

UC Santa Cruz

Archived UCSC Economics Department Seminars

Title

A Theory and Experiments of Learning in Social Networks

Permalink

<https://escholarship.org/uc/item/8853k4jd>

Author

Kariv, Shachar

Publication Date

2004-08-29

Learning in Networks: An Experimental Study*

Syngjoo Choi[†]
New York University

Douglas Gale[‡]
New York University

Shachar Kariv[§]
UC Berkeley

August 29, 2004

Abstract

Individuals living in society are bound together by a social network, the complex of relationships that brings them into contact with other agents. In many social and economic situations, individuals learn by observing the behavior of others in their local environment. This process is called *social learning*. Learning in incomplete networks, where different agents have different information sets, is especially challenging: because of the lack of common knowledge individuals must draw inferences about the actions others have observed as well as about their private information. Whether individuals can rationally process the information available in a network is ultimately an empirical question. This paper reports an experimental investigation of learning in three-person networks and uses the theoretical framework Gale and Kariv (2003) to interpret the data generated by the experiments. The family of three-person networks includes several non-trivial architectures, each of which gives rise to its own distinctive learning patterns. We find that the theory can account for the behavior observed in the laboratory in variety of networks and informational settings. To account for errors in subjects' behavior, we adapt the model of Quantal Response Equilibrium of McKelvey and Palfrey (1995, 1998) and find that its restrictions are also confirmed. The 'goodness of fit' is better for the QRE model than for the game-theory model. This provides important support for the use of QRE to interpret experimental data.

*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 Colin Camerer, Gary Charness, Xiaohong Chen, Charlie Holt, Tom Palfrey, and Andrew Schotter for helpful discussions. This paper has also benefited from suggestions by the participants of 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).

[‡]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

Social learning occurs when economic agents learn by observing the behavior of others. Whether choosing a restaurant, adopting a new technology, or investing in a portfolio, an individual's actions can reveal useful private information. So, in social settings, where agents can observe one another's actions, it is rational for them to try to learn from one another.

Individuals living in society are bound together by a *social network*, the complex of relationships that brings them into contact with other agents, such as neighbors, co-workers, family, and so on. An individual agent's ability to observe other agents is limited so, in practice, each agent has *imperfect information* about the actions of agents in the same network.

Gale and Kariv (2003) study a model of Bayesian learning in social networks. The social network is represented by a directed graph. Each agent is located at a node of the graph and agent i can observe agent j if there is an edge leading from node i to node j . Note that the links need not be symmetric: the fact that i can observe j does not necessarily imply that j can observe i . Agents make repeated choices from a finite set of actions. Information percolates through the network as one agent after another changes his action in response to what he observes his neighbors doing.

A simple example may clarify the setup in Gale and Kariv (2003). Many of us are repeatedly in the situation where we have to choose the contribution to a retirement plan such as a 401(k). It is natural for an individual to observe his colleagues' choices before making his own decision. Some of our colleagues, having worked in the department for a long time, will have friends in other departments and may know them well enough to have observed their 401(k) choices as well. So, indirectly, one's choice of 401(k) plan may be influenced by individuals one does not know and cannot observe. Furthermore, the network formed by these relationships may be large ("small-worlds" and "six degrees of separation") and complex. Nonetheless, a *rational* individual must take account of the network architecture in order to draw correct conclusions about the information revealed by the choices he observes.

Bayesian learning requires an agent to assign the correct probabilities to a potentially infinite set of states, implicitly constructing in his mind an infinite hierarchy of beliefs. In the Gale-Kariv model, agents can revise their action as more information becomes available. In this setting, the complexity of an agent's decision-problem increases over time. At the first date, an agent only has to interpret his private information. At the second date, he has to interpret his neighbors' actions and try to infer the private information on which it was based. At the third date, because of the lack of common knowledge about actions, an agent is forced think about his neighbors' knowledge of other their neighbors' actions and the private information they reveal. For example, consider a three-

person network in which agent A observes agent B , agent B observes agent C , and agent C observes agent A . Agent A , in interpreting B 's actions in the preceding period, has to think about the action B observed C choose in the period before that, what private information B thought C had, and what effect it had on B 's actions. Even in this three-person network, the exploitation of this information requires subtle reasoning because actions are not common knowledge. Lack of common knowledge forces agents to think about hierarchies of beliefs.

Whether individuals can rationally process the information available in a network is ultimately an empirical question. There is a large empirical literature which shows evidence of learning in social networks in many areas. Among others, Foster and Rosenzweig (1995) analyze technology adoption in developing countries, and Duflo and Saez (2002, 2003) use a quasi-experimental setting to show that the information transmission through social interactions affects retirement-plan decisions. However, these *observational* studies are subject to identification problems. Manski (1993, 1995) provides a formal exposition of the issues involved in identifying social effects. In the laboratory, by contrast, we can control subjects' neighborhoods and their private information. This provides an opportunity to test the model's predictions and, at the same time, study the effects of variables about which our existing theory has little to say. In the present paper, we report on a series of laboratory experiments based on a version of the Gale-Kariv model. Although the experimental setup is quite simple, the analysis of the game is sometimes complex. To draw correct inferences the subject must consider several levels of his hierarchy of beliefs.

The main conclusion of the paper is that the theoretical approach developed here is indeed relevant for the interpretation of the experimental data. In particular, it shows that a parsimonious model does a good job of explaining the data from a variety of networks and informational settings and provides a consistent explanation of apparently irrational behavior. To test the usefulness of the theory in interpreting the data, we first estimate a theoretically grounded, structural model based on the theory of Gale and Kariv (2003), modified to allow for the possibility of occasional mistakes. Then we show that the structural model does quite well in explaining the experimental data generated in a variety of networks and treatments.

The data generated by these experiments can also be used to address a variety of important and interesting questions about individual and group behavior. A related paper, Choi, Gale and Kariv (2004), uses the same data set to investigate behavioral aspects of individual and group behavior, including comparisons across networks and information treatments.

The experiments reported here involve three-person, connected social networks. We restrict attention to connected networks since obviously disconnected agents cannot learn from others. The case of three-person networks has several non-trivial architectures, each of which gives rise to its own distinctive learning patterns, and the Gale-Kariv model suggests that even in the three-person case the process of social learning in networks can be complicated. The complete set of networks is illustrated in Figure 1. A line segment between any two types

represents that they are connected and the arrowhead points to the participant whose action can be observed. Three representative networks are used in the experimental design:

- the *complete network*, in which each agent observes the actions chosen by all the other agents;
- the *circle network*, in which each agent observes the actions chosen by exactly one other agent and each agent is observed by someone;
- the *star network*, in which one agent (the center) observes the other two agents and the two (peripheral) agents only observe the center.

[Figure 1 here]

We chose these networks because they illustrate the main features of the complete set of networks — the excluded networks can each be obtained by adding a single link to one of the three chosen networks. For practical purposes, these three networks “span” the set of networks and provide a reasonable test of the theory.

In the experimental design, there are two equally likely events (states of nature). We allow subjects to be of two types: *informed agents*, who receive a private signal that is correlated with the unknown events, 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 two events has taken place, basing his forecast on a private signal and the history of his neighbors’ past decisions. Each experimental session, consisting of 15 rounds, used a single network and a single information treatment and a single group of subjects.

We begin our analysis of the experimental data by calculating the equilibrium strategies predicted by the Gale-Kariv model and use these to compute error rates (the percentage of times subjects deviate from the equilibrium strategy). Although there is some variation across decision-turns, networks and treatments, the error rates are uniformly fairly low. Nonetheless, mistakes are made and this should be taken into account in any theory of rational behavior.

This leads us to adopt the Quantal Response Equilibrium (QRE) model of McKelvey and Palfrey (1995, 1998). We extend the basic model of Gale-Kariv (2003) to allow for idiosyncratic preference shocks, which can be interpreted, following Harsanyi and Selten, as the effect of a “trembling hand”. More precisely, the payoff from a given action in the perturbed game is assumed to be a weighted average of the theoretical payoff and a logistic disturbance. The “weight” placed on the theoretical payoff is determined by a regression coefficient. This coefficient will be positive if the theory has any predictive power, and will approach infinity if the theory predicts subjects’ behavior exactly. For any finite value of the coefficient, there is a positive probability that the optimal action predicted by the perturbed model will be different from the prediction

of the Gale-Kariv model. This allows us to account for “mistakes” in subjects’ behavior. It also shows us how subjects should rationally take into account the mistakes of others when drawing inferences from their behavior.

The Gale-Kariv model has a natural recursive structure. At the first decision-turn in any game, an agent makes a decision based on his private signal (if he is informed) or his prior (if his uninformed). After he has made his decision, he observes the actions chosen by his neighbors and updates his beliefs. At the second turn, he chooses a new action based on his updated beliefs, observes the actions chosen by his neighbors and the second turn, and updates his beliefs again. At the third turn, he chooses a new action based on his information from the second turn, and so on. Thus, at each turn, his decision is backward-looking (based on past information).

The recursive structure of the model allows us to estimate the coefficients of the QRE model for each decision-turn sequentially. For each network and treatment, we begin by estimating a QRE using the data from the first turn. Then we use the estimated coefficient from the first turn to calculate the theoretical payoffs from the actions at the second turn. 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. 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 network and 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 conduct a series of specification tests to see whether the restrictions of the QRE model are confirmed by the data and the results are strikingly in conformity with the theory. The decision rules of the QRE model are qualitatively very similar to the empirical choice probabilities. In particular, the data confirms the prediction of the logistic model that errors are more likely when there is little at stake (payoff differences are small).

Our paper contributes to the large and growing body of work which studies the influence of the network structure on economic outcomes. Goyal (2003) and Jackson (2003) provide recent surveys of theoretical work in economics focusing on social and economic networks and Kosfeld (2004) surveys the experimental work. The paper most closely related to Gale and Kariv (2003) is Bala and Goyal (1998). The models differ in two ways. First, Bala and Goyal (1998) examines the decisions of *boundedly rational* agents, who try to extract information from the behavior of the agents they observe, but without taking account of the fact that those agents also observe other agents. Second, in Bala and Goyal (1998), agents observe payoffs as well as actions. In other words, it is a model of *social*

experimentation rather than social learning.

The paper also contributes to a large literature on social learning. Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992) (BHW) introduced the basic concepts and their work was extended by Smith and Sørensen (2000). These models show that social learning can easily give rise to herd behavior or informational cascades, phenomena that have elicited particular interest and can arise in a wide variety of social and economic circumstances. This is an important result and it helps us understand the basis for uniformity of social behavior. At the same time, these models 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. In other words, it is a game of perfect information.

Anderson and Holt (1997) investigate the social learning model of BHW experimentally and replicate informational cascades in the laboratory. Following Anderson and Holt (1997), a number of experimental papers analyzed different aspects of social learning. Among others, Hung and Plott (2001), Kübler and Weizsäcker (2004) and Çelen and Kariv (2004a, 2004b) extend Anderson and Holt (1997) to investigate other possible explanations for informational cascades. This growing body of experimental work in the social learning literature has also successfully utilized QRE models.

The rest of the paper is organized as follows. The next section illustrates some features of the underlying theory. Section 3 describes the experimental design and procedures. Section 4 summarizes some important features of the data. Section 5 provides the econometric analysis and Section 6 concludes.

2 Some theoretical examples

In this section we discuss briefly the theoretical implications of the model tested in the laboratory. Gale and Kariv (2003) provide an extensive analysis of a general version of the model.

A network consists of three agents indexed by $i = A, B, C$. Each agent i has a set of neighbors, that is, agents whose actions he can observe. Let N_i denote the neighbors for agent i . The neighborhoods $\{N_A, N_B, N_C\}$ completely define a three-person network. These networks are illustrated in Figure 1 above.

There are two equally likely events (states of nature) denoted by $\omega = -1, 1$. With probability q an agent is informed and receives a private signal at the beginning of the game. Signals take two values $\sigma = -1, 1$ and the probability that the signal σ equals the true state ω is $2/3$. By convention, we assume an uninformed agent receives the signal $\sigma = 0$ in each state. The agent's signals are assumed to be independently distributed conditional on the true state.

Time is divided into a finite set of dates indexed by $t = 1, 2, \dots, T$. At the beginning of each date t , agents are simultaneously asked to guess the true state. Agent i 's action at date t is denoted by $a_{it} = -1, 1$. Agent i receives a positive payoff if his action a_{it} equals the true state ω and zero otherwise. Then each

agent i observes the actions a_{jt} chosen by the agents $j \in N_i$ and updates his beliefs accordingly. Thus, agent i 's information set at date t consists of his private signal, if he observed one, and the history of neighbors' actions.

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 t . There are several reasons for focusing on these equilibria. First, we want to stay close to the existing social-learning literature, in which myopic behavior is usually optimal. Secondly, in the absence of forward-looking, strategic considerations, the equilibrium has a recursive structure that simplifies the theoretical and econometric analysis. Thirdly, our econometric results strongly suggest that myopic behavior is consistent with the experimental data. Finally, a careful analysis shows that the tie-breaking assumption that agents switch actions whenever they are indifferent between continuing to choose the same action in the next period and switching to the other action is fully revealing. Thus, there is no incentive to sacrifice short-run payoffs in any period in order to influence the future play of the game.

Because of the symmetry of the example and the fact that signals take only discrete values, an agent is often indifferent between choosing $a_{it} = -1$ and $a_{it} = 1$, in which case some tie-breaking rule has to be chosen. It is important to note that the nature of the equilibrium play depends on the tie-breaking assumption. Here we assume that, whenever an agent has no signal, he chooses each action with probability 0.5 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 and it also eases the exposition and analysis. The other advantage of this approach is that agents' actions are also optimal given perturbed beliefs that take into account the possibility that others make mistakes. We will point out and discuss alternatives whenever our tie-breaking assumption becomes relevant. Note, however, that for the purpose of estimating the QRE model, the tie-breaking rule is irrelevant because the "trembling hand" ensures that ties are probability zero events.

An agent's equilibrium behavior in this game is conceptually very simple: at each date he chooses the state he thinks is most likely and then he updates his beliefs, using Bayes' rule, based on what he observes the other agents do. Bayesian updating is conceptually simple but it is computationally very difficult because of the large number of information sets and the lack of common knowledge. So we do not believe that subjects perform this calculation. Instead, they use heuristics which mimic the effect of rational Bayesian maximizing behavior.

Gale and Kariv (2003) describe agents' behavior formally and discuss the essential elements of the weak perfect Bayesian equilibrium, so we skip the model development and analysis and instead illustrate how the dynamics of actions and learning differ across networks and information structures. In order to get a sense of the challenges of substantive rationality in different settings, as well as the implications for equilibrium behavior of the different networks and information treatments, we consider a series of theoretical examples of the underlying game. We begin with the complete network.

2.1 The complete network

A network is *complete* if each agent can observe the actions of all the other agents in the network. Otherwise the network is called incomplete. There is a unique complete network, in which $N_A = \{B, C\}$, $N_B = \{A, C\}$, and $N_C = \{A, B\}$. The experimental design uses three information treatments, corresponding to different values of the probability of being informed. We refer to these as *full information* ($q = 1$), *high information* ($q = 2/3$), and *low information* ($q = 1/3$), respectively.

Full-information ($q = 1$) When every agent is informed, the equilibrium behavior is particularly simple and closely resembles the herd behavior found in BHW. At the first date, each agent's information consists of his private signal σ_i . The true state is more likely to be σ_i so agent i puts $a_{i1} = \sigma_i$. At the second date, each agent's first-period action has revealed his signal and so the signals are common knowledge. Since there must be at least two signals with the same value, from date 2 onwards all agents agree on the most likely state and will choose the same action at date 2 and every following period.

So, in this case, the equilibrium behavior is very simple. An informational cascade at date 2 causes a herd that continues until the end of the game. Further, the herd chooses the efficient action, based on the sum of agents' information, unlike the model of BHW. Here a rule of thumb that says "follow the majority" would lead to both a rational and efficient outcome.

High-information ($q = 2/3$) Equilibrium behavior is slightly more complicated when there is high information, because agents have to take account of the possibility that some other agents are uninformed. In this case, information revelation may continue after date 2. Suppose, for example, that agent A receives the signal $\sigma_A = 1$, agent B receives the signal $\sigma_B = -1$, and agent C is uninformed $\sigma_C = 0$. At date 2, agent C observes that the actions of agents A and B at date 1 do not match, so he is indifferent between the two actions. If agent C takes action -1 at date 1 and switches to action 1 at date 2, he reveals that he is uninformed. At date 2, agent A observes that B and C chose 1 in the previous period, so he switches to action -1 at date 2. This can be confirmed with a simple calculation using Bayes' rule. However, at date 3, he realizes that C is uninformed and since he is still not sure whether B is informed (B might have chosen -1 two times in a row by chance), it is rational for him to switch back to 1.

This example shows that the possibility of uninformed agents changes the qualitative features of the equilibrium. First, we no longer necessarily get a herd at date 2 and learning continues after date 2. Secondly, learning continues even if there is no change in an agent's actions: the longer B persists in choosing action -1 , the more confident A is that B is informed; but there is always some positive probability that B is uninformed and chose consistently by accident. Note that this aspect of the equilibrium depends on our tie-breaking rule. If the uninformed agent chooses the same action as last period when he is indifferent,

then the distribution of signals assumed above implies a herd on action -1 starts at date 2.

Unlike the full-information case, the dynamics of actions and beliefs in the high-information example above are complex and do not correspond to any simple heuristics. The greater complexity of behavior stems from the fact that agents have different amounts of information and the ability to revise decisions reveals this asymmetry over time.

Low-information ($q=1/3$) Qualitatively, the low-information case is like the high information case. The possible existence of uninformed agents allows learning to continue after date 2. The main difference lies in the fact that agents think it is much less likely that their opponents will be informed and hence have less incentive to imitate them. Suppose, for example, that all the agents are informed and that $\sigma_A = 1$, $\sigma_B = 1$, and $\sigma_C = -1$. A simple calculation shows that agent C will continue to choose action -1 at date 2, because he thinks it quite likely that A and B are uninformed. At date 3, agent C observes that A and B chose action 1 again at date 2, which reinforces C 's belief that A and B are informed. So here learning continues but the actions do not change. If the game continues long enough (i.e., T is large) C will eventually switch. This conclusion depends on our tie-breaking rule that indifferent uninformed agents randomize.

2.2 The star network

The first incomplete network we examine is the star, in which $N_A = \{B, C\}$, $N_B = \{A\}$, and $N_C = \{A\}$. The most interesting feature of this network is its asymmetry: agent A can observe both B and C and thus has more information than either. In fact, agent A is informed about the entire history of actions that have already been taken, whereas B and C have imperfect information. So here we can see the impact of both lack of common knowledge and asymmetry on the dynamics of social learning.

Because of the imperfection of information, learning continues after date 2 even in the full information case. Suppose then that there is full information ($q = 1$) and suppose the realizations of the signals are $\sigma_A = 1$, $\sigma_B = 1$, and $\sigma_C = -1$. Now, at date 2, agent C only observes that his action at date 1 does not match A 's action, so our tie-breaking assumption becomes relevant. The tie-breaking rule requires that agent C continue to choose action -1 at date 2. Agent B , on the other hand, sees that agent A has chosen the same action and this merely increases B 's belief that the true state is 1. From agent A 's perspective, agent C 's signal cancels out agent B 's, so agent A 's belief about the true state is unchanged. At date 2, each agent will make the same choice as at date 1.

Although the actions do not change between dates 1 and 2, information is revealed. In particular, agent C knows that since A did not change his action at date 2, A must have observed B choose 1 at date 1. Thus, C knows that A

and B both received the signal 1. Thus, it is optimal for C to switch to action 1 at date 3. We have again reached an absorbing state.

This example shows both the complexity of behavior under full information and the subtlety of the reasoning that may be required to draw correct inferences from the observed actions. Here, agent A serves as a communication channel between agents B and C as well as a potential source of private information. It can be shown by example that actions and beliefs may continue to evolve after the third date.

2.3 The circle network

The second incomplete network is the circle, in which each agent observes one other agent: $N_A = \{B\}$, $N_B = \{C\}$, and $N_C = \{A\}$. In the circle, every agent has imperfect information about the history of actions chosen in the game. Further, each agent is forced to make inferences about what the others have seen. In this network, the equilibrium reasoning required to identify the optimal strategy is subtle, but the equilibrium strategy itself is quite simple: an informed agent should always follow his own signal and an uninformed agent should imitate the one other agent he can observe. This reminds us that substantive rationality can be simpler than procedural rationality. It does not imply that behavioral dynamics are simple. For example, if all agents are uninformed, it may take a long time for the agents to discover this fact. Both beliefs and actions will continue to evolve until this fact is revealed, after which our tie-breaking rule implies that the agents' behavior is random.

2.4 Takeaways

The preceding examples have illustrated several features of the theory:

- Perhaps the most important point is that, in spite of the simplicity of the game, the inferences agents must draw in order to make rational decisions are quite subtle. In particular, because of the lack of common knowledge, agents have to think about a large number of possible situations that are consistent with their limited information.
- Even though the reasoning required to identify them is quite complex, optimal strategies may be quite simple.
- Significant differences can be identified in the equilibrium behavior of agents in different networks. We saw that in the complete network learning stops almost immediately if there is full information, whereas the existence of asymmetrically informed agents is consistent with a longer period of learning and more complex strategies.
- Similarly, different information treatments lead to different dynamics of beliefs and actions. For example, comparing the full-information and high-information treatments, we see that less time is required for beliefs and actions to converge when information is full.

We have focused on examples that reveal some of the unexpected features of the model. One must remember, however, that in many situations the outcome is much simpler. As a general rule, we can say that initial diversity of private information causes diversity of actions but that, as agents learn from each other, diversity is replaced by uniformity (barring cases of indifference). Convergence to a uniform action tends to be quite rapid, typically occurring within two to three periods. Thus, what happens in those first few periods is important for the determination of the outcome. Note, however, that the converse of the convergence result — if all agents choose the same action, they have reached an absorbing state and will continue to choose that action at every subsequent date — is not true in general.

Finally, we note 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 potential complexity of equilibrium strategies and the complexity of the reasoning typically required for substantive rationality confirm the importance of verifying the relevance of the theory empirically.

3 Experimental design

The experiment was run at the Experimental Economics Laboratory of the Center for Experimental Social Sciences (C.E.S.S.) at New York University. The subjects in this experiment were recruited from undergraduate classes at New York University and had no previous experience in network or social-learning experiments. After subjects read the instructions (see Appendix II), the instructions were read aloud by an experimental administrator¹. A \$5 participation fee and subsequent earnings for correct decisions were paid in private at the end of the experimental session. Throughout the experiment we ensured anonymity and effective isolation of subjects in order to minimize any interpersonal influences that could stimulate uniformity of behavior².

We studied three connected, three-person network structures (the complete, star, and circle networks) and three different information treatments (full, high, and low information). 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

¹At the end of the first round subjects were asked if there were any misunderstandings. No subject reported any difficulty understanding the procedures or using the computer program.

²Participants' work-stations were isolated by cubicles making it impossible for participants to observe other screens or to communicate. At the end of a session, participants were paid in private according to the number of their work-stations.

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. Urn R contained 2 red balls, and 1 white ball. Urn W contained 1 red ball and 2 white balls. The urn remained constant throughout the round. The choice of urn was independent across networks and across rounds. In each decision-turn, subjects were asked to predict which of the two urns had been chosen in that round.

To help subjects determine which urn had been selected, with probability $q = 1, 2/3, 1/3$ each subject was allowed to observe one ball, drawn at random with replacement, from the urn. Before subjects were called to make their first decision, each was informed whether the computer had drawn a ball for him and whether it was white or red. After everyone had seen his draw, each subject was asked to input the letter of the urn, W or R , that he thought was most likely to have been chosen by the computer. When all subjects in the session had made a decision, each subject observed the choices of the subjects to whom he was connected in his network. This completed the first of six decision-turns in a round. Next, subjects were asked to make their second decision, without observing a new draw from the urn. This process was repeated until six decision-turns were completed. At each date, the information available to subjects included the actions they had observed at every previous date.

When the first round ended, the computer informed subjects which urn had actually been chosen and their individual earnings. 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 ensured that at each decision-turn subjects would make their best guess as to which urn had been chosen. After subjects learned the true urn and their earnings, the second round started by having the computer randomly forming new groups of participants in networks and selecting an urn for each group. This process was repeated until all the 15 rounds were completed.

Note that the experiment is different from the standard social-learning experiments paradigm of Anderson and Holt (1997) in two important ways. First, subjects can only observe the actions of subjects to whom they are connected by a social network. Thus, actions are not public information and subjects can observe the actions of some, but not necessarily all, of their neighbors. Second, subjects make decisions simultaneously, rather than sequentially, and can revise their decisions rather than making a single, irreversible decision.

The experiments provide us with a rich set of data. Each of the nine sessions (a single network and a single information treatment) comprised 18 subjects (or in two cases, 15 subjects). A session consists of 15 rounds and each round consists of six decisions. In each round, the subjects were randomly formed into six (respectively, five) networks. So for each session we have observations on $6 \times 15 = 90$ (respectively, $5 \times 15 = 75$) different rounds and a total of

$18 \times 90 = 1520$ (respectively $18 \times 75 = 1330$) individual decisions. 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).

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

We use a variety of different treatments and network architectures to generate a variety of different outcomes which are representative of the theory. More importantly, the variety of different outcomes provides a serious test of the ability of a structural econometric model based on the theory to interpret the data.

4 Overview of experimental data

In this section, we provide an overview of some important features of the experimental data. We summarize these features of the data using three measures: *stability*, *efficiency* and *rationality*. We explain the three measures and their motivations and report averages across different treatments. We do not describe the data at the level of the individual subject. In a related paper, Choi, Gale and Kariv (2004), we use the same data set to investigate more thoroughly the behavioral aspects of individual and group behavior.

4.1 Herd behavior

One of the most important contributions of the social learning literature is to provide a model of rational herd behavior. When rational individuals ignore their own private information and “follow the herd,” information is lost, with the result that inefficient actions may be chosen. Herd behavior arises in the laboratory when, from some decision-turn on, all subjects take the same action. We begin by with the positive question of existence of herd behavior and later consider the normative question of efficiency. Herd behavior can be characterized in terms of two related phenomena, *uniformity* and *stability* of actions. 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. Stability is measured by the proportion of subjects who continue to choose the same action they chose at turn $t - 1$. 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 turns $t < s \leq T$. In such a case, we say that a herd of length $T - t + 1$ has occurred in the laboratory. Our first observation describes the evolution of herd behavior.

Observation 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 this observation is presented in Table 1 which shows, turn by turn, the average level of stability and uniformity and the percentage of rounds in which subjects followed a herd from that turn on (Table 1A). For comparison purposes, the experimental results are presented along with the theoretical predictions which are derived with the help of simulations (Table 1B).³ Note that, 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 1 here]

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

Observation 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.*

Observation 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.*

The first evidence about the frequencies of herd behavior is provided in Table 1. The relevant support for Observations 2 and 3 comes from Figure 2 which presents, in graphical form, the data from Table 1 on herd behavior in each network under all information treatments (Figure 2A), and for all networks under each information treatment (Figure 2B). A set of binary Wilcoxon tests indicates that the differences are significant at the 5 percent level.

[Figure 2 here]

The right column of Table 1 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

³We compute the measures with the help of simulations that were carried out by MatLab. An experiment starts by drawing a vector of private signals. We then collect the actions generated by this vector according to the model. Each experiment was repeated 10^5 times.

increases with the probability q that an individual subject receives a signal. Also, under full and high information, the expected length of herd behavior in the complete network is greater than in the star, and is greater in the star than in the circle. Under low-information, the expected length of herds in the circle is greater than in the star.

Convergence to a uniform action is more rapid in the complete network under full-information, where it is common knowledge that all subjects are informed and all actions are common knowledge. In contrast, diversity continues for a longer time under high and low information and in the star and circle networks. We conjecture that 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 more complex calculations.

Recall that in the specific parametric model underlying our experimental design, except the complete network under high-information, herds always entail correct decisions in all treatments. Thus, it is particularly interesting that almost all herds longer than three turns selected the right action, but some differences can be identified in the behavior of different networks. These facts are summarized in the following observation.

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

Evidence for this observation is also provided by Table 1. The numbers in parentheses are the fractions of herds that choose the wrong action, defined relative to the information available.

4.2 Informational efficiency

One of the central concerns in the study of social learning is the efficiency of information aggregation. 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 information 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 signals; define the *private-information expected payoff* to be the expected earnings of an 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* 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.

In order to assess the quality of aggregation and use of information within a network, the efficiency of decisions is measured separately for informed and

uninformed individuals in two ways:

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

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

Thus, the actual (private-information) efficiency is the difference between the actual (private-information) expected payoff π_a (π_p) and the random-choice expected payoff π_r as a fraction of the difference between the efficient expected payoff π_e and the random-choice expected payoff π_r . Note that actual efficiency varies by decision turn whereas private-information efficiency is constant across all turns in a round. Notice also that efficient decisions have an efficiency of one and random decisions have an efficiency of zero. The comparison of actual and private-information efficiencies is useful in determining the extent to which informed and uninformed subjects use the information revealed by their neighbors' actions. Table 2 summarizes the actual and private-information efficiencies in all networks and information structures (Table 2A), and the theoretical predictions which are derived with the help of simulations (Table 2B).

[Table 2 here]

We next turn our attention to analyze how efficient our subjects were in using the information revealed by their neighbors' actions. The next two observations report average actual-efficiency calculations to measure the informational efficiency within a given network, information treatment, and turn.

Observation 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.*

Observation 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.*

The support for Observations 5 and 6 comes from Figure 3, which presents the data from Table 2 by comparing the total actual efficiency in each network under all information treatments (Figure 3A), and for all networks under each information treatment (Figure 3B). Figure 3B 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 significant at the 5 percent level.

[Figure 3 here]

4.3 Rationality

Rationality is measured by the percentage of times subjects follow an equilibrium strategy. At the first and second decision-turns, the data supports the following observation.

Observation 7 *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 Observation 7 is given by the diagram 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 percentages 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)

4.4 Summary

In summary, 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 average subject behavior among the three networks and three information treatments. We have suggested that these differences might be explained by differences in the amount of common knowledge and the symmetry or asymmetry of the network or the information treatment. The reader is referred to Choi, Gale, and Kariv (2004) for a fuller discussion of these issues. Finally, we have noted the high degree of rationality in subject behavior. Nonetheless, mistakes are made. We take explicit account of the possibility of mistakes in the next section where we introduce a model of Quantal Response Equilibrium.

5 Econometric Analysis

To explain the agent's propensity to choose an action different from the one predicted by the basic theory, we assume that each agent's payoff is perturbed by an idiosyncratic preference shock that has a logistic distribution. The logit equilibrium can be summarized by a choice probability function following a

binomial logit distribution :

$$\Pr(a_{it} = 1|I_{it}) = \frac{1}{1 + \exp(-\beta_{it}x_{it})},$$

where a_{it} is the action of agent i at date t , I_{it} is agent i 's information set at date t , β_{it} is a coefficient, and x_{it} is the difference between the expected payoffs from actions $a = 1$ and $a = -1$, respectively. The choice of action becomes purely random as β_{it} goes to zero, whereas the action with the higher expected payoff is chosen for sure as β_{it} goes to the infinity. For positive values of β_{it} , the choice probability is increasing in x_{it} .

In the QRE model, a rational subject must predict his neighbors' choice probabilities correctly to calculate the posterior probabilities correctly. In effect, we assume that subjects have rational expectations about their neighbors' true error rates (determined by the true value of beta) and use the estimated beta coefficients to approximate the true beta. Thus subjects use the estimated betas from the prior decision turn $t - 1$ to update their posterior beliefs and expected payoffs at any decision-turn $t > 1$. These in turn determine the choice probabilities via the logistic response function given above.

We use repeatedly the standard maximum likelihood (ML) method for the estimation of the logistic random-utility models. The data employed to implement the ML estimation for betas at each turn are the current actions and the implied expected payoffs for the current period. Taking into account the influence of the networks and information treatments on the calculation of expected payoffs, we pool homogeneous data at each turn to reduce sampling errors in the estimation of betas. At the first decision-turn, in any network and information treatment, decisions are based only on private information. So all the data from the first turn of the experiment were pooled to provide a unique beta estimate. The information treatment and the number of neighbors matter in the computation of expected payoffs at the second turn. So we pooled the data of subjects who observed the same number of neighbors in the same information treatment to estimate a set of second-turn beta estimates: betas were estimated separately for each information treatment and for each of two groups of subjects, (a) all subjects in the complete network and type-A subjects in the star network and (b) all subjects in the circle network and type-B and type-C subjects in the star network. From the third turn on, we estimate betas separately for each network and information treatment and, in the case of the star network, distinguished the betas for the center (type A) and the periphery (types B and C).

We can illustrate the recursive estimation procedure with reference to the circle network. At the first decision-turn, we calculate the difference in expected payoffs, x_{i1} , conditional on the private signals for $i = A, B, C$. Then the beta for the first decision-turn is estimated via the ML logit estimation. Then the beta estimate for the first turn, $\hat{\beta}_1$, is used to determine the choice probabilities of each subject's neighbor j , $\Pr(a_{j1}|I_{j1})$, for each possible I_{j1} . These choice probabilities, together with Bayes' rule, are used to calculate the posterior probability that the state is $\omega = 1$ conditional on subject i 's information set, $\Pr(\omega = 1|I_{i2})$, which in turn determines the difference in expected

payoffs, $x_{i2}(\hat{\beta}_1)$. Analogously, the beta estimate for the second turn, $\hat{\beta}_2$, can be obtained. Note that the estimation procedure follows precisely each subject's inference problem in the theory and it becomes more involved at later decision-turns. At the third turn, the incomplete structure of the circle network requires each subject to make inferences about the behavior of his neighbor's neighbor k . Thus, the beta estimates for the first and second turns are used to determine the choice probabilities of his neighbor j at the first and second turn, $\Pr(a_{j1}|I_{j1})$ and $\Pr(a_{j2}|I_{j2})$, and the choice probabilities of his neighbor's neighbor k at the first turn, $\Pr(a_{k1}|I_{k1})$. Again, together with Bayes' rule, these probabilities are used to compute the posterior probabilities and thus the difference in expected payoffs at the third turn, $x_{i3}(\hat{\beta}_1, \hat{\beta}_2)$, which serves as the independent variable in the estimation of the beta for the third turn. Continuing in this manner, we can estimate the entire logit equilibrium models for the circle network and each information treatment. The procedure is analogous for the other networks. The details for the inference problem in the QRE model for each network are relegated to Appendix I. Table 3 presents the results of the ML logit equilibrium estimation. Standard errors are given in parentheses.

[Table 3 here]

All the beta estimates are significantly positive. This implies that, under the specification of the logistic distribution, the behavior of subjects is not entirely random and the model of logit equilibrium has some predictive power in interpreting their behavior in the laboratory. Although it seems difficult to identify any marked behavioral differences of beta estimates across networks and information treatments, we found at least one apparent cross-sectional feature of the beta series: for each decision turn up to and including the fifth, the estimate beta coefficients from the circle network are monotonic with respect to information treatment. That is, for a fixed decision turn t the beta coefficient is lowest for the full-information treatment, higher for the high-information treatment, and highest for the low information treatment. Figure 4 provides a graphical re-presentation of the beta series in the complete, star (type A and types B and C) and circle networks.

[Figure 4 here]

Figure 4 indicates that subjects in the circle network are more sensitive to the difference in (theoretical) expected payoffs in the high and low information treatments. Recall that, in this network, the reasoning required to identify optimal strategies is complex, but the strategies themselves are quite simple: an informed subject should always follow his own signal and an uninformed subject should imitate the one other subject he can observe. We conclude that, overall, subjects were more likely to follow these strategies in lower information treatments. However, the differences may be explained by compositional differences resulting from the changes in the proportion of informed and uninformed subjects.

Although the results of the logit analyses show some power in predicting the behavior observed in the laboratory, further investigation is needed to determine whether this parametric specification of QRE fits the data well. In particular, the parametric specification implies that the probability distribution of choices has the familiar logistic shape and that subjects are more likely to make “mistakes” when the differences in expected payoffs are small. To test the predictions of the model, we first perform a series of graphical comparisons between predicted logit choice probabilities and empirical choice probabilities. The predicted logit choice probabilities across networks and treatments are graphed using the corresponding beta estimates. We use the method of nonparametric regression estimation to represent the empirical choice probabilities. Specifically, define $y_{it} = 1_{\{a_{it}=1\}}$, where $1_{\{\cdot\}}$ is an indicator function. Assume that the true relation between y_{it} and x_{it} may be expressed in terms of the conditional moment $E[y_{it}|x_{it}] = G(x_{it})$, where $G : \mathbb{R} \rightarrow [0, 1]$. Then given a data set $\{(y_{it}, x_{it})\}_{i=1}^n$ we employ the Nadaraya-Watson estimator with a Gaussian kernel function for the choice probability associated with each of the parametric cases.⁴ Note that we construct the data of expected payoffs x_{it} , for $t \geq 2$, using the logistic distribution specification. The bandwidth is chosen to be $n^{-1/5}$.⁵ In all cases the selected bandwidths provided properly smoothed kernel regression estimates. Figure 5 shows a set of comparisons between these two choice probabilities.

[Figure 5 here]

In each of the graphs in Figure 5, a solid line represents the nonparametrically estimated choice probability of action 1 and a dashed line represents the parametrically estimated logit choice probability for the same action. A beta estimate and a selected bandwidth are reported at the top of each panel. These graphical comparisons presents a rough indication for goodness of fit. Somewhat surprisingly, the fits are generally good except for the cases of type-A subjects in the star network with full information. In particular, the empirical data confirm the main prediction of the QRE model that errors are more likely when payoff differences are small. We investigated the irregularity in the case of type-A subjects in the star network with full information and found it was caused

⁴The Nadaraya-Watson estimator for $G(\cdot)$ is given by

$$\hat{G}(x) = \left[\sum_{i=1}^n K\left(\frac{x_{it} - x}{h}\right) y_{it} \right] / \sum_{i=1}^n K\left(\frac{x_{it} - x}{h}\right),$$

where h is a bandwidth and $K(\cdot)$ is a kernel function. The Gaussian kernel function is given by $K(u) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}u^2)$ for $u \in \mathbb{R}$.

⁵The optimal bandwidth in the nonparametric kernel regression with a single independent variable is proportional to $n^{-1/5}$. We tried several methods of automatic bandwidth-selection such as Generalized Cross Validation. However, the bandwidth yielded by those methods resulted in a kernel regression estimate that was too irregular to be plausible. It is interesting to note that the literature of bandwidth selection in nonparametric regression indicates that automatic bandwidth selection is not always preferable to graphical methods with a trial and error approach. See Pagan and Ullah (1999, p.120).

by a combination of the small-sample problem and one subject’s “irrational” behavior.⁶

The graphical comparison is highly suggestive but a formal test is more convincing, so we performed specification tests for the functional-form assumption of the logistic random-utility model using Zheng’s (1996) test. Given the unknown relation between y_{it} and x_{it} for any decision-turn t , we test the null hypothesis that the logit equilibrium model is correct:

$$H_0 : \Pr [\mathbb{E}(y_{it}|x_{it}) = G(\beta_t^0 x_{it})] = 1 \text{ for some } \beta_t^0 \in \mathbb{R},$$

where $G(\beta x) = 1/(1 + \exp(-\beta x))$. The alternative hypothesis is, without a specific alternative model, that the null is false:

$$H_1 : \Pr [\mathbb{E}(y_{it}|x_{it}) = G(\beta_t x_{it})] < 1 \text{ for all } \beta_t \in \mathbb{R}.$$

Note that the alternative includes all the possible departures from the null model.⁷ The results of the series of specification tests are reported in Table 4. P -values are reported in parentheses.

[Table 4 here]

The bandwidth is selected to be $cn^{-1/5}$, where c is equal to 1. The test results in Table 4 confirm the previous graphical comparisons. In most of the cases p -values are fairly high and support strongly the parametric specifications. As seen in the graphs, we reject the null in the case of type- A subjects in the star network with full information at decision-turns $t = 3, 5, 6$ with 5% significance level. Interestingly, we also reject the null at the third turn, in the high-information, circle-network treatment, with 5% significance level.⁸

6 Conclusion

Many economic decision problems involve incomplete and asymmetric information. That is, agents are uncertain about some underlying decision-relevant

⁶The subject played the following strategy $(a_1, a_2, a_3, a_4, a_5, a_6) = (1, 0, 1, 0, 1, 1)$ in 12 out of 15 rounds. Further, in 9 rounds out of 12, the optimal strategy required the subject to choose action 0 for all decision-turns. Most of the time, he did not even coordinate with his own signal.

⁷The test statistic T_n is given by

$$T_n = \frac{\sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n K\left(\frac{x_i - x_j}{h}\right) e_i e_j}{\left\{ \sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n 2K^2\left(\frac{x_i - x_j}{h}\right) e_i^2 e_j^2 \right\}^{1/2}},$$

where $K(\cdot)$ is a kernel function, h is a bandwidth, and $e_i = y_{it} - G(\hat{\beta} x_{it})$ with beta estimate $\hat{\beta}$ under the null. Under some mild conditions, the asymptotic distribution of T_n under the null hypothesis is the standard normal (Theorem 1 in Zheng, 1996).

⁸To investigate whether the test results are sensitive to the choice of bandwidth, we also calculated the test statistics when c is equal to 0.5 and 2. On the whole, we obtained the quite similar results with a small variation.

event and the information about it is shared asymmetrically among them. Consequently, agents have a very strong incentive to learn by observing the behavior of others. In social settings, agents are part of a social network and can only observe the actions of agents to whom they are connected through the network. Thus, networks are natural tools for understanding the social learning phenomenon.

Whether agents can rationally process the information available in a network is ultimately an empirical question. To test the relevance of the theory, 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. The family of three-person networks includes several non-trivial architectures, each of which 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. This confirms the importance of verifying the relevance of the theory empirically.

Somewhat surprisingly we find that the theory, modified to include the possibility of errors, does a good job of interpreting the subjects' behavior. Despite the complexity and sophistication of the decision-making required by the theory, the decision rules of the QRE model appear to be qualitatively very similar to the data. The series of specification tests we conducted to see whether the restrictions of the QRE model are confirmed by the data and the results are strikingly in conformity with the theory. This provides strong support for the use of theoretical models as the basis for structural estimation and the use of QRE to interpret experimental data.

The model and results that we have developed provide a foundation for future theoretical and experimental research of social learning in networks and the techniques can be applied to other setups such as random graphs, and dynamic graphs where the set of neighbors observed changes over time. We can also use the same methodology to examine more complex network architectures. Obviously, different network architectures and information structures may lead to different outcomes. This remains a subject for further theoretical and experimental research.

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] Çelen, B. and S. Kariv (2004a) "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review*, 94(3), pp. 484-497.
- [6] Çelen, B. and S. Kariv (2004b), "An Experimental Test of Observational Learning under Imperfect Information." *Economic Theory*, forthcoming.
- [7] Choi, S., D. Gale and S. Kariv (2004), "Behavioral Aspects of Learning in Social Networks: An Experimental Study." Forthcoming in *Advances in Behavioral and Experimental Economics* (in the *Advances in Applied Microeconomics* series), John Morgan editor, JAI Press.
- [8] Duflo, E. and E. Saez (2002), "Participation and Investment Decisions in a Retirement Plan: the Influence of Colleagues' choices." *Journal of Public Economics*, 85, pp. 121-148.
- [9] Duflo, E. and E. Saez (2003), "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence From a Randomized Experiment." *Quarterly Journal of Economics*, 118, pp. 815-842.
- [10] Foster, A. and M. Rosenzweig (1995), "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture." *Journal of Political Economy*, 103(6), pp. 1176-1209.
- [11] Gale, D. and S. Kariv (2003) "Bayesian Learning in Social Networks." *Games and Economic Behavior*, 45(2), pp. 329-346.
- [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] 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. 1508-1520.
- [14] 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.
- [15] Kosfeld, M. (2004), "Economic Networks in the Laboratory: A Survey." *Review of Network Economics*, 3, pp. 20-41
- [16] Kübler, D. and G. Weizsäcker (2004) "Limited depth of reasoning and failure of cascade formation in the laboratory." *Review of Economic Studies*, Review of Economic Studies, 71, pp. 425-441.

- [17] Manski, C. (1993) "Identification of Exogenous Social Effects: The Reaction Problem." *Review of Economic Studies*, 60, pp. 531–542.
- [18] Manski, C. (1995), "Identification Problems in the Social Sciences." Cambridge: Harvard University Press.
- [19] McKelvey, R.D. and T.R. Palfrey (1995), "Quantal Response Equilibria for Extensive Form Games." *Games and Economic Behavior*, 10, pp. 6-38.
- [20] McKelvey, R.D. and T.R. Palfrey (1998), "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics*, 1, pp. 9-41.
- [21] Pagan, A. and A. Ullah (1999). *Nonparametric Econometrics*. Cambridge University Press.
- [22] Smith, L. and P. Sørensen (2000) "Pathological Outcomes of Observational Learning." *Econometrica*, 68(2), pp. 371-398.
- [23] Zheng, J.X. (1996), "A Consistent Test of Functional Form via Nonparametric Estimation Techniques", *Journal of Econometrics*, 75, pp. 263-289.

7 Appendix I

The Gale-Kariv model is extended to allow for the possibility of errors in the behavior of subjects, which leads us to the QRE version of the model. The QRE model assumes that agents receive idiosyncratic preference shocks. Formally, for agent i at turn $t = 1, 2, \dots, T$, the random utility from a binary action $a \in \{-1, 1\}$ is given by

$$U_{it}^a = \beta_{it} \pi_{it}^a + \varepsilon_{it}^a, \quad \text{for } a \in \{-1, 1\},$$

where π_{it}^a represents an observed (theoretical) expected payoff from action a and coefficient β_{it} parametrizes the sensitivity of choices to such observed expected payoffs. Random variable ε_{it}^a represents agent i 's preference shock for action a , which is assumed to be privately observed only by agent i . The choice probability for action $a \in \{-1, 1\}$ can be obtained, given agent i 's information set I_{it} at turn t , by

$$\begin{aligned} \Pr(a = 1 | I_{it}) &= \Pr\{U_{it}^1 > U_{it}^{-1}\} \\ &= \Pr\{\varepsilon_{it}^{-1} - \varepsilon_{it}^1 < \beta_{it} x_{it}\}, \end{aligned}$$

where $x_{it} := \pi_{it}^1 - \pi_{it}^{-1}$ denotes the difference in expected payoffs between action 1 and -1 given information set I_{it} ⁹. For tractability, we adopt a parametric version of the QRE model called the *logit-equilibrium* model where ε_{it}^a is assumed to be independently and identically distributed according to the type I extreme value with cumulative distribution $F(\varepsilon) = \exp(-e^{-\varepsilon})$ for any $a \in \{-1, 1\}$, each agent i and all $t = 1, 2, \dots, T$.¹⁰ The assumption of error structures implies no serial correlation of errors across turns for an individual and no cross-sectional correlation of errors across individuals.

In computing the expected payoffs at each turn, an agent uses different levels of hierarchies of beliefs to infer his neighbors' signals through their observed actions, depending on the structure of networks. The hierarchies of beliefs are mainly grouped into three categories: beliefs about (i) *his neighbor's private signal*, (ii) *his neighbor's trembling in actions*, and (iii) *his neighbor's neighbor's actions at previous turns* if hidden to the subject. To see how such hierarchies of beliefs are utilized in updating beliefs, consider type- i agent in a network for $i \in \{A, B, C\}$ with his neighbors $N_i \subset \{A, B, C\}$. At turn t , type- i agent's information set is given by $I_{it} = \left\{ \sigma_i, (a_{js})_{s=1}^{t-1} \mid j \in \overline{N}_i \right\}$,¹¹ where $\overline{N}_i := \{i\} \cup N_i$. The posterior belief that the true state is $\omega = 1$ conditional on I_{it} is given, via

⁹In our experimental setting, the difference in two payoffs is written as $x_{it} = 2[2\Pr(\theta = \mathbf{1} | I_{it}) - 1]$. Varying earnings does not change any qualitative results in the estimation of beta coefficients but only affects the scale of them.

¹⁰The variance of this distribution is normalized to be $\pi^2/6$. It is well known that beta coefficients and the variance in the error term can not be separately identified.

¹¹If $t < 2$, then the set $(a_{js})_{s=1}^{t-1}$ is empty. In general, we adopt the convention throughout this appendix that a set is considered as empty whenever it is the case.

Bayes' rule, by

$$\begin{aligned}\Pr(\omega = 1|I_{it}) &= \frac{\Pr(I_{it}|\omega = 1)}{\Pr(I_{it}|\omega = 1) + \Pr(I_{it}|\omega = -1)} \\ &= \frac{\Pr(\sigma_i|\omega = 1) \prod_{s=1}^{t-1} \prod_{j \in N_i} \Pr(a_{js}|I_{is}, \omega = 1)}{\sum_{\omega} \Pr(\sigma_i|\omega) \prod_{s=1}^{t-1} \prod_{j \in N_i} \Pr(a_{js}|I_{is}, \omega)},\end{aligned}$$

where the second equality comes from the assumptions of distributions of errors and signals. The formula says that an agent processes information by forming a belief about new observation at each turn given information set available up to that turn as well as conditional on the state of the world. Thus, the focus is, in what follows, on how to decompose such beliefs about new observation into various hierarchies of beliefs, depending on the structure of networks.

The Complete Network We only consider type- A 's inference problem, whose information set at $t = 1, 2, \dots, T$ is given by $I_{At} = \{\sigma_A, (a_{As}, a_{Bs}, a_{Cs})_{s=1, \dots, t-1}\}$, because of the homogeneity of the positions in the complete network. Due to the common knowledge of the history of the play, type- A agent only needs to infer his neighbors' private signals while considering the possibility of their errors. Thus, for instance, the belief about type- B 's action at turn t conditional on I_{At} and state ω is decomposed into

$$\begin{aligned}\Pr(a_{Bt}|I_{At}, \omega) &= \Pr(a_{Bt} | \{a_{As}, a_{Bs}, a_{Cs}\}_{s=1}^{t-1}, \omega) \\ &= \sum_{\sigma_B} \Pr(a_{Bt} | \sigma_B, \{a_{As}, a_{Bs}, a_{Cs}\}_{s=1}^{t-1}, \omega) \Pr(\sigma_B | \{a_{As}, a_{Bs}, a_{Cs}\}_{s=1}^{t-1}, \omega) \\ &= \sum_{\sigma_B} \Pr(a_{Bt} | I_{Bt}) \Pr(\sigma_B | \{a_{As}, a_{Bs}, a_{Cs}\}_{s=1}^{t-2}, a_{Bt-1}, \omega),\end{aligned}$$

where $\Pr(a_{Bt}|I_{Bt})$ contains different values of σ_B in the different summands. Note that the first term in each summand represents type- B 's choice probability, which is independent of the state of the world, and the second term does type- A 's belief about type- B 's signal conditional on relevant information and state ω . Each second term in the summands can be further decomposed into, for $t \geq 3$,

$$\begin{aligned}&\Pr(\sigma_B | \{a_{As}, a_{Bs}, a_{Cs}\}_{s=1}^{t-2}, a_{Bt-1}, \omega) \\ &= \frac{\Pr(a_{Bt-1} | I_{Bt-1}) \Pr(\sigma_B | \{a_{As}, a_{Bs}, a_{Cs}\}_{s=1}^{t-3}, a_{Bt-2}, \omega)}{\sum_{\sigma'_B} \Pr(a_{Bt-1} | I'_{Bt-1}) \Pr(\sigma'_B | \{a_{As}, a_{Bs}, a_{Cs}\}_{s=1}^{t-3}, a_{Bt-2}, \omega)}.\end{aligned}$$

The Circle Network Type- A 's information set at turn t is given by $I_{At} = \{\sigma_A, (a_{As}, a_{Bs})_{s=1, \dots, t-1}\}$. The inference problem becomes more interesting from $t \geq 3$ due to the lack of common knowledge of the history: type- A agent needs to consider type- C 's action in processing information from type- B 's

actions. For any $s \geq 2$,

$$\begin{aligned}
\Pr(a_{Bs}|I_{As}, \omega) &= \Pr\left(a_{Bs} | \{a_{Ap}, a_{Bp}\}_{p=1}^{s-2}, a_{Bs-1}, \omega\right) \\
&= \sum_{\sigma_B, (a_{Cm})_{m=1}^{s-1}} \Pr\left(a_{Bs} | \sigma_B, \{a_{Ap}, a_{Bp}, a_{Cp}\}_{p=1}^{s-2}, a_{Bs-1}, a_{Cs-1}, \omega\right) \\
&\times \Pr\left(\sigma_B | \{a_{Ap}, a_{Bp}, a_{Cp}\}_{p=1}^{s-2}, a_{Bs-1}, a_{Cs-1}, \omega\right) \\
&\times \prod_{k=1}^{s-1} \Pr\left(a_{Ck} | \{a_{Ap}, a_{Bp}\}_{p=1}^{s-2}, a_{Bs-1}, \{a_{Cl}\}_{l=1}^{k-1}, \omega\right) \\
&= \sum_{\sigma_B, (a_{Cm})_{m=1}^{s-1}} \Pr(a_{Bs}|I_{Bs}) \Pr\left(\sigma_B | \{a_{Bp}, a_{Cp}\}_{p=1}^{s-2}, a_{Bs-1}, \omega\right) \\
&\times \prod_{k=1}^{s-1} \Pr\left(a_{Ck} | \{a_{Ap}, a_{Cp}\}_{p=1}^{k-1}, \omega\right).
\end{aligned}$$

Note that type- A 's belief about new observation entails beliefs about type- C 's actions at all previous turns because they affect beliefs about type- B 's signal and trembling. And those beliefs are, further, decomposed into

$$\begin{aligned}
&\Pr\left(a_{Ck} | \{a_{Ap}, a_{Cp}\}_{p=1}^{k-1}, \omega\right) \\
&= \sum_{\sigma_C} \Pr(a_{Ck}|I_{Ck}) \Pr\left(\sigma_C | \{a_{Cp}\}_{p=1}^{k-1}, \{a_{Ap}\}_{p=1}^{k-2}, \omega\right).
\end{aligned}$$

The Star Network The interaction between heterogeneous agents in the star network has also a salient feature in updating beliefs. First, consider type- A agent who has the perfect knowledge over the history of the play. Just as agents in the complete network, type- A agent only needs to infer his neighbors' signals with the consideration of trembling. However, the nature of forming beliefs is quite different because both agents on the periphery can only interact through type- A agent. For any $s \geq 2$,

$$\begin{aligned}
\Pr(a_{Bs}|I_{As}, \omega) &= \Pr\left(a_{Bs} | \{a_{Ap}, a_{Bp}\}_{p=1}^{s-1}, \omega\right) \\
&= \sum_{\sigma_B} \Pr\left(a_{Bs} | \sigma_B, \{a_{Ap}, a_{Bp}\}_{p=1}^{s-1}, \omega\right) \Pr\left(\sigma_B | \{a_{Ap}, a_{Bp}\}_{p=1}^{s-1}, \omega\right) \\
&= \sum_{\sigma_B} \Pr(a_{Bs}|I_{Bs}) \Pr\left(\sigma_B | \{a_{Ap}, a_{Bp}\}_{p=1}^{s-2}, a_{Bs-1}, \omega\right).
\end{aligned}$$

Consider the inference problem of an agent on the periphery, for example, type- B . Just as agents in the circle network, type- B agent should consider the impact of type- C 's unobserved actions on type- A 's observed actions. But, the nature of inference is also different because his action does not directly influence

type- C 's decision problem: for any $s \geq 2$,

$$\begin{aligned}
\Pr(a_{As}|I_{Bs}, \omega) &= \Pr\left(a_{As} | \{a_{Ap}, a_{Bp}\}_{p=1}^{s-1}, \omega\right) \\
&= \sum_{\sigma_A, (a_{Cm})_{m=1}^{s-1}} \Pr(a_{As}|I_{As}) \Pr\left(\sigma_A | \{a_{Ap}, a_{Bp}, a_{Cp}\}_{p=1}^{s-2}, a_{As-1}, \omega\right) \\
&\quad \times \prod_{k=1}^{s-1} \Pr\left(a_{Ck} | \{a_{Ap}, a_{Cp}\}_{p=1}^{k-1}, \omega\right).
\end{aligned}$$

8 Appendix II

This appendix provides the experimental instructions for the circle network high-information treatment. All instructions used in the experiments are available upon request.

Introduction

This is an experiment in the economics of decision-making in networks. A research foundation has provided funds for conducting this research. Your earnings will depend partly on your decisions and partly on chance.

If you follow the instructions and make careful decisions, you may earn a considerable amount of money. At this point, take a minute to write down the number of the computer you are using as appears on the top of the monitor. At the end of the experiment, you should use your computer number to claim your earnings. Details of how you will make decisions will be provided below.

In this experiment, you will participate in 15 independent and identical (of the same form) rounds each divided into six decision-turns. In each round you will be asked to predict which of two urns has been randomly chosen.

In each round, you will be assigned to a position in a three-person network. You will only be able to observe the choices of the other participants to whom you are connected in this network.

Before the first round, you will be randomly assigned to one of the network positions labeled *A*, *B*, or *C*. A third of the participants in the room will be designated as type-*A* participants, one third as type-*B* participants and one third as type-*C* participants.

Your type depends solely upon chance and will remain constant in all rounds throughout the experiment. When you are called to make your first decision you will be informed of your type on the left hand side of the dialog window (was attached to the instructions). The network is given to you in the window that appears just to the right of it and is illustrated in the scheme below (similar to Figure 1 above). A line segment between any two types represents that they are connected and the arrowhead points to the participant whose action can be observed.

Note that in the network used in this experiment, each type can observe the choices of one other type – type-*A* participants observe type-*B*, type-*B* observe type-*C* and type-*C* observe type-*A*.

A decision round

Next, we will describe in detail the process that will be repeated in all 15 rounds. Each round starts by having the computer randomly form 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 depend solely upon chance and are independent of the networks formed in any of the other rounds. That is, in

any network each participant of type-*A* is equally likely to be chosen for that network, and similarly with participants of type-*B* and type-*C*.

In each round you will be asked to predict which of two urns, labeled *R* and *W*, has been randomly chosen. For each group, it is equally likely that urn *R* or urn *W* will be chosen. **Urn *R* contains 2 red balls, and 1 white ball. Urn *W* contains 1 red ball and 2 white balls.**

To help you determine which urn has been selected, with some chance you will be allowed to observe one ball, drawn at random, from the urn. **Each participant will observe a private draw with 1/3 chance and will not observe a private draw with 2/3 chance.**

Whether you observe a draw is randomly determined in each round and is independent of whether others in your group observe a draw and independent of whether you observe a draw in any of the other rounds. The result of your draw will be your private information and should not be shared with any of the other participants.

After each draw, the ball will be returned to the urn before making the next private draw. This is done by the computer.

Before you are called to make your first decision, you will be informed on the left hand side of the program dialog window whether the computer has randomly drawn a ball for you and if it is white or red.

The window will read “Your draw in this round is *w*” if the ball the computer has randomly drawn for you is white, “Your draw in this round is *r*” if the ball the computer has randomly drawn for you is red, and “Your draw in this round is *N/A*” if the ball the computer has not drawn a ball for you.

After everyone for whom the computer has randomly drawn a ball has seen his or her draw, each participant will be asked to input the letter of the urn, *W* or *R*, that he or she thinks is more likely to have been used.

When you are ready to make your decision, all you need to do is just use the mouse to click on the *W*-Button or *R*-Button on the lower left hand corner of the dialog window. After that, confirm your decision by clicking on the Submit button.

When everyone in your group has made a decision, you will observe the choices of the participants to whom you are connected by the network. For example, if you are a type-*A* participant, you will only be informed what letter the type-*B* participant has chosen. This information is given to you in the large window that appears at the bottom of the dialog window.

Note that each of the other participants will also observe the letters chosen by the participants to whom he or she is connected by the network. This completes the first of six decision-turns in this round.

Next, you will be called to make your second decision without observing a new draw from the urn. You will again be asked to click on the letter of the urn, *W* or *R*, that you think is more likely to have been used, and after everyone in your group has made his or her decisions, you will observe the choices in the second decision-turn of the participants to whom you are connected by the network. This process will be repeated until six decision turns are completed.

Note, again, that the urn used does not change in all decision-turns throughout a given round.

When the first round ends, the computer will inform everyone of the urn that was actually used (see attachment). After letting you observe the results of the first round, the second round will start by having the computer randomly forming new groups of participants in networks and selecting an urn for each group.

This process will be repeated until all the 15 independent and identical rounds are completed. At the end of the last round, you will be informed the experiment has ended.

Your earnings at each round are determined as follows. At the end of the round, the computer will randomly select one of the six decision-turns. Everyone whose choice in this decision-turn matches the letter of the urn that was actually used earns \$2. All others earn nothing.

The turn is selected in random to insure that at every decision-turn your choice is your best guess as to what the urn is. Hence, the best way to maximize your earnings is to give your best guess about which urn is chosen.

Your final earnings in the experiment will be the sum of your earnings over the 15 rounds. The decision-turn selected in each round, the actual urn used, your choice and your payoff for the round, and your cumulative payoff, will be recorded in the large window that appears at the center of the program dialog window.

Rules

Please do not talk with anyone during the experiment. We ask everyone to remain silent until the end of the last round.

Your participation in the experiment and any information about your earnings will be kept strictly confidential. Your payments receipt and participant form are the only places in which your name and social security number are recorded.

If there are no further questions, you are ready to start. An instructor will approach your desk and activate your program.