w_{12} = w_{13} = 1/2$ and $\varepsilon_1(0) = \varepsilon_1(1) = 0$. Consider how individual 1 will play in period 1, as a function of how we choose the starting state. In one case, pick the initial choices of individuals 2 and 3 in an independent manner, selecting 1 with probability p. In that case, individual 1 will choose 1 with probability p. But pick the initial actions of individuals 2 and 3 to be the same, so with probability p they are both 1, and with probability $1 - p$ they are both 0. This choice does not make any difference in calculating the probability that individual 1 chooses action 1 in the first period. It is still p. The joint distribution

TABLE 9.2
Probability of action in Krackhardt's network of advice among managers

Agent	Probability of taking action 1	Level	Department	Age (years)	Tenure (years)
1	0.667	3	4	33	9.3
2	0.842	2	4	42	19.6
3	0.690	3	2	40	12.8
4	0.666	3	4	33	7.5
5	0.690	3	2	32	3.3
6	0.585	3	1	59	28
7	0.771	1	0	55	30
8	0.676	3	1	34	11.3
9	0.681	3	2	62	5.4
10	0.660	3	3	37	9.3
11	0.656	3	3	46	27
12	0.585	3	1	34	8.9
13	0.680	3	2	48	0.3
14	0.821	2	2	43	10.4
15	0.687	3	2	40	8.4
16	0.651	3	4	27	4.7
17	0.671	3	1	30	12.4
18	0.737	2	3	33	9.1
19	0.685	3	2	32	4.8
20	0.685	3	2	38	11.7
21	0.755	2	1	36	12.5

Note: Level indicates the level of hierarchy in the firm (e.g., 1 denotes the highest, 3 the lowest level).

over all players' actions will change, but the probability that any given individual chooses action 1 is unaffected.

This special property that leads to such power in calculating the steady-state probabilities of actions is not always satisfied. In fact, many applications of interest have a more complex structure to the incentives. For example, suppose that individual 1 would like to choose 1 if both of the others choose action 1, but not otherwise. These incentives would arise if action 1 has a cost associated with it (slightly higher than the cost of taking action 0), but agent 1 wants to choose an action that is compatible with as many other agents as possible. Then if we pick 2's and 3's actions independently, there is only a p^2 chance that both 2 and 3 will choose action 1 and that individual 1 will choose action 1 in period 1. If instead we pick 2's and 3's actions to be the same, then there is a p chance that they will both choose action 1 and that individual 1 will then choose action 1 in period 1.

Although we can still write the system as a Markov chain, it is not quite as powerful a tool now, as the transition probabilities and evolution of behavior are more complicated. To get a better handle on such more complicated interaction structures, let us turn to game-theoretic reasoning.

9.2 ▪ Graphical Games

Individual decisions often depend on the relative proportions of neighbors taking actions, as in deciding on whether to buy a product, change technologies, learn a language, smoke, engage in criminal behavior, and so forth. This can result in multiple equilibrium points: for instance, some people may be willing to adopt a new technology only if others do, and so it would be possible for nobody to adopt it, or for some nontrivial fraction to adopt it. A way of introducing such reactive or strategic behavior into the analysis of social interaction is to model the interaction as a game.

A useful class of such interactions was introduced by Kearns, Littman, and Singh [385] as what they called *graphical games*.[10]

More formally, there is a set N of players, with cardinality n, who are connected by a network (N, g). Each player $i \in N$ takes an action in $\{0, 1\}$. The payoff of player i when the profile of actions is $x = (x_1, \ldots, x_n)$ is given by:

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

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

There is nothing about this definition that precludes the network from being directed. For instance, player i may care about how player j acts, but not the reverse. Most of the definitions that follow work equally well for directed and undirected cases. I note points at which the analysis is special in the directed case. Most of the examples examine the undirected case for ease of exposition.

In a graphical game, players' payoffs depend on the actions taken by their neighbors in the network. Nevertheless, a player's behavior is related to that of indirect neighbors, since a player's neighbors' behavior is influenced by their neighbors, and so forth, and equilibrium conditions tie together all the behaviors in the network. Note that we can view this formulation as being without loss of generality, because we can define the network to include links to all of the players that affect a given player's payoff. As an extreme case, if a player cares about everyone's behavior, we have the complete network, which is a standard game with n players.

---

10. These can also be viewed as a special case of multi-agent influence diagrams, often referred to as MAIDs, which are discussed by Koller and Milch [408]. Earlier discussion of the possibility of using MAIDs to model strategic interactions in which players only respond to other subsets of agents' actions dates to Shachter [590], but the first fuller analysis appears in Koller and Milch [408]. Even though the general MAID approach allows for the encoding of multiple types of problems, information structures, and complex interaction structures, and hence includes graphical games as a special case, most of the analysis examines special cases that preclude the graphical games structure discussed here.

9.2.1 Examples of Graphical Games

To fix ideas, consider a couple of examples.

Example 9.2 (Threshold Games of Complements) *Many of the applications mentioned so far involve strategic complements, such that a player has an increasing incentive to take a given action as more neighbors take the action. In particular, consider situations in which the benefit to a player from taking action 1 compared to action 0 (weakly) increases with the number of neighbors who choose action 1, so that*

$$u_i(1, x_{N_i(g)}) \geq u_i(0, x_{N_i(g)}) \text{ if and only if } \sum_{j \in N_i(g)} x_j \geq t_i,$$

where t_i is a threshold. In particular, if more than t_i neighbors choose action 1, then player i should choose 1, and if fewer than t_i neighbors choose action 1 then it is better for player i to choose action 0.

A special case occurs if action 1 is costly (e.g., investing in a new technology or product) but the benefit of that action increases as more neighbors undertake the action:

$$\begin{aligned} u_i(1, x_{N_i(g)}) &= a_i \left(\textstyle\sum_{j \in N_i(g)} x_j \right) - c_i, \\ u_i(0, x_{N_i(g)}) &= 0, \end{aligned}$$

for some $a_i > 0$ and $c_i > 0$. Here the threshold is such that if at least $t_i = \frac{c_i}{a_i}$ neighbors choose action 1, then it is better for player i to choose 1, and otherwise player i should choose action 0.

Example 9.3 (A "Best-Shot" Public Goods Game) *Another case of interest has the opposite incentive structure. For example, if a player or any of the player's neighbors take action 1, then the player obtains a benefit of 1. For instance, the action might be learning how to do something, where that information is readily communicated; or buying a book or other product that is easily lent from one player to another.*[11] *Taking action 1 is costly, and a player would prefer that a neighbor take the action rather than having to do it himself or herself; but taking the action and paying the cost is better than having nobody take the action. This scenario is known as a* best-shot public goods game *(e.g., see Hirshleifer [329]), where*

$$\begin{aligned} u_i(1, x_{N_i(g)}) &= 1 - c. \\ u_i(0, x_{N_i(g)}) &= 1 \qquad \text{if } x_j = 1 \text{ for some } j \in N_i(g), \textit{ and} \\ u_i(0, x_{N_i(g)}) &= 0, \qquad \text{if } x_j = 0 \text{ for all } j \in N_i(g), \end{aligned}$$

where $1 > c > 0$.

---

11. The distinction between private and public goods is a standard one in economics. The term *public good* refers to the fact that one player might acquire the good, information or the like, that can be consumed by others and hence is not private to that player but is publicly available. The term *local* refers to the benefits of a given player's action being public only in the player's neighborhood.

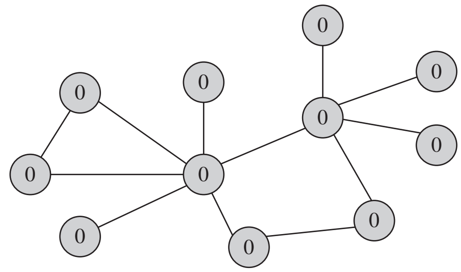

FIGURE 9.2 An equilibrium in a game of complements with threshold 2.

9.2.2 Equilibrium

Given the graphical game structure, we can use game theory to make predictions about players' behavior and how it depends on the network structure. I strongly recommend that those not familiar with game theory read Section 9.9, which provides a quick and basic tutorial in game theory, before proceeding with the remainder of this chapter.

In a graphical game, a *pure strategy Nash equilibrium* is a profile of strategies $x = (x_1, \ldots, x_n)$ such that

$$\begin{aligned} u_i(1, x_{N_i(g)}) &\geq u_i(0, x_{N_i(g)}) \quad \text{if } x_i = 1, \text{ and} \\ u_i(0, x_{N_i(g)}) &\geq u_i(1, x_{N_i(g)}) \quad \text{if } x_i = 0. \end{aligned}$$

So the equilibrium condition requires that each player choose the action that offers the highest payoff in response to the actions of his or her neighbors: no player should regret the choice that he or she has made given the actions taken by other players.

Figures 9.2–9.4 show some pure strategy equilibria for the threshold game of complements outlined in Example 9.2 when the threshold is 2 for all players and the network is undirected. That is, a player prefers to buy a product if at least two of his or her neighbors do, but prefers not to otherwise. Note that the case pictured in Figure 9.2, in which all players take action 0, is an equilibrium for any game in which all players have a threshold of at least 1 for taking action 1.

There is generally a multiplicity of equilibria in such threshold games. For instance, the configuration pictured in Figure 9.3 is also an equilibrium. An equilibrium also occurs when a maximal configuration of players takes action 1. The configuration pictured in Figure 9.4 shows each player taking the maximal action that he or she can in any equilibrium. While these figures show that multiple equilibria can exist in graphical games, it is also possible for none to exist. This is illustrated in Example 9.4.

Example 9.4 (A Fashion Game and Nonexistence of Pure Strategy Equilibria) *Consider a graphical game in which there are two types of players. One type consists of "conformists" and the others are "rebels." Conformists wish to take an action that matches the majority of their neighbors, while rebels prefer to take an action that matches the minority of their neighbors. This game is a variation on a classic one called "matching pennies," in which one player wishes to choose*

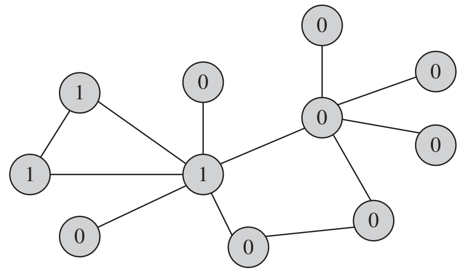

FIGURE 9.3 Another equilibrium in a game of complements with threshold 2.

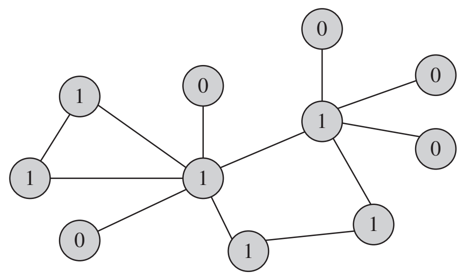

FIGURE 9.4 Maximal equilibrium in a game of complements with threshold 2.

the same action as the other, while the second player wishes to mismatch (see Section 9.4). It is easy to check that pairing one conformist and one rebel in a dyad results in no pure strategy equilibrium, as in the game of matching pennies. More generally, some graphical game structures have no pure strategy equilibrium. One class of networks with pure strategy equilibria is that in which all players have more than half of their neighbors being conformists: there is a pure strategy equilibrium where all the conformists take one action, and all the rebels take the other.

In light of Example 9.4, it is useful to define mixed strategy equilibria for a graphical game. Denote a mixed strategy of i by σ_i, where $\sigma_i \in [0, 1]$ is the probability that player i chooses $x_i = 1$ and $1 - \sigma_i$ is the probability that the player chooses $x_i = 0$. Let σ_{-i} denote a profile of mixed strategies of the players other than i, and let $u_i(x_i, \sigma_{N_i(g)})$ denote the expected utility of player i who plays x_i and whose neighbors play $\sigma_{N_i(g)}$. Let $u_i(\sigma_i, \sigma_{N_i(g)})$ be the corresponding expected utility when i plays a mixture σ_i.[12] Then a mixed strategy equilibrium is a profile

12. Thus

$$u_i(x_i, \sigma_{N_i(g)}) = \sum_{x_{N_i(g)} \in \{0,1\}^{d_i(g)}} u_i(x_i, x_{N_i(g)}) \Pr(x_{N_i(g)} | \sigma_{N_i(g)}),$$

and

$$u_i(\sigma_i, \sigma_{N_i(g)}) = \sigma_i u_i(1, \sigma_{N_i(g)}) + (1 - \sigma_i) u_i(0, \sigma_{N_i(g)}).$$

$\sigma = (\sigma_1, \ldots, \sigma_n)$ of mixed strategies such that for every i,

$$u_i(\sigma_i, \sigma_{N_i(g)}) = \max\left[u_i(1, \sigma_{N_i(g)}), u_i(0, \sigma_{N_i(g)})\right].$$

As graphical games have a finite set of strategies and players, they always have at least one equilibrium, which may be in mixed strategies (see Section 9.9.4). Many graphical game settings always have pure strategy equilibria, including the examples of best-shot and threshold games mentioned above. But we have also seen some natural situations that will only have mixed strategy equilibria, such as the fashion game (Example 9.4).

Let us turn to a class of graphical games that covers many applications of interest and also is nicely behaved, allowing for a tractable analysis of how social structure relates to behavior.

9.3 ▪ Semi-Anonymous Graphical Games

There are many situations in which a player's choice is influenced mainly by the relative popularity of a given action among his or her neighbors and is not dependent on the specific identities of the neighbors who take the action. While an approximation, this type of game can be very useful and has many natural applications. A class of such graphical games is examined by Galeotti et al. [274], and I refer to these as semi-anonymous graphical games.

These games are not quite anonymous (e.g., see Kalai [375]), in which a player is affected by the actions of all other players in a symmetric way; since in a graphical game a player cares only about a subset of the other players' actions. But it is anonymous in the way that a player is influenced by his or her neighbors. That is, the player cares only about how many of the neighbors take action 0 versus 1, but not precisely which of the neighbors take action 1 versus action 0. So semi-anonymity refers to this anonymity on a neighborly level. In addition, another aspect of anonymity is invoked here: players have similar payoff functions, so that differences between players arise from the network structure and not some other innate characteristics. A player's utility function depends on his or her degree and not on his or her label.

9.3.1 Payoffs and Examples

Formally, a *semi-anonymous graphical game* is a graphical game such that the payoff to player i with a degree d_i who chooses action x_i is described by a function $u_{d_i}(x_i, m)$, where m is the number of players in $N_i(g)$ taking action 1.

Thus the payoff function is dependent on the player's degree, the player's own action, and the number of neighbors who take each action. Note that since the function depends on both the degree d_i and the number of neighbors choosing action 1, m, we could equivalently have defined it to be a function of the degree and the number of players taking action 0 (which is simply $d_i - m$) or of the degree and the fraction of players taking action 1 (or 0). Note also that the best-shot game of Example 9.3 is a semi-anonymous graphical game, as is the fashion game in Example 9.4. The threshold games in Example 9.2 are semi-anonymous when each player's threshold depends only on his or her degree. Here are some other examples.

Example 9.5 (A Local Public Goods Game) *Consider a game in which each player's action contributes to some local public good: that is, an action by a given player provides some local benefits to all neighbors. This example generalizes the best-shot public goods graphical game to allow for situations in which having multiple players take action 1 is better than having just one player take the action. For example, having each player study a given candidate's record and then share that information with his or her neighbors could lead players to be more informed about how to vote in an election than having just one player study a candidate's record. In this case, a player who has m neighbors take action 1 gets a payoff of*

$$u_{d_i}(x_i, m) = f(x_i + \lambda m) - cx_i,$$

where f is a nondecreasing function, and $\lambda > 0$ and $c > 0$ are scalars. The case where $\lambda = 1$, $f(k) = 1$ for all $k \geq 1$, and $f(0) = 0$ is the best-shot public good graphical game.

Example 9.6 (A "Couples" Game) *Imagine learning a skill that is most easily enjoyed when there is at least one friend to practice it with, such as playing tennis, some video games, or gin rummy. In this situation, a player prefers to take action 1 if at least one neighbor takes action 1, but prefers to take action 0 otherwise. We can think of this as having a cost of investing in the skill of c, along with a benefit of 1 if there is a partner to participate with. Here*

$$\begin{aligned} u_{d_i}(1, m) &= 1 - c \quad \text{if } m \geq 1, \\ u_{d_i}(1, 0) &= -c, \qquad \text{and} \\ u_{d_i}(0, m) &= 0. \end{aligned}$$

This game is a special case of a threshold game of complements with a threshold of 1.

Example 9.7 (A Coordination Game) *Consider a situation in which a player prefers to coordinate his or her action with other players, and his or her payoff is related to the fraction of neighbors who play the same strategy:*

$$\begin{aligned} u_{d_i}(1, m) &= a\frac{m}{d_i}, \text{ and} \\ u_{d_i}(0, m) &= b\frac{d_i - m}{d_i}. \end{aligned}$$

This is a special case of a threshold game of complements, with the threshold for player i of degree d_i being $t_i = d_i \frac{b}{a+b}$.

9.3.2 Complements and Substitutes

In the examples, as in many applications, the incentives of a player to take an action either increase or decrease as other players take the action. Distinguishing between these cases is important, since they result in quite different behaviors. These two broad cases are captured by the following definitions.

A semi-anonymous graphical game[13] exhibits *strategic complements* if it satisfies the property of increasing differences; that is, for all d and $m \geq m'$:

$$u_d(1, m) - u_d(0, m) \geq u_d(1, m') - u_d(0, m').$$

A semi-anonymous graphical game exhibits *strategic substitutes* if it satisfies decreasing differences; that is, for all d and $m \geq m'$:

$$u_d(1, m) - u_d(0, m) \leq u_d(1, m') - u_d(0, m').$$

These notions are said to apply strictly if the inequalities above are strict when $m > m'$.

The best-shot public goods game is one of strategic substitutes, as is the local public goods game when f is concave. In that case, higher levels of actions by neighbors (i.e., higher m) lead to an incentive for a given player to take the lower action or to free-ride. Other examples are items that can be shared, such as products and information, as well as other situations with externalities, such as pollution reduction and defense systems (where the network could be one of treaties and the players are countries).

The couples game is one of strategic complements. Local public goods games when f is convex, so there are increasing returns to action 1, also exhibit strategic complementarities. Other examples of strategic complementarities include situations in which peer effects are important. It is important to emphasize that complementarities can exist for many different reasons. For instance, in decisions of whether to pursue higher education, neighbors may serve as role models for a given individual. They may also provide information about the potential benefits of higher education, or serve as future contacts in relaying job information. An individual may have an incentive to conform to the patterns of behavior of his or her peers. The critical common feature is that increased levels of activity among a given player's neighbors increase the incentives or pressures for that player to undertake the activity.

9.3.3 Equilibria and Thresholds

The nice aspect of semi-anonymous graphical games is that the behavior of a given individual can be succinctly captured by a threshold. In the case of strategic complements, there is a threshold $t(d)$, which can depend on a player's degree d such that if more than $t(d)$ neighbors choose action 1, then the player prefers action 1, while if fewer than $t(d)$ neighbors choose 1, then the player prefers 0. It is possible to have situations in which an individual is completely indifferent at the threshold. For the case of strategic substitutes, there is also a threshold, but the best response of the player is reversed, so that he or she prefers to take action 0 if more than $t(d)$ neighbors take action 1, and prefers action 1 if fewer than $t(d)$

13. The definitions extend readily to settings beyond the semi-anonymous case, by working with set inclusion. That is, complements are such that if the actions of all of a player's neighbors do not decrease and some increase, then the player's gain in payoffs from taking action 1 compared to action 0 increases; substitutes are the reverse case. See Exercise 9.8.

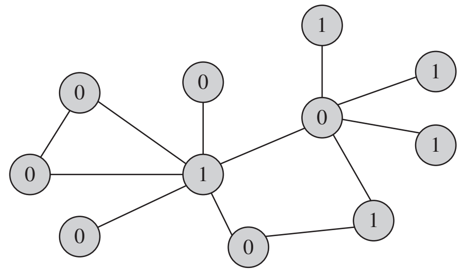

FIGURE 9.5 An equilibrium in a best-shot public goods graphical game.

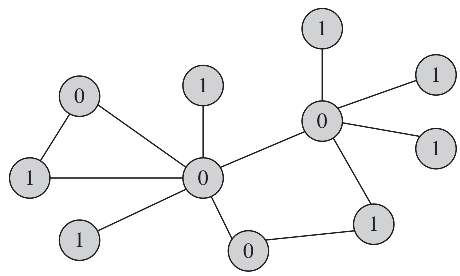

FIGURE 9.6 Another equilibrium in a best-shot public goods graphical game.

neighbors take action 1. The best-shot public goods game and the couples game are games of strategic substitutes and strategic complements, respectively, where the threshold is 1 (irrespective of degree).

As discussed in more detail in Section 9.8, semi-anonymous graphical games with strategic complementarities always have a pure strategy equilibria. In fact, the set of pure strategy equilibria has a nice structure and is what is known as a *complete lattice,* as outlined in Exercise 9.5. This structure implies that there exists a maximum equilibrium such that each player's action is at least as high as in every other equilibrium, and similarly a minimum equilibrium where actions take their lowest values out of all equilibria. For the network pictured in Figures 9.2–9.4, the minimum equilibrium is in Figure 9.2 and the maximum is in Figure 9.4. (There is one other pure strategy equilibrium for this network not pictured in Figures 9.2–9.4, which is the topic of Exercise 9.4.)

Semi-anonymous graphical games of strategic substitutes do not always have a pure strategy equilibrium, but they always have at least one equilibrium in mixed strategies. There are games of strategic substitutes that have pure strategy equilibria, as we have seen in the case of the best-shot public goods game, and they can have multiple equilibria. Figures 9.5 and 9.6 exhibit two different pure strategy equilibria in a best-shot public goods game for the same network as in Figures 9.2–9.4.

There are other equilibria for this network (see Exercise 9.4), but the structure of equilibria in the best-shot public goods game does not exhibit a lattice structure. Nevertheless, there still always exists at least one pure strategy equilibrium for a best-shot public goods graphical game when the network is undirected. Exercise 9.6 shows that an equilibrium may not exist in directed networks.

9.3.4 Comparing Behavior as the Network Is Varied

With some examples and definitions in hand, let us examine a few basic properties of how behavior changes as we vary the structure of a network. Such comparisons show how social structure influences behavior.

First, it is easy to see that in games of complements in which the threshold for taking action 1 is nonincreasing in degree, adding links will lead to (weakly) higher actions, as players will have more neighbors taking action 1.

Proposition 9.3 *Consider a semi-anonymous graphical game of strategic complements on a network* (N, g) *such that the threshold for taking action 1 is nonincreasing as a function of degree, so that* $t(d+1) \leq t(d)$ *for each* d*. If we add links to the network to obtain a network* g' *(so that* $g \subset g'$*), then for any pure strategy equilibrium* x *under* g*, there exists an equilibrium* x' *under* g' *such that all players play at least as high an action under* x' *as under* x*.*

Proposition 9.3 notes that if incentives to take an action increase as more neighbors take an action, then denser networks lead to higher numbers of players choosing the action. This result is not dependent on the 0-1 action space we have been considering, but extends to more general action spaces as outlined in Exercise 9.8. This conclusion requires that players react to the absolute level of activity by neighbors rather than to a proportion.

The more subtle case is that of strategic substitutes. Adding links can change the structure of payoffs in unpredictable ways, as illustrated in the following example of adding a link in a best-shot graphical game, which is drawn from an insight of Bramoullé and Kranton [103]. One might expect, reversing the intuition from the complements case, that adding links leads to new equilibria at which all agents take less action than they did before. However this is not quite right, as decreasing actions for some agents can lead to increasing actions for others in the case of strategic substitutes, and so changing network structure leads to more complex changes in behavior, as pictured in Figure 9.7.

The top panel of Figure 9.7 shows an equilibrium in which both players who are the centers of their respective stars provide the public good and other players free-ride. This configuration is the overall cheapest way of providing the public good to all players and so has a strong efficiency property. However, when we add a link between these two center agents, it is no longer an equilibrium for both of them to provide the public good. Thus adding a link can change the structure of equilibria in complicated ways.

Nevertheless, there is still a well-defined way in which the equilibrium adjusts so that, despite Figure 9.7, adding links still decreases actions if tracking how all equilibria change.

Proposition 9.4 (Galeotti et al. [274]) *Consider a best-shot graphical game on a network* (N, g) *and any pure strategy equilibrium* x *of* $g + ij$*. Either* x *is also an equilibrium of* g*, or there exists an equilibrium under* g *in which a strict superset of players chooses 1.*

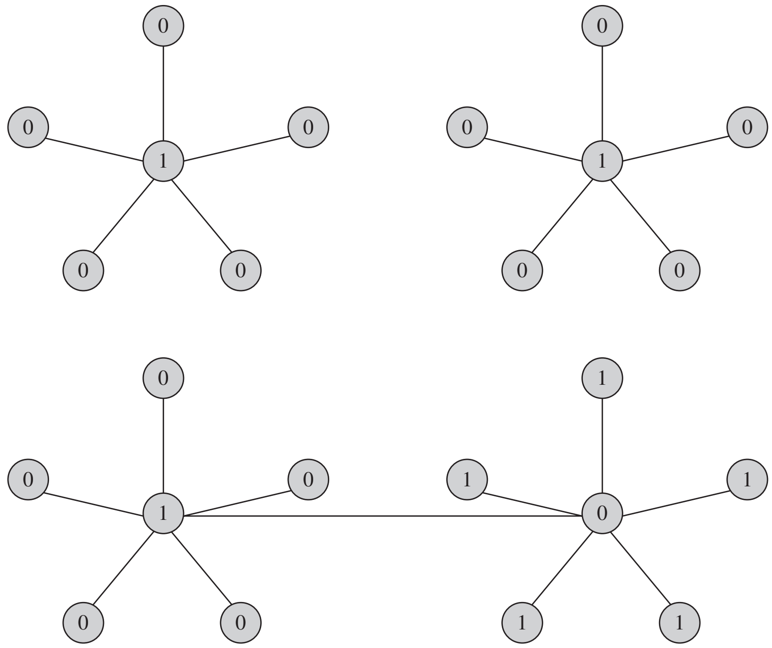

FIGURE 9.7 Adding a link changes the equilibrium structure in a best-shot graphical game.

Thus, although there are generally multiple equilibria and a particular equilibrium might not be an equilibrium when a link is added, any new equilibrium in the network with a new link has a subset of players who take action 1 compared to some equilibrium of the old network. Indeed, in Figure 9.7 the equilibrium in the bottom half of the figure is also an equilibrium without the link being present. It is just that the equilibrium in the top panel no longer survives. Thus in some sense adding links implies fewer players providing the public good. The proof of the proposition follows easily from the structure of maximal independent sets (e.g., see Observation 2.1).

The proposition is illustrated by examining all pure strategy equilibria on the networks of three individuals, as pictured in Figure 9.8. In Figure 9.8, we see how the equilibria vary as links are added or deleted, and although a particular equilibrium might not survive when links are deleted, it can be compared to some other equilibrium of the resulting network.

Making changes beyond the addition of links (e.g., moving links or changing the degree distribution but keeping the mean degree constant) can alter the landscape of equilibria in more complicated ways. The sensitivity of behavior to network changes leaves the graphical games model without sharp comparisons of behavior resulting from changes in network structure. However, there is a variation on graphical games in which behavior varies in predictable ways in response to general changes in network structure. This is not to say that one or the other is a better model, as they fit different situations, and the difference in the sharpness of their predictions is reflective of the differences across those settings.

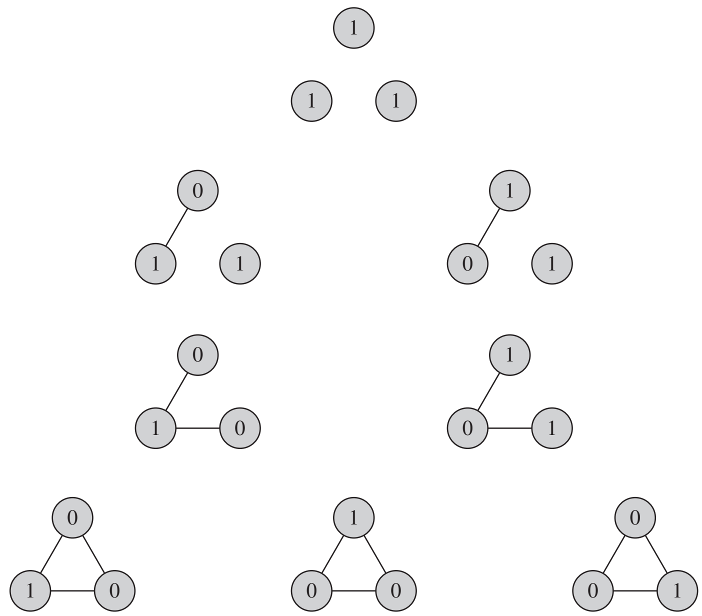

FIGURE 9.8 Equilibria in various three-player best-shot graphical games.

9.4 ▪ Randomly Chosen Neighbors and Network Games

While graphical games nicely model a number of networked interactions in which players have a good idea of their neighbors' actions when choosing their own (or they can adjust their behavior), there are also many situations in which players choose actions in at least partial ignorance of what their neighbors will do, or even in ignorance of who their neighbors will be. This limited knowledge applies when players are learning a skill or making some investment and are unsure of their future interactions. For example, they might be choosing majors in college, which will eventually be very important in interactions with their employers, their colleagues, and so forth. In choosing what major to undertake, they might know something about the number of other people choosing that action and the job market for different majors, without knowing with whom they will be interacting.[14] These ideas are formalized in a setting from Galeotti et al. [274], Jackson and Yariv [362], [364], and Sundararajan [618]. For expositional purposes, I use a setting with just two actions, but the analysis extends to richer settings.

A player knows his or her own degree as well as the distribution over the likely degrees of his or her neighbors, but nothing more about the network structure when choosing an action. Degree can be thought of as the number of interactions that a player is likely to have in the future.

14. See Pasini, Pin, and Weidenholzer [528] for some discussions of another application (to buyer-seller markets) where this network games formulation is appropriate.

Define a strategy to indicate which action is chosen as a function of a player's degree. In particular, let $\sigma(d) \in [0, 1]$ be the probability that a player of degree d chooses action 1. For most degrees (d) will be either 0 or 1, but in some cases there might be some mixing. This definition implicitly builds in a symmetry such that players of the same degree follow the same strategy.[15] In many cases, the symmetry holds without loss of generality, as players with the same degree face the same payoffs as a function of their actions and will often have a unique best response.

The degrees of a player's neighbor are drawn from a degree distribution $\widetilde{P}$. Recall that $\widetilde{P}(d) = \frac{P(d)d}{\langle d \rangle}$ approximates the distribution over a neighbor's degree from the configuration model with respect to a degree sequence represented by P. Under $\widetilde{P}$ there is a well-defined probability that a neighbor takes action 1:

$$p_\sigma = \sum_d \sigma(d)\widetilde{P}(d).$$

Thus the probability that exactly m out of the d_i neighbors of player i choose action 1 is given by the binomial formula $\binom{d_i}{m} p_\sigma^m (1 - p_\sigma)^{d_i - m}$. The expected utility of a player of degree d_i who takes action x_i is then

$$U_{d_i}(x_i, p_\sigma) = \sum_{m=0}^{d_i} u_{d_i}(x_i, m) \binom{d_i}{m} p_\sigma^m (1 - p_\sigma)^{d_i - m}, \tag{9.4}$$

where $u_{d_i}(x_i, m)$ is the payoff corresponding to an underlying graphical game. One can then think of this as a sort of graphical game, in which players choose their strategies knowing how many links they will have but not knowing which network will be realized. However, the above formulation does not require a specification of the precise set of players or even how many players there will be. Players just need to know their own degrees and have beliefs about their neighbors' behavior; they do not need to have a fully specified model of the world.

The formulation above presumes independence of neighbors' degrees. However, the results extend to allow for correlation among neighbors' degrees, which is important since, as we have seen, many networks exhibit such correlations. Here I limit the discussion to the independent case, since it makes the exposition more transparent. See Section 4.5.7 and Galeotti et al. [274] for details on the appropriate definitions for extensions of these results.[16]

A specification of a utility function u_d for each d and a distribution of neighbors' degrees $\widetilde{P}$ is referred to as a *network game*. It is now easy to define a (Nash)

---

15. The equilibrium definition below allows any player to deviate in any way, and so this symmetry in behavior is an equilibrium phenomenon and is payoff-maximizing for the players; it is not a constraint on behavior.

16. Effectively the generalization allows each degree to have a different anticipated distribution over vectors of neighbors' degrees. What is required in the case of strategic complements is that higher-degree players have a distribution over neighbors' degrees that leads them to expect (weakly) higher degrees among their neighbors than a lower-degree player would. This situation is reversed for substitutes. Comparing joint distributions over different-sized vectors is based on concepts discussed in Section 4.5.7.

equilibrium of a network game.[17] An equilibrium in a network game is a strategy σ such that for each d,

- if $\sigma(d) > 0$, then $U_d(1, p_\sigma) \geq U_d(0, p_\sigma)$, and
- if $\sigma(d) < 1$, then $U_d(1, p_\sigma) \leq U_d(0, p_\sigma)$.

9.4.1 Degree and Behavior

Players' strategies can be ordered as a function of their degrees. The idea is that players with higher degrees have more neighbors, and hence an expectation of having more neighbors choosing 1. In games of strategic complements, if having more total activity among one's neighbors leads one to prefer the higher action, then higher-degree players prefer a higher action compared to the preferences of a lower-degree player. This process reverses itself for substitutes.

The conclusion that players with higher degree have more of an incentive to take higher actions is not guaranteed simply by having strategic complementarities, since that condition examines a given player's incentives as his or her neighbors' behavior is changed. It does not make comparisons of how incentives vary with degree. To make comparisons across degrees, let us focus on the case in which payoffs depend on absolute numbers of neighbors taking action 1, so that

$$u_d(x_i, m) = u_{d+1}(x_i, m) \tag{9.5}$$

for each $m \leq d$ and x_i. Equation (9.5) is not necessary for the results that follow, which hold for much more general payoff settings, including those in which players are influenced by the percentage of neighbors taking a given action rather than by the absolute number of neighbors doing so. Exercise 9.9 addresses this payoff scheme. However, using (9.5) simplifies the exposition and conveys the basic ideas. What is needed for the following results is a payoff structure such that if a higher-degree individual is faced with the same typical behavior by any given neighbor, then he or she prefers to choose action 1 over 0 when a lower-degree player would prefer to choose action 1 over 0. Under such conditions, it is straightforward to deduce the existence of an equilibrium where higher types take higher actions, and similarly for substitutes and lower actions, which leads to the following proposition. Proposition 9.5 is quite useful in deducing how behavior varies with network structure.

Proposition 9.5 (Galeotti et al. [274]) *Consider a network game in which payoff functions satisfy (9.5). If it is a game of strategic complements, then there exists an equilibrium that is nondecreasing in degree;*[18] *and if it is a game of strategic substitutes, then there exists an equilibrium that is nonincreasing in degree. If the game is one of strict strategic complements, then all equilibria are*

17. Such an equilibrium, where players' strategies depend on a "type" (here, their degree) and players are not sure of the other players' types when they choose their action, is also known as a *Bayesian equilibrium*.

18. *Nondecreasing* refers to $\sigma(d)$ being a nondecreasing function of d.

nondecreasing in degree; analogously if it is of strict strategic substitutes then all equilibria are nonincreasing in degree.

The proof of Proposition 9.5 follows the logic of a variety of game-theoretic analyses in the presence of strategic complementarities (e.g., see Topkis [630], Vives [642], and Milgrom and Roberts [469]), here adapted to the network setting. The idea is that if we begin with some σ that is nondecreasing in degree, then under (9.5) there is a best response for the players that is nondecreasing in degree. An equilibrium is a fixed point of the best response correspondence. The set of nondecreasing strategies is convex and compact (with appropriate definitions), and so a fixed point exists by any of a variety of theorems on fixed points; the same holds for substitutes. Let me sketch a more direct and intuititive proof. Let σ^t be such that $\sigma^t(d) = 1$ if $d \geq t$ and $\sigma^t(d) = 0$ for $d < t$ (allowing for $t = \infty$ to allow all players to play 0). Consider the case of strategic complements. Note that if action 1 is a best response to some strategy σ for a player of degree d, then it is a best response for all higher-degree players; and similarly if action 0 is a best response to some σ for a player of degree d, then it is a best response for all lower-degree players. Begin with σ^1. If this is an equilibrium, then stop. Otherwise, a degree-1 player must prefer to play action 0 in response to σ^1. So consider σ^2. Now action 0 must still be a best response for the degree-1 players, as there is less aggregate action by other players. So if this configuration is not an equilibrium, degree-2 players must prefer to play 0 to 1. Iterating on this logic, either the process eventually stops at some σ^t, where the degree-t players do not wish to change from action 1 to 0, or else this process continues for all degrees, in which case all players playing 0 (σ^∞) is an equilibrium.

Note that this logic shows that for the case of complements, there is actually an equilibrium of pure strategies. The case of strategic substitutes is slightly more complicated, as lowering the action of a given type of player might actually reverse the incentives for that type. For instance, consider a situation in which all players have degree d in a best-shot public goods game. The equilibrium will then involve mixing, since if players were to choose 1, then any given player would prefer to play 0, and vice versa.[19] Nevertheless, there is still nice structure to incentives across degrees, so that if 1 is a best response to some strategy σ for a player of degree d, then action 1 is a best response for all lower-degree players; similarly if 0 is a best response to some σ for a player of degree d then it is a best response for all higher-degree players. One can then follow a similar algorithm as above, but starting at σ^∞. As a first step examine whether degree-1 players would prefer to change to action 1. If they do, then raise their action but this time raise the action continuously from $\sigma(1) = 0$ to $\sigma(1) = 1$. Given the continuity of preferences, the difference in utility for such players between action 0 and 1 will change continuously. Either at some point degree-1 players are indifferent between the two actions in response to this mixture and action 0 by other degree players, or else action 1 is their strict best response to a situation in which they play action 1 and all others play action 0. Continue in this manner.

19. Recall that strategies are specified as a function of degree, and so in a regular setting all players must take the same action.

The claim that all equilibria are nondecreasing when the strategic complements or substitutes are strict follows from noting that players with higher degrees expect to have more neighbors choosing action 1 (in terms of first-order stochastic dominance) and is the subject of Exercise 9.11.

The fact that players' incentives vary monotonically with their degree does not necessarily guarantee that their payoffs vary monotonically with their degree. That relation depends on how the actions of others affect a given player. Note that just because the incentives of a player to choose action 1 increase when more neighbors choose action 1 does not mean that the players are better off making this choice. For example, consider a game that involves athletes' choices of whether to engage in doping (e.g., taking illegal drugs or undergoing blood transfusions) to give them an advantage in competition. A player's neighbors are other athletes against whom the player competes. The strategy in the game is either to dope or not to dope. Doping improves an athlete's performance, but also has ethical costs, health costs, and potential costs of detection and punishment. Regardless of exactly how these different factors weigh on a given player's payoffs, as more neighbors are doping, a player is faced with tougher competition and has greater incentives to dope himself or herself just to maintain competitiveness. This game will generally be one of (strict) strategic complements. Nevertheless, all players would be better off if nobody doped compared to everyone doping. So conditions on incentives, such as strategic complementarities, do not necessarily imply orderings of overall payoffs without knowing more about the structure of the game. If the game is such that increased choices of action 1 by neighbors lead to higher payoffs, then indeed, higher-degree players will get higher payoffs (see Exercise 9.10); but the situation could also be reversed, as in the doping example, so that increased choices of action 1 by neighbors decrease payoffs.

While Proposition 9.5 is relatively straightforward to prove and quite intuitive, it concludes that the more-connected members of a society take higher action in situations with complementarities and lower action in those with substitutes. This observation is consistent, for instance, with the data of Coleman, Katz, and Menzel [165] discussed in Section 3.2.10. In the case of strict strategic complements, it also means that the equilibrium can be characterized in terms of a threshold degree, such that all players with degree above the threshold take action 1, and players with degree below it take action 0; the reverse holds for substitutes. It could be that players right at the threshold degree randomize. Figure 9.9 illustrates this threshold for a particular degree distribution. The figure shows a possible frequency distribution of neighbors' degrees, $\widetilde{P}$, as a function of degree. The threshold degree is such that players with higher degree take action 1 and those with lower degree take action 0. Thus in Figure 9.9, the probability of a neighbor taking action 1 is just the sum of the distribution of neighbors' degrees under $\widetilde{P}$ to the right of the threshold degree (adjusting for any mixing by the players exactly at the threshold).

9.4.2 Changes in Networks and Changes in Behavior

This analysis also leads to predictions of how behavior changes as we modify the distribution of neighbors' degrees. Suppose that we compare the equilibrium

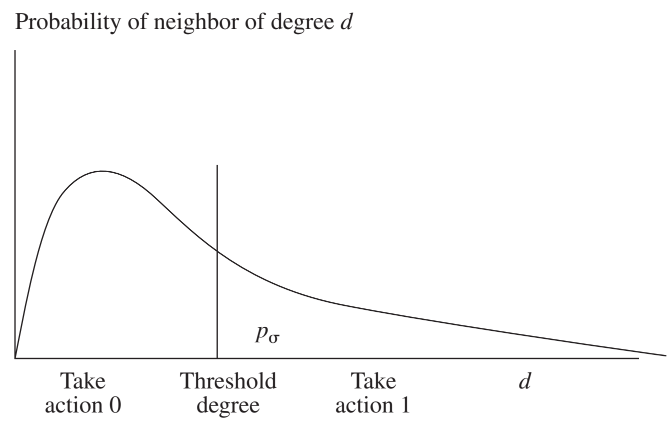

FIGURE 9.9 Behavior as a function of degree with complementarities.

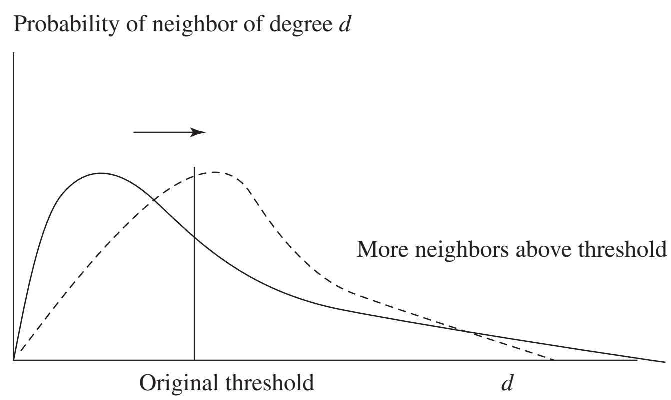

FIGURE 9.10 A shift in the degree distribution leads to more action.

behavior for the distribution pictured in Figure 9.9 with that for a different distribution, such as the dashed distribution in Figure 9.10. The new (dashed) degree distribution places more weight above the original threshold degree. If the equilibrium strategy does not change, then the shift would lead to a higher probability that any given neighbor plays action 1 and results in a higher expected number of neighbors taking action 1. Thus in the case of complements, players of any given degree now have a weakly higher incentive to play action 1 versus 0 than they did before. The threshold should then move down. As it decreases even more players have an incentive to play action 1, and so we move to a new equilibrium where even more neighbors play action 1.

This intuition is formalized in the following proposition. It is stated for the case of strategic complements, but also holds for the case of strategic substitutes, with an appropriate reversal of the direction of the shifts of thresholds and probabilities of action.

Proposition 9.6 (Galeotti et al. [274]) *Consider a network game of strict strategic complements that satisfies (9.5) and has a distribution of neighbors' degrees given by $\widetilde{P}$ and an equilibrium with threshold t. If the distribution of neighbors' degrees is changed to $\widetilde{P}'$ such that $\sum_{d\le t}\widetilde{P}'(d) \le \sum_{d\le t-1}\widetilde{P}(d)$, then there is an equilibrium threshold under $\widetilde{P}'$ that is at least as low as t, and the probability that any given neighbor chooses action 1 (weakly) increases. If instead $\sum_{d\ge t}\widetilde{P}'(d) \ge \sum_{d\le t+1}\widetilde{P}(d)$, then there is an equilibrium threshold under $\widetilde{P}'$ that is at least as high as t, and the probability that any given neighbor chooses action 1 (weakly) decreases.*

Note that Proposition 9.6 effectively allows us to compare any two degree distributions. The only complication is if the two distributions are very close and differ only at the threshold.[20]

Proof of Proposition 9.6. Let σ denote the equilibrium under $\widetilde{P}$, and consider the case $\sum_{d\le t}\widetilde{P}'(d) \le \sum_{d\le t-1}\widetilde{P}(d)$, as the other case is analogous. If σ is played under $\widetilde{P}'$, then there is a new probability $p'_\sigma \ge p_\sigma$ that any given neighbor will choose action 1. It is then easily verified that for any given degree d, this new distribution leads to a first-order stochastic dominance shift in the distribution of m, the number of neighbors who choose action 1. Given strict strategic complements, $u_d(1, m) - u_d(0, m)$ is an increasing function of m, and so given the first-order stochastic dominance shift, $U_d(1, p_\sigma) - U_d(0, p_\sigma)$, is at least as large as it was before for any d under the new distribution of neighbors' degrees. Thus $\sigma(d)$ is still a best response to σ for all $d > t$. If action 1 is a best response for degree-t players, then set their strategy to action 1. So, following the notation of the discussion after Proposition 9.5, we have strategy σ^t. Note also that if 1 is a best response to some strategy for a player of degree d, then action 1 is a best response for all higher-degree players; similarly if 0 is a best response for a player of degree d, then it is a best response for all lower-degree players. Then consider the best response of players of degree $t-1$ to σ^t under $\widetilde{P}'$. If it is action 0, then σ^t is an equilibrium and the conclusions of the proposition hold. Otherwise, move to strategy σ^{t-1}, and then consider the best responses of players of degree $t-2$. Continue in this manner until either stopping at some $\sigma^{t'}$ with $t' < t$, or hitting σ^1, in which case all players choosing action 1 is an equilibrium. In either case, the conclusions of the proposition hold. ■

Proposition 9.6 concludes that the probability of a neighbor choosing action 1 increases when the distribution of neighbors' degrees is shifted to place more weight above the threshold, but it does not claim that the probability that the overall fraction of players choosing action 1 increases at the new equilibrium. There is an important distinction between fractions of neighbors and fractions of players, which stems from the distinction between neighbors' degrees and players' degrees. Neighbors are more likely to be higher-degree players. The conclusion that the overall fraction of players choosing action 1 increases is valid if it is also

---

20. If the starting equilibrium involves no mixing by the threshold-degree players, then the conclusion also holds under the weaker conditions that $\sum_{d\le t-1}\widetilde{P}' \le \sum_{d\le t-1}\widetilde{P}$, or $\sum_{d\ge t}\widetilde{P}' \le \sum_{d\ge t}\widetilde{P}$, which then covers all possible comparisons.

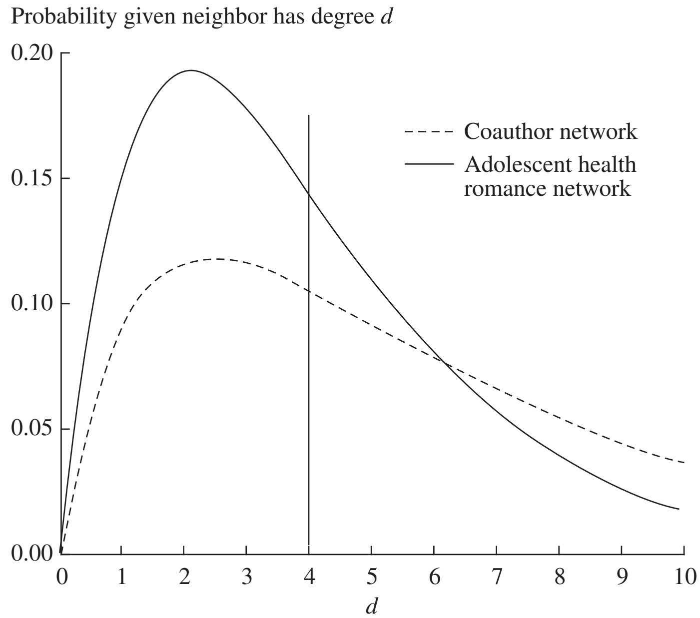

FIGURE 9.11 Distribution of neighbors' degrees for the romance network of Bearman, Moody, and Stovel [51] is first-order stochastically dominated by the distribution of neighbors' degrees for the coauthorship network of Goyal, van der Leij, and Moraga-González [304].

true that the weight that P' places below t is less than the weight that P places below $t-1$, where P' and P are the degree distributions corresponding to $\widetilde{P}'$ and $\widetilde{P}$, respectively. In many instances this condition holds, but one can find counter examples (e.g., see Galeotti et al. [274]).

To see the potential usefulness of Proposition 9.6 consider Figure 9.11. This figure shows two degree distributions from empirical studies. The first degree distribution is that of neighbors' degrees for the romance network of Bearman, Moody and Stovel [51] (from Figure 1.2; see Section 1.2.2). The second is a distribution of neighbors' degrees for a coauthorship network studied by Goyal, van der Leij, and Moraga-González [304]. This second distribution first-order stochastically dominates the first distribution and thus has greater weight above any threshold degree. While these distributions are from different applications, they show that examining empirically generated degree distributions allows for comparisons of the type treated in Proposition 9.6.

9.5 ▪ Richer Action Spaces

The analysis to this point has focused on situations with two actions. While two-action scenarios capture many applications and offer broad insights, there are settings in which the intensity of activity plays a substantial role. I present two such models. The first is a public goods model in which there are interesting implications for specializing in activities. The second is a model with complementarities that exhibits an interesting relationship between levels of activity and network centrality. For this section I return to the graphical games formulation.

9.5.1 A Local Public Goods Model

The following model analyzed by Bramoullé and Kranton [103] is a variation on a local public goods graphical game, like the one in Example 9.5, but the action space for each player is $X_i = [0, \infty)$. Players benefit from their own action plus the actions of their neighbors with payoffs described by

$$u_i(x_i, x_{N_i(g)}) = f(x_i + \sum_{j \in N_i(g)} x_j) - cx_i,$$

where f is a continuously differentiable, strictly concave function, and $c > 0$ is a cost parameter.

The interesting case occurs when $f'(0) > c > f'(x)$ for some large enough x, as otherwise optimal action levels are 0 or ∞. In this case, in every equilibrium each player's neighborhood has some production of the public good (at least with positive probability). Letting x^* be such that $f'(x^*) = c$, it is easy to see that any pure strategy equilibrium must have at least x^* produced in each player's neighborhood (so that $x_i + \sum_{j \in N_i(g)} x_j \geq x^*$ for each i); otherwise a player could increase his or her payoff by increasing his or her action. In fact, a strategy profile $(x_1, \ldots, x_n)$ is an equilibrium if and only if the following holds for each i:

- If $x_i > 0$, then $x_i + \sum_{j \in N_i(g)} x_j = x^*$; and
- If $x_i = 0$, then $\sum_{j \in N_i(g)} x_j \geq x^*$.

So a pure strategy equilibrium is any profile of actions such that every player's neighborhood produces at least x^*, a player only chooses a positive activity level if his or her neighbors produce less than x^* in aggregate, and in that case the player produces just enough to bring his or her aggregate neighborhood activity level to x^*. Figure 9.12 pictures some pure strategy equilibria for three-person networks in which $x^* = 1$.

There is a class of equilibria, which Bramoullé and Kranton [103] refer to as *specialized equilibria,* where players either choose an action level of x^* or 0. Thus there are players who specialize in providing the information or public good (e.g., the opinion leaders of Katz and Lazarsfeld discussed in Section 8.1), and others who free-ride on their neighbors. Even though the action spaces are richer, the specialized equilibria have the same structure as those of the pure strategy equilibria in the best-shot public goods graphical games. That is, the specialized equilibria are precisely those where the players who specialize in providing the local public good at the level x^* form a maximal independent set and the remaining players choose an action of 0.

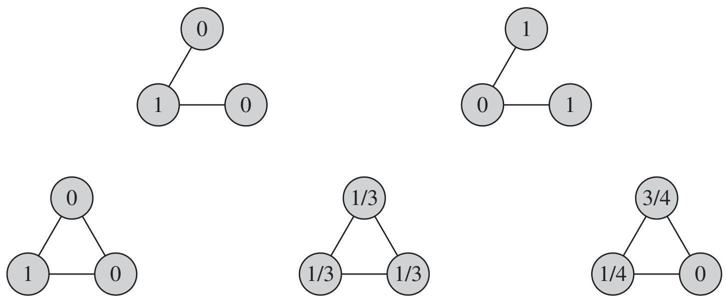

FIGURE 9.12 Examples of equilibrium local public good provision levels of $x^* = 1$. The numbers in each node indicate the node's action in Figures 9.12 and 9.13.

There is a sense in which the specialized equilibria are more robust than other equilibria. As Bramoullé and Kranton [103] point out, only specialized equilibria (and in fact only a subset of them) satisfy the following concept of stability:[21]

- Start with a pure strategy equilibrium profile $(x_1, \ldots, x_n)$.
- Perturb it slightly by adding some small perturbation ε_i to each x_i, with a requirement that $x_i + \varepsilon_i \geq 0$. Denote this perturbation by $x^1 = (x_1 + \varepsilon_1, \ldots, x_n + \varepsilon_n)$.
- Consider the best responses to x^1. That is, for each i, find a level of action that maximizes u_i, presuming that x^1_{-i} will be played by the other players. Let this profile of best-responses be denoted $x^2 = (x^2_1, \ldots, x^2_n)$, where x^2_i is the best response to x^1_{-i}.
- Iterate on the best responses, at each step examining the best responses x^k of the players to the previous step's strategies x^{k-1}.

If there is some $\overline{\varepsilon} > 0$ for which this process always converges back to $(x_1, \ldots, x_n)$ starting from any admissible perturbations such that $|\varepsilon_i| < \overline{\varepsilon}$ for all i, then the original equilibrium is said to be stable.

In this setting, the best responses to some x_{-i} take a simple form: if $\sum_{j \in N_i(g)} x_j \geq x^*$, then the best response is an action of 0; otherwise a best response is the action that raises the neighborhood production to x^*: $x^* - \sum_{j \in N_i(g)} x_j$. It follows fairly directly that the only stable equilibria, if they exist, are specialized equilibria such that each nonspecialist player has at least two specialists in his or her neighborhood, and each specialist has no neighbors choosing x. With at least two specialists in every nonspecialist's neighborhood, even a slight perturbation leads to a best-response dynamic in which the nonspecialists return to an action of 0 and the specialists return to an action of x^*. So such equilibria are stable. If there are fewer than two specialists in some nonspecialist's neighborhood, then the equilibria are unstable. The proof of this takes a bit more argument, but

21. This is a classic notion of stability that has been used in a variety of settings. See Chapter 1 in Fudenberg and Tirole [263] for more discussion and references.

to see the basic idea, consider a dyad for which there is no equilibrium with two specialists. It is easy to see that there are no stable equilibria: consider any pure strategy equilibrium, which must be such that $x_1 + x_2 = x^*$. At least one of the two strategies is larger than 0, so suppose that $x_2 > 0$. Then perturb the strategies to $x_1 + \varepsilon$, $x_2 - \varepsilon$. This configuration is also an equilibrium for any $\varepsilon \leq x^2$, and so the best responses do not converge back to the original point and no equilibrium is stable. In Figure 9.12, only the upper right-hand equilibrium is a stable one, and there are no stable equilibria for the complete network.

It is also worth noting that all equilibria in this public goods game are inefficient, in the sense that they stop short of maximizing total utility. Each player is maximizing his or her own payoff, and yet his or her action also benefits other agents. For example, in a dyad the total production is x^*, which maximizes $f(x) - cx$, while overall societal utility is $2f(x) - cx$, which will generally have a higher maximizer if f is smooth and strictly concave. This inefficiency is endemic to public goods provision, and players generally underprovide public goods, because they do not fully account for the benefits that their actions bestow on others. In this setting, there are also some differences across equilibria in terms of the total utility they generate. For example, if we consider the two different equilibria in the two top networks in Figure 9.12, they result in different aggregate payoffs. The one with one specialist on the left results in a payoff of $3f(1) - c$, while the one with two specialists on the right results in a payoff of $f(2) + 2f(1) - 2c$. Here we can rank these two equilibria, as the one on the left generates more total utility. We see this ranking by noting that the difference between the one on the left and that on the right is $c - f(2) + f(1)$. Since $x^* = 1$ and f is strictly concave, it follows that $c > f(2) - f(1)$ (otherwise, a player would prefer to increase the production to 2 even by himself or herself), and so this difference is positive. For more general networks, the comparisons across equilibria depend on the specific configurations and payoffs, but we can conclude that equilibria are generally inefficient, so that the stable equilibria are not always the most efficient equilibria.

This analysis of local public goods, although stylized, provides us with some basic insights into the emergence of individuals who provide local public goods, such as information, and who might act as opinion leaders, while other individuals free-ride, benefiting from this activity while not providing any benefit themselves. Significantly, the only stable equilibria are actually the asymmetric ones, even in very symmetric networks, and so this heterogeneity among individuals emerges because of the network interactions, even when there is no other a priori difference among individuals. While there can exist multiple equilibria, this analysis does not always offer useful predictions as to who will become providers or opinion leaders. Introducing heterogeneity into costs and benefits across individuals can help cut down the multiplicity of equilibria.

Before leaving this model, I comment on an aspect of the predictions of specialization and equilibria in graphical games more generally. Consider a simple example with four players and $x^* = 1$, and the individuals are connected in a circle network, as in Figure 9.13. In this case, there are only three pure strategy equilibria. According to the stability notion in Section 9.5.1, the two specialist equilibria are the only stable ones. However, if there is any cost to maintaining a link, neither of those networks would be pairwise stable in the sense of Section 6.1, as a player would not maintain a link to a neighbor who provides no public good. In contrast,

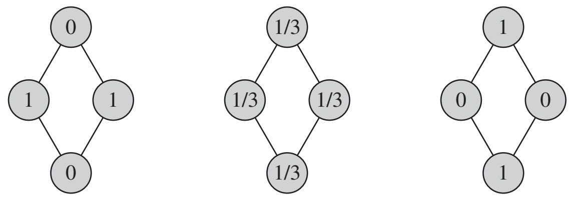

FIGURE 9.13 Examples of equilibrium public good provision choices.

there are situations in which the middle network in Figure 9.13 is pairwise stable (in particular, when the cost of a link is less than $f(1) - f(2/3)$ and greater than $f(4/3) - f(1)$).

The implication is that the graphical games analyses that we have been conducting are affected by considering the network to be endogenous. Indeed, people form relationships with others who provide local public goods, and they are expected to reciprocate in some fashion.[22] This is not to say that the insights behind specialization drawn from the analysis above are flawed, but that they need to be explored in a larger context. For example, if players were involved in two separate local public goods problems at the same time, with some players specializing in becoming informed about political campaigns and others about good local restaurants, then it could be possible to support specialized equilibria in conjunction with each other and an endogenous network. Some of the interplay between network formation and behavior on networks is examined in Section 11.4.3, but is still a largely underexplored subject.

9.5.2 Quadratic Payoffs and Strategic Complementarities

The public goods model described in the previous section is one in which activities are strategic substitutes. Let us now examine a different model, in which actions are again continuously adjustable, but there are strategic complementarities between players' actions. The following is a variation of the model of Ballester, Calvó-Armengol, and Zenou [34], which admits strategic complementarities.

Each player chooses an intensity with which he or she undertakes an activity. Let $x_i \in \mathbb{R}_+$ indicate that intensity, so that higher x_i corresponds to greater action. A player i's payoff is described by

$$u_i(x_i, x_{-i}) = a_i x_i - \frac{b_i}{2} x_i^2 + \sum_{j \neq i} w_{ij} x_i x_j,$$

where $a_i \geq 0$ and $b_i > 0$ are scalars, and w_{ij} is the weight that player i places on j's action. If $w_{ij} > 0$, then i and j's activities are strategic complements, so that more activity by j leads to increased incentives for activity by i; while if $w_{ij} < 0$,

22. One can also consider explicit transfers as a means for sustaining specialist equilibria. For more on transfers and stability, see Section 11.6.

then i and j's activities are strategic substitutes and increased activity by j lowers i's activity. The expression $\frac{b_i}{2}x_i^2$ leads to diminishing returns from the activity for player i, so that player i sees some trade-off to taking the action, ensuring a well-defined optimal strategy.

The payoff-maximizing action for player i is found by setting the derivative of the payoff $u_i(x_i, x_{-i})$ with respect to the action level x_i equal to 0, which leads to a solution of

$$x_i = \frac{a_i}{b_i} + \sum_{j \neq i} \frac{w_{ij}}{b_i} x_j. \tag{9.6}$$

The interdependence between the players' actions is quite evident.

Let $g_{ij} = w_{ij}/b_i$ (and set $g_{ii} = 0$). We can think of g as a weighted and directed network.[23] This variable captures the relative dependence of i's choice of action on j's choice. The vector of actions that satisfy (9.6) is described by[24]

$$x = \alpha + gx, \tag{9.7}$$

where x is the $n \times 1$ vector of x_i and α is the $n \times 1$ vector of a_i/b_i. If $a_i = 0$ for each i, then (9.7) becomes $x = gx$, so that x is a unit right-hand eigenvector of g. Otherwise,

$$x = (\mathbb{I} - g)^{-1}\alpha, \tag{9.8}$$

where $\mathbb{I}$ is the identity matrix, provided $\mathbb{I} - g$ is invertible and the solution is nonnegative. These two conditions hold if the b_is are large enough so that the entries of g are small.[25] From (9.7), by substituting for x repeatedly on the right-hand side, we also see that

$$x = \sum_{k=0}^{\infty} g^k \alpha. \tag{9.9}$$

Equations (9.8) and (9.9) have nice interpretations. They are variations on the centrality indices discussed in Section 2.2.4. The intuition is similar. Being linked to players who are more active (have higher levels of x_i) leads a player to want to increase his or her level of activity, presuming g nonnegative. Correspondingly, the more active a player's neighbors' neighbors are, the more active the player's neighbors are, and so forth. Thus, the activity levels in the system depend on activity levels. The fact that the payoffs are quadratic in the Ballester, Calvó-Armengol, and Zenou [34] model leads to a precise relationship to centrality measures, but even more generally, we would expect similar effects to be present.

---

23. It could even allow for some negative weights, depending on the values of w_{ij}.

24. Finding solutions to this problem is related to what is known as the *linear complementarity problem.* See Ballester and Calvó-Armengol [33] for a discussion of the relation and more general formulations of such games.

25. A sufficient condition is that the sum of the entries of each row of g be less than 1 and the sum of entries in each column of g be less than 1, in the case in which they are nonnegative.

To develop this concept a bit further, Ballester, Calvó-Armengol, and Zenou [34] also examine the case in which $a_i = a$ and $b_i = b$ for all i, so that the only heterogeneity in the society comes through the weights w_{ij} in the network of interactions. In that case, the equilibrium levels of actions in (9.8) can be written as

$$x = (I\!I - \frac{1}{b}w)^{-1}\frac{a}{b}\mathbb{1}, \tag{9.10}$$

where $\mathbb{1}$ is a vector of 1s. Equation (9.10) looks very much like the equations for Katz prestige-2 (2.9) and Bonacich centrality (2.10). In fact, we can write[26]

$$x = \frac{a}{b}\left(\mathbb{1} + P^{K2}\big(w, \frac{1}{b}\big)\right), \tag{9.11}$$

where $P^{K2}(w, \frac{1}{b})$ is the Katz prestige-2 from (2.9) (which is the same as the Bonacich centrality $Ce^B(w, \frac{1}{b}, \frac{1}{b})$ from (2.10)). To ensure that x is well defined, the term $1/b$ has to be small enough so that the Katz prestige-2 measure is well defined and nonnegative. There are various sufficient conditions, but ensuring that the rows (or columns) of w/b each sum to less than 1 (presuming they are all nonnegative) is enough to ensure convergence.

There are some clear comparative statics. If we decrease b or increase a, then the solution x increases, and the action of every player increases. There is a direct effect of making higher levels of x_i more attractive, fixing the level of the other player's actions, which then feeds back to increase other neighbors' actions, which further increases incentives to increase player i's action, and so forth. Presuming that w is nonnegative, increasing an entry of w, say w_{ij}, increases the equilibrium actions of all players who have directed paths to i in w. This observation takes a bit of proof and can be shown via different methods. A direct technique is to start at an equilibrium x, increase w, and then consider each player's best response to x at the new w. Player i's best response will be higher, as he or she has an increased benefit from neighbors' actions. Iterating on the best responses, any player ℓ such that $w_{\ell i} > 0$ will increase his or her action in response to i's higher actions, and those having links to ℓ will increase their actions, and so forth. Actions will only move upward, and so convergence is monotone upward to a new equilibrium, provided that an equilibrium is still well defined.[27]

---

26. To see (9.11) note that (9.9) implies that

$$x = \left(1 + (\frac{1}{b})w + (\frac{1}{b})^2 w^2 \ldots\right)\frac{a}{b}\mathbb{1},$$

whereas the corresponding Katz prestige-2 from (2.7) is

$$P^{K2}(w, \frac{1}{b}) = \left(\frac{1}{b}w + (\frac{1}{b})^2 w^2 \ldots\right)\mathbb{1}.$$

27. Another way to see the increase in the equilibrium action levels is to examine (9.9), noting that the entries of g^k increase in some row j for some large enough k if and only if there is a directed path from j to i in g; note that no entries decrease.

This model provides a tractable formulation that shows how actions relate to network position in a very intuitive manner. Its tractability also allows the equilibrium to be studied in more detail, given its closed form. For example, one interpretation of the above model that Ballester, Calvó-Armengol, and Zenou [34] pursue is that players are choosing levels of criminal activity. A player sees direct benefits ($a_i x_i$) and costs ($-b_i x_i^2$) to crime, and there are also interactive effects if there are complementarities with one's neighbors ($\sum_{j \neq i} w_{ij} x_i x_j$), so that more criminal activity by player's neighbors leads to greater benefits from crime to that player. This enhancement can be due to coordination, if cooperation results in more effective criminal activity, or it might reflect activities such as learning from neighbors. In the context of criminal activity, a natural question is which player should be removed to have the maximal impact on actions? For instance, if some police authority wants to lower criminal activity and can remove a single player, which player should it target? If the a_i and b_i terms coincide, then the interactive effects boil down to the centrality measure and the structure of w. The problem is then to compare players' activity levels when all players are present to equilibrium activity levels when a player is removed. By (9.11), this problem is equivalent to asking how removing one player affects the Katz prestige-2 measures. Ballester, Calvó-Armengol, and Zenou [34] show that the largest reduction in total activity comes from removing the player with the highest value of a variation on the Katz measure, which adjusts for the extent to which a player's prestige comes from paths back to himself or herself.

9.6 ▪ Dynamic Behavior and Contagion

The analyses of behavior up to this point are static in that they examine equilibrium behavior. In many situations, we are interested in the extent to which a new behavior diffuses through a society. For instance, if a new movie opens or a new product becomes available, how many people will take advantage of it? If there are complementarities in the product, so that a person is more likely to want to purchase it if another does, then the system may well have multiple equilibria, but simply examining them does not give us a full picture of which behavior is likely to emerge.

A powerful way of answering such questions is by examining the best-response behavior of a society over time. We have already seen some uses of iterating on best responses in checking for stability and identifying equilibria, but such iteration has been a prominent dynamic for more general analyses. That is, start by having some small portion of the population adopt a new action, say action 1. Then in a situation with complementarities, we can see how their neighbors respond. How many of them buy the product in response? This response then leads to further waves of adoption or diffusion. This process has been examined in variations on network settings by Ellison [221] and Blume [81], among others. An analysis that ties directly to the graphical games setting is one by Morris [487], which helps illustrate some interesting ideas.

Morris [487] considers a semi-anonymous graphical game with strategic complements. He examines a case in which each player responds to the fraction of

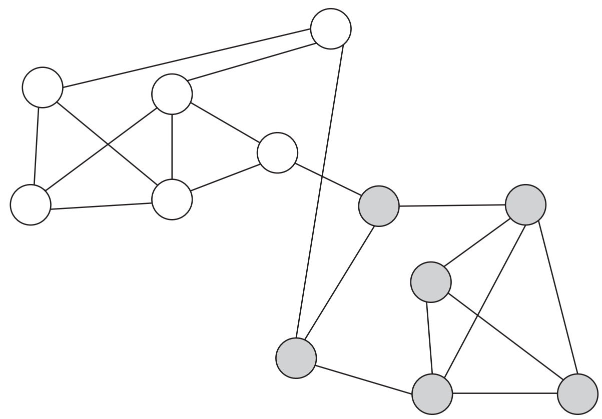

FIGURE 9.14 The sets of light and shaded nodes are each 2/3-cohesive.

neighbors playing action 1 versus 0. There is a threshold fraction q such that action 1 is a best response for a given player if and only if at least that fraction q of the player's neighbors choose 1. This fraction is the same for all players, independent of their degrees.[28] In the nontrivial case for which q lies strictly between 0 and 1, this is effectively a coordination game, and there are at least two equilibria, one where all players choose action 0 and the other where all players choose action 1.

What else can we deduce about equilibrium structure? For example, when is it possible that both actions coexist in a society, so that there is an equilibrium where some nonempty strict subset of the society plays action 1 and the rest plays action 0? Let S be the subset of the society that plays action 1. Each player in S must have at least a fraction q of his or her neighbors in the set S. It must also be that every player outside S has a fraction of no more than q of his neighbors in S, or equivalently, has a fraction of at least $1-q$ of his neighbors outside S.

To capture these conditions, given $1 \geq r \geq 0$, Morris [487] defines the set of nodes S to be r-*cohesive* with respect to a network g if each node in S has at least a fraction r of its neighbors in S. That is (recalling (2.15)), S is r-cohesive relative to g if

$$\min_{i \in S} \frac{|N_i(g) \cap S|}{d_i(g)} \geq r,$$

where $0/0$ is set to 1. Figure 9.14 illustrates this definition with disjoint sets of nodes that are each $2/3$-cohesive.

If a set is such that each player has at least some fraction r of his or her neighbors within the set, then it is easy to see that each player must have at least a fraction r' of his or her neighbors within that set when $r' < r$. So define the *cohesiveness*

28. This is a special case of complements games in which the threshold in terms of numbers of neighbors of degree d is simply qd.

of a given set S relative to a network (N, g) to be the maximum r such that S is r-cohesive. Then we have the following proposition, for which the proof is direct.

Proposition 9.7 (Morris [487]) *Consider a network (N, g) and a coordination game such that action 1 is a best response for any player if and only if a fraction of at least q of his or her neighbors play action 1. Both actions are played by different subsets of the society in some pure strategy equilibrium if and only if there exists some nonempty and strict subset of players S that is q-cohesive and such that its complement $N \setminus S$ is $(1-q)$-cohesive.*

An obvious sufficient condition for both actions to be played in equilibria is to have at least two separate components, as then different actions can be played on each component. The cohesiveness of a component is 1, and thus it is also q- and $(1-q)$-cohesive for any q.

Beyond components, cohesiveness provides enough of a separation in a network for different equilibria to exist adjacent to one another. For example, in Figure 9.14, the gray and white sets of nodes are connected to each other, but both are sufficiently inward-looking so that they can each sustain different equilibria in any game with q between 1/3 and 2/3.

Morris [487] also asks the following question.[29] Consider a given network (N, g) and start with all nodes playing action 0. "Infect" some number m of the nodes by switching them to play action 1 (and they can never switch back). Let players (other than the initially infected) best respond to the current action of their neighbors, switching players to action 1 if their payoffs are at least as good with action 1 as with action 0 against the actions of the other players. Repeat this process starting from the new actions, and stop when no new players change to action 1. If there is some set of m nodes whose initial infection leads to all players taking action 1 under the best response dynamic, then we say that there is *contagion from m nodes*. Let us say that a set S is *uniformly no more than* r-cohesive if there is no nonempty subset of S that is more than r-cohesive. We then have the following proposition.

Proposition 9.8 *Consider a network (N, g) and a coordination game such that action 1 is a best response for any player if and only if a fraction of at least q of his or her neighbors play action 1. Contagion from m nodes occurs if and only if there exists a set of m nodes such that its complement is uniformly no more than $(1-q)$-cohesive.*

Proof of Proposition 9.8. Consider a set S of m nodes. If its complement has some subset A that is more than $1-q$-cohesive, then that set A of nodes will all play 0 under the process above, at every step. Thus it is necessary for the complement to be uniformly no more than $(1-q)$-cohesive to have contagion to all nodes. Next let us show that this condition is sufficient. Since the complement is uniformly no more than $(1-q)$-cohesive, then it is no more than $(1-q)$-cohesive. Thus there must be at least one player in the complement who has at least a fraction of q of his or her neighbors in S. So, at the first step, at least one player changes

29. Morris [487] works with infinite networks. I have adapted his formulation and results to a finite setting to be comparable to graphical games.

strategies. Subsequently, at each step, the set of players who have not yet changed strategies is no more than $(1-q)$-cohesive, and so some player must have at least q neighbors who are playing 1 and will change. Thus as long as some players have not yet changed, the process will have new players changing, and so every player must eventually change. ■

A set is uniformly no more than $(1-q)$-cohesive if every subset of at least one node has more than q of its neighbors outside that subset. Thus such a network is quite dispersed in terms of its connections and does not have any highly segregated groups.

Corollary 9.1 *Consider a network (N, g) and a coordination game such that action 1 is a best response for any player if and only if a fraction of at least q of his or her neighbors play action 1. If a set S of nodes is uniformly no more than r-cohesive, then there will be contagion starting from the complement of that set, provided $q \leq 1-r$.*

While cohesion is an easy concept to state and provides for compact and intuitive characterizations of contagion, it is not always an easy condition to verify. Part of this difficulty is simply because there are 2^n different subsets of players in any given network of n players, and so checking the cohesion of each subset becomes impractical even with relatively few players. Thus verifying whether a given network is r-cohesive needs to take advantage of some structural characteristics of the network. To obtain some feel for the cohesion of different network structures, let us examine the cohesion of a few types of networks, beginning with the simplest case of the complete network. A subset of players S of a complete network is $\frac{|S|-1}{n-1}$-cohesive, since each player in S has $|S|-1$ neighbors in S and the remaining $n-1$ neighbors outside S. As the size of the set S grows, so does its cohesion. As the number of players becomes large and the S becomes large relative to N, the cohesion approaches 1, which makes contagion impossible except for q approaching 0. Clearly, this type of network is extraordinary in at least two ways: the degree of each player is large, and every subset of agents forms a clique so that the network is highly clustered. Let us examine the opposite extreme of a tree network in which all agents have degree of no more than some d. To keep things simple, consider a tree in which all agents have degree d or 1. Here it is easy to hit the upper bound of a strict subset being $\frac{d-1}{d}$-cohesive. To see this, simply pick a subtree, as in Figure 9.15, so that there is only one link from one player to the rest of the network. That player has $\frac{d-1}{d}$ of his or her neighbors in the subtree, and the other players in the subtree have all their neighbors in the subtree. As we know that many random networks have some subsets of nodes that are nearly tree-like in their structure when n is large (e.g., see Exercise 1.2), there are many networks in which the cohesion of some subsets is quite high, and so contagion requires a low threshold.[30]

30. Morris [487] examines only connected networks on an infinite set of nodes, and for those he shows that an upper bound on the contagion threshold is 1/2 (by showing that every co-finite set—a set with a finite complement—contains an infinite subset that is 1/2 cohesive). These networks are not good approximations for (even very) large finite networks, as we see from the high cohesiveness of various finite networks.

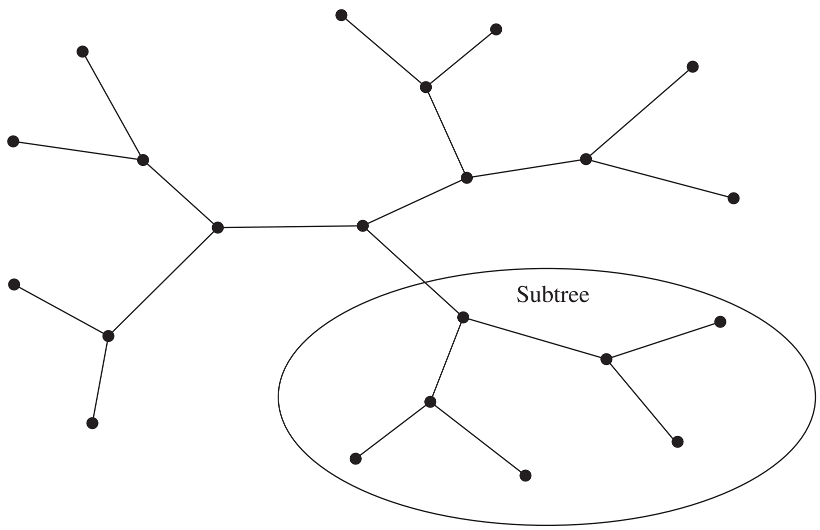

FIGURE 9.15 A subtree that is 2/3-cohesive.

When bridge links are present in a network (so that deleting that link would lead to a new component structure), then it is clear that the two sets of nodes that are bridged each have relatively high cohesion of $\frac{d_i-1}{d_i}$, where i is the bridging node. This cohesion makes contagion difficult and the support of different actions in equilibrium relatively easier.

Contagion is demanding in that it requires that all nodes be reached. Thus a network that has even a few players who are very cohesive among themselves will fail to be susceptible to contagion using the above definition, even though action 1 might spread to almost the entire network. The ideas behind cohesion still provide partial characterizations, as it must be that there is some large set that has low cohesion in nodes to have substantial contagion. A precise characterization of partial contagion is the subject of Exercise 9.17. To better grasp the multiplicities of equilibria and the partial diffusion of actions, let us return to the network games setting.

9.7 ▪ Multiple Equilibria and Diffusion in Network Games*

The various models we have explored often have multiple equilibria, and some of the analyses consider aspects of stability and contagion such as movement from one equilibrium to another. Indeed, the multiplicity of equilibria is important in many applications, and there are many case studies (e.g., see Rogers [564]) in which for some cases a product of behavior diffuses broadly while in other cases it does not. The analysis of contagion in Section 9.6 gives us some feel for the notion but is extreme in requiring that an action be adopted by the entire population. As we have seen in the network games setting and many case studies, there is often some heterogeneity in a population, with some people adhering to one behavior and others adopting a different behavior. Introducing some heterogeneity among players, beyond their degrees, can actually help in producing a more tractable analysis of the structure and stability of multiple equilibria and the diffusion of behavior.

9.7.1 Best-Response Dynamics and Equilibria

An analysis of such diffusion is performed by Jackson and Yariv [364] (see also [362]) in the following context. The setting is similar to the network game setting with a couple of modifications.

Players all begin by taking action 0, which can be thought of as a status quo, for instance, not having bought a product, not having learned a language, or not having become educated. Players are described by their degrees, which indicate the number of future interactions they might undertake. Players' preferences are as in network games, described by (9.4), with one variation. The players also can have some idiosyncratic cost of taking action 1, which is described by c_i. This cost captures the idea that some players might have a personal bias toward buying a given product, or a proclivity or aversion to learning a language, or becoming educated, or the like. If the c_is are all 0, the model reduces to the network games that we considered before. But in the more general model, the payoff to a player is

$$U_{d_i}(0,\ p)$$

if the player stays with action 0, where p is the probability that any given neighbor chooses action 1, U is the network game payoff as described by (9.4), and

$$U_{d_i}(1,\ p) - c_i$$

if the player switches to adopt action 1.

Without loss of generality, normalize the payoff to adhering to action 0 to be 0, so that $U_{d_i}(1,\ p) - c_i$ captures the change in payoffs from switching to action 1 for a given player of degree d_i with idiosyncratic cost c_i and faced with a probability of neighbors' adoptions of p. Thus player i prefers to switch to action 1 when

$$U_{d_i}(1,\ p) \geq c_i.$$

Let us focus on the case of strategic complements, so that the player's payoff from switching to action 1, $U_d(1,\ p)$, increases with the probability p of neighbors taking action 1. The case of strategic substitutes is examined in Exercise 9.18.

Let F describe the distribution function of costs, so that $F(c)$ is the probability that any given player's cost c_i is less than or equal to c. Then the probability that a player of degree d prefers action 1 is the probability that his or her cost is less than the benefit from adopting action 1 of $U_d(1,\ p)$ and so that probability is

$$F(U_d(1,\ p)).$$

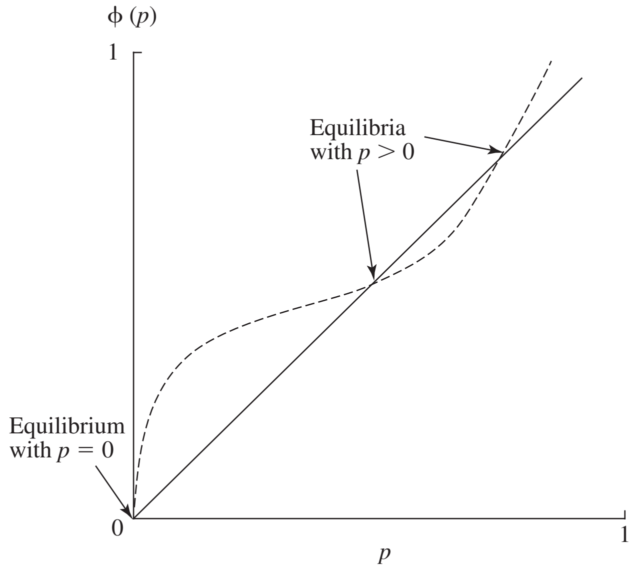

FIGURE 9.16 Resulting fraction of neighbors who choose action 1 ($\phi(p)$) as a best response to a fraction of neighbors who choose action 1 (p).

Now consider the following dynamic.[31] Start with some beginning probability that a player's neighbors will choose action 1, say, p^0. Players then best respond to p^0, which results in a new fraction of players who wish to adopt action 1, p^1, and so forth. Under strategic complements, this process is monotone, so that players never wish to switch back to action 0 as the adoption increases over time. In particular, given a probability that a neighbor chooses action 1, p^t, the new probability that a neighbors chooses action 1 at $t+1$ is

$$p^{t+1} = \phi(p^t) \equiv \sum_d \widetilde{P}(d) F(U_d(1, p^t)). \tag{9.12}$$

An equilibrium corresponds to a probability p of neighbors' choosing action 1 such that $p = \phi(p)$. Figure 9.16 pictures a hypothetical function $\phi(p)$, indicating the best-response levels of action 1 as a function of the starting level of action 1.[32] Figure 9.16 shows three different equilibria. There are also situations, as in Figure 9.17, in which a unique equilibrium occurs. In that figure, $\phi(0) > 0$, so that there are some players who choose action 1 regardless of whether any other players do.

---

31. This "dynamic" has various interpretations. It can explicitly be a dynamic, or it might also simply be a tool to define stability of equilibria and study the properties of various equilibria.
32. This sort of analysis of the multiplicity of, dynamics leading to, and stability of equilibria draws from quite standard techniques for analyzing the equilibria of a system. For example, see Fisher [250] for a survey of the analysis of the equilibria of economic systems and Granovetter [306] for an application of such techniques to a social setting.

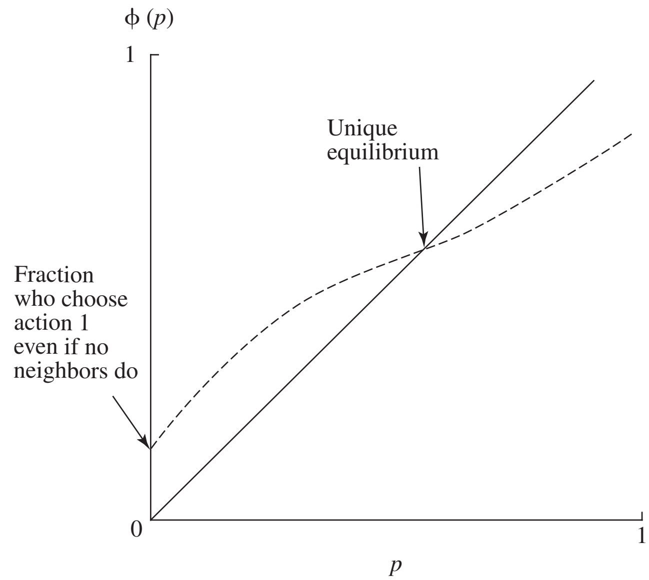

FIGURE 9.17 A unique equilibrium.

9.7.2 Stability

As discussed above, there are various notions of the stability of equilibria. Using dynamics described by best responses and the function ϕ, we can define stability as in Section 9.5.1. That is, start at some equilibrium $p = \phi(p)$, and then perturb p to $p + \varepsilon$ or $p - \varepsilon$ for some small ε, with the constraint that the perturbed probability lie in [0,1]. If iterating on ϕ from this point always converges back to p for small enough ε, then the equilibrium is stable; if it does not for arbitrarily small ε, then it is unstable. A stable equilibrium is pictured in Figure 9.18.

Generally, if ϕ cuts the 45-degree line from above, then the equilibrium is stable; if it cuts the 45-degree line from below, then it is unstable.[33] Figure 9.19 shows a multiplicity of equilibria, some stable and others not. The figure shows some interesting aspects of equilibria. The equilibrium at 0 is unstable, and the next higher one is stable. There is a tipping phenomenon such that if p is pushed above 0, then the best-response dynamics leads behavior upward to the higher stable equilibrium. Thus if some initial adoption occurs, then behavior diffuses up to the higher stable equilibrium.

33. It is possible to have ϕ be tangent to the 45-degree line, in which case the equilibrium is unstable.

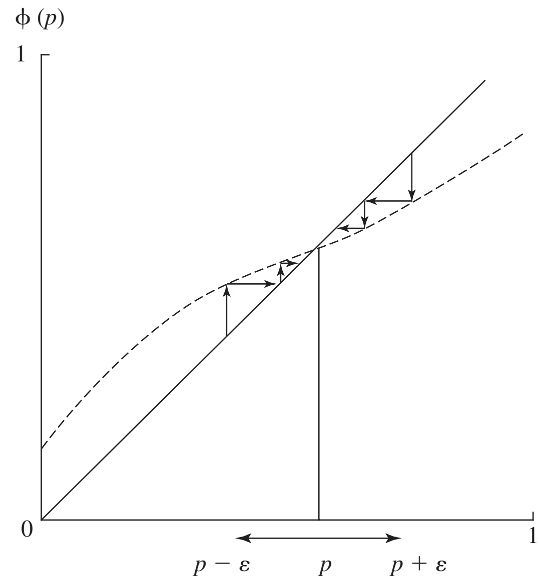

FIGURE 9.18 A stable equilibrium.

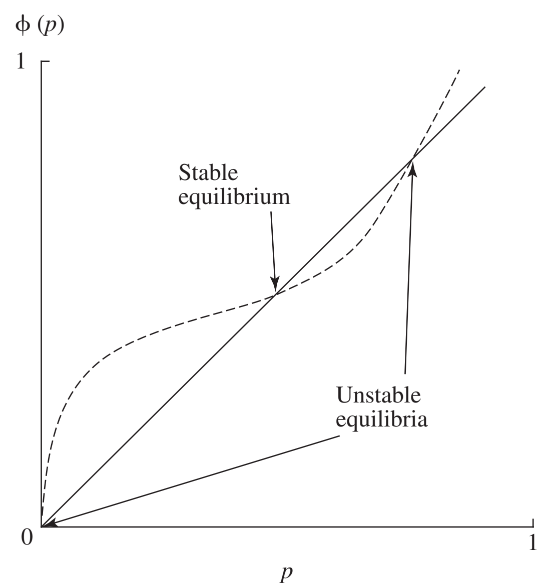

FIGURE 9.19 Multiple equilibria: some stable, some not.

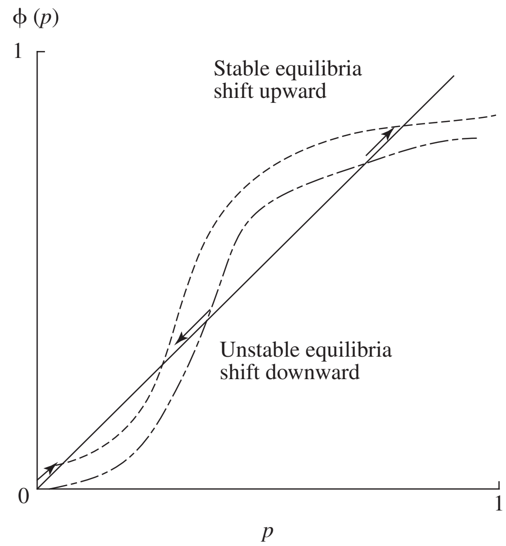

FIGURE 9.20 Change in equilibria due to a shift in best responses.

9.7.3 Equilibrium Behavior and Changes in the Environment

A change in the environment, such as an alteration to the cost of adopting action 1, to the relative attractiveness of the two actions, or to the network structure, leads to a change in the best-response function ϕ.

If the shift is systematic, so that $\phi(p)$ shifts up (or down) for every p, then we can deduce how equilibria change. For instance, in Figure 9.20 the best response to any p is higher under the dashed curve. The figure shows that as ϕ shifts upward, the unstable equilibria move down and the stable equilibria move up. This reconfiguration makes it easier to reach tipping points (the unstable equilibrium) and leads dynamics to reach higher equilibria, so that the diffusion of behavior is more prevalent in a well-defined way. Thus any changes in the setting that result in systematic shifts in ϕ lead to concrete conclusions about how equilibrium behavior will respond, even in the presence of multiple equilibria.

To be careful, these conclusions about changes in equilibria due to an upward shift in ϕ hold for sufficiently small shifts in ϕ and presume that ϕ is a continuous function, which is true if F and U are continuous. Given the continuity of ϕ, any stable equilibrium is locally unique; however, unstable equilibria may not be unique and may even be such that ϕ is tangent to the 45-degree line at some point. If ϕ is continuous, then if we shift $\phi(p)$ up at every point, then for every stable equilibrium there is a new equilibrium that is higher than the old one. In addition, for an unstable equilibrium p, it is possible that there is no longer any equilibrium at or below p.[34] In what follows, the conclusions based on shifts in the best-response function ϕ should be interpreted with these consequences in mind.

---

34. When $p = 1$ is initially an equilibrium, then $\phi(1)$ cannot increase, but any shift up of ϕ at other points would still leave 1 as an equilibrium. Note also the new equilibrium above a given

From (9.12) the best response to a given p is described by

$$\phi(p) = \sum_d \widetilde{P}(d) F(U_d(1, p)). \tag{9.13}$$

As Jackson and Yariv [364] point out, (9.13) makes clear that some types of changes systematically shift ϕ up. Let us examine such changes.

Lowering the Cost of Changing Actions Changing F, so that costs of adopting action 1 are lower (e.g., increasing $F(c)$ for each c), leads to a shift up in ϕ and so to lower tipping points and higher stable equilibrium choices of action 1.[35] Lowering the costs of taking action 1 corresponds to increasing the probability that the cost of choosing action 1 is below its benefit $U_d(1, p)$, so this change corresponds to increasing $F(U_d(1, p))$ for any given d and p. The result is an increase on the right-hand side of (9.13), so that indeed ϕ shifts up at every point.

Changes in Network Structure Beyond lowering costs, other alterations can lead to the same sorts of systematic changes observed above. For example, consider a first-order stochastic dominance shift from $\widetilde{P}$ to some new distribution of neighbors' degrees $\widetilde{P}'$. If the payoff to choosing action 1, $U_d(1, p)$, is increasing in degree d for any positive proportion $p > 0$ of neighbors taking action 1 and the distribution of costs F is increasing (which is true when the distribution corresponds to a continuous density function on the relevant range), then the result is again an increase on the right-hand side of (9.13) at every positive p and an upward shift of ϕ at every positive p. This shift then leads to an increase in stable equilibria and a lowering (or disappearance) of the unstable equilibria or tipping points, and similarly to the case of lowering of costs, we should expect higher overall diffusion in the sense that there are lower thresholds for diffusion and higher equilibria.

Note that we can compare this situation to Proposition 9.6, which did not distinguish between stable and unstable equilibria but did deal with shifts. Here, by accounting for all equilibria and seeing how ϕ adjusts, we have a more complete picture of how equilibria change with modifications to the environment. Proposition 9.6 only concluded that for any equilibrium under the old distribution, there is at least one that has moved up under the new distribution, so that the highest equilibrium is (weakly) higher than it was before.

As Jackson and Yariv [362] point out, we can also examine changes in terms of mean-preserving spreads of the degree distribution, recalling how such spreads affect the expectation of a convex function (see Section 4.5.5). The impact of such a change depends on the convexity of $F(U_d(1, p))$ as a function of d. If F is convex and increasing in d, then the change in equilibria can be well ordered, as $\phi(p)$

---

stable equilibrium might no longer be stable, as the new ϕ may be tangent to the 45-degree line at the higher equilibrium.

35. This statement presumes that ϕ is continuous after the change.

increases for every p. That is, a mean-preserving spread in the degree distribution, such as a change from a Poisson degree distribution to a scale-free degree distribution, leads to more diffusion of action 1 (in the sense of shifting stable equilibria to be higher and unstable equilibria to be lower). This increased diffusion might apply in the case of strategic complements, provided there are sufficient complementarities. In contrast, if the compound function $F(U_d(1, p))$ is concave and decreasing in d, then the shift is reversed. Thus if $F(U_d(1, p))$ is increasing and convex in d, then power, Poisson, and regular degree distributions with identical means generate corresponding best-response functions ϕ^{power}, ϕ^{Poisson}, and ϕ^{regular} such that $\phi^{\text{power}}(p) \geq \phi^{\text{Poisson}}(p) \geq \phi^{\text{regular}}(p)$ for all p.

The various graphical game and network game models examined provide a basis for understanding how equilibria behave as a function of payoffs, player degree or position in the network, and network characteristics. The multiplicity of equilibria make systematic conclusions a challenge, and we also see some sensitivity of the conclusions to the fine details of the setting. Nonetheless, such concepts as strategic complementarities provide powerful tools that allow us to draw some fairly specific conclusions about equilibrium properties and how they vary with network structure.

9.8 ▪ Computing Equilibria*

Beyond the analysis of graphical games and network games, it is also important to know how to compute an equilibrium and to appreciate the difficulty of finding one. This knowledge is not only useful for researchers or scientists exploring the behavior of a society but also important in determining whether the society will reach an equilibrium. There are a variety of ways that we might posit that players adjust their behavior, including deductive reasoning, communicating with others, updating strategies over time in response to what actions have been played in the past, or responding to evolutionary or other selective pressures over time. If computing an equilibrium using the full description of the game is so complicated that it is not expected to be done in finite time, it is hard to expect a system to lead to equilibrium (as otherwise we could mimic that system to compute equilibria). Also, knowing something about the multiplicity of equilibria is also important, for at least two reasons. Having many equilibria lowers predictive power, as more profiles of behavior are consistent with the model. And having many equilibria can make it difficult for players to coordinate or reach equilibrium, even if they can communicate with one another.

Obviously, these issues have not escaped game theorists' attention, and a good deal of attention has been devoted to understanding the multiplicities of equilibria; computing equilibria; modeling how societies might learn or evolve to play equilibria, refining equilibrium predictions; and studying focalness, social norms, and other methods of coordination. I will not try to distill such a breadth of material here and instead I refer the reader to standard game-theory texts, such as Binmore [70], Fudenberg and Tirole [263], Myerson [497], and Osborne and Rubinstein [514], to learn more about these issues. It is important, however, to see how these issues manifest themselves in graphical games. So let us begin by computing equilibria.

Computing equilibria for the threshold games described in Example 9.2 is quite easy and takes advantage of the strategic complementarities of the game, as we have already seen in special cases. Here is a method that takes at most $(n+1)n/2$ steps for any such threshold game.

Set all players' actions to 1. Now consider player 1. If player 1 would improve by changing to action 0, then do so and otherwise leave the profile of actions as it is. Do the same for player 2, given the new profile of actions. Continue iterating in this manner. After passing through all players, repeat the procedure. Stop adjusting actions when all players have been considered and no actions have changed.

This algorithm takes advantage of the fact that players' actions only change from 1 to 0, and given the strategic structure of the game, such changes can only lead other players to change from 1 to 0, but never to reverse their decisions. Actually, this algorithm finds the maximum equilibrium in the sense that there is no other equilibrium where any player ever takes a higher action (see Exercise 9.22). Moreover, the technique also finds the maximum equilibrium in a wider class of games in which strategies are ordered and there are such complementarities between strategies.[36]

Next, consider the best-shot public goods game in the case of an undirected network.[37] As noted, the pure strategy equilibria occur when the players who take action 1 form a maximal independent set. Maximal independent sets are easy to find. For instance, some of the maximal independent sets on a tree can be found as follows. The following set A is a maximal independent set:

$$A = \bigcup_{m: m \text{ is even}} D_i^m(g),$$

where $D_i^m(g)$ are the nodes at distance m from some i in g (with $D_i^0(g) = \{i\}$ being one of the sets where m is even).

Even when the network is not a tree, there still exist obvious (and fast) methods of finding maximal independent sets. Here is an algorithm that finds an equilibrium for a connected network (N, g). By applying it to each component, this method can be used for any network. At step k, the algorithm constructs sets A_k and B_k; the eventual maximal independent set will be the final A_k. In terms of finding an equilibrium for the best-shot game, the final A_k is the list of players who take action 1, and the final B_k is the set of players who take action 0:

Step 1: Pick some node i and let A_1 be i and B_1 be i's neighbors ($A_1 = \{i\}$ and $B_1 = D_i^1(g)$).

Step 2: Pick some node j at distance 2 from i ($j \in D_i^2(g)$), and let $A_2 = \{i, j\}$ and $B_2 = B_1 \cup D_j^1(g)$.

36. For more background on games with strategic complementarities, see Topkis [630], Vives [642], and Milgrom and Roberts [469].

37. For a directed network in the best-shot public goods game the analysis is a bit different, as outlined in Exercise 9.6.

Step k: Iterate by selecting one of the players j' who has a minimal distance to i out of those players not yet assigned to a set A_{k-1} or B_{k-1}. Let $A_k = A_{k-1} \cup \{j'\}$ and $B_k = B_{k-1} \cup D^1_{j'}$.[38]

End: Stop when $A_k \cup B_k = N$.

Although varying the starting player i and the order in which new players are chosen in each step k results in finding several different equilibria, there can be many more that are not found by this algorithm (see Exercise 9.23).

For the general class of graphical games, computing equilibria becomes more challenging. For example, consider a game in which players have thresholds for choosing action 1 but are also concerned about congestion, so that they do not take action 1 if too many neighbors choose action 1. In particular, a player has a lower threshold and an upper threshold, so that the player prefers action 1 if and only if the number of neighbors playing 1 lies between the two thresholds. Moreover, allow the lower and upper thresholds to differ across players. In such a setting, which might not admit any pure strategy equilibria, it could be hard to find even one equilibrium pure or mixed. Each time one player's strategy is adjusted, we may have to adjust the strategies of previously considered players in response, as these can feed back on one another. In the best-shot and threshold examples, changing a player's strategy in one direction had clear implications for how others respond, but more generally, as in the example with multiple thresholds, the feedback and interaction can be complex. The details of defining what is meant by a "hard to find" equilibrium are beyond the scope of this text, but here I sketch the basic ideas.

For any algorithm that computes equilibria, there are some inputs that describe the game. In the case of a graphical game, the inputs are the number of players, the network that connects them, and each player's payoff function. The number of players is simply n, and the information about the network can be coded in an $n \times n$ matrix, so that there are n^2 bits of information (although this number can be lowered in some classes of games). If each player has a degree of at most d, then each player's payoff matrix has 2^{d+1} entries, indicating the payoff as a function of each vector of choices of 0 or 1 for each neighbor and the player. Thus the full description of the game involves on the order of n^2 plus n times a constant (related to the maximum d) bits of information. Given this information, we construct an algorithm for finding an equilibrium. This involves prescribing a series of steps that use the information about the game to do some calculations and eventually spit out a list of strategies for each player. How many steps will it take to terminate with the determination of an equilibrium? The method of counting steps that is generally followed is to look at the upper bound, or worst possible performance. So the performance measure is to find a game and payoff structure that would lead to the most steps before the given algorithm finds an equilibrium and to keep track of this number as a function of n. Algorithms are considered to be relatively quick if the upper bound on the number of steps needed is at most some polynomial function of n. Slow algorithms require more steps than polynomial in n for some games, for example, when worst-case scenarios use a number of

38. Note that these sets are well defined, since no neighbors of j' can be in A_{k-1}, as otherwise j' would have been in B_{k-1}.

steps that grows exponentially with n. It is difficult to show that a problem is such that all algorithms sometimes require more than a polynomial number of steps. There is a deep and long-standing open question on which hinges the answer to complexity for a number of problems, including equilibrium computation in a class of graphical games.[39]

There are several side issues of interest here. For instance, is it reasonable to measure the performance of an algorithm for computing equilibria based on worst cases? Is an exponential number of steps really that much more than a high-order polynomial for some n? How large does n have to be before there is a serious distinction between polynomial and exponential time for the problem in question? Is this the right accounting for complexity, given that we are not considering the complexity of the calculation at each step? Are there large classes of graphical games for which the task is much easier, and are the graphical games for which computing an equilibrium is difficult very interesting? What happens when we look for approximate equilibria (so that players nearly maximize their payoff) instead of exact equilibria? Are there other definitions of equilibria for which computation is easy? These are all difficult and open questions that have received attention. There are also other difficulties that we face with graphical games. For instance, if players do not have some maximal degree, but the maximal degree grows with n, then describing payoffs could take up to 2^n bits of information. Part of the reason that the threshold and best-shot games were easy to handle is that the payoffs were quite simple to describe.[40]

Here I summarize what is known about finding equilibria of graphical games. When the network is a tree, there are algorithms that involve a polynomial in n number of steps and find an equilibrium of any graphical game on a tree, as shown by Littman, Kearns, and Singh [437]. Once we venture beyond trees, however, the strong conjecture is that no such algorithm exists with a number of steps that is always polynomial in n (see Daskalakis, Goldberg, and Papadimitriou [185]).[41] Although this conjecture is somewhat pessimistic with regards to being able to make predictions of behavior in the broad class of graphical games, there is often much more structure to the games of interest compared to the worst-case scenarios. As we have seen, strategic complementarities make finding equilibria easy and fast.[42]

---

39. For more background on algorithms and complexity, see Papadimitriou [523] and Roughgarden [570].

40. For more about the complexity of describing payoffs and representing such games see Daskalakis and Papadimitriou [183].

41. The strong conjecture is based on the fact that the problem of computing Nash equilibria in a graphical game has been shown to be equivalent to a problem (lying in a class called *PPAD-complete*) that is conjectured to have no polynomial time algorithm for finding a solution. That conjecture is among a class of long-standing open problems regarding the complexity of algorithms, which have received a great deal of attention.

42. Note also that while computing Nash equilibria in general graphical (and other large) games can be hard, there are polynomial-time algorithms for finding correlated equilibria (which are a generalization of Nash equilibria that admit correlation in the players' strategies) in certain graphical games that have nice representations. See Kakade et al. [373] and Papadimitriou and Roughgarden [525].

9.9 ▪ Appendix: A Primer on Noncooperative Game Theory

This appendix discusses what is known as *noncooperative game theory,* in which agents act in self-interested ways to maximize their own payoffs and equilibrium notions are applied to predict outcomes. Cooperative game theory examines coalitions and how payoffs might be allocated within coalitions. It is examined in Section 12.1.

The basic elements of performing a noncooperative game-theoretic analysis are (1) framing the situation in terms of the actions available to players and their payoffs as a function of actions, and (2) using various equilibrium notions to make either descriptive or prescriptive predictions. In framing the analysis, a number of questions become important. First, who are the players? They may be people, firms, organizations, governments, ethnic groups, and so on. Second, what actions are available to them? All actions that the players might take that could affect any player's payoffs should be listed. Third, what is the timing of the interactions? Are actions taken simultaneously or sequentially? Are interactions repeated? The order of play is also important. Moving after another player may give player i an advantage of knowing what the other player has done; it may also put player i at a disadvantage in terms of lost time or the ability to take some action. What information do different players have when they take actions? Fourth, what are the payoffs to the various players as a result of the interaction? Ascertaining payoffs involves estimating the costs and benefits of each potential set of choices by all players. In many situations it may be easier to estimate payoffs for some players (such as yourself) than others, and it may be unclear whether other players are also thinking strategically. This consideration suggests that careful attention be paid to a sensitivity analysis.

Once we have framed the situation, we can look from different players' perspectives to analyze which actions are optimal for them. There are various criteria we can use.

9.9.1 Games in Normal Form

Let us begin with a standard representation of a game, which is known as a *normal form* game, or a game in *strategic form:*

The set of players is $N = \{1, \ldots , n\}$.

Player i has a set of actions, X_i, available. These are generally referred to as *pure strategies*. This set might be finite or infinite.

Let $X = X_1 \times \cdots \times X_n$ be the set of all profiles of pure strategies or actions, with a generic element denoted by $x = (x_1, \ldots , x_n)$.

Player i's payoff as a function of the vector of actions taken is described by a function $u_i : X \to \mathbb{R}$, where $u_i(x)$ is i's payoff if the x is the profile of actions chosen in the society.

Normal form games are often represented by a table. Perhaps the most famous such game is the *prisoners' dilemma*, which is represented in Table 9.3. In this game

TABLE 9.3
A prisoners' dilemma game

		Player 2	
		C	D
Player 1	C	−1, −1	−3, 0
	D	0, −3	−2, −2

Note: C, cooperate; D, defect.

TABLE 9.4
A rescaling of the prisoners' dilemma game

		Player 2	
		C	D
Player 1	C	4, 4	0, 6
	D	6, 0	2, 2

Note: C, cooperate; D, defect.

there are two players who each have two pure strategies, where $X_i = \{C, D\}$, and C stands for "cooperate" and D stands for "defect." The first entry indicates the payoff to the row player (or player 1) as a function of the pair of actions, while the second entry is the payoff to the column player (or player 2).

The usual story behind the payoffs in the prisoners' dilemma is as follows. The two players have committed a crime and are now in separate rooms in a police station. The prosecutor has come to each of them and told them each: "If you confess and agree to testify against the other player, and the other player does not confess, then I will let you go. If you both confess, then I will send you both to prison for 2 years. If you do not confess and the other player does, then you will be convicted and I will seek the maximum prison sentence of 3 years. If nobody confesses, then I will charge you with a lighter crime for which we have enough evidence to convict you and you will each go to prison for 1 year." So the payoffs in the matrix represent time lost in terms of years in prison. The term *cooperate* refers to cooperating with the other player. The term *defect* refers to confessing and agreeing to testify, and so breaking the (implicit) agreement with the other player.

Note that we could also multiply each payoff by a scalar and add a constant, which is an equivalent representation (as long as all of a given player's payoffs are rescaled in the same way). For instance, in Table 9.4 I have doubled each entry and added 6. This transformation leaves the strategic aspect of the game unchanged.

There are many games that might have different descriptions motivating them but have a similar normal form in terms of the strategic aspects of the game. Another example of the same game as the prisoners' dilemma is what is known as a *Cournot duopoly.* The story is as follows. Two firms produce identical goods. They each have two production levels, high or low. If they produce at high production, they will have a lot of the goods to sell, while at low production they have less to sell. If they cooperate, then they agree to each produce at low production. In this case, the product is rare and fetches a very high price on the market, and they each make a profit of 4. If they each produce at high production (or defect), then they will depress the price, and even though they sell more of the goods, the price drops sufficiently to lower their overall profits to 2 each. If one defects and the other cooperates, then the price is in a middle range. The firm with the higher production sells more goods and earns a higher profit of 6, while the firm with the lower production just covers its costs and earns a profit of 0.

9.9.2 Dominant Strategies

Given a game in normal form, we then can make predictions about which actions will be chosen. Predictions are particularly easy when there are dominant strategies. A dominant strategy for a player is one that produces the highest payoff of any strategy available *for every possible action by the other players.*

That is, a strategy $x_i \in X_i$ is a *dominant* (or weakly dominant) strategy for player i if $u_i(x_i, x_{-i}) \geq u_i(x_i', x_{-i})$ for all x_i' and all $x_{-i} \in X_{-i}$. A strategy is a *strictly dominant strategy* if the above inequality holds strictly for all $x_i' \neq x_i$ and all $x_{-i} \in X_{-i}$.

Dominant strategies are powerful from both an analytical point of view and a player's perspective. An individual does not have to make any predictions about what other players might do, and still has a well-defined best strategy.

In the prisoners' dilemma, it is easy to check that each player has a strictly dominant strategy to defect—that is, to confess to the police and agree to testify. So, if we use dominant strategies to predict play, then the unique prediction is that each player will defect, and both players fare worse than for the alternative strategies in which neither defects. A basic lesson from the prisoners' dilemma is that individual incentives and overall welfare need not coincide. The players both end up going to jail for 2 years, even though they would have gone to jail for only 1 year if neither had defected. The problem is that they cannot trust each other to cooperate: no matter what the other player does, a player is best off defecting.

Note that this analysis presumes that all relevant payoff information is included in the payoff function. If, for instance, a player fears retribution for confessing and testifying, then that should be included in the payoffs and can change the incentives in the game. If the player cares about how many years the other player spends in jail, then that can be written into the payoff function as well.

When dominant strategies exist, they make the game-theoretic analysis relatively easy. However, such strategies do not always exist, and then we can turn to notions of equilibrium.

9.9.3 Nash Equilibrium

A pure strategy Nash equilibrium[43] is a profile of strategies such that each player's strategy is a best response (results in the highest available payoff) against the equilibrium strategies of the other players.

A strategy x_i is a *best reply*, also known as a *best response,* of player i to a profile of strategies $x_{-i} \in X_{-i}$ for the other players if

$$u_i(x_i, x_{-i}) \geq u_i(x_i', x_{-i})$$

for all x_i'. A best response of player i to a profile of strategies of the other players is said to be a *strict best response* if it is the unique best response.

A profile of strategies $x \in X$ is a *pure strategy Nash equilibrium* if x_i is a best reply to x_{-i} for each i. That is, x is a Nash equilibrium if

$$u_i(x_i, x_{-i}) \geq u_i(x_i', x_{-i})$$

for all i and x_i'. This definition might seem somewhat similar to that of dominant strategy, but there is a critical difference. A pure strategy Nash equilibrium only requires that the action taken by each agent be best against the actions taken by the other players, and not necessarily against all possible strategies of the other players.

A Nash equilibrium has the nice property that it is stable: if each player expects x to be the profile of strategies played, then no player has any incentive to change his or her action. In other words, no player regrets having played the strategy that he or she played in a Nash equilibrium.

In some cases, the best response of a player to the strategies of others is unique. A Nash equilibrium such that all players are playing strategies that are unique best responses is called a *strict* Nash equilibrium. A profile of dominant strategies is a Nash equilibrium but not vice versa.

To see a Nash equilibrium in action, consider the following game between two firms that are deciding whether to advertise. Total available profits are 28, to be split between the two firms. Advertising costs a firm 8. Firm 1 currently has a larger market share than firm 2, so it is seeing 16 in profits while firm 2 is seeing 12 in profits. If they both advertise, then they will split the market evenly 14 each but must bear the cost of advertising, so they will see profits of 6 each. If one advertises while the other does not, then the advertiser captures three-quarters of the market (but also pays for advertising) and the nonadvertiser gets one-quarter of the market. (There are obvious simplifications here: just considering two levels of advertising and assuming that advertising only affects the split and not the total profitability.) The net payoffs are given in the Table 9.5.

To find the equilibrium, we have to look for a pair of actions such that neither firm wants to change its action given what the other firm has chosen. The search

43. On occasion referred to as *Cournot–Nash equilibrium,* with reference to Cournot [179], who first developed such an equilibrium concept in the analysis of oligopoly (a set of firms in competition with one another).

TABLE 9.5
An advertising game

		Firm 2	
		Not	Adv
Firm 1	Not	16, 12	7, 13
	Adv	13, 7	6, 6

Note: Adv, advertise; Not, do not advertise.

TABLE 9.6
A coordination game

		Player 2	
		A	B
Player 1	A	5, 5	0, 3
	B	3, 0	4, 4

is made easier in this case, since firm 1 has a strictly dominant strategy of not advertising. Firm 2 does not have a dominant strategy; which strategy is optimal for it depends on what firm 1 does. But given the prediction that firm 1 will not advertise, firm 2 is best off advertising. This forms a Nash equilibrium, since neither firm wishes to change strategies. You can easily check that no other pairs of strategies form an equilibrium.

While each of the previous games provides a unique prediction, there are games in which there are multiple equilibria. Here are three examples.

Example 9.8 (A Stag Hunt Game) *The first is an example of a coordination game, as depicted in Table 9.6. This game might be thought of as selecting between two technologies, or coordinating on a meeting location. Players earn higher payoffs when they choose the same action than when they choose different actions. There are two (pure strategy) Nash equilibria: (A, A) and (B, B).*

This game is also a variation on Rousseau's "stag hunt" game.[44] *The story is that two hunters are out, and they can either hunt for a stag (strategy A) or look for hares (strategy B). Succeeding in getting a stag takes the effort of both hunters, and the hunters are separated in the forest and cannot be sure of each other's behavior. If both hunters are convinced that the other will hunt for stag, then hunting stag is a strict or unique best reply for each player. However, if one turns out to be mistaken and the other hunter hunts for hare, then one will go hungry. Both hunting for hare*

44. To be completely consistent with Rousseau's story, (B, B) should result in payoffs of (3, 3), as the payoff to hunting for hare is independent of the actions of the other player in Rousseau's story.

TABLE 9.7
A battle of the sexes game

		Player 2	
		A	B
Player 1	A	3, 1	0, 0
	B	0, 0	1, 3

is also an equilibrium and hunting for hare is a strict best reply if the other player is hunting for hare. This example hints at the subtleties of making predictions in games with multiple equilibria. On the one hand, (A, A) (hunting stag by both) is a more attractive equilibrium and results in high payoffs for both players. Indeed, if the players can communicate and be sure that the other player will follow through with an action, then playing (A, A) is a stable and reasonable prediction. However, (B, B) (hunting hare by both) has properties that make it a useful prediction as well. It does not offer as high a payoff, but it has less risk associated with it. Here playing B guarantees a minimum payoff of 3, while the minimum payoff to A is 0. There is an extensive literature on this subject, and more generally on how to make predictions when there are multiple equilibria (see, e.g., Binmore [70], Fudenberg and Tirole [263], Myerson [497], and Osborne and Rubinstein [514]).

Example 9.9 (A "Battle of the Sexes" Game) *The next example is another form of coordination game, but with some asymmetries in it. It is generally referred to as a "battle of the sexes" game, as depicted in Table 9.7. The players have an incentive to choose the same action, but they each have a different favorite action. There are again two (pure strategy) Nash equilibria: (A, A) and (B, B). Here, however, player 1 would prefer that they play equilibrium (A, A) and player 2 would prefer (B, B). The battle of the sexes title refers to a couple trying to coordinate on where to meet for a night out. They prefer to be together, but also have different preferred outings.*

Example 9.10 (Hawk-Dove and Chicken Games) *There are also what are known as anticoordination games, with the prototypical version being what is known as the hawk-dove game or the chicken game, with payoffs as in Table 9.8. Here there are two pure strategy equilibria, (Hawk, Dove) and (Dove, Hawk). Players are in a potential conflict and can be either aggressive like a hawk or timid like a dove. If they both act like hawks, then the outcome is destructive and costly for both players with payoffs of 0 for both. If they each act like doves, then the outcome is peaceful and each gets a payoff of 2. However, if the other player acts like a dove, then a player would prefer to act like a hawk and take advantage of the other player, receiving a payoff of 3. If the other player is playing a hawk strategy, then it is best to play a dove strategy and at least survive rather than to be hawkish and end in mutual destruction.*

TABLE 9.8
A hawk-dove game

		Player 2	
		Hawk	Dove
Player 1	Hawk	0, 0	3, 1
	Dove	1, 3	2, 2

TABLE 9.9
A penalty-kick game

		Goalie	
		L	R
Kicker	L	−1, 1	1, −1
	R	1, −1	−1, 1

Note: L, left; R, right.

9.9.4 Randomization and Mixed Strategies

In each of the above games, there was at least one pure strategy Nash equilibrium. There are also simple games for which no pure strategy equilibrium exists. To see this, consider the following simple variation on a penalty kick in a soccer match. There are two players: the player kicking the ball and the goalie. Suppose, to simplify the exposition, that we restrict the actions to just two for each player (there are still no pure strategy equilibria in the larger game, but this limitation makes the exposition easier). The kicking player can kick to the left side or to the right side of the goal. The goalie can move to the left side or to the right side of the goal and has to choose before seeing the kick, as otherwise there is too little time to react. To keep things simple, assume that if the player kicks to one side, then she scores for sure if the goalie goes to the other side, while the goalie is certain to save it if the goalie goes to the same side. The basic payoff structure is depicted in Table 9.9. This is also the game known as "matching pennies." The goalie would like to choose a strategy that matches that of the kicker, and the kicker wants to choose a strategy that mismatches the goalie's strategy.[45]

It is easy to check that no pair of pure strategies forms an equilibrium. What is the solution here? It is just what you see in practice: the kicker randomly picks left versus right, in this particular case with equal probability, and the goalie does the same. To formalize this observation we need to define randomized strategies, or

---

45. For an interesting empirical test of whether goalies and kickers on professional soccer teams randomize properly, see Chiappori, Levitt, and Groseclose [150]; and see Walker and Wooders [646] for an analysis of randomization in the location of tennis serves in professional tennis matches.

what are called *mixed strategies.* For ease of exposition suppose that X_i is finite; the definition extends to infinite strategy spaces with proper definitions of probability measures over pure actions.

A mixed strategy for a player i is a distribution μ_i on X_i, where $\mu_i(x_i)$ is the probability that x_i is chosen. A profile of mixed strategies $(\mu_1, \ldots, \mu_n)$ forms a mixed-strategy Nash equilibrium if

$$\sum_x \left(\prod_j \mu_j(x_j) \right) u_i(x) \geq \sum_{x_{-i}} \left(\prod_{j \neq i} \mu_j(x_j) \right) u_i(x_i', x_{-i})$$

for all i and x_i'.

So a profile of mixed strategies is an equilibrium if no player has some strategy that would offer a better payoff than his or her mixed strategy in reply to the mixed strategies of the other players. Note that this reasoning implies that a player must be indifferent to each strategy that he or she chooses with a positive probability under his or her mixed strategy. Also, players' randomizations are independent.[46] A special case of a mixed strategy is a pure strategy, where probability 1 is placed on some action.

It is easy to check that each mixing with probability 1/2 on L and R is the unique mixed strategy of the matching pennies game above. If a player, say the goalie, places weight of more than 1/2 on L, for instance, then the kicker would have a best response of choosing R with probability 1, but then that could not be an equilibrium as the goalie would want to change his or her action, and so forth.

There is a deep and long-standing debate about how to interpret mixed strategies, and the extent to which people really randomize. Note that in the goalie and kicker game, what is important is that each player not know what the other player will do. For instance, it could be that the kicker decided before the game that if there was a penalty kick then she would kick to the left. What is important is that the kicker not be known to always kick to the left.[47]

We can begin to see how the equilibrium changes as we change the payoff structure. For example, suppose that the kicker is more skilled at kicking to the right side than to the left. In particular, keep the payoffs as before, but now suppose that the kicker has an even chance of scoring when kicking right when the goalie goes right. This leads to the payoffs in Table 9.10. What does the equilibrium look like? To calculate the equilibrium, it is enough to find a strategy for the goalie that

46. An alternative definition of correlated equilibrium allows players to use correlated strategies but requires some correlating device that only reveals to each player his or her prescribed strategy and that these are best responses given the conditional distribution over other players' strategies.
47. The contest between pitchers and batters in baseball is quite similar. Pitchers make choices about the location, velocity, and type of pitch (e.g., whether various types of spin are put on the ball). If a batter knows what pitch to expect in a given circumstance, that can be a significant advantage. Teams scout one another's players and note any tendencies or biases that they might have and then try to respond accordingly.

TABLE 9.10
A biased penalty-kick game

		Goalie	
		L	R
Kicker	L	−1, 1	1, −1
	R	1, −1	0, 0

Note: L, left; R, right.

makes the kicker indifferent, and a strategy for the kicker that makes the goalie indifferent.[48]

Let μ_1 be the kicker's strategy and μ_2 be the goalie's strategy. It must be that the kicker is indifferent. The kicker's payoff from L is $-\mu_2(L) + \mu_2(R)$ and the payoff from R is $\mu_2(L)$, so that

$$-\mu_2(L) + \mu_2(R) = \mu_2(L),$$

or $\mu_2(L) = 1/3$ and $\mu_2(R) = 2/3$. For the goalie to be indifferent, it must be that

$$\mu_1(L) - \mu_1(R) = -\mu_1(L) + \mu_1(R),$$

and so the kicker must choose $\mu_1(L) = 1/2 = \mu_2(R)$.

Note that as the kicker gets more skilled at kicking to the right, it is the *goalie's* strategy that adjusts to moving to the right more often! The kicker still mixes evenly. It is a common misconception to presume that it should be the kicker who should adjust to using his or her better strategy with more frequency.[49]

While not all games have pure strategy Nash equilibrium, every game with a finite set of actions has at least one mixed strategy Nash equilibrium (with a special case of a mixed strategy equilibrium being a pure strategy equilibrium), as shown in an important paper by Nash [498].

48. This reasoning is a bit subtle, as we are not directly choosing actions that maximize the goalie's payoff and maximize the kicker's payoff, but instead are looking for a mixture by one player that makes the other indifferent. This feature of mixed strategies takes a while to grasp, but experienced players seem to understand it well (e.g., see Chiappori, Levitt, and Groseclose [150] and Walker and Wooders [646]).

49. Interestingly, there is evidence that professional soccer players are better at playing games that have mixed strategy equilibria than are people with less experience in such games, which is consistent with this observation (see Palacios-Huerta and Volij [521]).

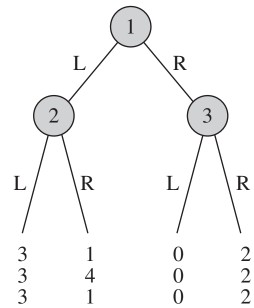

FIGURE 9.21 A game tree with three players and two actions each.

9.9.5 Sequentiality, Extensive Form Games, and Backward Induction

Let us now turn to the question of timing. In the above discussion it was implicit that each player was selecting a strategy with beliefs about the other players' strategies but without knowing exactly what they were.

If we wish to be more explicit about timing, then we can consider what are known as games in *extensive form,* which include a complete description of who moves in what order and what they have observed when they move.[50] There are advantages to working with extensive form games, as they allow for more explicit treatments of timing and for equilibrium concepts that require credibility of strategies in response to the strategies of others.

Definitions for a general class of extensive form games are notationally intensive. In this book, we mainly look at a special class of extensive form games—finite games of perfect information—which allows for a treatment that avoids much of the notation. These are games in which players move sequentially in some prespecified order (sometimes contingent on which actions have been chosen), each player moves at most a finite number of times, and each player is completely aware of all moves that have been made previously. These games are particularly well behaved and can be represented by simple trees, where a node is associated with the move of a specified player and an edge corresponds to different actions the player might take, as in Figure 9.21. I will not provide formal definitions, but simply refer directly to games representable by such finite game trees.

Each node has a player's label attached to it. There is an identified *root node* that corresponds to the first player to move (player 1 in Figure 9.21) and then subsequent nodes that correspond to subsequent players who make choices. In Figure 9.21, player 1 has a choice of moving either left or right. The branches in the tree correspond to the different actions available to the player at a given node.

---

50. One can collapse certain types of extensive form games into normal form by simply defining an action to be a complete specification of how an agent would act in all possible contingencies. Agents then choose these actions simultaneously at the beginning of the game. But the normal form becomes more complicated than the two-by-two games in Sections 9.9.3 and 9.9.4.

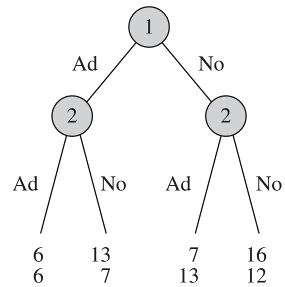

FIGURE 9.22 Advertising choices of two competitors.

In this game, if player 1 moves left, then player 2 moves next; while if player 1 moves right, then player 3 moves next. It is also possible to have trees in which player 1 chooses twice in a row, or no matter what choice a given player makes it is a certain player who follows, and so forth. The payoffs are given at the end nodes and are listed for the respective players. The top payoff is for player 1, the second for player 2, and the bottom for player 3. So the payoffs depend on the set of actions taken, which then determines a path through the tree. An equilibrium provides a prediction about how each player will move in each contingency and thus makes a prediction about which path will be taken; we refer to that prediction as the *equilibrium path.*

We can apply the concept of a Nash equilibrium to such games, which here is a specification of what each player would do at each node with the requirement that each player's strategy be a best response to the other players' strategies. Nash equilibrium does not always make sensible predictions when applied to the extensive form. For instance, reconsider the advertising example discussed above (Table 9.5). Suppose that firm 1 makes its decision prior to firm 2, and that firm 2 knows firm 1's choice before it chooses. This scenario is represented in Figure 9.22. To apply the Nash equilibrium concept to this extensive form game, we must specify what each player does at each node. There are two Nash equilibria of this game in pure strategies. The first is where firm 1 advertises, and firm 2 does not (and firm 2's strategy conditional on firm 1 not advertising is to advertise). The other equilibrium corresponds to the one identified in the normal form: firm 1 does not advertise, and firm 2 advertises regardless of what firm 1 does. This is an equilibrium, since neither wants to change its behavior, given the other's strategy. However, it is not really credible in the following sense: it involves firm 2 advertising even after it has seen that firm 1 has advertised, and even though this action is not in firm 2's interest in that contingency.

To capture the idea that each player's strategy has to be credible, we can solve the game backward. That is, we can look at each decision node that has no successor, and start by making predictions at those nodes. Given those decisions, we can roll the game backward and decide how player's will act at next-to-last decision nodes, anticipating the actions at the last decision nodes, and then iterate. This is called *backward induction.* Consider the choice of firm 2, given that firm 1 has decided not to advertise. In this case, firm 2 will choose to advertise, since 13 is larger than 12. Next, consider the choice of firm 2, given that firm 1 has decided to advertise. In this case, firm 2 will choose not to advertise, since 7 is larger than 6. Now we can collapse the tree. Firm 1 will predict that if it does not advertise, then firm 2 will advertise, while if firm 1 advertises then firm 2 will not. Thus when making its choice, firm 1 anticipates a payoff of 7 if it chooses not to advertise and 13 if it chooses to advertise. Its optimal choice is to advertise. The backward induction prediction about the actions that will be taken is for firm 1 to advertise and firm 2 not to.

Note that this prediction differs from that in the simultaneous move game we analyzed before. Firm 1 has gained a first-mover advantage in the sequential version. Not advertising is no longer a dominant strategy for firm 1, since firm 2's decision depends on what firm 1 does. By committing to advertising, firm 1 forces firm 2 to choose not to advertise. Firm 1 is better off being able to commit to advertising in advance.

A solution concept that formalizes the backward induction solution found in this game and applies to more general classes of games is known as *subgame perfect equilibrium* (due to Selten [586]). A subgame in terms of a finite game tree is simply the subtree that one obtains starting from some given node. Subgame perfection requires that the stated strategies constitute a Nash equilibrium in every subgame (including those with only one move left). So it requires that if we start at any node, then the strategy taken at that node must be optimal in response to the remaining specification of strategies. In the game between the two firms, it requires that firm 2 choose an optimal response in the subgame following a choice by firm 1 to advertise, and so it coincides with the backward induction solution for such a game.

I close this appendix by noting that moving first is not always advantageous. Sometimes it allows one to commit to strategies which would otherwise be untenable, which can be advantageous; but in other cases the information that the second mover gains from knowing which strategy the first mover has chosen may be a more important consideration. For example, when playing the matching pennies game sequentially, it is clearly not good for a player to move first.

9.10 ▪ Exercises

9.1 ***Fashionable Ants*** Consider the model described in Section 9.1.2. Suppose that a player has a probability $\varepsilon > 0$ of flipping a coin to choose an (binary) action, and a probability of $1 - \varepsilon$ of matching the action being taken by the majority of other individuals. Taking n to be even, so that the number of other individuals is always odd, describe p_s for any s. Next, pick a value of n, and for several values of ε plot the steady-state probability of there being s individuals taking action 1 as a function of s (similar to Figure 9.1).

9.2 ***Proof of Proposition 9.1*** Prove Proposition 9.1.

9.3 ***Steady-State Probabilities of Action*** Consider the following variation on a model of social interaction by Calvó-Armengol and Jackson [131] and discussed in Section 9.1.2. Let $p_s = q$ for some $1 > q > 0$ when $s \geq \tau$ and $p_s = 1 - q$ when $s < \tau$, where $\tau \in \{0, \dots, n-1\}$ is a threshold. Thus individuals choose action 1 with probability q if at least τ others have, and with probability $1 - q$ otherwise. Solve for the steady-state probability μ_s as a function of μ_0.

9.4 ***Another Pure Strategy Equilibrium for the Game in Figures 9.2–9.4*** Find a pure strategy equilibrium of the game in Figures 9.2–9.4 that is not pictured there.

9.5 ***The Lattice Structure of Equilibria in Semi-Anonymous Games of Complementaries*** Show that if $x = (x_1, \ldots, x_n)$ and $x' = (x_1', \ldots, x_n')$ are pure strategy equilibria of a semi-anonymous graphical game with strategic complementarities, then there exists a pure strategy equilibrium $\overline{x}$ such that

$$\overline{x}_i \geq \max\left(x_i, x_i'\right)$$

for all i, as well as a pure strategy equilibrium such that

$$\overline{x}_i \leq \min\left(x_i, x_i'\right)$$

for all i. This property means that the set of equilibria form a lattice. (In fact, the lattice is complete, so that for any set of pure strategy equilibria we can find a pure strategy equilibrium which is greater than or equal to each set member and another pure strategy equilibrium that is less than or equal to each member.)

9.6 ***Possible Nonexistence of Pure Strategy Equilibria in Best-Shot Graphical Games on a Directed Network*** Provide an example of a directed network with three players for which the only equilibria to a best-shot game played on that network are mixed strategies. Identify a mixed strategy equilibrium.

9.7 ***Existence of Pure Strategy Equilibria in Semi-Anonymous Graphical Games of Strategic Complements with Infinite Action Spaces**** Consider a graphical game on a network (N, g) in which player i has a compact action space $X_i \subset [0, M]$. Let $u_i(x_i, x_{N_i(g)})$ be continuous for each i. A graphical game exhibits *strategic complements* if

$$u_i(x_i', x_{N_i(g)}') - u_i(x_i, x_{N_i(g)}') \geq u_i(x_i', x_{N_i(g)}) - u_i(x_i, x_{N_i(g)}) \qquad (9.14)$$

for every i, $x_i' > x_i$ and $x_{N_i(g)}' \geq x_{N_i(g)}$.[51]

Show that there exists a pure strategy equilibrium in such a game. Show that the set of pure strategy equilibria form a complete lattice (see Exercise 9.5). Show that there are examples for which each of these conclusions fails if we set $X_i = \mathbb{R}_+$.

9.8 ***Graphical Games of Complements**** Consider a graphical game as in Exercise 9.7 but that is also semi-anonymous, so that all agents have the same action space and payoffs depend only on the vector $x_{N_i(g)}'$ up to a relabeling of the agents.[52]

---

51. The inequality $x_{N_i(g)}' \geq x_{N_i(g)}$ indicates that each coordinate of $x_{N_i(g)}'$ is at least as large as the corresponding coordinate of $x_{N_i(g)}$.

52. That is, if there exists a bijection π from $N_i(g)$ to $N_i(g)$ such that jth coordinate of $x_{N_i(g)}'$ is equal to the $\pi(j)$th coordinate of $x_{N_i(g)}$ for each j, then $u_i(\cdot, x_{N_i(g)}) = u_i(\cdot, x_{N_i(g)}')$.

Suppose also that for any z in $[0, M]^{d_i}$

$$u_{d_i}(x_i, z) = u_{d_i+1}(x_i, (z, 0)). \tag{9.15}$$

Thus if we add a link from one player to a second player who is choosing action 0, then the payoff is as if the second player were not there.

(a) Show that if $g \subset g'$, then for every pure strategy equilibrium x under g there exists an equilibrium $x' \geq x$ under g'.

(b) Consider a case in which the inequality in (9.14) is strict when $x'_{N_i(g)} \neq x_{N_i(g)}$ and X_i is connected for each i. Show that if an equilibrium x relative to g is such that $x_i < M$ for each i, then there exists an equilibrium x' under $g + ij$ in which all players in the component of i and j play strictly higher actions.

(c) Show that the conclusions of (a) and (b) can fail if (9.15) is violated.

9.9 ***Payoffs Increase with Degree**** Galeotti et al. [274] state that a network game exhibits *degree complementarity* if

$$U_d(1, \sigma) - U_d(0, \sigma) \geq U_{d'}(1, \sigma) - U_{d'}(0, \sigma) \tag{9.16}$$

when $d > d'$. Equation (9.16) states that facing the same behavior by other players, a player with a higher degree has at least as big an incentive to take action 1 compared to a player with a lower degree.[53]

(a) Show that if (9.5) holds and the network game is one of strategic complements, then degree complementarity holds. Show that degree complementarity also holds in the case of strategic complements when a player cares about the fraction of neighbors taking action 1, so that $u_d(1, m) = \frac{m}{d} - c$ and $u_d(0, m) = b\frac{m}{d} - a$, with $b \leq 1$.

(b) Show that a network game that satisfies the condition in (a) has an equilibrium that is nondecreasing in degree.

9.10 ***Payoffs Increase with Degree*** Consider the setting of Proposition 9.5. Suppose that there are *positive externalities*, so that for each d and x_i, $u_d(x_i, m)$ is nondecreasing in m.[54] Show that for every equilibrium in a game with either strategic complements or substitutes, the payoff to a player with degree d', where $d' > d$, is at least as high as the payoff to a player with degree d.

9.11 ***All Equilibria of a Network Game Are Monotone**** Consider the setting of Proposition 9.5. Prove the last claim that if the game is one of strict strategic complements, then all equilibria are nondecreasing in degree.

9.12 ***A Local Public Goods Graphical Game in Which Players Have Heterogeneous Costs*** Bramoullé and Kranton [103] consider a variation on the model of Section

53. Analogously, payoffs exhibit *degree substitution* if the inequality above is reversed, and the following statements hold as well.

54. This setting is different from that of network formation, for which externalities were defined relative to network structure. Here the actions considered are those in the graphical game.

9.5.1 in which players can have different costs of providing the public good. That is, payoffs are given by

$$u_i(x_i, x_{N_i(g)}) = f(x_i + \sum_{j \in N_i(g)} x_j) - c_i x_i,$$

where c_i can differ across players. Let the function f be increasing and strictly concave, let x_i^* denote the maximizer of $f(x) - c_i x$, and suppose that x_i^* is well defined and nonzero for every player. Provide an algorithm that finds an equilibrium in a setting with $c_1 < c_2 < \cdots < c_n$. Provide an example in which there is more than one stable equilibrium.

9.13 ***Convex Costs in a Local Public Goods Graphical Game*** Consider the following variation on the local public goods graphical game of Bramoullé and Kranton [103] from Section 9.5.1. Payoffs are given by

$$u_i(x_i, x_{N_i(g)}) = f(x_i + \sum_{j \in N_i(g)} x_j) - c(x_i),$$

where f is strictly concave and c is strictly convex, and there exists $x^* > 0$ such that $f'(x^*) = c'(x^*)$, which is the action level that an individual chooses if he or she is the only provider.

(a) Find a pure strategy equilibrium on a complete network and show that it is the unique pure strategy equilibrium and that all players choose positive actions.

(b) Consider a circle network with an even number of players, and suppose that $f'(2x^*) < c'(0)$. Describe a specialized equilibrium where only some players choose positive actions.

9.14 ***Cohesiveness*** Find a partition of the set of nodes in Figure 9.14 into two sets such that one set is 2/3-cohesive and its complement is 3/4-cohesive.

9.15 ***Labelings of Nodes and Cohesion*** Consider a network (N, g) and a coordination game such that action 1 is a best response for any player if and only if a fraction of at least q of his or her neighbors play action 1.

(a) Show the following result from Morris [487]. Let a labeling of nodes be a bijection (a one-to-one and onto function) ℓ from N to N. Let $\alpha_\ell(i)$ be the fraction of $\ell(i)$'s neighbors who have labels less than $\ell(i)$. Show that there is a contagion from m nodes if and only if there exists a labeling ℓ such that $\alpha_\ell(i) \geq q$ for all $\ell(i) \geq m + 1$.

(b) From (a) show that there exists a set S that is uniformly no more than r-cohesive if and only if there is a labeling ℓ such that $\alpha_\ell(i) \geq 1 - r$ for all $\ell(i) \geq m + 1$, where m is the cardinality of the complement of S.

9.16 ***A Sufficient Condition for the Failure of Contagion*** Consider a network (N, g) and a coordination game such that action 1 is a best response for any player if and only if a fraction of at least q of his or her neighbors play action 1. Show that a sufficient condition for never having contagion from any group of m nodes is to have at least $m + 1$ separate groups that are each more than $(1 - q)$ cohesive.

9.17 ***Contagion to a Subset of Nodes*** Consider a network (N, g) and a coordination game such that action 1 is a best response for any player if and only if a fraction of at least q of his or her neighbors play action 1. Show that $B \cup A$ is the eventual set of nodes playing 1 under the contagion system described in Section 9.6 if and only if the complement of $B \cup A$, denoted by C, is more than q-cohesive and for every nonempty subset D of B, $D \cup C$ has a cohesiveness of no more than q.

9.18 ***Diffusion of Behavior in Network Games of Strategic Substitutes*** Consider the network games setting from Section 9.7.1 and suppose that $U_d(1, p)$ is decreasing and continuous in p for each d, and suppose that F is increasing and continuous on the entire range of U_d for each d. Show that there is a unique equilibrium p and that it is a stable equilibrium.

9.19 ***Adoption Patterns by Degree: Diffusion of Behavior in Network Games*** Consider the network games setting from Section 9.7.1 in a case such that $U_d(1, p) = pd$ and F is uniform on [0, 5], so that $F(U_d(1, p)) = \min\left[pd, 5\right]/5$.

(a) Suppose that the network game is regular so that all players have degree d. What is the unique equilibrium p for $d < 5$? What are the two equilibria p when $d > 5$? What are the equilibrium ps when $d = 5$?

(b) Consider a degree distribution that has equal weights on degrees $\{1, 2, \ldots, 10\}$ (so you need to use the corresponding $\widetilde{P}$ that is biased toward higher degrees with weight $d/55$ on degree d to obtain the distribution of neighbors' degrees). Using a simple spreadsheet or other program of your choosing, start with an initial $p^0 = .1$ and trace the evolution of the proportion of degree-d types that have chosen action 1 at a sequence of dates $t = 1, 2, \ldots$ until you have some sense of convergence. Plot the resulting adoption curves for $d = 1$, $d = 5$, and $d = 10$ versus time.

9.20 ***S-Shaped Adoption Curves: Diffusion of Behavior in Network Games***
S-shaped adoption curves have been found in a variety of studies of diffusion. For such curves, adoption starts slowly, then increases its rate of diffusion, and then eventually slows down again.[55] In terms of diffusion in the network games setting from Section 9.7.1, we can track $p^{t+1} - p^t = \phi(p^t) - p^t$ as a proxy for the rate of diffusion as in Jackson and Yariv [362].

Let $H(d, p) = F(U_d(1, p))$, which lies between 0 and 1 for every d and p, since F is a distribution function. Suppose that $H(d, 0) > 0$ for some d such that $\widetilde{P}(d) > 0$, and that H is twice continuously differentiable and increasing in both variables and strictly concave in p. Show that ϕ will be S-shaped. That is, show that there exists $p^* \in [0, 1]$ such that $\phi(p) - p$ is increasing when $p < p^*$ and then decreasing when $p > p^*$ (when $\phi(p) < 1$).

---

55. See Bass [46] for a discussion of this behavior, Rogers [564] for more detailed references, and Young [669] for alternative learning-based models.

9.21 ***The Expected Number of Equilibria in a Generic Graphical Game*** *[56] Consider an arbitrary network (N, g) as the basis for a graphical game. Define the payoffs for players as a function of their actions as follows. For each player i and configuration of strategies $(x_i, x_{N_i(g)}) \in \{0, 1\}^{d_i(g)+1}$, assign the payoff $u_i(x_i, x_{N_i(g)})$ according to an atomless distribution F on $\mathbb{R}$. Do this independently for each player and profile of strategies. Once we have specified every u_i, the graphical game is well defined. It might have one pure strategy Nash equilibria, it might have several, or it might not have any, depending on the values of the u_is. Show that the expected number of pure strategy Nash equilibria is 1.

Hint: What is the probability that $x_i = 1$ is a best reply to some $x_{N_i(g)} \in \{0, 1\}^{d_i(g)}$? Then what is the probability that some profile of actions $(x_1, \ldots, x_n)$ is an equilibrium?

9.22 ***Finding Equilibria in Graphical Games of Strategic Complements*** (a) Show that the algorithm for threshold games described in Section 9.8 finds the maximal equilibrium x, in the sense that $x_i \geq x'_i$ for all other equilibrium x' and all i.

(b) Describe an algorithm for finding the minimal equilibrium x such that $x_i \leq x'_i$ for all other equilibrium x' and all i.

(c) Argue that the claims in (a) (and (b)) are true even when considering mixed strategy equilibria so that x_i is at least (atmost) as large (small) as the maximum (minimum) of the support of the strategy of player i in any alternative equilibrium.

(d) Show that your algorithm also works for any graphical game of strategic complements with an action space of $\{0, 1\}$.

9.23 ***Finding All Equilibria in Best-Shot Graphical Games*** Provide an example of an equilibrium in a best-shot public goods graphical game that would not be found by the algorithm for best-shot games described in Section 9.8.[57]

9.10.1 Exercises on Games

9.G1 ***Product Choices*** Two electronics firms are making product-development decisions. Each firm is choosing between the development of two alternative computer chips. One system has higher efficiency but requires a larger investment and is more costly to produce. Based on estimates of development costs, production costs, and demand, the present-value calculations shown in Table 9.11 represent the value of the alternatives (high-efficiency chips or low-efficiency chips) to the firms. The first entry in each cell in Table 9.11 is the present value to firm 1 and the second entry is the present value to firm 2. The payoffs in the table are not symmetric. Firm 2 has a cost advantage in producing the higher-efficiency chip, while firm 1 has a cost advantage in producing the lower-efficiency chip. Overall profits are largest when the firms choose different chips and do not compete head to head.

(a) Firm 1 has a dominant strategy. What is it?
(b) Given your answer to (a), what should firm 2 expect firm 1's choice to be? What is firm 2's optimal choice, given what it anticipates firm 1 will do?

---

56. This exercise is based on a result by Daskalakis, Dimakis, and Mossel [184].
57. For an algorithm that finds all maximal independent sets, see Johnson, Papadimitriou, and Yannakakis [368].

TABLE 9.11
A production-choice game

		Firm 2	
		High	Low
Firm 1	High	1, 2	4, 5
	Low	2, 7	5, 3

(c) Do firm 1's strategy (answer to (a)) and firm 2's strategy (answer to (b)) form an equilibrium? Explain.
(d) Compared to (c), firm 1 would make larger profits if the choices were reversed. Why don't those strategies form an equilibrium?
(e) Suppose that firm 1 can commit to a product before firm 2. Draw the corresponding game tree and describe the backward induction/subgame perfect equilibrium.

9.G2 ***Hotelling's Hotels*** Two hotels are considering a location along a newly constructed highway through the desert. The highway is 500 miles long with an exit every 50 miles (including both ends). The hotels may choose to locate at any exit. These will be the only two hotels available to any traveler using the highway. Each traveler has a most-preferred location along the highway (at some exit) for a hotel, and will choose to go to the hotel closest to that location. Travelers' most-preferred locations are distributed evenly, so that each exit has the same number of travelers who prefer that exit. If both hotels are the same distance from a traveler's most-preferred location, then that traveler flips a coin to determine which hotel to stay at. Each hotel would like to maximize the number of travelers who stay at it.

(a) If hotel 1 locates at the 100-mile exit, where should hotel 2 locate?
(b) Given hotel 2's location in (a) where would hotel 1 prefer to locate?
(c) Which pairs of locations form Nash equilibria?

9.G3 ***Backward Induction*** Find the backward induction solution to Figure 9.21 and argue that there is a unique subgame perfect equilibrium. Provide a Nash equilibrium of that game that is not subgame perfect.

9.G4 ***The Colonel Blotto Game*** Two armies are fighting a war. There are three battlefields. Each army consists of six units. The armies must each decide how many units to place on each battlefield. They do this without knowing the number of units that the other army has committed to a given battlefield. The army who has the most units on a given battlefield wins that battle, and the army that wins the most battles wins the war. If the armies each have the same number of units on a given battlefield, then there is an equal chance that either army wins that battle. A pure strategy for an army is a list (u_1, u_2, u_3) of the number of units it places on battlefields 1, 2, and 3, respectively, where each u_k is in $\{0, 1, \ldots, 6\}$ and the sum of the u_ks is 6. For example, if army A allocates its units (3,2,1), and army B allocates its units (0,3,3), then army A wins the first battle, army B wins the second and third battles, and army B wins the war.

(a) Argue that there is no pure strategy Nash equilibrium in this game.
(b) Show that mixing uniformly at random over all possible configurations of units is not a mixed strategy Nash equilibrium. (Hint: placing all units on one battlefield is not a good idea).
(c) Show that each army mixing with equal probability between (0,3,3), (3,0,3), and (3,3,0) is not an equilibrium.[58]

9.G5 ***Divide and Choose*** Two children must split a pie. They are gluttons and each prefers to eat as much of the pie as possible. The parent tells one child to cut the pie into two pieces and then allows the other child to choose which piece to eat. The first child can divide the pie into any multiple of a tenth (e.g., 1/10 and 9/10, 2/10 and 8/10). Show that there is a unique backward induction solution.

9.G6 ***Take It or Leave It Bargaining*** Two players are bargaining over a pie. The first player can suggest a division x_1, x_2 such that $x_1 + x_2 = 1$ and each share is nonnegative. Thus the game is in extensive form with an infinite number of strategies for the first player. The second player can then say either "yes" or "no." If the second player says "yes," then they each get their proposed share. If the second player says "no," then they both get nothing. Player's payoffs are their share of the pie. Argue that there is a *unique* subgame perfect equilibrium to this game.

9.G7 ***Information and Equilibrium*** Each of two players receives an envelope containing money. The amount of money has been randomly selected to be between 1 and 1,000 dollars (inclusive), with each dollar amount equally likely. The random amounts in the two envelopes are drawn independently. After looking in their own envelope, the players have a chance to trade envelopes. That is, they are simultaneously asked if they would like to trade. If they both say "yes," then the envelopes are swapped and they each go home with the new envelope. If either player says "no," then they each go home with their original envelope. What is the highest amount for which either player says "yes" in a Nash equilibrium? (Hint: Should a player say "yes" with 1,000 dollars in his or her envelope?)

58. Finding equilibria in Colonel Blotto games is notoriously difficult. One exists for this particular version, but finding it will take you some time.