

Social Learning in Networks: A Quantal Response Equilibrium Analysis of Experimental Data*

Syngjoo Choi[†]
UCL

Douglas Gale[‡]
NYU

Shachar Kariv[§]
UC Berkeley

February 21, 2012

Abstract

Individuals living in society are bound together by a social network and, 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 individuals have different information, 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. This paper reports an experimental investigation

*The results reported here were previously distributed in a paper titled “Learning in Networks: An Experimental Study.” 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 thank Colin Camerer, Boğaçhan Çelen, Gary Charness, Xiaohong Chen, Charlie Holt, John Morgan, Tom Palfrey, Matthew Rabin, Andrew Schotter, and Georg Weizsäcker 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 the University of California, Berkeley (COR Grant).

[†]Department of Economics, University College London, Gower Street, London WC1E 6BT, UK (Email: syngjoo.choi@ucl.ac.uk, URL: <http://www.homepages.ucl.ac.uk/~uctpsc0>).

[‡]Department of Economics, New York University, 19 W. 4th Street, New York, NY, 10012 (E-mail: douglas.gale@nyu.edu, URL: <http://www.econ.nyu.edu/user/galed>).

[§]Department of Economics, University of California, Berkeley, 530 Evans Hall # 3880, Berkeley, CA 94720, USA (E-mail: kariv@berkeley.edu, URL: <http://emlab.berkeley.edu/~kariv/>).

of learning in three-person networks and uses 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. To test the usefulness of the theory in interpreting the data, we adapt the Quantal Response Equilibrium (QRE) model of McKelvey and Palfrey (1995, 1998). We find that the theory can account for the behavior observed in the laboratory in a variety of networks and informational settings. This provides important support for the use of QRE to interpret experimental data.

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

Key Words: Social networks, social learning, Quantal Response Equilibrium (QRE), experiment.

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 observes a subset of the network, the members with whom he is in direct contact, but his ability to observe the network is limited. He has *imperfect information* about the actions of most agents in the same network.

In the present paper, we analyze data from a series of laboratory experiments of social learning in networks. The experimental design is based on a model of Bayesian learning in social networks developed by Gale and Kariv (2003), henceforth GK; the experimental data is described in Choi, Gale and Kariv (2005), henceforth CGK. Even in the laboratory, the process of social learning in networks can be complicated. In incomplete networks, the absence of common knowledge requires subjects to interpret the information contained in the actions of others by forming higher order beliefs and performing complex calculations. It is not at all clear whether experimental subjects are able to draw correct, Bayesian inferences in this setting. For this reason, it is important to extend the GK model of Bayesian learning to allow for mistakes.

The Quantal Response Equilibrium (QRE) model of McKelvey and Palfrey (1995, 1998) allows 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 predictive power and approaches infinity if subjects are perfect Bayesians; for any finite value of the coefficient, there will be a positive probability that “mistakes” are made. The QRE model has two distinct advantages. First, the probability of mistakes depends on the payoff differences between actions, so mistakes are less likely when there is a lot at stake. Second, it is a consistent equilibrium theory in the sense that subjects’ responses take into account the mistakes of others.

We estimate a structural QRE model and find that subjects are highly rational in two senses. First, their behavior is predicted by a parsimoniously parameterized model. Second, their predicted behavior is highly sensitive to the correct (rational expectations) payoffs. We also test the model specification and find that the prediction of the QRE model—that mistakes are more likely when payoff differences are small—is confirmed by the data. Thus, both the rationality of the subjects’ behavior and the model’s qualitative predictions find strong support in the data.

The experiments reported in CGK involve three-person, connected social networks. Attention is restricted to connected networks since obviously disconnected agents cannot learn from others. The set of three-person networks contains several non-trivial architectures, each of which gives rise to its own distinctive learning patterns. The GK 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, where a line segment between any two types represents that they are connected and the arrowhead points to the subject whose action can be observed. Note that the links need not be symmetric: the fact that i can observe j does not necessarily imply that j can observe i .

[Figure 1 here]

Three representative networks are used in the experimental design: the *complete network*, in which each subject observes the actions chosen by all the other subjects; the *circle network*, in which each subject observes the actions chosen by exactly one other subject and each subject is observed by

someone; and the *star network*, in which one subject (the center) observes the other two subjects and the two (peripheral) subjects only observe the center. We chose these networks because they illustrate the main features of the complete set of networks. For practical purposes, these three networks “span” the set of networks—the excluded networks can each be obtained by adding a single link to one of the three chosen networks—and provide a reasonable test of the theory.

In the CGK experimental design, there are two equally likely events (states of nature). Subjects are 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 consists 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, uses a single network, a single information treatment and a single group of subjects.

The GK 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 at 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 QRE has a similar recursive structure that 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, sec-

ond, because it assumes that agents take into account the possibility that others are making mistakes when drawing inferences from their actions. 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).

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 GK model, agents can revise their actions 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 to think about his neighbors' knowledge of his neighbors' actions and the private information they revealed.

In the laboratory, lack of common knowledge forces subjects to think about hierarchies of beliefs. For example, in the circle network in which subject A observes subject B , subject B observes subject C , and subject C observes subject A . Subject 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. The fact that subjects' decisions increase in difficulty at each decision turn suggests that mistakes are more likely to occur, other things being equal, at later stages of the game. We find that the estimated coefficients of the QRE model decline at each decision turn, confirming the hypothesis that complexity leads to mistakes.

The rich data set generated by these experiments has been used to address a variety of important and interesting questions about individual and group behavior. A related paper, CGK, uses the same data set to investigate behavioral aspects of individual and group behavior, including comparisons across networks and information treatments. They show 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 average subject behavior among the three networks and three information treatments. These differences can be

explained by differences in the amount of common knowledge and the symmetry or asymmetry of the network or the information treatment. In this paper, we extend the analysis by providing a more systematic comparison of the theory with the data using structural methods.

In a paper published in this volume, Choi (2012) explores an *alternative* behavioral approach that can be brought to the interpretation of our data set. The econometric specification is based on a *type-mixture model* in which subjects randomly draw cognitive types from a common distribution. Each type corresponds to a different level of cognitive ability and determines how much information he can usefully process. Choi (2012) assumes further that the choice of each type is stochastic. This is done by combining the QRE approach with the cognitive hierarchy model (CH-QRE). In equilibrium, the value of information processing and the cognitive type distribution are *endogenously* determined. The estimation procedure finds the distribution of types that maximizes the likelihood of the empirical data. Choi (2012) finds that the proportion of rational types is very high and their behavior fits the predictions of the Gale-Kariv model quite well, taking into account the existence of boundedly rational subjects in the population. Choi (2012) compares the QRE model and CH-QRE model and discusses their goodness of fit.

The rest of the paper is organized as follows. The next section describes the related literatures. Section 3 describes the experimental design and procedures. Section 4 illustrates the main features of the Gale-Kariv model, and Section 5 provides the QRE analysis. Section 6 contains some concluding remarks. Technical details are gathered in Section 7.

2 Related literature

Our paper contributes to the large and growing body of work which studies the influence of the network structure on economic outcomes. Jackson (2008) provide a recent survey of the work in economics focusing on social and economic networks. Goyal (2005) and Jackson (2005) provide excellent, if now already somewhat dated, surveys of the theoretical work and Kosfeld (2004) surveys the experimental work. The preceding paper most closely related to GK 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. DeMarzo, Vayanos, and Zwiebel (2003) also assume boundedly rational agents. Rosenberg, Solan and Vieille (2009) and Mueller-Frank (2011) also examine the decisions of rational agents and extend the analysis of GK.

Whether agents 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. Griliches (1957) first studied the gradual adoption of a new agricultural technique, and observed that farmers learned from salespersons and from their neighbors. Among others, Foster and Rosenzweig (1995), Conley and Udry (2001) and Munshi (2005) examine how agents in developing countries learn from their social contacts in various contexts. 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.

The paper also contributes to a large literature on social learning. Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992), henceforth 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), Çelen and Kariv (2004, 2005), and Goeree, McKelvey, Palfrey, and Rogers (2007) extend Anderson and Holt (1997) to investigate other possible explanations for informational cascades.

Weizsäcker (2010) presents a meta data analysis of 13 different experiments of the BHW model. This large and growing body of experimental work in the social learning literature has also successfully utilized QRE models. 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.

3 Design and procedures

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, the instructions were read aloud by an experimental administrator.¹ Sample experimental instructions are reproduced in Online Appendix I.² 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 behavior.³

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 networks are illustrated in Figure 1 above. 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 subject's type, *A*, *B*, or *C*, remained constant throughout

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.

²Online Appendix I: http://emlab.berkeley.edu/~kariv/CGK_I_A1.pdf.

³Subjects' 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.

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 subject 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 subjects in networks and selecting an urn for each group. This process was repeated until all the 15 rounds were completed.

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. More importantly, we use a variety of different network architectures and information treatments to generate a variety of different outcomes which are representative of the theory. 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 Some theory

In this section we discuss briefly the theoretical implications of the model tested in the laboratory. GK provide an extensive analysis of a general version of the model.

4.1 The game

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

$$I_{it} = \left\{ \sigma_i, (a_{js})_{s=1}^{t-1} \mid j \in N_i \right\}.$$

4.2 Equilibrium

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 equilibrium is fully revealing under the tie-breaking assumption that agents switch actions whenever they are indifferent between choosing the same action in the next period and switching to the other action. 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 *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 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 zero probability events.

GK 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.

4.3 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 chooses $a_{i1} = \sigma_i$. 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 (T is large) C will eventually switch. This conclusion depends on our tie-breaking rule that indifferent uninformed agents randomize.

4.4 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.

4.5 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.

4.6 Summary

The preceding examples have illustrated several features of the theory:

- 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.
- 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 ($q = 1$), whereas the existence of asymmetrically informed agents ($q < 1$) 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 ($q = 1$).

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.

5 Quantal Response Equilibrium (QRE)

5.1 Specification

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. To this end, we extend the Gale-Kariv model 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. The 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 $a \in \{-1, 1\}$, conditional on agent i 's information set I_{it} at turn t , is given 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

$$I_{it} = \left\{ \sigma_i, (a_{js})_{s=1}^{t-1} \mid j \in N_i \right\}.$$

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 with cumulative distribution $F(\varepsilon) = \exp(-e^{-\varepsilon})$ for any $a \in \{-1, 1\}$, each agent i and all $t = 1, 2, \dots, T$.⁴ The assumption about error structures implies no serial correlation of errors across turns for an agent and no correlation of errors across agents.

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

⁴The 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.

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 from the subject. To see how such hierarchies of beliefs are utilized in updating beliefs, consider a type-*i* agent, for $i \in \{A, B, C\}$, with neighbors $N_i \subset \{A, B, C\}$. The posterior belief that the true state is $\omega = 1$ conditional on I_{it} is given, via 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 on the distributions of errors and signals. The formula says that an agent processes information by forming a belief about new observation at each turn given the information set available up to that turn as well as conditional on the state of the world.

In the QRE model, a rational agent must predict his neighbors' choice probabilities correctly to calculate the posterior probabilities correctly. In effect, we assume that agents 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 agents 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 a logistic response function. The logit equilibrium can thus 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 $a_{it} = 1$ and $a_{it} = -1$, 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} .

5.2 Estimation

We use repeatedly the standard maximum likelihood (ML) method for the estimation of the logistic random-utility models. The data employed to im-

plement 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 vari-

able 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 Section 7.

5.3 Results

Table 1 below presents the results of the ML logit equilibrium estimation. Standard errors are given in parentheses. 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 estimated 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.

[Table 1 here]

Figure 2 provides a graphical re-presentation of the beta series in the complete, star (type A and types B and C) and circle networks. Figure 2 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.

[Figure 2 here]

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. The 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 $u \in \mathbb{R}$. 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.

Online Appendix II presents a set of comparisons between these two choice probabilities.⁶ In each of the figures, a solid line represents the non-parametrically 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

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).

⁶Online Appendix II: http://emlab.berkeley.edu/~kariv/CGK_I_A2.pdf. The figures are difficult to see in the small black and white format required in the printed version.

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 by a combination of the small-sample problem and one subject’s “irrational” behavior.⁷

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 that, without a specific alternative model, 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. The test statistic T_n is given by

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

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).

The results of the series of specification tests are reported in Table 2. The bandwidth is selected to be $cn^{-1/5}$, where c is equal to 1. The test results in Table 2 confirm the previous graphical comparisons. In most of the cases p -values (reported in parentheses) 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

⁷The 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.

at decision turns $t = 3, 5, 6$ with a 5% significance level. Interestingly, we also reject the null at the third turn, in the high-information, circle-network treatment, with a 5% significance level.⁸

[Table 2 here]

6 Conclusion

Many economic decision problems involve incomplete and asymmetric information. That is, agents are uncertain about some underlying decision-relevant 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 GK to interpret the data generated by the experiments. 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 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. Thus, the experimental design offers an elegant setting for further experimental

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.

investigation of social learning in networks. Perhaps, the most important subject for future research is to identify the impact of network architecture on the efficiency and dynamics of social learning: How does architecture affect the dynamics of social learning? What architectures facilitate or hinder the choice of an optimal action? Obviously, different network architectures and information structures may lead to different outcomes.

7 Technical appendix

7.1 The Complete Network

We only consider the inference problem of a type- A agent, 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 symmetry of the complete network. Due to the common knowledge of the history of play, a 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

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

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 represents type A 's belief about type B 's signal conditional on relevant information and state ω . The second term in each summand can be further decomposed into

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

for $t \geq 3$.

7.2 The Star Network

The interaction between heterogeneous agents in the star network has also a salient feature in updating beliefs. First, consider a type- A agent who has the perfect knowledge over the history of play. Just like agents in the complete network, a type- A agent only needs to infer his neighbors' signals, taking into account the probability of trembles. However, the formation of beliefs is different because the peripheral agents can only interact through the 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 like agents in the circle network, a type- B agent should consider the impact of a type C 's unobserved actions on a type A 's observed actions. But, the nature of inference is also different because his action does not directly influence a type C 's decision problem: for any $s \geq 2$,

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

7.3 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 for $t \geq 3$ due to the lack of common knowledge of the history: a type- A agent needs to consider type C 's action in processing information from type B 's actions. For any $s \geq 2$,

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

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

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

References


- [1] Anderson, L. and C. Holt (1997), "Information Cascades in the Laboratory." *American Economic Review*, 87, 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, 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, pp. 992-1026.
- [5] Çelen, B. and S. Kariv (2004), “Distinguishing Informational Cascades from Herd Behavior in the Laboratory.” *American Economic Review*, 94, pp. 484-497.
- [6] Çelen, B. and S. Kariv (2005), “An Experimental Test of Observational Learning under Imperfect Information.” *Economic Theory*, 26, pp. 677-699.
- [7] Choi, S. (2012), “A Cognitive Hierarchy Model of Learning in Networks.” *Review of Economic Design*, forthcoming.
- [8] Choi, S., D. Gale and S. Kariv (2005), “Behavioral Aspects of Learning in Social Networks: An Experimental Study.” In *Advances in Behavioral and Experimental Economics* (the Advances in Applied Microeconomics series), John Morgan editor, JAI Press.
- [9] Conley, T. and C. Udry (2001), “Social learning through networks: the adoption of new agricultural technologies in Ghana.” *American Journal of Agricultural Economics*, 83, pp. 668-673.
- [10] DeMarzo, P., D. Vayanos and J. Zwiebel (2003), “Persuasion Bias, Social Influence, and Unidimensional Opinions.” *The Quarterly Journal of Economics*, 118, pp. 909-968.
- [11] 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.
- [12] 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.
- [13] 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, pp. 1176-1209.
- [14] Gale, D. and S. Kariv (2003), “Bayesian Learning in Social Networks.” *Games and Economic Behavior*, 45, pp. 329-346.

- [15] Goeree, J., R. McKelvey, T. Palfrey, and B. Rogers (2007), "Self-Correcting Information Cascades." *Review of Economic Studies*, 74, pp. 733-762.
- [16] Goyal S. (2005), "Learning in Networks: A Survey." In *Group Formation in Economics; Networks, Clubs and Coalitions*, ed. G. Demange and M. Wooders. Cambridge University Press.
- [17] Griliches, Z. (1957), "Hybrid corn: an exploration in the economics of technological change." *Econometrica*, 25, pp. 501-522.
- [18] Hung, A. and C. Plott (2001), "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions." *American Economic Review*, 91, pp. 1508-1520.
- [19] Jackson, M. (2005), "A Survey of Models of Network Formation: Stability and Efficiency." In *Group Formation in Economics: Networks, Clubs, and Coalitions* edited by G. Demange and M. Wooders. Cambridge University Press.
- [20] Jackson, M. (2008), *Social and Economic Networks*. Princeton University Press.
- [21] Kosfeld, M. (2004), "Economic Networks in the Laboratory: A Survey." *Review of Network Economics*, 3, pp. 20-41.
- [22] Kübler, D. and G. Weizsäcker (2003), "Limited depth of reasoning and failure of cascade formation in the laboratory." *Review of Economic Studies*, 71, pp. 425-441.
- [23] Manski, C. (1993) "Identification of Exogenous Social Effects: The Reaction Problem." *Review of Economic Studies*, 60, pp. 531-542.
- [24] Manski, C. (1995) "Identification Problems in the Social Sciences." Cambridge: Harvard University Press.
- [25] McKelvey, R. and T. Palfrey (1995), "Quantal Response Equilibria for Extensive Form Games." *Games and Economic Behavior*, 10, pp. 6-38.
- [26] McKelvey, R. and T. Palfrey (1998), "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics*, 1, pp. 9-41.
- [27] Mueller-Frank, M. (2011), "A General Framework for Rational Learning in Social Networks." *Theoretical Economics*, forthcoming.

- [28] Munshi, K. (2004), “Social learning in a heterogeneous population: technology diffusion in the Indian green revolution.” *Journal of Development Economics*, 73, pp. 185-213.
- [29] Pagan, A. and A. Ullah (1999), *Nonparametric Econometrics*. Cambridge University Press.
- [30] Rosenberg D., E. Solan and N. Vieille (2009), “Informational externalities and emergence of consensus.” *Games and Economic Behavior*, 66, pp. 979-994.
- [31] Smith, L. and P. Sørensen (2000), “Pathological Outcomes of Observational Learning.” *Econometrica*, 68, pp. 371-398.
- [32] Weizsäcker, G. (2010), “Do we follow others when we should? A simple test of rational expectations,” *American Economic Review*, 100, pp. 2340-2360.
- [33] Zheng, J.X. (1996), “A Consistent Test of Functional Form via Nonparametric Estimation Techniques”, *Journal of Econometrics*, 75, pp. 263-289.


Figure 1: Three-person connected networks




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

Figure 2: The beta time-series in the each network


Complete



Circle



Star type-A



Star type-B and -C




Table 1: ML estimates.

Turn 1 - 4.171 (0.160)

Full-information

Turn	Complete	Star (type A)	Star (types B & C)	Circle
2	2.43 (0.22)		2.78 (0.28)	
3	2.73 (0.28)	1.23 (0.27)	3.93 (0.55)	2.19 (0.26)
4	2.60 (0.27)	2.17 (0.39)	3.06 (0.42)	1.84 (0.22)
5	2.58 (0.27)	1.31 (0.27)	2.90 (0.38)	2.23 (0.25)
6	2.74 (0.28)	1.29 (0.26)	3.17 (0.42)	2.34 (0.26)
# of obs.	270	90	180	270

High-information

Turn	Complete	Star (type A)	Star (types B & C)	Circle
2	2.42 (0.27)		2.84 (0.25)	
3	2.56 (0.31)	3.12 (0.66)	2.93 (0.41)	2.72 (0.32)
4	2.59 (0.31)	2.85 (0.59)	2.09 (0.29)	2.20 (0.26)
5	2.59 (0.31)	1.94 (0.41)	2.92 (0.39)	2.71 (0.31)
6	2.33 (0.28)	2.17 (0.44)	1.50 (0.23)	1.74 (0.21)
# of obs.	225	90	180	270

Low-information

Turn	Complete	Star (type A)	Star (types B & C)	Circle
1	2.83 (0.35)		3.19 (0.36)	
2	2.98 (0.40)	2.42 (0.59)	2.70 (0.48)	4.30 (0.54)
3	2.21 (0.32)	3.14 (0.66)	2.73 (0.45)	3.42 (0.44)
4	2.11 (0.30)	1.56 (0.47)	2.14 (0.39)	3.38 (0.43)
5	2.73 (0.34)	2.97 (0.63)	2.13 (0.38)	2.53 (0.34)
6	270	90	180	225
# of obs.	270	90	180	225

Standard errors are given in parentheses.

of obs. - the number of individual decisions per type and turn.

Table 2: Specification tests

Full-information

Turn	Complete	Star (type A)	Star (types B & C)	Circle
2	0.83 (0.409)		-0.47 (0.640)	
3	-0.46 (0.644)	3.37 (0.001)*	-0.08 (0.935)	-0.13 (0.896)
4	0.01 (0.993)	1.42 (0.154)	0.94 (0.345)	0.73 (0.466)
5	-0.67 (0.504)	4.53 (0.000)*	0.14 (0.888)	-0.68 (0.494)
6	-0.31 (0.756)	6.96 (0.000)*	-0.62 (0.538)	-0.78 (0.435)

High-information

Turn	Complete	Star (type A)	Star (types B & C)	Circle
2	-1.12 (0.264)		0.94 (0.346)	
3	-0.55 (0.580)	-0.21 (0.837)	-0.86 (0.388)	2.10 (0.036)*
4	-0.32 (0.746)	-0.34 (0.737)	-0.63 (0.531)	-0.47 (0.638)
5	0.72 (0.475)	0.11 (0.909)	-0.18 (0.856)	1.01 (0.312)
6	-0.40 (0.693)	-0.35 (0.727)	-0.03 (0.979)	-0.02 (0.986)

Low-information

Turn	Complete	Star (type A)	Star (types B & C)	Circle
2	-0.94 (0.345)		0.33 (0.743)	
3	-0.91 (0.363)	0.27 (0.784)	-0.59 (0.557)	1.69 (0.091)**
4	-0.91 (0.364)	0.98 (0.329)	-0.94 (0.350)	1.87 (0.061)**
5	-0.09 (0.929)	0.14 (0.886)	-0.12 (0.908)	-0.12 (0.902)
6	-0.14 (0.886)	-0.80 (0.424)	-0.74 (0.459)	-0.37 (0.714)

* The null hypothesis can be rejected with 5% significance level.

** The null hypothesis can be rejected with 10% significance level.