

Behavioral Aspects of Learning in Social Networks: An Experimental Study*

Syngjoo Choi[†]
New York University

Douglas Gale[‡]
New York University

Shachar Kariv[§]
UC Berkeley

October 10, 2004

Abstract

Networks are natural tools for understanding social and economic phenomena. For example, all markets are characterized by agents connected by complex, multilateral information networks, and the network structure influences economic outcomes. In an earlier study, we undertook an experimental investigation of learning in various three-person networks, each of which gives rise to its own learning patterns. In the laboratory, learning in networks is challenging and the difficulty of solving the decision problem is sometimes massive even in the case of three persons. We found that the theory can account surprisingly well for the behavior observed in the laboratory. The aim of the present paper is to investigate important and interesting questions about individual and group behavior, including comparisons across networks and information treatments. We find that in order to explain subjects' behavior, it is necessary to take into account the details of the network architecture as well as the information structure. We also identify some "black spots" where the theory does least well in interpreting the data.

Journal of Economic Literature Classification Numbers: D82, D83, C92.

*This research was supported by the Center for Experimental Social Sciences (C.E.S.S.) and the C. V. Starr Center for Applied Economics at New York University. We are grateful to John Morgan for his comments and suggestions. This paper has benefited from suggestions by the participants of SITE 2004 Summer Workshop, the New and Alternative Directions for Learning Workshop at Carnegie Mellon University and seminars at several universities. For financial support, Gale acknowledges the National Science Foundation Grant No. SBR-0095109) and the C. V. Starr Center for Applied Economics at New York University, and Kariv thanks University of California Berkeley (COR Grant).

[†]Department of Economics, New York University, 269 Mercer St., 7th Floor, New York, NY, 10003 (e-mail: sc743@nyu.edu, url: <http://homepages.nyu.edu/~sc743/>).

[‡]Department of Economics, New York University, 269 Mercer St., 7th Floor, New York, NY, 10003 (e-mail: douglas.gale@nyu.edu, url: <http://www.econ.nyu.edu/user/galed>).

[§]Department of Economics, University of California, Berkeley, Evans Hall # 3880, Berkeley, CA 94720 (e-mail: kariv@berkeley.edu, url: <http://socrates.berkeley.edu/~kariv/>).

1 Introduction

In a modern economy, an individual knows only a small fraction of the information distributed throughout the economy as a whole. Consequently, he has a very strong incentive to try to benefit from the knowledge of others before making an important decision. Sometimes he learns from public sources, books, newspapers, the Internet, etc.; at other times he needs information that is not available from these public sources and then he must try to find the information in his local environment. In social settings, where an individual can observe the choices made by other individuals, it is rational for him to assume that those individuals may have information that is not available to him and then to try to infer this information from the choices he observes. This process is called *social learning*. The literature on social learning contains numerous examples of social phenomena that can be explained in this way. In particular, it has been argued that the striking uniformity of social behavior is an implication of social learning.

Much of the social-learning literature has focused on examples of inefficient information aggregation. The seminal papers of Bikchandani, Hirshleifer and Welch (1992) (BHW) and Banerjee (1992) show that social learning can easily give rise to herd behavior and informational cascades. Herds or cascades can be started by a small number of agents who choose the same action. Subsequently, other agents ignore their own information and join the herd. Once an agent decides to join the herd, his own information is suppressed. Since only a small amount of information is revealed by the agents who started the herd, the herd is likely to have chosen a sub-optimal action. Smith and Sørensen (2000) extend the basic model to allow for richer information structures and to provide a more general and precise analysis of the convergence of actions and beliefs¹.

The models of BHW and Banerjee (1992) are special in several respects. They assume that each agent makes a once-in-a-lifetime decision and the decisions are made sequentially. Further, when each agent makes his decision, he observes the decisions of all the agents who have preceded him. An alternative model, described in Gale and Kariv (2003), assumes that agents are part of a social network and can only observe the actions of agents to whom they are connected through the network². Information percolates through the network as an agent's action is first observed by his neighbors and then (indirectly) by his neighbors' neighbors and so on. In order to model the diffusion of information through the network, it is natural to assume that agents can revise their decisions as more information becomes available. More precisely, Gale and Kariv (2003) assume that all agents make simultaneous and repeated decisions.

Whereas herd behavior arises quickly in the sequential model of BHW and

¹For excellent surveys see, Gale (1996) and Bikchandani, Hirshleifer and Welch (1998), which also provide examples and applications of observational learning in economic contexts. Among others, Lee (1993), Chamley and Gale (1994), Gul and Lundholm (1995), Moscarini, Ottaviani and Smith (1998), and Çelen and Kariv (2004a) provide further extensions of the theory.

²There is a large and growing body of work which studies the influence of the network structure on economic outcomes. For recent surveys see Goyal (2003) and Jackson (2003).

Banerjee (1992), in the social-network model learning may continue for some time as information diffuses through the network. Paradoxically, in spite of the agents' limited powers of observation, the informational efficiency of the network model may be greater than the sequential decision model.

Another difference between the network model and the sequential model is related to the complexity of decision making. Because the history of actions is not common knowledge, the agents in the network model have to make inferences not just about their neighbors' private signals, but also about their neighbors' observations and inferences about *their* neighbors. The greater complexity of the learning process raises questions about the plausibility of a rational learning model. The absence of common knowledge of the history of actions requires agents to hold beliefs about beliefs about beliefs ... about the actions and information of agents they cannot observe directly. This has led some authors, e.g., Bala and Goyal (1998), to suggest that models of bounded rationality are more appropriate for describing learning in networks.

Whether individuals can rationally process the information available in a network is ultimately an empirical question. To test the relevance of the theory, Choi, Gale and Kariv (2004) (CGK) examined the behavior of subjects in a variety of three-person networks based on the model of Gale and Kariv (2003). The information structure for the experiments was adapted from BHW and the experiment utilized the procedures of Anderson and Holt (1997)³. The family of three-person networks includes several non-trivial architectures, each of which gives rise to its own distinctive learning patterns. CGK study three of these networks: the complete network, in which each agent can observe the other two agents, and two incomplete networks, the circle, in which each agent observes one other agent, and the star, in which the agent in the center of the star is connected to the two agents at the periphery.

In the experimental design, there are two decision-relevant events equally likely to occur *ex ante* and two corresponding signals. Signals are informative in the sense that there is a probability higher than 1/2 that a signal matches the label of the realized event. We allow subjects to be of two types: informed agents who receive a private signal, and uninformed agents who know the true prior probability distribution of the states but do not receive a private signal. Each experimental round consisted of six decision-turns. At each decision turn, the subject is asked to predict which of the events has taken place, basing his forecast on a private signal and the history of his neighbors' past decisions.

CGK find that the theory can account for the behavior observed in the laboratory in most of the networks and informational treatments. In fact, the rationality of behavior is striking. The error rates calculated using deviations from the equilibrium strategies implied by the Gale-Kariv model are positive but moderate. To account for these errors, CGK adapt the model to allow for the effect of the "trembling hand" and estimate a recursive QRE model. They find that the QRE model appears to account for the large-scale features

³Anderson and Holt (1997) investigate the model of BHW experimentally. Among others, Hung and Plott (2001), Kübler and Weizsäcker (2004) and Çelen and Kariv (2004b, 2004c) analyze several aspects of sequential social learning.

of the data with surprising accuracy. Of course, there may be other ways of accounting for the data and one must consider whether the apparent success of the theory may be due to special features of the experimental design, such as the simplicity of the networks chosen or the fact that the optimal strategies are well approximated by simple heuristics.

The data generated by the CGK experiments can also be used to address a variety of important and interesting questions about individual and group behavior. In this paper, we use the same data set to investigate behavioral aspects of individual and group behavior, including comparisons across networks and information treatments. We also look more closely at the data in order to identify the “black spots” where the theory does least well in interpreting the data and ask whether additional “behavioral” explanations might be needed to account for the subjects’ behavior.

Much of the theoretical and experimental literature on social learning has focused on the phenomenon of herd behavior which is said to occur when every agent acts like others do, even though he would have chosen a different action on the basis of her own information alone. In this sense, individuals rationally ‘ignore’ their own information and ‘follow the herd’. Furthermore, since actions aggregate information poorly, despite the available information, herds need not adopt an optimal action. Therefore, the efficiency of information aggregation is one of the main concerns in the study of social learning.

We find that the experimental data exhibit a strong tendency toward herd behavior and a marked efficiency of information aggregation. The data also suggest that there are significant and interesting differences in individual and group behavior among the three networks and three information treatments. We argue that these differences can be explained by and the symmetry or asymmetry of the network or the information treatment and the resulted differences in the amount of common knowledge.

We first provide information about the evolution of herd behavior. First, diversity of private information initially leads to diversity of actions, which then gives way to uniformity as subjects learn by observing the actions of their neighbors (Result 1). Second, although convergence to a uniform action is quite rapid, frequently occurring within two to three turns, there are significant differences between the behavior of different networks (Result 2) and information treatments (Result 3). Finally, most herds tend to entail correct decisions (Result 4), which is consistent with prediction of the parametric model underlying our experimental design.

Next, we use expected payoff calculations to measure the efficiency of the decisions made by our subjects in the laboratory. We compare the levels of efficiency across networks (Result 5) and treatments (Result 6). We then provide information as to why the evolution of actual-efficiency depends on the information treatment (Result 7). We also discuss the behavioral regularities at the individual level and how they are affected by the network (Result 8) and information treatment (Result 9). Comparing individual behavior indicates that there is indeed high variation in individual behavior across subjects.

Finally, we examine how well the theory approximates the actual behavior

observed in the laboratory. We begin by computing the optimal strategies as predicted by the theoretical model and use these to compute the level of rationality (Result 10). At the first and second turns, the error rates are uniformly fairly low although there are significant differences across information sets. In the sequel, we identify some “black spots” in which there are sharp drops in rationality and discuss the departures from the predictions of the theory that ought to be considered in future work.

The rest of the paper is organized as follows. We describe the theoretical model and the experimental design in Section 2. In Section 3 we introduce three summary measures of subject behavior that are used in the sequel. The results are contained in Section 4. We group our results under three headings, relating to group behavior, efficiency, and rationality. The important features of the QRE analysis are summarized in Section 5. Some concluding remarks and important topics for further research are contained in Section 6.

2 Theory, predictions and design

In this section, we describe the theory on which the experimental design is based and the design itself. Gale and Kariv (2003) provide a more extensive description and analysis of the model and CGK provide a fuller description of the experimental design (the experimental instructions are available upon request).

2.1 The model

We restrict attention to three-person networks. Each network has three locations, indexed by $i = A, B, C$, and, at each location i , a single (representative) agent i who maximizes his short-run payoff in each period. The network is a *directed graph* represented by a family of sets $\{N_i : i = A, B, C\}$, where N_i denotes i 's neighbors, i.e., the set of agents $j \neq i$ who can be observed by agent i .

We study three networks, the *complete network*, in which each agent can observe the actions chosen by the other agents, the *star*, in which one agent, the center, can observe the actions of the other two peripheral agents, and the peripheral agents can only observe the center, and the *circle*, in which each agent can observe only one other agent and each agent is observed by one other agent. The networks are illustrated in Figure 1, where an arrow pointing from agent i to agent j indicates the $j \in N_i$.

[Figure 1 here]

There are two equally likely *states of nature* represented by two urns, a red urn (R) and a white urn (W). The red urn contains two red balls (r) and one white ball (w); the white urn contains two white balls and one red ball. One of these urns is randomly selected by nature before the start of the game. This

is the ball-and-urn social learning experiments paradigm of Anderson and Holt (1997).

Once the urn is chosen, each agent receives a private signal with probability q . In this experiment, the signal consists of seeing the color of a ball that is randomly drawn from the urn (with replacement). Signals are informative in the sense that there is a probability $2/3$ that a signal matches the label of the realized event. With probability $1 - q$ the agent does not receive a signal. An agent who receives a signal is called *informed*; otherwise he is called *uninformed*. The information structure is summarized in the diagram below.

	State	
Signal	R	W
\emptyset	$1 - q$	$1 - q$
r	$2q/3$	$q/3$
w	$q/3$	$2q/3$

An uninformed agent has a uniform prior across the two states. An informed agent has a posterior probability that depends on his signal. For example, if he sees a red ball, he believes the true state is red with probability $2/3$ and, if he sees a white ball, he believes the true state is white with probability $2/3$.

2.2 The decision problem

After the state of nature has been determined and some agents are informed, a simple guessing game is played. There are six stages or *decision turns* in the game. At each decision turn, each agent is asked to guess the true state of nature based on the information they have at that turn. The choice of agent i at date t is denoted by $x_{it} \in \{R, W\}$. An agent receives a positive payoff for guessing the correct state and zero for guessing the wrong state. The payoff is received at the end of the game, after all the decisions are made, so there is no learning from payoffs.

At the first turn, the agent's information consists of his private signal, if he has one, and the structure of the game, which is common knowledge. An informed agent will maximize his expected payoff by choosing the state he thinks is more likely. An uninformed agent regards each state as equally likely and so is indifferent between them. We assume that, whenever an agent has no signal, he chooses each action with equal probability and, when an agent is indifferent between following his own signal and following someone else's choice, he follows his own signal. One may assume different tie-breaking rules, but our experimental data supports this specification.

At the end of the first turn, after all the agents have made their decisions, agents are allowed to observe the decisions of all the agents to whom they are connected by the network. At the second turn, the agents update their probability beliefs about the true state of nature based on the information obtained at the first turn, and are again asked to make their best guess of the true state.

After all the agents have made their decisions, they observe what their neighbors have chosen. This procedure is repeated until the agents have made six decisions.

Note that we restrict attention to equilibria in which myopic behavior is optimal, that is, it is rational for agents in equilibrium to choose the actions that maximize their short-run payoffs at each date. A careful analysis shows that in our setting there is no incentive to sacrifice short-run payoffs in any period in order to influence the future play of the game because full revelation obtains if agents switch actions whenever they are indifferent between continuing to choose the same action in the next period.

2.3 The complete network

We illustrate the play of the game using the complete network, in which each agent can observe the other two, as an example. The complete network is defined by the conditions:

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

There are three different information treatments, corresponding to different values of the probability q that an individual agent receives a signal. We refer to them as *full information* ($q = 1$), *high information* ($q = 2/3$) and *low information* ($q = 1/3$).

Full-information In the case of full information ($q = 1$), all three agents are informed. At the first turn, each agent chooses the state he thinks is more likely, that is, R if he draws r and W if he draws w . Each agent's action reveals his private information and, since each agent can observe the actions of the other two, the private information becomes common knowledge at the end of the first round.

This means that, at the beginning of the second round, all the agents have the same information, they all choose the same action, and no further information is revealed. Both the actions and the beliefs of the agents will remain constant from the second turn onwards. This is a simple example of herd behavior. Although the decisions from the second turn on are based on all the information available, the herd will be incorrect with positive probability.

High-information The game changes in two ways when $q = 2/3$. First, an agent may be informed or uninformed and, secondly, the other agents do not know whether he is informed or not. Obviously, the informational value of observing another agent's action is smaller than in the full-information case.

Suppose, for example, that the pattern of signals is given by the diagram below.

		Agent/Signal		
		<i>A</i>	<i>B</i>	<i>C</i>
Period		<i>r</i>	<i>w</i>	\emptyset
1		<i>R</i>	<i>W</i>	<i>W</i>
2		<i>W</i>	<i>W</i>	<i>W</i>
3		<i>W</i>	<i>W</i>	<i>W</i>
...	

At the first decision turn, agents *A* and *B* follow their signals, while *C*, being uninformed, randomizes and ends up choosing *W*. At the second turn, *A* sees that two others have chosen *W*, but he knows they are informed with probability $2/3$ and must take this into account in updating his belief of the true state. If *B* and *C* had observed exactly one *w* signal, that would make *A* indifferent, so it becomes a question of whether they observed two *w* signals or no signals. When $q = 2/3$, two *w* signals is more likely than no signals, so *A* will switch. By similar reasoning, at the second turn *B* will not switch and *C* will be indifferent. Assuming that *C* does not switch when indifferent, we have reached an absorbing state at the second decision turn.

If *C* were to switch to *R* at the second decision turn, this would signal that he is uninformed. At that point, *A* should switch back to *R* at the third turn, thus revealing that he is informed. Then the fact that *B* continues to choose *W* at the fourth turn reveals that he is informed. At the fifth turn, everyone is indifferent, knowing that there is one *w* and one *r* signal, and they can continue to choose different actions in the remainder of the game. The lack of common knowledge (of private signals) postpones the development of herd behavior and allows information to be revealed over a longer period of time. Whether this results in better decision-making overall depends on the particular realization of the signals.

Low-information With low information, the probability that the other agents are uninformed increases. In that case, an informed agent will continue to follow his own information at the second date. Even if *A* observes the other two agents choosing *W*, the possibility that they are both uninformed is so high that he would rather ignore their actions and follow his own signal. This will reveal that *A* is informed at the second turn. At this point, *C* will imitate *A*, because of the possibility that *B* is uninformed, and switch to *R*. This reveals *C* to be uninformed.

At the third decision turn, *B* will be indifferent. If we assume that he follows his own signal when indifferent, he will be revealed to be informed and from that point onwards, everyone is indifferent between the two states and may continue to choose different actions. This pattern is given by the diagram below. Comparing the high- and low-information examples, we can see that one effect

of reducing q is to make it clearer who is informed and who is uninformed.

		Agent/Signal		
		A	B	C
Period	r	w	\emptyset	
1	R	W	W	
2	R	W	W	
3	R	W	R	
...	

2.4 The star

The first incomplete network we examine is the star, which is defined by the following conditions:

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

At each decision turn, agent A is informed about the entire history of actions taken, whereas B and C have imperfect information. The asymmetry of this network gives rise to effects not found in the other networks. The central agent A plays an important role because it is only through him that information can be transmitted between the other two.

Suppose that $q = 2/3$ and the signals are shown in the diagram below. At the first decision turn, A and B follow their signals and C randomizes and chooses R , say. At the second turn, A observes the complete history of actions but B and C only observe A 's action. B will continue to choose W because he is informed, whereas A might be uninformed, and C will continue to choose R , because he is uninformed, whereas A might be informed. A will continue to choose R because, from his point of view, B and C cancel each other out.

At the third decision turn, A knows that B is informed and C may be informed or uninformed, so A continues to choose R . B now knows that A is either informed or observed C choose R at the first turn. Eventually, if A and C continue to choose R , their information overwhelms B and B will switch to R .

		Agent/Signal		
		A	B	C
Period	r	w	\emptyset	
1	R	W	R	
2	R	W	R	
3	R	W	R	
...	

2.5 The circle

The second incomplete network we examine is the circle, in which each agent observes exactly one of the others. The circle is defined by the conditions:

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

The peculiarity of the circle is that, while the equilibrium strategies are very simple, the analysis is quite subtle because of the lack of common knowledge. Assuming once again that $q = 2/3$ and the signals are given in the diagram below, we can trace out one possible evolution of play.

		Agent/Signal		
		<i>A</i>	<i>B</i>	<i>C</i>
Period	<i>r</i>	<i>w</i>	\emptyset	
1	<i>R</i>	<i>W</i>	<i>W</i>	
2	<i>R</i>	<i>W</i>	<i>R</i>	
3	<i>R</i>	<i>W</i>	<i>R</i>	
4	<i>W</i>	<i>W</i>	<i>R</i>	
...	

At the first turn, the informed agents follow their signals and the uninformed agent randomizes (here we assume he chooses *W*). At the second turn, *C* switches to *R* because he is uninformed and observes *A* who might be informed. Imitating the behavior of the neighboring agent is always the optimal strategy for an uninformed agent. Conversely, it is always optimal for an informed agent to follow his signal, though this is far from obvious. Agent *B*, seeing *C* switch to *R*, will know that *C* is uninformed and has observed *A* choose *R*. If *C* continues to choose *R* this will tell *B* that *A* has continued to choose *R*, which means that *A* is informed.

At this point, *B* is indifferent between *R* and *W* and we will assume he continues to follow his signal. All that *A* knows is that *B* continues to choose *W*. He cannot infer what *C* is doing, so he simply has to rely on Bayes rule to take account of all the possibilities. A lengthy calculation would show that it is optimal for *A* to continue choosing *R*, but that eventually he will become indifferent. In the limit, everyone will be indifferent except for *C*.

2.6 Summary

In summary, the above examples have illustrated several features of the theory that can be tested in the laboratory. First, initially, diversity of private information leads to diversity of actions. But, as agents learn by observing the actions of their neighbors diversity is over time replaced by uniformity. Second, convergence to a uniform action is quite rapid, typically occurring within two to three periods. Thus, what happens in those first few periods is crucial for the determination of the social behavior. Third, significant differences can be identified in the behavior of different networks. In particular, in the complete network learning stops almost immediately, while in the incomplete networks learning can continue for a longer time. Finally, despite the fact that agents suppress their own information and follow a herd, a careful analysis shows that in all treatments, except with very small probability in the complete network under high-information, herds always adopt an action that is optimal relative to the total information available to agents.

The theory clearly suggests that even in the three-person case the process of social learning in networks can be complicated, particularly in the incomplete networks. That is why we believe that insights obtained from an experiment may provide better understanding of social learning in networks.

2.7 Experimental design

We studied three different network structures (the complete network, the star and the circle) and three different information treatments ($q = 1, 2/3, 1/3$). The network structure and the information treatment were held constant throughout a given experimental session. In each session, the network positions were labeled A , B , or C . A third of the subjects were designated type- A participants, one third type- B participants and one third type- C participants. The participant's type, A , B , or C , remained constant throughout the session.

Each session consisted of 15 independent rounds and each round consisted of six decision-turns. The following process was repeated in all 15 rounds. Each round started with the computer randomly forming three-person networks by selecting one participant of type A , one of type B and one of type C . The networks formed in each round depended solely upon chance and were independent of the networks formed in any of the other rounds. The computer also chose one of two equally probable urns, labeled R and W , for each network and each round. The urn remained constant throughout the round. The choice of urn was independent across networks and across rounds.

When the first round ended, the computer informed subjects which urn had actually been chosen. Then, the second round started by having the computer randomly form new groups of participants in networks and select an urn for each group. This process was repeated until all the 15 rounds were completed. Earnings at each round were determined as follows: at the end of the round, the computer randomly selected one of the six decision-turns. Everyone whose choice in this decision-turn matched the letter of the urn that was actually used earned \$2. All others earned nothing. This procedure insured that at each decision-turn subjects would make their best guess as to what urn had been chosen.

The data was generated by experiments run during the Summer and Fall of 2003 at the Center for Experimental Social Science (C.E.S.S.) at New York University. In total, we have observations from 156 subjects who had no previous experience in networks or social learning experiments. Subjects were recruited from undergraduate classes at New York University and each subject participated in only one of the nine experimental sessions. The diagram below summarizes the experimental design (the entries have the form a / b where a is the number of subjects and b the number of observations per type and turn).

Network	Information		
	Full	High	Low
Complete	18 / 90	15 / 75	18 / 90
Star	18 / 90	18 / 90	18 / 90
Circle	18 / 90	18 / 90	15 / 75

3 Three measures

For a better understanding of the decision mechanism of our subjects we organize the data according to three measures: *stability*, *uniformity* and *efficiency*. Next, we explain the three measures and their motivations.

Much of the theoretical and experimental literature on social learning has focused on the related phenomena of informational herd behavior and informational cascades. Herd behavior is said to occur when, after some point, all agents choose the same action. A herd may arise even if agents would have chosen a different action on the basis of their own information alone. In this sense, agents rationally ‘ignore’ their own information and ‘follow the herd’. We characterize herd behavior by two related phenomena, stability and uniformity of actions.

Stability *At each turn t , stability is measured by the proportion of subjects who continue to choose the action they chose at turn $t - 1$. For each network a stability variable is denoted by S_t and defined by*

$$S_t = \frac{\#\{i : x_{it} = x_{it-1}\}}{n}.$$

We report averages of S_t across different networks.

Uniformity *At each turn t , uniformity is measured by a score function that takes the value 1 if all subjects act alike and takes the value 0 otherwise. For each network a uniformity variable is denoted by U_t and defined by*

$$U_t = \begin{cases} 1 & \text{if } x_{it} = x_{jt}, \forall i, j, \\ 0 & \text{otherwise.} \end{cases}$$

We report averages of U_t across different networks.

Herd behavior arises in the laboratory when, from some decision-turn on, all subjects take the same action. Notice that uniformity of actions at some date t will persist and lead to herd behavior if and only if stability takes the value 1 at all subsequent stages or decision-turns.

As the examples in the preceding section have illustrated, the theory predicts that convergence to a uniform action typically occurs within two to three periods. Furthermore, except with very small probability in the complete network under high-information, herds always adopt the optimal action. As a benchmark for our empirical analysis of stability and uniformity, we first calculated

the values of these measures predicted by the theory. The theoretical predictions are derived with the help of simulations, which are summarized in Table 1 and show, turn by turn, the average level of stability and uniformity and the percentage of herd behavior. The numbers in parentheses are the fractions of herds that choose the wrong action, defined relative to the information available.

[Table 1 here]

Informational efficiency of markets is a natural question for economists and the efficiency of information aggregation is one of the main concerns in the study of social learning. A central result of the literature is that herd behavior may result in most agents choosing the wrong action (where the right action is defined relative to the information available in the economy). This outcome is both informationally inefficient and Pareto inefficient. This failure of information aggregation is explained by two facts. First, an agent’s action is a coarse signal of his private information and, secondly, after some point, agents suppress their private information and join the herd, so that only a small fraction of the private information in the game may ever be revealed.

Like Anderson and Holt (1997), we use expected payoff calculations to measure the efficiency of the decisions made by our subjects in the laboratory. As a benchmark we use the payoff to a hypothetical agent who has access to all private signals in his network. Define the *efficient expected payoff* to be the expected earnings of an agent who makes his decision based on the entire vector of private signals; define the *private-information expected payoff* to be the expected earnings of an informed agent who makes his decision on the basis of his own private signal; and define the *random expected payoff* to be the expected earnings of an agent who randomizes uniformly between the two actions. Finally, for each turn, let the *actual expected payoff* be the expected earnings from the subject’s actual decision in the laboratory.

The *sum* (over agents) of the efficient, private-information, random, and actual payoffs, for all rounds, will be denoted by π_e , π_p , π_r and π_a , respectively. We use these payoff calculations to assess the quality of aggregation and use of information within a network.

Efficiency *The efficiency of decisions is measured in two ways:*

$$\text{actual efficiency} = \frac{\pi_a - \pi_r}{\pi_e - \pi_r},$$

which is calculated turn by turn, and

$$\text{private-information efficiency} = \frac{\pi_p - \pi_r}{\pi_e - \pi_r}.$$

There are two normalizations in our measure of actual efficiency. First, since even uninformed random choices will be correct half the time, we subtract random efficiency π_r from actual efficiency π_a in order to get a more accurate

measure of the benefit the subjects get from the information they use. Second, there is more information available in some treatments than in others, so we express the net actual efficiency $\pi_a - \pi_r$ as a fraction of the net efficiency $\pi_e - \pi_r$ that could be achieved if subjects pooled their information. Hence, efficient decisions have an efficiency of one and random decisions have an efficiency of zero. A similar rationale applies to the measure of private-information efficiency. The comparison of actual- and private-information efficiencies is useful in determining the extent to which subjects use the information revealed by their neighbors' actions, i.e., the extent to which they did worse than choosing according to all the information and the extent to which they did better than choosing only according to their private information.

Note that the prediction of the theory is that, except in the circle network under full-information in which agents can rationally choose different actions forever, complete learning occurs quite rapidly with the result that an efficient action, in the sense that the same action would be chosen if all the signals were public information, will be chosen. Thus, the theory predicts a marked efficiency of information aggregation. Table 2 summarizes the theoretical predictions, which are derived with the help of simulations, in all networks and information structures.

[Table 2 here]

4 Results

4.1 Group behavior

We characterize herd behavior in terms of stability and uniformity of actions. Our first result provides information about the evolution of herd behavior.

Result 1 *There is an upward trend in the degree of uniformity and a high and constant level of stability in all treatments, with the result that, over time, subjects tend to follow a herd more frequently.*

Support for Result 1 is presented in Table 3 which shows, turn by turn, the average level of stability and uniformity and the percent of rounds in which subjects followed a herd from that turn on. For comparison purposes, the experimental results presented in Table 3 are given in the same format as the theoretical predictions presented in Table 1. By definition, the number of herds is monotonically non-decreasing over time, but the increase in stability and uniformity is not implied by the definitions. It appears to be the result of learning and information aggregation.

[Table 3 here]

Result 1 confirms that over time subjects are increasingly persuaded by the observed actions and gradually build confidence in the information revealed by their neighbors' actions. In the incomplete networks, some of the information

of unobserved subjects is accumulated in the observed actions, so the fact that subjects tend to follow a herd more frequently indicates that they try to extract information of unobserved subjects from the actions they observe.

This is consistent with the prediction of the theory that over time more and more information is revealed. The theory, however, also predicts that once agents have chosen the same action and they are not indifferent between the two actions, they have reached an absorbing state and will continue to choose the same action in every subsequent period. In the laboratory, in contrast to this prediction, we sometimes observe deviations from a herd.

Next, we turn to the frequencies of herd behavior in different networks and treatments. Our next results report that, within a given decision-turn, in some treatments there is no significant difference between the frequencies of herd behavior, but the situation clearly reverses, particularly in early turns, in other treatments.

Result 2 (networks) *In the complete and star networks, the frequency of herds is highest under full-information and lowest under low-information; in the circle, the frequency of herds under low-information is the same as under high-information but lower than under full-information.*

Result 3 (information) *Under full-information, the frequency of herds in the complete network is the same as in the star but higher than in the circle; under high-information the frequency of herds in the circle network is the same as in the star but lower than in the complete network; under low-information, there are no significant differences between the frequencies of herd behavior in the different networks.*

Note that the behavior summarized in these two results is consistent with the theoretical results described in Section 2 and Section 3. For example, in the complete network under full information, a herd must start at the second decision turn for every realization of the private signals. By contrast, under the low-information treatment, subjects do not know who is informed and who is uninformed at the second turn, so learning continues after the second turn and herd behavior is delayed. Thus, herd behavior is more likely and will begin sooner in the full-information treatment. In the circle network, it is optimal for informed subjects to follow their own signals, regardless of the behavior they observe in others, and for uninformed subjects to imitate the behavior they observe. For this reason, we expect herd behavior to develop only if the informed subjects receive identical signals. A herd will develop sooner in the full-information treatment than in the high- and low-information treatments, because it takes time for the uninformed subjects to get on board.

The first evidence about the frequencies of herd behavior is provided in Table 3. The relevant support for Result 2 and Result 3 comes from Figure 2 which presents, in graphical form, the data from Table 3 on herd behavior in each network under all information treatments (left panel), and for all networks under each information treatment (right panel). A set of binary Wilcoxon tests indicates that the differences are highly significant.

[Figure 2 here]

The right column of Table 3 summarizes, treatment by treatment, the average level of stability and uniformity over all turns and the expected length of herd behavior. Note that in all networks the expected length of herd behavior is increasing in the probability that an individual subject receives a signal. Also, under full- and high-information, herd behavior in the complete network is longer than in the star, and in the star is longer than in the circle. Under low-information, there are longer herds in the circle than in the star.

Hence, convergence to a uniform action is more rapid in the complete network and under full-information because it is common knowledge that all subjects are informed and all actions are common knowledge. In contrast, diversity can continue for a longer time under high- and low-information and in the star and circle networks. Clearly, the absence of common knowledge makes it harder for subjects to interpret the information contained in the actions of others and requires them to perform complex calculations.

Herd behavior has elicited particular interest because erroneous outcomes may occur despite individual rationality, and they may in fact be the norm in certain circumstances. In the model underlying our experimental design, we note that, except with very small probability in the complete network under high-information, herds always adopt an action that is optimal relative to the total information available to agents. Thus, it is particularly interesting that, in the laboratory, almost all herds longer than three decision turns selected the right action, but some differences can be identified in the behavior of different networks. We can report the following result.

Result 4 *Relative to the information available, herds entail correct decisions. There are, however, significantly more incorrect herds in the complete network under high-information.*

Evidence for Result 4 is also provided by Table 3. The numbers in parentheses are the fractions of herds that choose the wrong action, defined relative to the information available. It is noteworthy that herds entail correct decisions even in the star and circle networks in which subjects had imperfect information about the history of decisions.

In summary, we observe several empirical regularities in the experimental data. First, diversity of private information initially leads to diversity of actions, which then gives way to uniformity as subjects learn by observing the actions of their neighbors. Second, convergence to a uniform action can be quite rapid, frequently occurring within two to three periods. Third, herd behavior develops frequently and most herds turned out to be correct.

4.2 Efficiency

We next turn our attention to analyze how efficient were our subjects in using the information revealed by their neighbors' actions. The next results report

average actual-efficiency calculations to measure the informational efficiency within a given network, information treatment, and turn.

Result 5 (network) *In the complete network, average actual-efficiency is highest under full-information and lowest under low-information; in the star, average actual-efficiency under full-information is the same as under high-information but higher than under low-information; in the circle, the levels of average actual-efficiency are the same under all information treatments.*

Result 6 (information) *Under full-information, average actual-efficiency is highest in the complete network and lowest in the circle; under high- and low-information, there are no significant differences between the levels of average actual-efficiency in the different networks.*

Table 4 which summarizes, turn by turn, the actual and private-information efficiencies in all networks and information treatments, provides a first indication. Under high and low information, Table 4 also provides actual efficiency for informed and uninformed individual subjects' decisions. For comparison purposes, the experimental results presented in Table 4 are given in the same format of the theoretical predictions presented in Table 2.

The support for Result 5 and Result 6 comes from Figure 3, which presents the data from Table 4 by comparing the total actual efficiency in each network under all information treatments (left panel), and for all networks under each information treatment (right panel). Figure 3 also depicts the average private-information efficiency over all subjects within each information treatment. A set of binary Wilcoxon tests indicates that all the differences above are highly significant.

[Table 4 here]

Result 5 and Result 6 emphasize the role of common knowledge in the laboratory. Under full information, it is common knowledge that all subjects are informed. In the complete network, the subjects' actions are also common knowledge, whereas in incomplete networks the actions are not common knowledge. Under full information, the greater efficiency in the complete network compared to incomplete networks can be attributed to the agents' ability to use the greater amount of information available to them. In the other information treatments, subjects are uncertain whether the other participants are informed or uninformed, and this uncertainty appears to prevent them from making use of the additional information available in the complete network.

[Figure 3 here]

Our previous results relate to the levels of actual efficiency across treatments but it appears that efficiency increases over time only in the complete and star networks under low information. Our next result provides information as to why the evolution of actual-efficiency depends on the information treatment, i.e., the

probability that an individual subject receives a signal. We use the same payoff calculations but sum π_e , π_p , π_r and π_a over informed and uninformed decision points and not over individual subjects.

Result 7 *Whereas the actual efficiency in informed decision points falls slightly from the first to the last decision turn, the actual efficiency in uninformed decision points increases over time. The rate of increase of actual efficiency is greater under low information than under high and full information.*

Figure 4 provides the support for Result 7 by comparing the efficiencies of informed and uninformed decisions in each of the networks under low- and high-information. Since there are more uninformed subjects in the low-information treatment and actual efficiency is increasing for uninformed subjects and decreasing for informed subjects, the rate of increase will be highest under low information. Figure 4 also depicts the average private-information efficiency over all informed individual subjects within each information treatment, and reveals that actual efficiency of informed decisions is higher than or roughly the same as private-information efficiency. Note also that Result 5 matches the prediction of the theory that, in many cases, efficiency is not expected to increase in late turns as beliefs reach an absorbing state in which agents either choose the same action or are indifferent between the two actions and no further learning occurs.

[Figure 4 here]

Our results so far deal only with average efficiency. We are also interested, however, in the behavioral regularities at the individual level and how they are affected by the network and information treatment. For a better understanding of the decision mechanism of the subjects, we next focus on the data at the level of the individual subject. We find strong evidence that there is a good degree of conformity with the theory in the aggregate data, which we sometimes fail to observe in the individual data.

Next, we look at actual efficiency and ask whether individual behavior is heterogeneous within a given network and information treatment and across networks and information treatments. To focus on the individual effects, Result 8 deals only with the complete and circle networks under full-information since all types have identical sets of neighbors and subjects are equally informed. Result 9 summarizes the behavioral regularities in this regard across networks and information treatments.

Result 8 (network) *Comparing individual behavior in the complete and circle networks under full information indicates that there is high variation in individual behavior across subjects.*

Result 9 (information) *There are significant differences between the distributions of actual efficiency across information treatments in the complete*

network and across network structures under full information. The distributions of actual efficiency are roughly the same in the other treatments.

Figure 5 provides the support for Result 8 and Result 9. It presents the actual-efficiency distributions in each network under all information treatments (left panel), and for all networks under each information treatment (right panel) in the form of histograms. The horizontal axis consists of the intervals of efficiency scores and the vertical axis measures the percentages of subjects corresponding to these efficiency scores.

Note, for example, that under full-information the histograms in Figure 5 show that subject behavior is more efficient in the complete network as the distribution of efficiency scores shifts considerably to the right when calculated using the complete network data. A Kolmogorov-Smirnov test confirms this observation at the 5 percent significance level.

[Figure 5 here]

4.3 Rationality

In the laboratory, learning in networks is challenging: because of the lack of common knowledge about the history of play subjects must draw inferences about the actions other subjects have observed as well as about their private signals. Moreover, not all of the decisions in the different networks are comparable in terms of their complexity or sophistication. For example, subjects have larger information sets in some networks than in others and sometimes there is more information to be gleaned from the actions of others. Also, differences in the amount of common knowledge will require different degrees of sophistication to discern the optimal strategies in different networks.

Hence, even in the three-person case, the difficulty of solving the problem of social learning in networks is sometimes massive. Therefore, one of our main goals is to examine, treatment by treatment, how well the theory approximates the actual behavior observed in the laboratory. We begin by computing the optimal strategies as predicted by the theoretical model and use these to compute the level of rationality in the first and second decision turns. Thus, rationality is simply measured by the percentage of times subjects follow an equilibrium strategy.

At the first decision-turn in any treatment, a subject should make a decision based on his private signal (if he is informed) or his prior (if his uninformed). At the second decision-turn, he should make a decision based on his private information and the actions taken at the first turn, and so on. Thus, the complexity of a subject's decision-problem increases over time. This leads us to examine how well Bayes rationality approximates the actual behavior observed in the laboratory. At the first and second decision-turns, the data supports the following result.

Result 10 *Over all treatments, only 5.8 percent of the first-turn actions in each round were inconsistent with the information implicit in the private*

signal. At the second turn, although there are significant differences across information sets, the error rates are uniformly fairly low.

Evidence for Result 10 is given by the table below, which reports the error rates, i.e., the percentage of times subjects deviate from the equilibrium strategy, at the second turn. The data is grouped according to the number of actions observed, i.e., all types in the complete network and type *A* (the center) in the star observe $N = 2$ actions, and all types in the circle network and types *B* and *C* in the star observe $N = 1$ actions. The numbers in parentheses are the percent of decisions in which subjects were indifferent between the two actions.

Information	$N = 2$	$N = 1$
Full	12.5 (0.00)	4.40 (44.4)
High	17.8 (16.2)	16.7 (0.00)
Low	20.3 (36.9)	26.9 (0.00)

The histograms in Figure 6 show this data across subjects. The horizontal axis measures the error-rates for different intervals and the vertical axis measures the percentage of subjects corresponding to each interval. Note that, in both the first and second decision-turns, the distribution is considerably skewed to the left, but the two distributions are significantly different at the 5 percent significance level using a Kolmogorov-Smirnov test.

[Figure 6 here]

To provide further evidence for Result 10, Table 5 summarizes the error rates in the second decision-turn in each treatment, including for informed and uninformed decisions. Table 5 clearly identifies some “black spots” in which there are sharp drops in rationality, especially in the star network under low- and high-information. Therefore, subjects must presumably estimate the errors of others and consider this in processing the information revealed by their neighbors’ actions.

[Table 5 here]

It is noteworthy that, under low-information, the error rates in uninformed peripherals’ decisions and in uninformed decisions in the circle network were very high and roughly the same (33.9 percent and 34.9 percent respectively). Error rates were very high in informed centers’ decisions compare to informed decisions in the complete network (22.2 percent and 13.8 percent respectively). Under high-information, the error rates in uninformed peripherals’ decisions were much lower than in uninformed decisions in the circle network (26.7 percent and 16.3 percent respectively). These differences are highly significant according to a Wilcoxon matched-pair test.

Note that the complexity of a subject’s decision problem increases over time. At the first turn, a subject only has to interpret his private information. At the second turn, he has to interpret his neighbors’ actions and try to infer the

private information on which it was based. At the third turn, because of the lack of common knowledge about actions in incomplete networks, a subject is forced to think about subjects' knowledge of other subjects' actions and the private information it reveals. Thus, in the laboratory, mistakes are inevitable and this should be taken into account by evaluating the degree to which the game-theoretic model explains behavior in the laboratory. In the sequel, we discuss a modification of the game-theory model that abandons the assumption of common knowledge of rationality.

5 Quantal Response Equilibrium (QRE)

Since mistakes are made, especially under low- and high-information, and this should be taken into account in any theory of rational behavior, in CGK, we attempt to formulate this by estimating a recursive model that allows for the possibility of errors in earlier decisions. We adapt the model of Quantal Response Equilibrium (QRE) of McKelvey and Palfrey (1995, 1998) which enables us to evaluate the degree to which the theory explains behavior in the laboratory. We skip the model development and instead briefly explain the analysis and discuss why the results of the QRE model differ from those of the basic game-theoretic model.

We first extend the basic model of Gale and Kariv (2003) to allow for idiosyncratic preference shocks, which can be interpreted, following Harsanyi and Selten, as the effect of a 'trembling hand'. The basic model has a natural recursive structure, which suggests a recursive estimation procedure for the logistic random-utility model. In effect, we are assuming subjects have rational expectations and use the true mean error rate when interpreting the actions they observe at the first turn. This is the behavioral interpretation of the recursive econometric method.

We begin by estimating the random-utility model using the data from the first turn. Then we use the estimated coefficient to calculate the theoretical payoffs from the actions at the second turn. We then estimate the random-utility model based on the perturbed payoffs and the observed decisions at the second turn. Continuing in this way, we estimate the entire QRE for each treatment. The parameter estimates are highly significant and positive, showing that the theory does help predict the subjects' behavior.

The predictions of the QRE model are different from those of the basic game-theoretic model for two reasons: first, because it allows agents to make mistakes and, secondly, because it assumes that agents take into account the possibility that others are making mistakes when drawing inferences from their actions. The 'goodness of fit,' as measured by the error rates, is better for the QRE model than for the game-theory model. We also conducted a series of specification tests and find that restrictions of the QRE model are confirmed by the data.

Among our conclusions in this paper, the following are particularly relevant to the evaluation of the QRE model.

- The model appears to fit the data best in the full information treatments, where every subject receives a private signal. In other information treatments, subjects are randomly informed and the subjects do not know whether another subject is informed or not. This asymmetry appears to reduce the efficiency and rationality of the subjects' behavior.
- The model appears to fit the data best in symmetric networks, where each subject observes the same number of other subjects. In asymmetric networks, where one subject is known to have an informational advantage over the others, the behavior of subjects is less close to the predictions of the model and the discrepancy is highest for the subjects with the informational advantage.

In evaluating these departures from the predictions of the theory, we have become aware of several factors that ought to be considered in future work.

- In order to explain subjects' behavior, it is necessary to take into account the details of the network architecture as well as the information treatment. Simple summary characteristics of the network, such as the average distance between subjects, do not account for the subtle and complicated behaviors that we observe.
- Because of the lack of common knowledge in the networks, the decision problems faced by subjects require quite sophisticated reasoning. At the same time, the optimal strategy is simple in some networks, so it is plausible that subjects following a simple heuristic might behave "as if" they were following the optimal strategy, even though it would be quite implausible to expect them to be able to "solve" for the best response.
- In other networks, by contrast, the optimal behavior is sometimes extremely complex, even though the networks themselves are quite simple. This may explain the failure of the QRE to fit some of the data generated by asymmetric networks, for example.

To determine which of these factors are important in explaining subject behavior in a variety of settings, it will be necessary to investigate a larger class of networks in the laboratory. This is perhaps one of the most important topics for future research.

6 Conclusion

We have undertaken an experimental investigation of learning in three-person networks and focus on using the theoretical framework of Gale and Kariv (2003) to interpret the data generated by the experiments. In our experimental design we used three networks, the complete network and two incomplete networks, the circle and the star, each of which theoretically gives rise to its own distinctive learning patterns.

The theory suggests that even in the three-person case the process of social learning in networks can be complicated. In particular, in the incomplete networks the absence of common knowledge makes it harder for agents to interpret the information contained in the actions of others and requires them to perform complex calculations.

Indeed, in the laboratory, we have seen that removing in three-person networks has a significant effect on social behavior, even if removing links does not have much effect on the degree of separation, i.e., the social distance within the network. The reason is the impact of lack of common knowledge on the dynamics of social learning and the efficiency of aggregation.

The presence of common knowledge makes it easier for subjects to agree on the interpretation of the information contained in the actions of any set of subjects. Lack of common knowledge forces a subject to think about hierarchies of beliefs: for example, subject *A*'s beliefs about subject *B*'s beliefs about subject *C*'s action and the private information it reveals. When links are removed, actions are not common knowledge. This uncertainty appears to prevent them from making use of the additional information available to them from others' actions.

We have identified some situations where the theory does less well in accounting for subjects' behavior. We conjecture that the theory fails in those situations because the complexity of the decision problem exceeds the bounded rationality of the subjects. There is convincing theoretical and empirical evidence that some of the decisions faced by subjects are quite "complex"; however, "complexity" is a very difficult idea to conceptualize in a precise or formal way. Also, because of the simplicity of the experimental design, it is difficult to distinguish among different concepts of complexity, and there is a danger that any attempt to measure complexity will be over-adapted to this particular application. Progress in this area may require both new theory and new experimental data.

A theory of complexity has to take into account the following points. First, subjects' success or failure in the experiment results from the appropriateness of the heuristics they use as much as the inherent difficulty of the decision-making. Second, the optimal strategy may be intuitively simple even though the analysis is complex, and a simple heuristic may work very well, even though the game or the analysis of the game is complex. Third, complexity is endogenous: in a complex extensive-form game subjects may be forced to adopt simple strategies; and the decision-making problem for each subject may be simplified as a result.

Another idea of complexity has to do with the equilibrium path: in the star network, for example, the equilibrium path is more complex than in the complete network. Hence, an interesting question is whether we can identify a sufficient statistic for the difficulty or complexity of decisions that will allow us to interpret variation in efficiency and rationality measures. If we can identify more information sets where there are sharp drops in efficiency or rationality ("black spots"), we may be able to come up with some hypotheses about why the decision is complex or difficult and suggest experiments to test this hypothesis.

Our results suggest that the theory adequately accounts for large-scale fea-

tures of the data. The models, and results that we have developed provide a foundation for future theoretical and experimental research and the techniques can be applied to other setups. For example, we can apply our theoretical model to random graphs, as long as connectedness is satisfied, and it could also be applied to dynamic graphs where the set of neighbors observed changes over time.

There are many more important questions that remain to be explored using our data set. Perhaps, the most important subject for future research is to identify the impact of network architecture on the efficiency and dynamics of social learning. Whether other network architectures will lead to sharply different results is not clear, since all the decision rules will have to be changed to reflect the new environment. It will probably require a more sophisticated analysis to detect the effect of differences in network architectures.

Also, we have not yet explored alternative approaches that might be brought to the interpretation of the data. Clearly, there is much to be done and the uses of this data set are far from exhausted.

References

- [1] Anderson, L. and C. Holt (1997) "Information Cascades in the Laboratory." *American Economic Review*, 87(5), pp. 847-62.
- [2] Bala, V. and S. Goyal (1998) "Learning from Neighbors," *Review of Economic Studies* 65, pp. 595-621.
- [3] Banerjee, A. (1992) "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107(3), pp. 797-817.
- [4] Bikhchandani, S., D. Hirshleifer and I. Welch (1992) "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascade." *Journal of Political Economy*, 100(5), pp. 992-1026.
- [5] Bikhchandani, S., D. Hirshleifer and I. Welch (1998) "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades." *Journal of Economic Perspective*, 12(3), pp. 151-70.
- [6] Çelen, B. and S. Kariv (2004a) "Observational Learning Under Imperfect Information," *Games and Economic Behavior*, 47(1), pp. 72-86.
- [7] Çelen, B. and S. Kariv (2004b) "Distinguishing Informational Cascades from Herd Behavior in the Laboratory," *American Economic Review*, 94(3), pp. 484-497.
- [8] Çelen, B. and S. Kariv (2004c) "An Experimental Test of Observational Learning under Imperfect Information," *Economic Theory*, forthcoming.
- [9] Chamley, C. and D. Gale (1994) "Information Revelation and Strategic Delay in a Model of Investment." *Econometrica*, 62(5), pp. 1065-85.

- [10] Gale, D. (1996) "What Have We Learned from Social Learning?" *European Economic Review*, 40(3-5), pp. 617-28.
- [11] Gale, D. and S. Kariv (2003) "Bayesian Learning in Social Networks." *Games and Economic Behavior*, 45(2), pp. 329-46.
- [12] Goyal S. (2003) "Learning in Networks: A Survey," Forthcoming in, *Group Formation in Economics: Networks, Clubs, and Coalitions* edited by G. Demange and M. Wooders. Cambridge University Press.
- [13] Gul, F. and R. Lundholm (1995) "Endogenous Timing and the Clustering of Agents' Decisions," *Journal of Political Economy*, 103(5), pp. 1039-66.
- [14] Hung, A. and C. Plott (2001) "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions." *American Economic Review*, 91(5), pp. 1508-20.
- [15] Jackson, M. (2003), "A Survey of Models of Network Formation: Stability and Efficiency," Forthcoming in, *Group Formation in Economics: Networks, Clubs, and Coalitions* edited by G. Demange and M. Wooders. Cambridge University Press.
- [16] Kübler, D. and G. Weizsäcker (2003) "Limited depth of reasoning and failure of cascade formation in the laboratory." *Review of Economic Studies*, 71(2), pp. 425-41.
- [17] Lee, I. H. (1993) "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 61(2), pp. 396-411.
- [18] McKelvey, R.D. and T.R. Palfrey (1995), "Quantal Response Equilibria for Extensive Form Games." *Games and Economic Behavior*, 10, pp. 6-38.
- [19] McKelvey, R.D. and T.R. Palfrey (1998), "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics*, 1, pp. 9-41.
- [20] Moscarini, G., M. Ottaviani and L. Smith (1998) "Social Learning in a Changing World," *Economic Theory*, 11, pp. 657-65.
- [21] Smith, L. and P. Sørensen (2000) "Pathological Outcomes of Observational Learning." *Econometrica*, 68(2), pp. 371-98.

Table 1. Theoretical results: uniformity, stability and herd behavior, turn by turn, under the different information structures and networks
Average level of uniformity and stability and the percent of rounds in which subjects followed a herd from that turn on.

Decision-turn								
Uniformity		1	2	3	4	5	6	Average
Full Information	Complete	0.33	1.00	1.00	1.00	1.00	1.00	0.89
	Star	0.33	0.56	1.00	1.00	1.00	1.00	0.81
	Circle	0.33	0.33	0.33	0.33	0.33	0.33	0.33
High Information	Complete	0.29	0.78	0.86	0.89	0.89	0.89	0.76
	Star	0.29	0.54	0.80	0.86	0.86	0.85	0.70
	Circle	0.29	0.47	0.58	0.58	0.58	0.58	0.51
Low Information	Complete	0.26	0.49	0.77	0.78	0.77	0.76	0.64
	Star	0.26	0.54	0.68	0.88	0.86	0.86	0.68
	Circle	0.26	0.43	0.65	0.71	0.71	0.71	0.58

Stability		1	2	3	4	5	6	Average
Full Information	Complete		0.78	1.00	1.00	1.00	1.00	0.96
	Star		0.93	0.85	1.00	1.00	1.00	0.96
	Circle		1.00	1.00	1.00	1.00	1.00	1.00
High Information	Complete		0.68	0.89	0.97	0.97	0.97	0.90
	Star		0.78	0.85	0.95	0.98	0.98	0.91
	Circle		0.83	0.94	0.99	0.99	0.99	0.95
Low Information	Complete		0.67	0.82	0.90	0.91	0.91	0.84
	Star		0.67	0.82	0.88	0.93	0.92	0.84
	Circle		0.67	0.78	0.93	0.91	0.90	0.83

Herds		1	2	3	4	5	6	Length
Full Information	Complete	33.33 (0.00)	100.00 (0.00)	100.00 (0.00)	100.00 (0.00)	100.00 (0.00)	100.00 (0.00)	5.3
	Star	33.33 (0.00)	55.56 (0.00)	100.00 (0.00)	100.00 (0.00)	100.00 (0.00)	100.00 (0.00)	4.9
	Circle	33.33 (0.00)	33.33 (0.00)	33.33 (0.00)	33.33 (0.00)	33.33 (0.00)	33.33 (0.00)	2.0
High Information	Complete	28.77 (0.00)	78.35 (1.71)	85.40 (1.57)	88.14 (1.52)	88.19 (1.52)	88.57 (1.51)	4.6
	Star	28.75 (0.00)	54.21 (0.00)	79.55 (0.00)	85.12 (0.00)	85.15 (0.00)	85.45 (0.00)	4.2
	Circle	28.70 (0.00)	46.60 (0.00)	57.72 (0.00)	57.77 (0.00)	57.86 (0.00)	58.43 (0.00)	3.1
Low Information	Complete	25.80 (0.00)	48.68 (0.00)	71.69 (0.00)	74.18 (0.00)	74.37 (0.00)	76.21 (0.00)	3.7
	Star	25.99 (0.00)	54.33 (0.00)	66.77 (0.00)	82.95 (0.00)	83.23 (0.00)	85.58 (0.00)	4.0
	Circle	25.93 (0.00)	43.21 (0.00)	65.43 (0.00)	65.76 (0.00)	66.49 (0.00)	71.01 (0.00)	3.4

Table 3. Experimental results: stability and uniformity, turn by turn, under the different information structures and networks
 (average level of uniformity and stability and the percent of rounds in which subjects followed a herd from that turn on)

Decision-turn								
Uniformity		1	2	3	4	5	6	Average
Full Information	Complete	0.38	0.64	0.74	0.71	0.72	0.71	0.65
	Star	0.32	0.43	0.62	0.64	0.67	0.71	0.57
	Circle	0.31	0.44	0.54	0.53	0.58	0.68	0.51
High Information	Complete	0.32	0.55	0.64	0.67	0.75	0.69	0.60
	Star	0.33	0.40	0.61	0.59	0.62	0.52	0.51
	Circle	0.31	0.36	0.52	0.49	0.53	0.49	0.45
Low Information	Complete	0.23	0.36	0.46	0.44	0.51	0.63	0.44
	Star	0.20	0.36	0.46	0.38	0.40	0.51	0.38
	Circle	0.25	0.44	0.45	0.51	0.45	0.48	0.43
Stability		1	2	3	4	5	6	Average
Full Information	Complete		0.87	0.93	0.94	0.95	0.94	0.92
	Star		0.81	0.83	0.89	0.90	0.96	0.88
	Circle		0.75	0.78	0.77	0.80	0.83	0.79
High Information	Complete		0.76	0.84	0.84	0.87	0.86	0.84
	Star		0.71	0.76	0.83	0.83	0.79	0.78
	Circle		0.74	0.74	0.76	0.78	0.77	0.76
Low Information	Complete		0.67	0.71	0.72	0.75	0.79	0.73
	Star		0.71	0.73	0.70	0.74	0.77	0.73
	Circle		0.79	0.79	0.84	0.87	0.84	0.83
Herds		1	2	3	4	5	6	Length
Full Information	Complete	33.33 (0.00)	58.89 (0.00)	64.44 (0.00)	66.67 (0.00)	66.67 (0.00)	71.11 (0.00)	3.6
	Star	28.89 (0.00)	40.00 (0.00)	60.00 (2.92)	62.22 (8.67)	65.56 (8.65)	71.11 (11.33)	3.3
	Circle	20.00 (0.00)	31.11 (3.15)	35.56 (3.12)	45.56 (9.30)	51.11 (9.31)	67.78 (3.05)	2.5
High Information	Complete	25.33 (10.53)	46.67 (5.71)	49.33 (5.41)	58.67 (4.55)	62.67 (6.38)	69.33 (7.69)	3.1
	Star	12.22 (0.00)	24.44 (0.00)	32.22 (0.00)	35.56 (0.00)	41.11 (0.00)	52.22 (6.68)	2.0
	Circle	13.33 (0.00)	21.11 (0.00)	27.78 (0.00)	32.22 (0.00)	34.44 (0.00)	48.89 (10.21)	1.8
Low Information	Complete	7.78 (0.00)	14.44 (0.00)	20.00 (0.00)	34.44 (11.05)	43.33 (14.79)	63.33 (14.50)	1.8
	Star	6.67 (0.00)	14.44 (0.00)	21.11 (0.00)	22.22 (0.00)	31.11 (0.00)	51.11 (11.30)	1.5
	Circle	10.67 (0.00)	18.67 (0.00)	25.33 (0.00)	33.33 (0.00)	38.67 (0.00)	48.00 (2.78)	1.7

Table 4. Experimental results: actual and private-information efficiencies in all networks and information structures for informed and uninformed subjects.

Information	Network	Private-information	Actual							
			Turn 1	Turn 2	Turn 3	Turn 4	Turn 5	Turn 6	Average	
Full	Complete	0.716	0.683	0.793	0.829	0.818	0.828	0.839	0.798	
	Star	0.746	0.627	0.649	0.707	0.718	0.691	0.681	0.679	
	Circle	0.709	0.611	0.526	0.553	0.531	0.580	0.588	0.565	
High	Complete	All	0.620	0.487	0.617	0.625	0.713	0.692	0.653	0.631
		Informed	0.877	0.731	0.692	0.744	0.772	0.776	0.720	0.739
		Uninformed	0.000	-0.103	0.435	0.337	0.571	0.489	0.489	0.370
	Star	All	0.688	0.596	0.660	0.707	0.659	0.745	0.580	0.658
		Informed	0.854	0.767	0.765	0.741	0.738	0.816	0.612	0.740
		Uninformed	0.000	-0.118	0.220	0.565	0.329	0.450	0.450	0.316
	Circle	All	0.616	0.613	0.566	0.644	0.604	0.705	0.580	0.618
		Informed	0.831	0.767	0.581	0.637	0.610	0.728	0.561	0.647
		Uninformed	0.000	0.171	0.524	0.662	0.586	0.637	0.632	0.535
Low	Complete	All	0.461	0.271	0.458	0.495	0.414	0.460	0.523	0.437
		Informed	0.975	0.719	0.706	0.783	0.563	0.660	0.701	0.689
		Uninformed	0.000	-0.130	0.236	0.236	0.280	0.280	0.364	0.211
	Star	All	0.440	0.177	0.331	0.389	0.467	0.463	0.545	0.395
		Informed	0.952	0.625	0.543	0.581	0.725	0.673	0.804	0.658
		Uninformed	0.000	-0.207	0.149	0.224	0.245	0.282	0.324	0.169
	Circle	All	0.494	0.452	0.469	0.579	0.590	0.601	0.592	0.547
		Informed	0.957	0.935	0.853	0.836	0.793	0.771	0.689	0.813
		Uninformed	0.000	-0.065	0.060	0.304	0.373	0.419	0.488	0.263

Table 5. Error rates at the second turn under the different information structures and networks

$N=2$

		Complete		
		Full	High	Low
All		13.0	14.2	11.9
Informed			17.7	13.8
Uninformed			8.3	10.8

		Star - center		
		Full	High	Low
All		11.1	16.7	15.6
Informed			18.8	22.2
Uninformed			11.5	12.7

$N=1$

		Circle		
		Full	High	Low
All		13.0	15.9	25.3
Informed			15.8	7.6
Uninformed			16.3	34.9

		Star - peripherals		
		Full	High	Low
All		1.1	17.8	28.9
Informed			14.8	18.6
Uninformed			26.7	33.9




 - "black spots"


Figure 1. The complete, circle and star networks with three agents



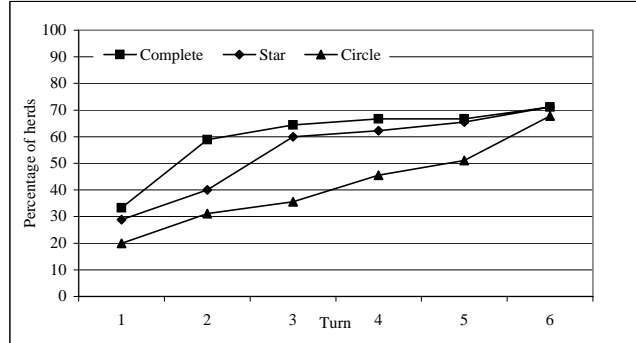
A line segment between any two types represents that they are connected and the arrowhead points to the participant whose action can be observed

Figure 2. Herd behavior in each network under all information treatments (left panel), and for all networks under each information treatment (right panel).
 (the percent of rounds in which subsequently all subjects acted alike)


Complete network



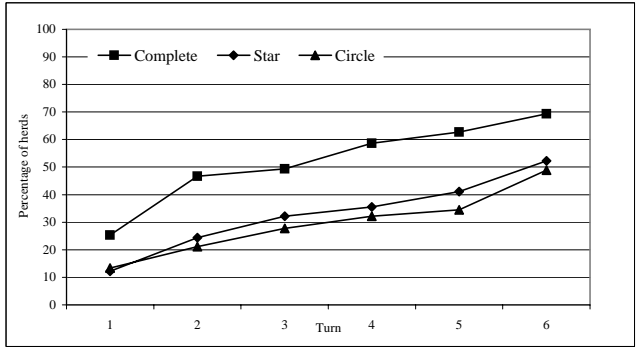
Full-information




Star network



High-information



Circle network



Low-information

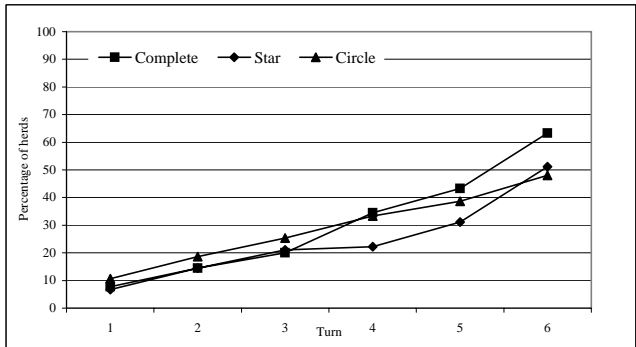




Figure 3. Actual efficiency in each network under all information treatments (left panel), and for all networks under each information treatment (right panel), and average private-information efficiency over all subjects within each information treatment.


Complete network




Full-information




Star network



High-information



Circle network



Low-information


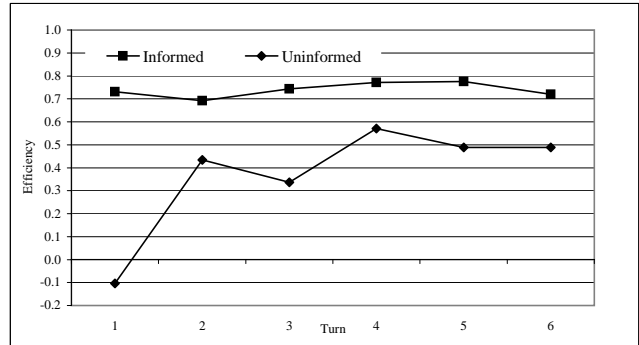





Figure 4. Actual efficiencies of informed and uninformed decisions in each of the networks under low- (left panel) and high-information (right panel).

Complete network



Star network



Circle network







Figure 5. Actual-efficiency distributions in each network under all information treatments (left panel), and for all networks under each information treatment (right panel).


Complete network




Full-information




Star network



High-information



Circle network



Low-information




Figure 6. Error-rates in the first and second decision turn

