

Bayesian Learning in Social Networks¹

Douglas Gale (*Corresponding Author*)

Department of Economics, New York University
269 Mercer St., 7th Floor, New York, NY, 10003-6687.

E-mail: douglas.gale@nyu.edu

Url: <http://www.econ.nyu.edu/user/galed/>

Phone: (212) 998-8944

Fax: (212) 995-3932

and

Shachar Kariv

Department of Economics, New York University
269 Mercer St., 7th Floor, New York, NY, 10003-6687.

E-mail: sk510@nyu.edu

Url: <http://home.nyu.edu/~sk510>

Version: March 13, 2003.

We extend the standard model of social learning in two ways. First, we introduce a social network and assume that agents can only observe the actions of agents to whom they are connected by this network. Secondly, we allow agents to choose a different action at each date. If the network satisfies a connectedness assumption, the initial diversity resulting from diverse private information is eventually replaced by uniformity of actions, though not necessarily of beliefs, in finite time with probability one. We look at particular networks to illustrate the impact of network architecture on speed of convergence and the optimality of absorbing states. Convergence is remarkably rapid, so that asymptotic results are a good approximation even in the medium run.

Journal of Economic Literature Classification Numbers: D82, D83

Key Words: Networks, Social learning, Herd behavior, Informational cascades.

Running Title: Bayesian Learning in Social Networks.

¹One of us discussed this problem with Bob Rosenthal several years ago, when we were both at Boston University. At that time, we found the problem of learning in networks fascinating but made no progress and were eventually diverted into working on boundedly rational learning, which led to our paper on imitation and experimentation. We thank seminar participants at NYU, DELTA, INSEAD, Cergy, Cornell and Iowa for their comments. The financial support of the National Science Foundation through Grant No. SES-0095109 is gratefully acknowledged.

1. INTRODUCTION

The canonical model of social learning comprises a set of agents I , a finite set of actions \mathcal{A} , a set of states of nature Ω , and a common payoff function $U(a, \omega)$, where a is the action chosen and ω is the state of nature. Each agent i receives a private signal $\sigma_i(\omega)$, a function of the state of nature ω , and uses this private information to identify a payoff-maximizing action.

This setup provides an example of a *pure information externality*. Each agent's payoff depends on his own action and on the state of nature. It does *not* depend directly on the actions of other agents. However, each agent's action reveals something about his private signal, so an agent can generally improve his decision by observing what others do before choosing his own action. In social settings, where agents can observe one another's actions, it is rational for them to learn from one another.

This kind of social learning was first studied by Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992). Their work was extended by Smith and Sørensen (2000). These models of social learning assume a simple *sequential* structure, in which the order of play is fixed and exogenous. They also assume that the actions of all agents are public information. Thus, at date 1, agent 1 chooses an action a_1 , based on his private information; at date 2, agent 2 observes the action chosen by agent 1 and chooses an action a_2 based on his private information and the information revealed by agent 1's action; at date 3, agent 3 observes the actions chosen by agents 1 and 2 and chooses an action a_3 ...; and so on. In what follows we refer to this structure as the *sequential social-learning model* (SSLM).

One goal of the social learning literature is to explain the striking uniformity of social behavior that occurs in fashion, fads, "mob psychology", and so forth. In the context of the SSLM, this uniformity takes the form of herd behavior.² Smith and Sørensen (2000) have shown that, in the SSLM, herd behavior arises in finite time with probability one. Once the proportion of agents choosing a particular action is large enough, the public information in favor of this action outweighs the private information of any single agent. So each subsequent agent "ignores" his own signal and "follows the herd".

This is an important result and it helps us understand the basis for uniformity of social behavior.³ At the same time, the SSLM has several

²A *herd* occurs if, after some finite date t , every agent chooses the same action. An *informational cascade* occurs if, after some finite date t , every agent finds it optimal to choose the same action regardless of the value of his private signal. An informational cascade implies herd behavior, but a herd can arise without a cascade.

³The most interesting property of the models of Bikhchandani, Hirshleifer and Welch (1992) and Banerjee (1992) is that informational cascades arise very rapidly, before much information has been revealed. For example, in these models if the first two agents make the same choice, all subsequent agents will ignore their information and imitate the first two. The behavior of a potential infinity of agents is determined by the behavior of the first two. This is both informationally inefficient and Pareto inefficient.

special features that deserve further examination: (i) each agent makes a single, irreversible decision; (ii) the timing of the agent’s decision (his position in the decision-making queue) is fixed and exogenous; (iii) agents observe the actions of *all* their predecessors; and (iv) the number of signals, like the number of agents, is infinite, so once a cascade begins the amount of information lost is large. These features simplify the analysis of the SSLM, but they are quite restrictive.

In this paper, we study the uniformity of behavior in a framework that allows for a richer pattern of social learning. We depart from the SSLM in two ways. First, we drop the assumption that actions are public information and assume that agents can observe the actions of some, but not necessarily all, of their neighbors. Second, we allow agents to make decisions simultaneously, rather than sequentially, and to revise their decisions rather than making a single, irreversible decision. We refer to this structure as the *social network model* (SNM). For empirical examples that illustrate the important role of networks in social learning, see Bikhchandani, Hirshleifer and Welch (1998).

On the face of it, uniform behavior seems less likely in the SNM, where agents have very different information sets, than in the SSLM. However, uniformity turns out to be a robust feature of *connected* social networks.⁴

The following results are established for any connected network:

Uniformity of behavior: Initially, diversity of private information leads to diversity of actions. Over time, as agents learn by observing the actions of their neighbors, some convergence of beliefs is inevitable. A central question is whether agents can rationally choose different actions forever. Disconnected agents can clearly ‘disagree’ forever. Also, there may be cases where agents are indifferent between two actions and disagreement of actions is immaterial. However, apart from cases of disconnectedness and indifference, all agents must eventually choose the same action. Thus, learning occurs through diversity but is eventually replaced by uniformity.

Optimality: We are interested in whether the common action chosen asymptotically is optimal, in the sense that the same action would be chosen if all the signals were public information. In special cases, we can show that asymptotically the optimal action is chosen but, in general, there is no reason why this should be the case.

Although the process of learning in networks can be very complicated, the SNM has several features that make the asymptotic analysis tractable. The first is the *welfare-improvement principle*. Agents have perfect recall, so expected utility is non-decreasing over time. This implies that equilibrium

⁴A network is a directed graph in which the nodes correspond to representative agents. Agent i can observe the actions of agent j if i is connected to agent j . A network is *connected* if, for any two agents i and j , there is a sequence i_1, \dots, i_K such that $i_1 = i$, $i_K = j$ and i_k is connected to i_{k+1} for $k = 1, \dots, K - 1$.

payoffs form a submartingale. We use the martingale convergence theorem to establish that an agents' (random) payoff converges almost surely to a constant.

Another useful property of the model is the *imitation principle*. If agent i can observe the actions of agent j , then one strategy available to him is to imitate whatever j does. Since i and j have different information sets, their conditional payoffs under this strategy may be different. However, on average, i must do as well as j .

The imitation principle, together with the connectedness of the network, is used to show that, asymptotically, i and j must get the same average (unconditional) payoffs. It turns out that this is only possible if agents choose the same actions. More precisely, agents choose different actions only if they are indifferent.

While the convergence properties of the model are quite general, other properties have only been established for particular networks:

Convergence in finite time: In special cases, we can rule out the possibility of indifference between actions with probability one. In that case, all agents choose the same action in finite time with probability one.

Speed of convergence: In two- and three-person networks, we can show that convergence to a uniform action is extremely rapid, typically occurring within five or six periods with probability close to 1. What happens in those first few periods is important for the determination of the asymptotic state.

Network architecture: Systematic differences can be identified in the behavior of different networks. For example, in three-person complete networks (where each agent observes all the others), learning stops almost immediately and the probability of an incorrect action in the long run is high. In three-person incomplete networks, learning continues for a longer time and the probability of choosing an incorrect action in the long run is lower.

The rest of the paper is organized as follows. In Section 2 we define the model and the equilibrium concept more precisely. In Section 3 we use the case of two-person networks to illustrate the working of the general model and some special features of complete networks. In Section 4 we derive the asymptotic properties of the general model. In Section 5 we study a selection of three-person graphs. Here we see the impact of lack of common knowledge on the dynamics of social learning and the efficiency of aggregation. We also compare the dynamic and asymptotic properties of different networks. The results are discussed in Section 6. Proofs are gathered in Section 7.

2. THE MODEL

The social learning literature ignores the complications of strategic behavior in order to focus on pure Bayesian learning. Non-strategic behavior

is simpler to analyze and it also seems more appropriate for a model of social behavior. However, special assumptions are needed to rationalized non-strategic behavior. In the SSLM, for example, an agent is assumed to make a once-in-a-lifetime decision. Because his payoff is independent of other agents' actions, it is rational for him to behave myopically and ignore the affect of his action on the agents who follow him. In the SNM, an agent's payoff is independent of other agents' actions but, unlike the SSLM, agents make repeated decisions. In order to eliminate strategic behavior, we assume that the economy comprises a large number of individually insignificant agents and that agents only observe the distribution of actions at each date. Since a single agent cannot affect the distribution of actions, he cannot influence the future play of the game. This allows us to ignore "strategic" considerations and focus on the pure Bayesian-learning aspect of the model.

The agents

Formally, we assume there is a finite set of locations indexed by $i = 1, \dots, n$. At each location, there is a non-atomic continuum of identical agents. In the sequel, the continuum of agents at location i is replaced by a single representative agent i who maximizes his short-run payoff in each period.

Uncertainty is represented by a probability measure space $(\Omega, \mathcal{F}, \mathbf{P})$, where Ω is a compact metric space, \mathcal{F} is a σ -field, and \mathbf{P} a probability measure. Time is represented by a countable set of dates indexed by $t = 1, 2, \dots$

Let $\mathcal{A} \subset \mathbf{R}$ be a finite set of actions and let $U : \mathcal{A} \times \Omega \rightarrow \mathbf{R}$ be the common payoff function, where $U(a, \cdot)$ is a bounded, measurable function for every action a . Each (representative) agent i receives a private signal $\sigma_i(\omega)$ at date 1, where $\sigma_i : \Omega \rightarrow \mathbf{R}$ is a random variable.

The network

A *social network* is represented by a family of sets $\{N_i : i = 1, \dots, n\}$, where

$$N_i \subseteq \{1, \dots, i-1, i+1, \dots, n\}.$$

For each agent i , N_i denotes the set of agents $j \neq i$ who can be observed by agent i . We can think of N_i as representing i 's "neighborhood". The sets $\{N_i : i = 1, \dots, n\}$ define a *directed graph* with nodes $N = \{1, \dots, n\}$ and edges $E = \cup_{i=1}^n \{(i, j) : j \in N_i\}$. The social network determines the information flow in the economy. Agent i can observe the action of agent j if and only if $j \in N_i$. Agents have perfect recall so their information set at each date includes the actions they have observed at every previous date.

For any nodes i and j , a *path* from i to j is a finite sequence i_1, \dots, i_K such that $i_1 = i$, $i_K = j$ and $i_{k+1} \in N_{i_k}$ for $k = 1, \dots, K-1$. A node i is *connected* to j if there is a path from i to j . The network $\{N_i\}$ is *connected* if every pair of nodes i and j is connected. Connectedness is essential for uniformity of behavior, but not for other results.

Equilibrium

At the beginning of each date t , agents choose actions simultaneously. Then each agent i observes the actions a_{jt} chosen by the agents $j \in N_i$ and updates his beliefs accordingly. Agent i 's information set at date t consists of his signal $\sigma_i(\omega)$ and the history of actions $\{a_{js} : j \in N_i, s \leq t-1\}$. Agent i chooses the action a_{it} to maximize the expectation of his short-run payoff $U(a_{it}, \omega)$ conditional on the information available.

An agent's behavior can be described more formally as follows. Agent i 's choice of action at date t is described by a random variable $X_{it}(\omega)$ and his information at date t is described by a σ -field \mathcal{F}_{it} . Since the agent's choice can only depend on the information available to him, X_{it} must be measurable with respect to \mathcal{F}_{it} . Since \mathcal{F}_{it} represents the agent's information at date t , it must be the σ -field generated by the random variables σ_i and $\{X_{js} : j \in N_i, s \leq t-1\}$. Note that there is no need to condition explicitly on agent i 's past actions because they are functions of the past actions of agents $j \in N_i$ and the signal $\sigma_i(\omega)$. Finally, since X_{it} is optimal, there cannot be any other \mathcal{F}_{it} -measurable choice function that yields a higher expected utility. These are the essential elements of our definition of equilibrium, as stated below.

DEFINITION 1. A weak perfect Bayesian equilibrium consists of a sequence of random variables $\{X_{it}\}$ and σ -fields $\{\mathcal{F}_{it}\}$ such that for each $i = 1, \dots, n$ and $t = 1, 2, \dots$,

- (i) $X_{it} : \Omega \rightarrow \mathcal{A}$ is \mathcal{F}_{it} -measurable,
- (ii) $\mathcal{F}_{it} = \mathcal{F}(\sigma_i, \{X_{js} : j \in N_i, s=1, \dots, t-1\})$, and
- (iii) $E[U(x(\omega), \omega)] \leq E[U(X_{it}(\omega), \omega)]$, for any \mathcal{F}_{it} -measurable function $x : \Omega \rightarrow \mathcal{A}$.

Note that our definition of equilibrium does not require optimality "off the equilibrium path". This entails no essential loss of generality as long as it is assumed that the actions of a single agent, who is of measure zero, are not observed by other players. Then a deviation by a single agent has no effect on the subsequent decisions of other agents.

3. LEARNING WITH TWO (REPRESENTATIVE) AGENTS AND TWO ACTIONS

To fix ideas and illustrate the workings of the basic model, we first consider the special case of two representative agents, A and B , and two actions, 0 and 1. There are three graphs, besides the empty graph $N_A = N_B = \emptyset$,

- (i) $N_A = \{B\}, N_B = \{A\}$;
- (ii) $N_A = \{B\}, N_B = \emptyset$;
- (iii) $N_A = \emptyset, N_B = \{A\}$.

Cases (ii) and (iii) are uninteresting because there is no possibility of mutual learning. For example, in case (ii), agent B observes a private signal and chooses the optimal action at date 1. Since he observes no further information, he chooses the same action at every subsequent date. Agent A observes a private signal and chooses the optimal action at date 1. At date 2, he observes agent B 's action at date 1, updates his beliefs and chooses the new optimal action at date 2. After that, A receives no additional information, so agent A chooses the same action at every subsequent date. Agent A has learned something from agent B , but that is as far as it goes. In case (i), on the other hand, the two agents learn from each other and learning can continue for an unbounded number of periods. We focus on the network defined in (i) in what follows.

For simplicity, we consider a special information and payoff structure. We assume that $\Omega = \Omega_A \times \Omega_B$, where Ω_i is an interval $[a, b]$ and the generic element is $\omega = (\omega_A, \omega_B)$. The signals are assumed to satisfy

$$\sigma_i(\omega) = \omega_i, \forall \omega \in \Omega, i = A, B,$$

where the random variables ω_A and ω_B are independently and continuously distributed, that is, $\mathbf{P} = \mathbf{P}_A \times \mathbf{P}_B$ and \mathbf{P}_i has no atoms. There are two actions $a = 0, 1$ and the payoff function is assumed to satisfy

$$u(a, \omega) = \begin{cases} 0 & \text{if } a = 0 \\ U(\omega) & \text{if } a = 1, \end{cases}$$

where the function $U(\omega_A, \omega_B)$ is assumed to be a continuous and increasing function. To avoid trivialities we assume that neither action is weakly dominated.

These assumptions are sufficient for the optimal strategy to have the form of a cutoff rule. To see this, note that for any history that occurs with positive probability, agent i 's beliefs at date t take the form of an event $\{\omega_i\} \times B_{jt}$, where the true value of ω_j is known to belong to B_{jt} . Then the payoff to action 1 is $\varphi_i(\omega_i, B_{jt}) = E[U(\omega_A, \omega_B) | \{\omega_i\} \times B_{jt}]$. Clearly, $\varphi_i(\omega_i, B_{jt})$ is increasing in ω_i , because the distribution of ω_j is independent of ω_i , so there exists a cutoff $\omega_i^*(B_{jt})$ such that

$$\begin{aligned} \omega_i > \omega_i^*(B_{jt}) &\implies \varphi_i(\omega_i, B_{jt}) > 0, \\ \omega_i < \omega_i^*(B_{jt}) &\implies \varphi_i(\omega_i, B_{jt}) < 0. \end{aligned}$$

We assume that when an agent is indifferent between two actions, he chooses action 1. The analysis is essentially the same for any other the tie-breaking rule.

The fact that agent i 's strategy takes the form of a cutoff rule implies that the set B_{it} is an interval. This can be proved by induction as follows. At date 1, agent j has a cutoff ω_{j1}^* and $X_{j1}(\omega) = 1$ if and only if $\omega_j \geq \omega_{j1}^*$.

Then at date 2 agent i will know that the true value of ω_j belongs to $B_j(\omega)$, where

$$B_{j2}(\omega) = \begin{cases} [\omega_{j1}^*, b] & \text{if } X_{j1}(\omega) = 1, \\ [a, \omega_{j1}^*] & \text{if } X_{j1}(\omega) = 0. \end{cases}$$

Now suppose that at some date t , the information set $B_{jt}(\omega) \subseteq [a, b]$ is an interval and agent j 's cutoff is $\omega_{jt}^*(B_{it}(\omega))$. Then at date $t + 1$, agent i knows that ω_j belongs to $B_{jt+1}(\omega)$, where

$$B_{jt+1}(\omega) = \begin{cases} B_{jt}(\omega) \cap [\omega_{jt}^*(B_{it}(\omega)), b] & \text{if } X_{jt}(\omega) = 1, \\ B_{jt}(\omega) \cap [a, \omega_{jt}^*(B_{it}(\omega))] & \text{if } X_{jt}(\omega) = 0. \end{cases}$$

Clearly, $B_{jt+1}(\omega)$ is also an interval. Hence, by induction, $B_{it}(\omega)$ is an interval for all t and the common knowledge at date t is $B_t(\omega) = B_{At}(\omega) \times B_{Bt}(\omega)$. By construction, $\omega \in B_{t+1}(\omega) \subseteq B_t(\omega)$ for every t . Then $B_t(\omega) \searrow B(\omega) = \cap_{t=1}^{\infty} B_t(\omega)$ and $\{B(\omega) : \omega \in \Omega\}$ defines a partition of Ω . Note that $\omega \in B(\omega)$ so $B(\omega) \neq \emptyset$.

In the limit, when all learning has ceased, agent A knows that ω_B belongs to a set $B_B(\omega)$ and agent B knows that ω_A belongs to $B_A(\omega)$. Furthermore, because the actions chosen at each date are common knowledge, the sets $B_A(\omega)$ and $B_B(\omega)$ are common knowledge.

An interesting question is whether, given their information in the limit, the two agents must agree which action is best. In the two-person case, we can show directly that both agents must eventually agree, in the sense that they choose different actions only if they are indifferent. The proof is by contradiction. Suppose, contrary to what we want to prove, that for some B and every ω such that $B(\omega) = B$,

$$E[U(\omega_A, \omega_B) | \{\omega_A\} \times B_B] > 0$$

and

$$E[U(\omega_A, \omega_B) | B_A \times \{\omega_B\}] < 0.$$

Then the same actions must be optimal for every element in the information set (otherwise, more information would be revealed) and this implies

$$E[U(\underline{\omega}_A, \omega_B) | \{\underline{\omega}_A\} \times B_B] \geq 0$$

and

$$E[U(\omega_A, \bar{\omega}_B) | B_A \times \{\bar{\omega}_B\}] \leq 0,$$

where $\underline{\omega}_A = \inf B_A(\omega)$ and $\bar{\omega}_B = \sup B_B(\omega)$. Then

$$U(\underline{\omega}_A, \bar{\omega}_B) \geq 0$$

and

$$U(\underline{\omega}_A, \bar{\omega}_B) \leq 0,$$

or $U(\underline{\omega}_A, \bar{\omega}_B) = 0$. If B_B is not a singleton,

$$E[U(\underline{\omega}_A, \omega_B) | \{\underline{\omega}_A\} \times B_B] < 0$$

a contradiction. Similarly, if B_A is not a singleton,

$$E[U(\omega_A, \bar{\omega}_B) | B_A \times \{\bar{\omega}_B\}] > 0,$$

a contradiction. Thus, B is a singleton and $U(\omega) = 0$ if $\omega \in B$.

The set $\{\omega : U(\omega) = 0\}$ has probability zero, so the probability of disagreeing forever is 0. In other words, both agents will choose the same action in finite time and once they have chosen the same action, they have reached an absorbing state and will continue to choose the same action in every subsequent period.

3.1. An example

To illustrate the short-run dynamics of the model, we can further specialize the example by assuming that, for each agent i , the signal $\sigma_i(\omega) = \omega_i$ is uniformly distributed on the interval $[-1, 1]$ and the payoff to action 1 is $U(\omega) = \omega_A + \omega_B$.

At date 1, each agent chooses 1 if his signal is positive and zero if it is negative. If both choose the same action at date 1, they will continue to choose the same action at each subsequent date. Seeing the other agent choose the same action will only reinforce each agent's belief that he has made the correct choice. No further information is revealed at subsequent dates and so we have reached an absorbing state, in which each agent knows his own signal and that the other's signal has the same sign, but nothing more. So interesting dynamics occur only in the case where agents choose different actions at date 1. The exact nature of the dynamics depends on the relative strength of the two signals, measured here by their absolute values. Without loss of generality, we assume that A has a negative signal, B a positive signal, and B 's signal is relatively the stronger, i.e., $|\omega_A| < |\omega_B|$.

Case 1: $\omega_A > -1/2$ and $\omega_B > 1/2$. In the first round at date 1, agent A will choose action 0 and agent B will choose action 1. At the second date, having observed that agent B chose 1, agent A will switch to action 1, while agent B will continue to choose 1. Thereafter, there is an absorbing state in which both agents choose 1 for ever and no further learning occurs.

Case 2: $3/4 < \omega_A < -1/2$ and $\omega_B > 3/4$. As before, A chooses 0 and B chooses 1 at date 1. At date 2, A observes that B chose 1 and infers that his signal has expected value $1/2$. Since $\omega_A < -1/2$, it is optimal for A to choose 0 again. Since B has an even stronger signal, he will continue to choose 1. At date 3, A observes that B chose 1, thus revealing that $\omega_B > 1/2$ so the expected value of B 's signal is $3/4$ and since $\omega_A > -3/4$ it is optimal for him to switch to 1, which then becomes an absorbing state.

Case 3: $-(t-1)/t > \omega_A > -(t-2)/t$ and $\omega_B > (t-1)/t$. By analogous reasoning, A chooses 0 and B chooses 1 until date t when A switches to 1.

The other interesting case to consider is when the signals are similar in strength. For example, suppose that $\omega_A = -\omega_B = x^*$ where x^* is the limit of the sequent $\{x_t\}_{t=1}^\infty$ defined by putting $x_1 = \frac{1}{2}, x_2 = \frac{1}{4}$, and

$$x_t = \frac{1}{2}(x_{t-1} + x_{t-2})$$

for $t = 3, 4, \dots$. Notice that if t is even then $x_t < x^* < x_{t-1}$.

As usual A chooses 0 and B chooses 1, at date 1. At date 2, A observes B 's choice in the previous period, realizes that the expected value of B 's signal is $1/2 > -x^*$ and switches to 1. By the symmetric argument, B switches to 0. At date 3, A observes B 's switch to 0 and realizes that $1/4 < \omega_B < 1/2$, that is, the expected value of ω_B is greater than $x^* = -\omega_A$. So it is optimal for A to choose 0 again. By a symmetric argument, B switches back to 1 at date 3.

At any even date t , A will choose 1 and B will choose 0 and at any odd date t , A will choose 0 and B will choose 1. B 's choice at an even date t reveals that $x_{t-2} < \omega_B < x_{t-1}$ and his choice at an odd date reveals $x_{t-1} < \omega_B < x_{t-2}$. By construction, at any odd date t , $-\omega_A = x^* < x_t = \frac{1}{2}(x_{t-1} + x_{t-2})$, so it is optimal for A to choose 1 at $t + 1$. Likewise, at any even date t , $-\omega_A = x^* > x_t = \frac{1}{2}(x_{t-1} + x_{t-2})$, so it is optimal for A to choose 0 at $t + 1$.

In fact, we can find a signal ω to rationalize any sequence of actions with the properties that for some T , $x_{At} \neq x_{Bt}$ for $t < T$ and $x_{At} = x_{Bt} = a$ for $t \geq T$. However, the sequences corresponding to $T = \infty$ occur with probability 0 and the sequences with $T < \infty$ occur with positive probability.

This example can also be used to illustrate the speed of convergence to uniformity of actions. In the first period, the probability that agents choose the same action is $1/2$. In the second period, it is $3/4$. In the third period, it is $7/8$, and so on. This is a very special example, but simulations of other examples confirm these results.

Finally, we note that in this simple example, where the signals of the two players are symmetrically distributed, the asymptotic outcome must be Pareto-efficient. This follows from the fact that the agent with the stronger signal, as measured by its absolute value, will ultimately determine the action chosen. However, a simple adjustment to this example shows the possibility of an inefficient outcome. Suppose that A has a signal uniformly distributed on $[0, 1]$ and B has a signal uniformly distributed on $[-\frac{1}{2}, 1]$. Then both A and B will choose action 1 at the first date and there will be no learning. However, if ω_A is close to 0 and ω_B is close to $-\frac{1}{2}$ then action 0 is clearly preferred conditional on the information available to the two agents.

4. ASYMPTOTIC PROPERTIES

Now we return to the general model described in Section 2 and study the asymptotic properties.

Although the process of learning in a social network is complicated, it has a number of features that make the characterization of the asymptotic outcomes tractable. The first is the Welfare-Improvement Principle: since an agent's information is non-decreasing over time, his payoff must also be non-decreasing. This allows us to apply the Martingale Convergence Theorem to show that equilibrium payoffs converge with probability one as time goes to infinity. Smith and Sørensen (2000) also use the Martingale Convergence Principle to show that Bayesian learning eventually leads to convergence of actions and beliefs.

The second useful feature is the Imitation Principle: since an agent can always imitate his neighbor, he must do at least as well as his neighbor on average (with a one-period lag). Similar ideas have been used by Banerjee and Fudenberg (1995) and Bala and Goyal (1998). We must be careful in exploiting this property, since it does not imply that an agent will do as well as his neighbor with probability one. Nonetheless, it turns out to be a useful property.

The Imitation Principle is particularly useful in a connected network. If i is connected to j then we can use the Imitation Principle recursively to argue that i does as well as j . If j is connected to i the same argument implies that i and j receive the same payoff on average. This fact is then used to show that i and j must choose the same action unless they are both indifferent. In other words, they essentially agree on the best action to take. In cases where indifference occurs with probability zero, we have uniformity in the limit.

Without this connectedness assumption, there is no reason to expect equal payoffs or uniformity. A trivial example, would be the two-person network $N_A = \{B\}, N_B = \emptyset$, where A observes B but B does not observe A . Clearly, A must do at least as well as B but may do better and there is no reason why A should always choose the same action as B . On the other hand, the complete network $N_A = \{B\}, N_B = \{A\}$ is connected and as we saw in the previous section, uniformity of actions arises with probability one in finite time.

4.1. Convergence

The first step in our analysis is to establish convergence of beliefs and payoffs. From the definition of equilibrium, we note that $\mathcal{F}_{it} \subseteq \mathcal{F}_{it+1} \subseteq \mathcal{F}$ for every i and t . In other words, an agent's information is non-decreasing over time. Then his equilibrium payoff must be non-decreasing over time and, since it is bounded, must converge. This property is established in the next theorem.

THEOREM 1. Let $\{X_{it}, \mathcal{F}_{it} : i = 1, \dots, n, t = 1, 2, \dots\}$ be an equilibrium. For each i , define $V_{it}^* : \Omega \rightarrow \mathbf{R}$ by

$$V_{it}^* = E[U(X_{it}, \cdot) | \mathcal{F}_{it}].$$

Then $\{V_{it}^*\}$ is a submartingale with respect to $\{\mathcal{F}_{it}\}$ and there exists a random variable $V_{i\infty}^*$ such that V_{it}^* converges to $V_{i\infty}^*$ almost surely.

4.2. The Imitation Principle

The next step is to establish the Imitation Principle, which states that asymptotically an agent must do at least as well as his neighbors. This follows directly from the fact that one strategy available to agent i is to imitate the actions of agent $j \in N_i$.

COROLLARY 1. Let $\{X_{it}, \mathcal{F}_{it}\}$ be the equilibrium in Theorem 1 and let V_{it}^* be the equilibrium payoffs. Then for any $j \in N_i$ and any t , $V_{it}^* \geq E[V_{jt-1}^* | \mathcal{F}_{it}]$. Furthermore, in the limit,

$$V_{i\infty}^* \geq E[V_{j\infty}^* | \mathcal{F}_{i\infty}],$$

where $\mathcal{F}_{i\infty}$ is the σ -field generated by $\cup\{\mathcal{F}_{i1}, \mathcal{F}_{i2}, \dots\}$.

4.3. Connectedness

In order to make use of the foregoing results to establish uniformity of actions, we need to make use of connectedness. We begin by studying the behavior of adjacent agents i and $j \in N_i$ and then extend the results to the entire network.

It is easy to see that if i is connected to j , then $E[V_{i\infty}^*] \geq E[V_{j\infty}^*]$ by induction. In particular, if $j \in N_i$ and j is connected to i then $V_{i\infty}^* \geq E[V_{j\infty}^* | \mathcal{F}_{it}]$ and $E[V_{j\infty}^*] \geq E[V_{i\infty}^*]$, which implies that $V_{i\infty}^* = E[V_{j\infty}^* | \mathcal{F}_{i\infty}]$.

COROLLARY 2. Let $\{X_{it}, \mathcal{F}_{it}\}$ be the equilibrium in Theorem 1 and let V_{it}^* be the equilibrium payoffs. If $j \in N_i$ and j is connected to i then $V_{i\infty}^* = E[V_{j\infty}^* | \mathcal{F}_{i\infty}]$.

Our next result concerns the possibility for agents to choose different actions in the long run. The fact that agents get the same payoff in the long run suggests that they must choose the same actions unless they are indifferent. This result requires a certain amount of care because their information sets are different, but the intuition is essentially correct as the next theorem shows.

Let $\{X_{it}, \mathcal{F}_{it} : i = 1, \dots, n, t = 1, 2, \dots\}$ be an equilibrium and define $V_{it}^a : \Omega \rightarrow \mathbf{R}$ by

$$V_{it}^a = E[U(a, \cdot) | \mathcal{F}_{it}]$$

for any agent i , date t , and action a . Then $\{V_{it}^a\}$ is a martingale with respect to $\{\mathcal{F}_{it}\}$ and V_{it}^a converges to a random variable $V_{i\infty}^a$ almost surely.

THEOREM 2. Let i and j be two agents such that $j \in N_i$ and j is connected to i . Let E^{ab} denote the measurable set on which i chooses a infinitely often and j chooses b infinitely often. Then $V_{i\infty}^a(\omega) = V_{i\infty}^b(\omega)$ for almost every ω in E^{ab} . That is, i is indifferent between a and b for almost every ω in E^{ab} .

Intuitively, if i always believes that he is getting the same payoff as agent j then i cannot believe that he is choosing a better action than j . In this sense, they cannot disagree forever.

Clearly, since the network is connected, every agent asymptotically has the same payoff and, indifference apart, all agents must choose the same actions.


The concept of connectedness used here is strong in the sense that there must be a path running both ways between any pair of nodes. If the network is not connected in this sense, one can still apply Theorem 2 to connected components of the graph, that is, maximal subsets of nodes such that every pair of nodes in the subset is connected.

5. SHORT-RUN DYNAMICS

To illustrate the short-run dynamics of the model, we adapt the example introduced in Section 3.1 by assuming there are three agents, A , B , and C , and two actions, 0 and 1. As before, the payoff from action 0 is identically 0 and the payoff from action 1 is $U(\omega) = \omega_A + \omega_B + \omega_C$, where, for each agent i , the signal $\sigma_i(\omega) = \omega_i$ is uniformly distributed on the interval $[-1, 1]$.

The three-person case, unlike the two-person case, has several non-trivial social networks, each of which gives rise to its own distinctive learning patterns. We refer to a network in which every agent directly observes every other agent as *complete*. Otherwise, the network is *incomplete*. The network studied in Section 3 is complete whereas most of the networks for the three-person case are incomplete. Several social networks are illustrated in Figure 1.

In a complete network, the entire history of past actions is common knowledge at each date. In an incomplete social network, past actions are typically not common knowledge at each date. This lack of common knowledge plays an important role in the learning process. It forces agents to make more or less complicated inferences about what other agents have seen, as well as about the inferences those agents have drawn, and changes the nature of the decision rules adopted by the agents.



Illustrations of Social Networks with Three Agents

- (i) The unique complete network
- (ii) The Circle (iii) The Star
- (iv) Two other connected but incomplete networks

FIG. 1

5.1. The complete network

As a benchmark, consider the (unique) complete network, in which each agent can observe the other two:

$$N_A = \{B, C\}, N_B = \{A, C\}, N_C = \{A, B\}.$$

If agents choose the same action at the first date, learning effectively ends there. For example, suppose that $\omega_i > 0$ for $i = A, B, C$. Agent i 's expected payoff from action 1 is $\omega_i > 0$, since $E[\omega_j] = 0$ for $j \neq i$. So each agent will choose action 1 at the first date. At the second date, seeing that the others have chosen the same action at date 1, agent i will infer that $\omega_j > 0$ and hence that $E[\omega_j | \omega_j > 0] = 1/2$ for $j \neq i$. This will increase agent i 's payoff from action 1 from ω_i to $\omega_i + 1$ and reinforce agent i 's preference for action 1. So each agent will continue to choose action 1 at date 2. At date 3 there is no change in actions or beliefs, so we have reached an absorbing state. Given the assumed values of the signals, the outcome is efficient.

A more interesting case is one in which there is diversity of actions at date 1. For example, suppose that $\omega_A > 0$, $\omega_B > 0$, and $\omega_C < 0$. At date 1, agents A and B will choose action 1 and agent C will choose action 0. At date 2 it becomes common knowledge that $\omega_A > 0$, $\omega_B > 0$, and $\omega_C < 0$.

The payoff from action 1 conditional on agent A 's information is

$$\omega_A + E[\omega_B | \omega_B > 0] + E[\omega_C | \omega_C < 0] = \omega_A + \frac{1}{2} - \frac{1}{2} = \omega_A.$$

Similarly, the payoff from action 1 conditional on agent B 's information is ω_B . So both A and B will continue to choose action 1. Conditional on agent C 's information, however, the payoff from action 1 is

$$E[\omega_A | \omega_A > 0] + E[\omega_B | \omega_B > 0] + \omega_C = \frac{1}{2} + \frac{1}{2} + \omega_C > 0$$

since $\omega_C > -1$. So C will switch to action 1 at date 2. At date 3, no further information is revealed, so actions and beliefs remain unchanged and once again we have reached an absorbing state. Agent C ignores his own information and joins the "herd" consisting of agents A and B . Clearly, the outcome will be inefficient if ω_A and ω_B are small relative to the absolute value of ω_C .

In either of the cases considered, the learning process comes to a halt by the end of the second period at the latest.

5.2. Incomplete networks

The first incomplete network we examine is the Star:

$$N_A = \{B, C\}, N_B = \{A\}, N_C = \{A\}.$$

Thus, at each date, A is informed about the entire history of actions that have already been taken, whereas B and C have imperfect information and thus have to form expectations about the actions of the unobserved third agent.

As with the complete network, if all the agents choose the same action at date 1, this is an absorbing state. So consider again the more interesting case where there is diversity at date 1, for example, $\omega_A > 0$, $\omega_B > 0$, and $\omega_C < 0$. The analysis of the decisions of the agents at date 1 is unchanged but now at date 2 agent C only observes that his action at date 1 does not match A 's action. Conditional on C 's information, the expected value of A 's signal is $1/2$ and the expected value of B 's signal is zero. Thus, at date 2, it is optimal for C to switch to 1 if $\omega_C \geq -1/2$ and to continue to choose 0 otherwise.

By the third round at date 3, agent C can draw some conclusions about the actions that B could have taken by observing the actions of agent A at dates 1 and 2. If A chooses 1 at both dates then it is revealed to C that B 's signal is positive; otherwise agent A would have switched to action 0 at date 2. Thus, a simple calculation shows that, having observed that A continues to choose 1, it is optimal for C to switch to action 1 for any realization of his private signal.

Even so, we have not necessarily reached an absorbing state, as agent A might himself switch at date 3. To see this, note that at date 3, A 's expected value of B 's signal is $1/2$ and A 's expected value of C 's signal is $-3/4$. Thus, it is optimal for A to choose 1 again if $\omega_A \geq 1/4$ and to switch to 0 otherwise. In case $\omega_A < 1/4$, it is common knowledge at date 4 that $0 \leq \omega_A < 1/4$, $0 \leq \omega_B \leq 1$ and $-1 \leq \omega_C < -1/2$. Table 1 summarizes the play of the game and shows that it can continue for quite a few periods.

t	(a_A, a_B, a_C)	$E_A[\omega_B], E_A[\omega_C]$	$E_B[\omega_A], E_B[\omega_C]$	$E_C[\omega_A], E_C[\omega_B]$
1	(1, 1, 0)	0, 0	0, 0	0, 0
2	(1, 1, 0)	$1/2, -1/2$	$1/2, 0$	$1/2, 0$
3	(0, 1, 1)	$1/2, -3/4$	$1/2, 0$	$1/2, 1/2$
4	(0, 1, 0)	$1/2, -3/4$	$1/8, -3/4$	$1/8, 1/2$

Table 1

As Table 1 illustrates, the dynamics of actions and learning are quite different in complete and incomplete social networks. First, in an incomplete network, learning does not end after two periods and more information may be revealed as a result. Secondly, for the diversity of actions to persist, the agent at the center of the Star must have a signal that is relatively weak (as measured by its absolute value) compared to the agents at the periphery. When his signal is relatively weak, the central agent changes his action frequently, thus transmitting information between the peripheral agents.

Alternating actions can also arise in the Star network when agent A has the negative signal and B and C have positive signals. At date 2, agent A observes that both B and C chose action 1 at date 1, so it is optimal for him to ignore his own information and to switch to action 1. However, at the same time, either agent B or agent C (or both) would switch from action 1 to action 0 if their signals are weak (less than $1/2$). Table 2 illustrates that alternating actions may continue beyond period 2 and that if the signals of agents B and C are relatively weak, say $\omega_B < 1/4$ and $\omega_C < 1/4$, and A 's signal is relatively strong, $\omega_A < -1/2$, there is an absorbing state in which all choose action 0.

t	(a_A, a_B, a_C)	$E_A[\omega_B], E_A[\omega_C]$	$E_B[\omega_A], E_B[\omega_C]$	$E_C[\omega_A], E_C[\omega_B]$
1	(0, 1, 1)	0, 0	0, 0	0, 0
2	(1, 0, 0)	$1/2, 1/2$	$-1/2, 0$	$-1/2, 0$
3	(0, 1, 1)	$1/4, 1/4$	$-1/2, 1/2$	$-1/2, 1/2$
4	(0, 0, 0)	$1/8, 1/8$	$-3/4, 1/2$	$-3/4, 1/2$

Table 2

A second example of an incomplete social network is provided by the Circle, in which each agent observes one other agent:

$$N_A = \{C\}, N_B = \{A\}, N_C = \{B\}.$$

In this case, generically no subset of the history of actions is shared as public information, and thus each agent makes different inferences about what others have learned. Thus, this network best illustrates how lack of common knowledge plays a crucial role.

Proceeding with the above example where $\omega_A > 0$, $\omega_B > 0$ and $\omega_C < 0$, suppose that $(t-2)/t > \omega_A > (t-1)/t$ and $\omega_C < -(t-1)/t$. As in the two-person case where agents have strong signals with opposite signs, so too in this situation despite the lack of common knowledge agents can agree to disagree. In fact, over time agent A and agent C learn that at least one of the other agents must have contrary private information which is stronger than they thought and thus they adjust their expectations towards -1 and 1 respectively. This continues until date t when agent A realizes that his own signal is weaker and will switch to action 0 .

Then, at date $t+1$, having observed that agent A switched to action 1 , agent B infers that agent A 's signal has expected value $(2t-3)/2t$. On the other hand, agent B cannot tell whether agent C has also switched, but given his beliefs about C 's information and strategy, he infers that agent C 's signal has expected value $(t-1)/t$. Thus, at date $t+1$, it is optimal for agent B to switch to 0 if $\omega_B < 1/2t$ and to continue to choose 1 otherwise.

6. DISCUSSION

There is a large literature on the economics of networks. The most closely related paper is by Bala and Goyal (1998), henceforth BG. In the BG model, at each date, an agent chooses one of several available actions with unknown payoff distributions. The agent observes his payoff from the action and uses this information to update his beliefs about the payoff distribution. Agents are organized in a network and can observe the actions and payoffs of their neighbors, that is, the agents with whom they are directly linked. This is a model of *social experimentation*, in the sense that it generalizes the problem of a single agent experimenting with a multi-armed bandit to a social setting, rather than social learning: agents learn by observing the outcome (payoff) of an experiment (choice of action) rather than by inferring another agent's private information from his action. A model of social experimentation is quite different from a model of social learning. In a model of social experimentation, there is an informational externality but there is no informational asymmetry.

There is private information in the BG model, but agents are assumed to ignore it. For example, suppose that agent A observes agent B 's action

and payoff but not agent C 's, whereas agent B observes agent C 's action and payoff. Then agent B has private information, not available to agent A , that is clearly payoff-relevant for agent A as well as agent B . However, A is assumed to ignore this linkage. A learns from B 's experiments (actions), but does not ask what information might have led B to choose those experiments (actions).

BG show that, in a connected network, in the long run, everyone adopts the same action and the action chosen can be sub-optimal.

Our model differs from BG in two ways. First, we examine the decisions of *fully rational* agents, who infer the information of unobserved agents from the actions they observe. Although the beliefs of agents are very complicated, it captures the idea that agents try to extract information about unobserved one, especially in small groups. Second, in our model agents observe only the actions of other agents, whereas in BG agents observe payoffs as well as actions. Obviously, there is a lot more information available in BG.

The techniques used in this paper can be applied to other models. For example, there is no difficulty applying them to random graphs, as long as connectedness is satisfied with probability one. They could also be applied to dynamic graphs where the set of neighbors observed changes over time.

Many questions about social learning in networks remain open. In special cases, we have established that uniformity arises in finite time with probability one. We conjecture that this result is true for all connected social networks, but we have yet to provide a general proof. This result would follow as a corollary of Theorem 2 if we could prove that the probability of indifference in the limit is zero, as we did in the two-person case. The impossibility of indifference is harder to establish for networks with more agents but we believe it must be true under some regularity conditions.

A second conjecture is that the result about asymptotic uniformity has a converse: if all agents choose the same action, they have reached an absorbing state and will continue to choose that action forever. This is true in the special cases we have looked at but we believe it must be true in general.

Speeds of convergence can be established analytically in simple cases. For more complex cases, we have been forced to use numerical methods. The computational difficulty of solving the model is massive even in the case of three persons. However, the results are sufficiently dramatic that they suggest the same might be true for more general cases. This is an important subject for future research.

7. PROOFS

7.1. Proof of Theorem 1

By definition (Billingsley, 1986, p. 480), the sequence $\{(V_{it}^*, \mathcal{F}_{it}) : t = 1, 2, \dots\}$ is a submartingale if the following four conditions hold:

- (i) $\mathcal{F}_{it} \subseteq \mathcal{F}_{it+1}$;
- (ii) V_{it}^* is measurable \mathcal{F}_{it} ;
- (iii) $E[|V_{it}^*|] < \infty$;
- (iv) with probability 1, $E[V_{it+1}^* | \mathcal{F}_{it}] \geq V_{it}^*$.

The first conditions follows directly from the definition of equilibrium. The second holds because $U(a, \cdot)$ is \mathcal{F}_{it} -measurable, X_{it} is \mathcal{F}_{it} -measurable, and

$$V_{it}^* = E[U(X_{it}(\cdot), \cdot) | \mathcal{F}_{it}].$$

The third condition follows because $\mathcal{A} \subset \mathbf{R}$ is finite and $U(a, \cdot)$ is bounded for each a . To establish the fourth condition, note that since $\mathcal{F}_{it} \subseteq \mathcal{F}_{it+1}$, X_{it} is \mathcal{F}_{it+1} -measurable and the equilibrium conditions imply that

$$\begin{aligned} E[U(X_{it}, \cdot) | \mathcal{F}_{it+1}] &\leq E[U(X_{it+1}, \cdot) | \mathcal{F}_{it+1}] \\ &= V_{it+1}^*. \end{aligned}$$

Then

$$\begin{aligned} V_{it}^* &= E[U(X_{it}, \cdot) | \mathcal{F}_{it}] \\ &= E[E[U(X_{it}, \cdot) | \mathcal{F}_{it+1}] | \mathcal{F}_{it}] \\ &\leq E[V_{it+1}^* | \mathcal{F}_{it}] \end{aligned}$$

and $\{(V_{it}^*, \mathcal{F}_{it}) : t = 1, 2, \dots\}$ is a sub-martingale.

From the martingale convergence theorem, there exists a random variable $V_{i\infty}^*$ such that $V_{it}^* \rightarrow V_{i\infty}^*$ almost surely.

7.2. Proof of Corollary 1

We note that for any $j \in N_i$, X_{jt-1} is \mathcal{F}_{it} -measurable so the equilibrium conditions imply that

$$\begin{aligned} V_{it}^* &\geq E[U(X_{jt-1}, \cdot) | \mathcal{F}_{it}] \\ &= E[E[U(X_{jt-1}, \cdot) | \mathcal{F}_{jt-1}] | \mathcal{F}_{it}] \\ &= E[V_{jt-1}^* | \mathcal{F}_{it}]. \end{aligned}$$

From this inequality it follows that $V_{i\infty}^* \geq E[V_{j\infty}^* | \mathcal{F}_{i\infty}]$, where $\mathcal{F}_{i\infty}$ is the σ -field generated by $\cup\{\mathcal{F}_{i1}, \mathcal{F}_{i2}, \dots\}$.

7.3. Proof of Theorem 2

Let i and j be two agents such that $j \in N_i$ and j is connected to i and let a and b be fixed but arbitrary actions. Define $E = \{\omega : X_{it} = a \text{ i.o.}, X_{jt} = b \text{ i.o.}\}$. Then E is a measurable set, in fact, $E \in \mathcal{F}_{i\infty}$. Let $\chi_E : \Omega \rightarrow \{0, 1\}$ denote the indicator function for the set E , that is,

$$\chi_E(\omega) = \begin{cases} 1, & \omega \in E, \\ 0, & \omega \notin E. \end{cases}$$

Since $V_{it}^a(\omega) = V_{it}^*(\omega)$ i.o. for every $\omega \in E$ and $V_{it}^a \rightarrow V_{i\infty}^a$ almost surely, we have $V_{i\infty}^a = V_{i\infty}^*$ for almost every $\omega \in E$. Similarly, $V_{j\infty}^b = V_{j\infty}^*$ for almost every $\omega \in E$. From Theorem 1, with probability one,

$$V_{it}^* \chi_E \rightarrow V_{i\infty}^* \chi_E$$

and

$$V_{jt}^* \chi_E \rightarrow V_{j\infty}^* \chi_E.$$

From Corollary 2, with probability one,

$$E[V_{i\infty}^* \chi_E | \mathcal{F}_{i\infty}] = E[V_{j\infty}^* \chi_E | \mathcal{F}_{i\infty}].$$

For any action a in \mathcal{A} let

$$V_{it}^a = E[U(a, \cdot) | \mathcal{F}_{it}].$$

Clearly, $\{V_{it}^a\}$ is a martingale and $V_{it}^a \rightarrow V_{i\infty}^a$ almost surely. Let

$$E_i^a = \{\omega : V_{i\infty}^a(\omega) \geq V_{i\infty}^b(\omega), b \neq a\}.$$

Then $E_i^a \cap E_j^b$ belongs to $\mathcal{F}_{i\infty} \cap \mathcal{F}_{j\infty}$. Suppose that $\mathbf{P}[E_i^a \cap E_j^b] > 0$ for some $a \neq b$. We conclude

$$\begin{aligned} E[V_{i\infty}^* | E_i^a \cap E_j^b] &= E[U(a, \cdot) | E_i^a \cap E_j^b] \\ &\geq E[U(b, \cdot) | E_i^a \cap E_j^b] \\ &= E[V_{j\infty}^* | E_i^a \cap E_j^b]. \end{aligned}$$

Since $V_{i\infty}^* = E[V_{j\infty}^* | \mathcal{F}_{i\infty}]$ we have $V_{i\infty}^a(\omega) = V_{i\infty}^b(\omega)$ for almost every $\omega \in E_i^a \cap E_j^b$. Thus, agents i and j can disagree (choose different optimal actions) in the limit only if i is indifferent.

REFERENCES

- Bala, V. and Goyal, S. (1998). "Learning from Neighbors," *Rev. Econ. Stud.* **65**, 595-621.
- Banerjee, A. (1992). "A Simple Model of Herd Behavior," *Quart. J. Econ.* **107**, 797-817.


Banerjee, A. and Fudenberg, D. (1995). "Word-of-Mouth Learning," MIT, mimeo.

Bikhchandani, S., Hirshleifer, D. and Welch, I. (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascade," *J. Polit. Economy* **100**, 992-1026.

Billingsley, P. (1986). *Probability and Measure, 2nd edition*. John Wiley and Sons, New York.

Bikhchandani, S., Hirshleifer, D. and Welch, I. (1998). "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades," *J. Econ. Perspect.* **12**, 151-170.

Smith, L. and Sørensen, P. (2000). "Pathological Outcomes of Observational Learning," *Econometrica* **68**, 371-398.



Illustrations of Social Networks with Three Agents

- (i) The unique complete network
- (ii) The Circle (iii) The Star
- (iv) Two other connected but incomplete networks