

A Cognitive Hierarchy Model of Learning in Networks*

Syngjoo Choi[†]

University College London

November 15, 2006

Abstract

This paper proposes a method for estimating a hierarchical model of bounded rationality in games of learning in networks. A cognitive hierarchy comprises a set of cognitive types whose behavior ranges from random to substantively rational. Specifically, each cognitive type in the model corresponds to the number of periods in which economic agents process new information. Using experimental data, we estimate type distributions in a variety of task environments and show how estimated distributions depend on the structural properties of the environments. The estimation results identify significant levels of behavioral heterogeneity in the experimental data and overall confirm comparative static conjectures on type distributions across task environments. Surprisingly, the model replicates the aggregate patterns of the behavior in the data quite well. Finally, we found that the dominant type in the data is closely related to Bayes-rational behavior.

Journal of Economic Literature Classification Numbers: C51, C92, D82, D83.

KeyWords: Cognitive hierarchy, Bounded rationality, Social learning, Social networks

“Human rational behavior ... is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor”, Herbert A. Simon (1990).

1 Introduction

A major controversy in economics and psychology surrounds the rationality of human behavior. While substantive rationality is considered unrealistic by many if not most economists, there ap-

*I thank Colin Camerer, Andrew Caplin, Xiaohong Chen, Guillaume Fréchet, Shachar Kariv, Donghoon Lee, and Andrew Schotter for comments and suggestions. I am especial grateful to Douglas Gale for his continued and invaluable discussion and advice. This paper has also benefitted from suggestions by the participants of seminars at National University of Singapore, New York University, University College London, University Western Ontario, Tinbergen Institute at Amsterdam, International ESA 2005 in Montreal, and World Cogress Econometric Society 2005 in London.

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

appears to be little consensus on how to model bounded rationality. As Simon’s scissors metaphor suggests, any serious attempt to study human rationality, either theoretically or empirically, should, on one hand, provide a variety of task environments and, on the other hand, be flexible enough to accommodate the full spectrum of economic agents’ cognitive abilities.

The existence of a vast diversity of cognitive abilities is a fact of life. Recent experimental studies on beauty-contest games¹ present well-documented evidence that individuals inside and outside the laboratory exhibit different levels of sophistication. A natural way of capturing such diversity is through a *cognitive hierarchy* consisting of a set of cognitive types whose behavior ranges from random to substantively rational.

The objective of this paper is to provide a method for estimating a hierarchical model of bounded rationality. Conceptually, we think of the distribution of cognitive types as being determined in an evolutionary equilibrium by the marginal costs and benefits of increasing one’s cognitive type. Using experimental data, we estimate the type distributions found in a variety of task environments and then show how estimated distributions depend on the structural properties of the task environments.

The opportunity to estimate a cognitive hierarchy model is provided by a rich data set produced in the laboratory by Choi, Gale and Kariv (2005, 2006; henceforth CGK05, CGK06). The data set has a couple of advantages from our point of view. First, the experimental design contains a number of different tasks that vary systematically in terms of their cognitive difficulty. Secondly, the data contain a large number of individual decisions of different levels of difficulty for each subject, thus allowing us to identify both the individual decision rules and the heterogeneity of those rules.

The experiments performed by CGK05 consisted of various games of social learning in networks. The following example is typical. Consider a network consisting of three players, A , B , and C . The network is *star-shaped*, with A at the center of the star and B and C at the periphery. A is linked to B and C and B and C are linked to A but not to each other. The network governs each player’s information flow: a player observes another player’s choices if and only if he is linked to that player by the network. Thus, in the star network, A observes the behavior of B and C , and B and C can observe the behavior of A , but B cannot observe C and C cannot observe B (See Figure 1).

At the beginning of the game, Nature randomly selects (i.e., with equal probabilities) one of two states of the world. With probability q each player independently receives a binary private signal that is correlated with the unknown state. With probability $1 - q$, no information is received. Time is divided into a finite number of periods, $t = 1, 2, \dots, T$. In each period, players

¹As part of testing the solution concept of dominance in game theory, many researchers have experimentally tested beauty-contest games and related behavioral heterogeneity in the laboratory to various cognitive limitations. To name a few, Nagel (1995), Ho, Camerer and Weigelt (1998), and Costa-Gomes and Crawford (2004). Camerer (2003) provided an excellent survey on this literature.

simultaneously choose which state is more likely to have occurred at the beginning of the game. Their payoffs, which are not revealed until the end of the game, are assumed to be an increasing function of the number of correct guesses.

This game has a very nice structure, in the sense that the cognitive difficulty of the players' decision tasks increases as the game proceeds. In period 1, a player bases his decision on his private signal (if he has one) or his prior (if he does not get a signal). In the second period, A observes the actions of B and C in the first period, whereas B and C can only observe A 's first-period action. Before making a decision in the second period, each player has to think about what information his neighbor's action reveals and update his beliefs accordingly. Then at the beginning of the third period each player observes what his neighbors did in the second period and reflects on what that implies about his neighbor's neighbor's action in the first period. For example, in the third period B should notice that A 's action in the second period is based on what B and C did in the first period as well as A 's private signal (if any). Similarly, A should know that B 's action at the second period is based on A 's action in period 1 as well as B 's private signal (if any). In each period, the inference problem becomes more demanding because it requires a player to consider higher order beliefs. In this sense, the sequence of tasks constitutes a cognitive hierarchy.

This cognitive hierarchy of *tasks* suggests a natural hierarchy of cognitive *types*. The lowest cognitive type would be someone who randomly guesses the state of nature, without processing any information. The second lowest type would be someone who could process his signal in period 1 but make no use of information obtained from observing his neighbor(s) at periods 2, 3, ..., T . The next lowest type would be someone who could process his signal in period 1 and make an inference about his neighbor(s)' signal(s) in period 2, but could not make any higher order inferences in periods 3, 4, ..., T . Note that each of these types corresponds to the number of periods in which the player processes new information and that is exactly how we shall define the cognitive hierarchy. It consists of T types $\tau = 0, \dots, T - 1$ where type τ processes information in the first τ periods and learns nothing from information received in the last $T - \tau - 1$ periods. The details of the type's information processing problem will be made clear later.

Higher types will clearly do at least as well as lower types since they have at least as much information. At the same time, since there is a fixed amount of information in the network (the private signals observed by the different players at the beginning of the game) we expect that the marginal benefit of processing information will decline over time. So the gain from being type $\tau + 1$ rather than type τ will be decreasing as τ increases. Depending on the cognitive costs of implementing a higher type of strategy, we may expect the tradeoff between costs and benefits would lead to a distribution of types concentrated on the lower types. This is an empirical question, of course, and it is this question we hope to answer.

Our approach of modeling a hierarchy of cognitive types can be illustrated as follows. A

group of agents are randomly selected from the population to play a game of social learning. Before playing a game, each of them compares the benefits and costs of being each cognitive type. The benefit of processing information in the first τ periods, V_τ , is captured by the sum of *ex ante* expected increments of maximized expected utilities over all time periods. And the cost of processing information in the first τ periods, c_τ , is related to individual-level cognitive costs. Thus, depending on the level of cognitive cost, different agents adopt different cognitive types whose behavior ranges from random to substantively rational. In equilibrium, the benefits of information processing V_τ and the type distribution π are endogenously determined. Since the different task environments relate to different values of V_τ , the equilibrium type distribution differs in different task environments.

The notion of optimality that underlies our approach raises an interesting question. How can agents choose decision rules “optimally,” taking into account the corresponding computation costs and value of information, without having solved the decision problem corresponding to each cognitive type? We adopt an evolutionary point of view, in which evolution or learning through lifetime experiences is assumed to reveal to agents the expected benefits associated with different cognitive types without the necessity of solving a formal optimization problem. Given this point of view, it is not unreasonable to think that individuals of different levels of cognitive ability will be led to choose the decision rules or cognitive types that are optimal for them. The specification of types from which the evolving subjects are assumed to choose is somewhat ad hoc, but without some such assumptions, it is impossible to open the black box that determines the relation between bounded rationality and task complexity, still less confronts the empirical data and suggests cross-environment comparisons.

In the cognitive hierarchy model an equilibrium distribution over cognitive types depends on the structure of the task environment. As the information parameter q and the network architecture change, the cognitive difficulty of the different tasks and the value of information, in other words, the benefits of being a higher type, will vary. For instance, when the value of information parameter q decreases, the marginal benefit of information processing in the first two periods becomes uniformly lower, holding constant the network structure, and it usually takes more periods of information processing to extract all valuable information. We discuss a number of comparative static conjectures in a later section and test them in the data.

Other approaches to modeling a cognitive hierarchy can be found in the literature on heterogeneity of strategic behavior in normal-form games such as beauty-contest games (e.g., see Camerer, Ho and Chong (2004)). This literature assumes that each agent believes that the rest of the population consists of agents who are less sophisticated than he is. Suppose one wants to analyze a one-shot, normal form game. A hierarchy can start from type 0 players who choose a strategy randomly. Type 1 players believe that all other players are type 0 and choose a best response to the distribution of type 0 strategies. Type 2 players believe that all other players

belong to types 0 and 1 and choose a best response to the distribution of those types. And so on. The highest type will be close to choosing a best response to the entire population and thus is approximately substantively rational. However, the beliefs of most players are not consistent with the actual type distribution or with the actual play of the game. Our notion of a cognitive hierarchy is somewhat different. While the hierarchy we use could be justified by the belief that each type of agent believes that the rest of the agents have a lower type, in practice we assume that agents are responding to the true type distribution (including higher types) and the true distribution of strategies. Another difference arises from the fact that we study a dynamic game which involves learning over time. Because we choose to identify types with the number of periods of information processing, this imposes a very particular structure on the strategies chosen. It also makes it possible to identify the types from the data generated over time, whereas in a one shot game it is very hard to distinguish a player's type from a simple mistake.

The approach we adopt here in order to explain behavioral heterogeneity in the data leads us to a natural specification for econometric analysis. The specification is a type-mixture model in which subjects randomly draw cognitive types from a common distribution. Then the estimation procedure searches for a type distribution that maximizes the likelihood of the empirical data.

This structural model and its empirical application provide new insights into the relationship between bounded rationality and task complexity. Among the main results of the paper are the following.

- First, we propose a hierarchical model of bounded rationality in which the equilibrium type distribution depends on the task environments.
- Secondly, the cognitive hierarchy model provides a set of testable conjectures regarding cross-environment comparisons of type distributions. The estimation results overall confirm the conjectures, suggesting the structural approach is a good tool for understanding behavior in the laboratory.
- Thirdly, we found that the dominant cognitive type across all treatments is closely related to Bayes-rational behavior. Despite the presence of multiple cognitive types in the subject pool, this finding tells us that the Bayesian paradigm has considerable explanatory power as far as the majority of subjects is concerned.
- Finally, we perform a goodness-of-fit test for the cognitive hierarchy model. Surprisingly, the cognitive hierarchy model replicates the empirical patterns of average herd behavior across networks and information treatments at the aggregate level. This sheds light on the potential importance of structural approaches allowing individual heterogeneity in explaining even aggregate patterns of data.

The rest of the paper is organized as follows. The next section introduces the model of

cognitive hierarchy and presents some theoretical predictions. Section 3 illustrates the experimental data briefly, presents the econometric method for estimating the CH model and reports the estimation results. Sections 4 and 5 include discussion and conclusions.

2 The Cognitive Hierarchy Model

In this section we sketch a hierarchical model of bounded rationality in the three-person networks studied by CGK05. 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 a set of neighbors for agent i . The neighborhood structure $\mathcal{G} = \{N_A, N_B, N_C\}$ defines a three-person network. The span of three-person, connected² networks is illustrated by Figure 1.

[Insert Figure 1]

There are two equally likely states of nature denoted by $\omega \in \Omega = \{-1, 1\}$. With probability q , an agent is *informed* and receives a private signal at the beginning of the game. Private signals take two values $\sigma = -1, 1$ and the probability that the signal σ equals the true state ω is $2/3$. For notational convenience, we assume that an uninformed agent receives a signal $\sigma = 0$. Agents' signals are assumed to be independently and identically distributed conditional on the true state.

Time is divided into a finite number of periods, indexed by $t = 1, 2, \dots, T$. At each period t , agents simultaneously choose the state they think is more likely to be the true state, denoted by $a_t \in A = \{-1, 1\}$. Agents have a common payoff function at each period, for $0 < M < \infty$,

$$u(a, \omega) = \begin{cases} M & \text{if } a = \omega \\ 0 & \text{otherwise} \end{cases}.$$

The true state is realized at the end of the final period T and the total payoff in the game is the sum of per-period payoffs. We assume that agents maximize per-period payoffs at each time period t conditional on available information.

2.1 Equilibrium

There are a large number of agents in the population who differ in their (constant) unit costs of cognition at each period, c , where $0 \leq c \leq \bar{c} < \infty$. Let F denote a cumulative distribution of cognitive costs on the interval $[0, \bar{c}]$, which is assumed to be common knowledge among agents. Agents with heterogeneous costs are randomly selected from the population to play a social learning game.

²A network is connected if and only for any pair of two agents (i, j) there is a path from agent i and agent j .

Let $I_{it} = \left\{ \sigma_i, \left((a_{js})_{s=1}^{t-1} \right)_{j \in N_i}, \pi \right\}$ denote the information set at location i in time period t and $\mathcal{F}(I_{it})$ be a σ -field generated by that information set. Let π denote a distribution over a set of cognitive types in which $\pi(\tau)$ is the relative frequency in the population of type τ which processes information in the first τ periods only. When information in the first t periods is processed in location i of a network \mathcal{G} with information parameter q , the *ex ante* expected increment of maximized expected utility at time period t is defined by

$$\begin{aligned} v_t((\mathcal{G}, i, q), \pi) &= E_0 \left[\max_{a \in A} E[u(a, \omega) | \mathcal{F}_{it}] \right] - \max_{a \in A} E_0 u(a, \omega) \\ &= M \left[\sum_{I_t} \Pr(I_t) \left\{ \max \{ \Pr(\omega = 1 | I_t), \Pr(\omega = -1 | I_t) \} - \frac{1}{2} \right\} \right], \end{aligned}$$

where $E_0(\cdot)$ denotes the expectation with respect to prior information over the states and thus $\max_{a \in A} E_0 u(a, \omega) = M/2$. Note that $v_t((\mathcal{G}, i, q), \pi) = 0$ for $t = 0$.

Then the value of processing information in the first τ periods, called *value of information for type τ* , is defined by the sum of *ex ante* expected increments of maximized expected utilities over all time periods:

$$V_\tau((\mathcal{G}, i, q), \pi) = \sum_{t=0}^{\tau} v_t((\mathcal{G}, i, q), \pi) + (T - \tau) v_\tau((\mathcal{G}, i, q), \pi).$$

An equilibrium, defined below, maps the distribution of cognitive costs F into the distribution of cognitive types π . Let $\mathbf{1}\{\cdot\}$ denote an indicator function.

Definition 1 *A weak perfect Bayesian equilibrium consists of the distribution π over cognitive types and a sequence of random variables $\{X_{it}\}$ and σ -fields $\{\mathcal{F}_{it}\}$ for an agent with cognitive cost c such that*

$$\begin{aligned} (i) \quad & \tau_c \in \arg \max_s \{V_s((\mathcal{G}, i, q), \pi) - s \cdot c\}, \\ (ii) \quad & \mathcal{F}_{it} = \begin{cases} \mathcal{F}_0 & \text{if } \tau_c = 0 \\ \mathcal{F} \left(\left\{ \sigma_i, \left((X_{js})_{s=1}^{t-1} \right)_{j \in N_i}, \pi \right\} \right) & \text{if } t \leq \tau_c \\ \mathcal{F} \left(\left\{ \sigma_i, \left((X_{js})_{s=1}^{\tau_c-1} \right)_{j \in N_i}, \pi \right\} \right) & \text{otherwise} \end{cases}, \\ (iii) \quad & X_{it} : \Omega \rightarrow A \text{ is } \mathcal{F}_{it} \text{ - measurable,} \end{aligned}$$

$$\begin{aligned} (iv) \quad & E[u(X_{it}(\omega), \omega) | \mathcal{F}_{it}] \geq E[u(x(\omega), \omega) | \mathcal{F}_{it}], \\ & \text{for any } \mathcal{F}_{it} \text{ - measurable function } x : \Omega \rightarrow A, \end{aligned}$$

and

$$(v) \quad \pi(\tau) = \int_0^{\bar{c}} \mathbf{1}\{\tau_c = \tau\} dF, \text{ for } \tau = 0, 1, \dots, T.$$

Note that the equilibrium concept requires agents to have rational expectations about the distribution of cognitive types and the value of information for each type is computed on the equilibrium path. The value of information depends on the network architecture and information parameter q as well as the distribution of cognitive types. It is thus fundamental to understand the relationship between the value of information on the one hand and the network architecture and information parameter on the other, in order to make comparative static conjectures on the relation between the distribution of cognitive types and the task environment.

2.2 Value of information

Let $MV_\tau((\mathcal{G}, i, q), \pi)$ denote the marginal gain in the value of information of being type τ rather than being type $\tau - 1$, for $\tau = 1, 2, \dots, T$. Then

$$\begin{aligned} MV_\tau((\mathcal{G}, i, q), \pi) &= V_\tau((\mathcal{G}, i, q), \pi) - V_{\tau-1}((\mathcal{G}, i, q), \pi) \\ &= (T - \tau + 1) [v_\tau((\mathcal{G}, i, q), \pi) - v_{\tau-1}((\mathcal{G}, i, q), \pi)]. \end{aligned}$$

We call $MV_\tau((\mathcal{G}, i, q), \pi)$ the *marginal value of information for type τ* .

As a benchmark, we compute the marginal value of information in the model where all agents are Bayes-rational and there is no cognitive cost. Because of the analytical difficulty of computing the marginal values, we rely on simulations with many replications of independent games³. Table 1 presents the simulated marginal values of information across time periods for three networks and information parameter $q = 1, 2/3, 1/3$, which were used in CGK05. The payoff for a correct action is chosen to be $M = \$2$. When individuals are indifferent between two actions, we use randomization with equal probabilities as the tie-breaking rule in simulations throughout the paper.

[Insert Table 1]

Let us first take a look at the time series patterns of marginal values of information for each network and each information parameter. The marginal values decline as time period increases and eventually they become zero. This simply reflects the fact that learning eventually ends and there is no valuable information from that point onwards. The marginal values decrease monotonically in the complete network and the center of the star network, whereas they do not necessarily decline monotonically in the periphery of the star network and in the circle network. For example, in the periphery (for example, agent B) of the star network with $q = 1$, the marginal value starts at 1.98 in the first period and then becomes zero in the second period. However, in the third period it again becomes positive with a value of 0.60. This reflects the fact that the information at the second period (agent A 's first-period action) makes agent B at most indifferent between two actions and thus there is no marginal gain from this information. However, agent

³The number of replications in the simulation throughout the paper is 10000.

A 's second-period action together with his first-period action reveals the information on his neighbor's (agent C 's) first-period action, which would make agent B switch to follow agent A at some history and thus create a positive marginal gain.

Now compare the patterns of marginal values of information as the network architecture and information parameter change. Figure 2 presents graphically the marginal values of information across networks and values of the information parameter.

[Insert Figure 2]

First of all, holding constant the network architecture, the marginal value in the first period falls, as q decreases, from 1.98 when $q = 1$ to 1.32 when $q = 2/3$, and 0.66 when $q = 1/3$. Note that the marginal value in the first period depends neither on the network architecture nor on the distribution of cognitive types, since the decision is only based on a private signal. In the complete network and the center of the star network, the marginal value at the second period becomes lower as well when q decreases (Figure 2A and 2C), whereas this is not the case in the circle network and the periphery of the star network (Figure 2B and 2D). Nonetheless, it takes at least as many periods to extract all valuable information when q decreases and the sum of all marginal values falls when q decreases. Therefore, agents either make more effort or get smaller benefits from extracting all valuable information when q decreases.

As we have seen, holding constant the information parameter, the marginal values in the first period are the same across all locations and networks. For sufficiently high q , the marginal value in the second period is higher when agents observe the behavior of the other two agents than when agents only observe the behavior of one other agent. Also, it takes at least as many periods to extract all valuable information and the sum of all marginal values is lower when agents only observe the behavior of one other agent. See Figures 2E, 2F and 2G. In the same vein, agents either make more effort or get smaller benefits in extracting all valuable information when they observe only one other agent.

The patterns of marginal values revealed by the simulation of the Bayesian model are interesting in their own right, but they also suggest comparative static conjectures regarding the equilibrium type distributions in the cognitive hierarchy model. These will be discussed in a later section.

2.3 Some predictions on learning dynamics

One main result about learning dynamics in the Bayesian model is that (asymptotic) uniformity of actions is a robust phenomenon in *connected* networks (Gale and Kariv, 2003). The cognitive hierarchy model encompasses the Bayesian model as a degenerate case. Hence, it can generate the same behavior but also allows for richer patterns of learning dynamics. First, the presence of multiple cognitive types may result in slower convergence to a uniform action, since agents may

have to wait longer in order to learn others' types as well as their private signals. Secondly, herd behavior, defined as the uniformity of actions from some time period on, may not arise either because of the inertia of lower cognitive types or because information transmission is blocked by lower cognitive types. Finally, a group of agents in a network may end up in an incorrect herd, relative to the distribution of private signals, more frequently in the cognitive hierarchy model than in the Bayesian model.

We illustrate these predictions below with the examples from the three networks.

Slower convergence Consider the complete network when $q = 1$. Suppose that the signal distribution is given by $(\sigma_A, \sigma_B, \sigma_C) = (-1, 1, 1)$. In the Bayesian model each agent follows his own signal in the first period, which results in $(a_{A1}, a_{B1}, a_{C1}) = (-1, 1, 1)$. Since each neighbor's first-turn action reveals perfectly his private signal, learning effectively ends at the second turn. Thus, herd behavior emerges from the second period onwards. On the other hand, the learning dynamics in the cognitive hierarchy model depends on the distribution of cognitive types. If the probability of cognitive type 0, $\pi(0)$, is sufficiently high, then higher cognitive types should wait longer before they join a herd. The following diagrams show a comparison of learning dynamic between Bayesian model and cognitive hierarchy model.

Bayesian model				$(\tau_A, \tau_B, \tau_C) = (4, 4, 4)$			
Period	Agent / Signal			Period	Agent / Signal		
	<i>A</i>	<i>B</i>	<i>C</i>		<i>A</i>	<i>B</i>	<i>C</i>
	-1	1	1		-1	1	1
1	-1	1	1	1	-1	1	1
2	1	1	1	2	-1	1	1
3	1	1	1	3	1	1	1
4	1	1	1	4	1	1	1
...

No convergence Consider the star network when $q = 1$. Suppose that the signal distribution is given by $(\sigma_A, \sigma_B, \sigma_C) = (-1, 1, 1)$. In the Bayesian model, agent *A* will switch to action 1 from the second period on. Thus, a herd will emerge at least from the third period onwards. On the other hand, suppose that the following types $(\tau_A, \tau_B, \tau_C) = (1, 1, 5)$ are randomly matched to play a game in the cognitive hierarchy model. Then, the two type-1 agents in locations *A* and *B* will follow their own private signals from the first period on, whereas agent *C* will eventually switch to follow agent *A*'s actions. Due to the persistent conflict in action choices between agents *A* and *B*, there is no chance that a herd arises. The following diagrams summarize a case of no

convergence in the cognitive hierarchy model.

Bayesian model				$(\tau_A, \tau_B, \tau_C) = (1, 1, 5)$			
Period	Agent / Signal			Period	Agent / Signal		
	<i>A</i>	<i>B</i>	<i>C</i>		<i>A</i>	<i>B</i>	<i>C</i>
	-1	1	1		-1	1	1
1	-1	1	1	1	-1	1	1
2	1	1	1	2	-1	1	1
3	1	1	1	3	-1	1	-1
4	1	1	1	4	-1	1	-1
...

Incorrect herds Consider again the star network when $q = 1$. Suppose that the signal distribution is given by $(\sigma_A, \sigma_B, \sigma_C) = (-1, 1, 1)$. Given the signal distribution, the optimal action is 1. In the Bayesian model agent *A* will switch to follow other two agents' actions from the second period on. On the other hand, since an agent in the center of the star network plays the role of an information channel between the two agents on the periphery, all agents may end up choosing action -1. Suppose that the following types $(\tau_A, \tau_B, \tau_C) = (1, 5, 5)$ are randomly selected for each location. Since agent *A* will only process the information from his private signal, he will continue to choose action -1. But the other two agents are sophisticated enough to eventually follow agent *A*'s action, i.e., to switch to -1. This is illustrated in the diagrams below.

Bayesian model				$(\tau_A, \tau_B, \tau_C) = (1, 5, 5)$			
Period	Agent / Signal			Period	Agent / Signal		
	<i>A</i>	<i>B</i>	<i>C</i>		<i>A</i>	<i>B</i>	<i>C</i>
	-1	1	1		-1	1	1
1	-1	1	1	1	-1	1	1
2	1	1	1	2	-1	1	1
3	1	1	1	3	-1	-1	-1
4	1	1	1	4	-1	-1	-1
...

3 Econometric Analysis

This section describes the experimental data briefly and then presents an econometric analysis of the data based on the cognitive hierarchy model.

3.1 Experimental Data

The experimental data we use were collected by CGK05. The experimental design focuses on three-person, connected networks: the *complete*, *star* and *circle* networks. In addition, there are three information treatments, corresponding to different values of the probability of receiving a private signal: *full information* ($q = 1$), *high information* ($q = 2/3$), and *low information* ($q = 1/3$).

There is a total of nine sessions. In each session, the network and information treatment are held constant. A session consists of 15 independent games and each game consists of 6 decision periods. At the beginning of a session, subjects were randomly assigned to one of three locations, which remained constant through the session. Subjects were randomly matched to form a group of three players at the beginning of each game. The computer randomly selected one of the six decision periods at the end of each game and each subject who guessed the correct state in the selected period earned \$2 and otherwise earned nothing. See CGK05 for more details.

As motivated in the previous section, the type structure is intuitively plausible but, more importantly, it is not too difficult to see that these types really do exist in the experimental data. In certain cases we can easily identify the decision rules corresponding to different numbers of periods of information processing. The following figure presents the relative frequencies of three decision rules that are observed in two different networks, the complete and the star networks, when every subject was informed ($q = 1$). We focus on histories of play where a player's neighbor(s) always choose(s) the state that is different from the one indicated by the player's own signal. So if the player received a signal indicating that state 1 was more likely, his neighbor(s) always choose state 2. The following three behaviors are clearly observed in the data. First, there is behavior that looks as if it never uses any information and corresponds to the random type $\tau = 0$. Secondly, there is behavior that corresponds to type $\tau = 1$, who follows his own private signal. Finally, there is a large proportion of subjects who respond to their neighbors' behavior by switching to follow them. This corresponds to a sophisticated cognitive type τ for some τ sufficiently large.

[Insert Figure 3]

The heterogeneity in the data has profound consequences for information aggregation in the experiments. Sometimes it caused subjects to choose an incorrect action and other times it might block completely the transmission of information. One consequence is that the overall frequency of herd behavior (uniformity of actions) is significantly lower than the Bayesian model predicts. Herd behavior is said to arise in time period t when, from that period on, all subjects take the same action. Table 2 reports the average herd behavior from the empirical data and data simulated from the Bayesian model.

[Insert Table 2]

The comparison between simulated and empirical data for each network and information treatment is quite informative. Figure 4 presents these comparisons graphically for the full information treatment.

[Insert Figure 4]

First, the average herd behavior from the data is uniformly lower than the Bayesian model predicts. Secondly, despite the gaps between the theoretical and empirical averages of herd behavior, the shapes of the two graphs are remarkably similar (see Figure 4). For instance, the theoretical herd behavior in the star network jumps in the second and third periods and remains constant afterwards. The increase of herd behavior in the second period is because agent *A* learns by observing the other two agents and joins a herd after some history, while the jump in the third period is due to learning by agents *B* and *C*. We find the same pattern in the empirical data. Given the presence of multiple decision rules in the data, we expect that these two features can be explained by the cognitive hierarchy. This can be confirmed after we estimate the cognitive hierarchy model.

3.2 Identification

From the experimental design, there are in principle 6 cognitive types, starting with type 0 and ending with type 5. Not all cognitive types are identifiable from the data, however. In order to separately identify two different types, we need to have some history in equilibrium where the two types predict different choices of actions. To fix ideas, consider the full information case. At the beginning of the game, there are two possible histories: one where every agent receives the same signal and the other where one agent receives a different signal. It is clear to see that type 0 and higher cognitive types can be identified from the play of the game starting at the first history. At the second history, when there is an initial diversity of signals, we are able to identify lower cognitive types, who stop processing information before it is optimal to switch to follow neighbors, and higher cognitive types, who process information long enough to switch to follow neighbors. Finally, once herd behavior emerges, all cognitive types who process information at least until that period would choose the same behavior. Of course, under high and low information, there are other histories at the beginning to be considered. Nonetheless, similar arguments still remain valid.

There is another reason to prefer a smaller number of types. A theory that requires fewer types to explain the data is more parsimonious and, since the types are not directly observable, parsimony should lead us to minimize the number of types used, other things being equal.

The concerns about identification and parsimony lead us to use the following specification of types in the econometric estimation. We include three basic types in the complete network and the star network: type 0 (called *T0*), type 1 (called *T1*), and some higher types who respond to neighbors' behavior (called *SB*). In order to make type 0 and type 1 distinguishable under

high and low information, we consider $T1$ a hybrid of type 1 if an agent receives a signal and type 2 if an agent does not receive a signal. Finally, in the circle network it is much harder to identify between $T1$ and SB . Hence, we only include $T0$ and SB in the circle network. Note that, depending the network and type distribution, $T1$ could incorporate types higher than type 1 as long as such types stop processing information before they would respond to neighbors' behavior.

3.3 Testable conjectures

In order to make predictions about cross-network and cross-information (q) comparisons of equilibrium type distributions, we have to take account of the endogenous determination of the value of information and the type distribution in equilibrium. Although it is very hard to pin down the precise equilibrium distribution of types, the series of simulation results can provide reasonable conjectures on the qualitative difference of equilibrium type distribution across networks and information parameters. Table 3 collects the simulated marginal values of information for three type distributions: $(\Pr(T0), \Pr(T1), \Pr(SB)) = (0.2, 0.2, 0.6), (0.4, 0.0, 0.6),$ and $(0.0, 0.4, 0.6)$.

[Insert Table 3]

Although the simulation results are delicate to interpret, what we want to see from the simulation results is whether the patterns of marginal values of information we have found in the Bayesian model remain valid when there is some proportion of lower cognitive types. Note that the type distribution has no influence on the first-period marginal values. First, holding constant the network architecture and the type distribution, it generally takes as many time periods to extract all valuable information when q decreases and the sum of the marginal values falls when q decreases. Secondly, holding the information treatment and the type distribution constant, it takes at least as many periods to extract all valuable information and the sum of all marginal values is lower when agents only observe the behavior of one other agent. This somewhat clarifies the general shift of the equilibrium type distributions when the task environment changes.

Given the types we include in the estimation, we expect that, holding constant a network, there would be higher proportion of $T0$ when q decreases. This reflects the fact that there is a significant drop of marginal value at the first period when q decreases. On the other hand, there would be lower proportion of SB when q decreases since the marginal values starting from the second period when $q = 1$ are higher than the marginal values starting from the third period when $q = 2/3$ or $1/3$. Secondly, although it is less clear, we expect that, holding constant q , there would be higher proportion of SB in the complete and star network than in the periphery of the star network. The conjectures are summarized below.

Conjecture I (Information) Given each location in a network, (i) $\Pr(T0; q)$ decreases in q .
That is, the probability of $T0$ is highest under low information and lowest under full

information. (ii) $\Pr(SB; q)$ increases in q . That is, the probability of SB is lowest under low information and highest under full information.

Conjecture II (Network) Given an information parameter q , (i) the probability of SB is higher in the center (location A) of the star network than in the periphery (location B and C) of the star network. (ii) The probability of SB is higher in the complete network than in the periphery of the star network.

3.4 Likelihood function

The econometric specification of the cognitive hierarchy model is a type-mixture model in which a subject's type is randomly drawn from a common distribution over three types: $T0$, $T1$ and SB . We assume that subjects' types are independently drawn from a common distribution at each game rather than that an initially drawn type remains constant across all 15 games.^{4, 5}

We observe the set of action choices and the associated evolution of information sets at each game, $\left\{ \{a_{nt}, I_{nt}\}_{t=1}^T \right\}_{n=1}^N$, where N denotes the number of total observations. Then the likelihood function can be constructed as follows:

$$\begin{aligned} \mathcal{L} & \left(\{ \beta_t \}_{t=1}^T, \{ \pi_k \}_{k \in \{T0, T1, SB\}} \mid \left\{ \{ a_{nt}, I_{nt} \}_{t=1}^T \right\}_{n=1}^N \right) \\ & = \prod_{n=1}^N \left\{ \sum_{k \in \{T0, T1, SB\}} \pi_k \Pr \left(a_{n1}, \dots, a_{nT} \mid \{ I_{nt} \}_{t=1}^T, k, \{ \beta_t \}_{t=1}^T, \{ \pi_k \}_{k \in \{T0, T1, SB\}} \right) \right\}. \end{aligned}$$

In order to have a well-defined likelihood function, we assume further that the choice from each type is stochastic. This is done by combining the Quantal response equilibrium (QRE) approach by McKelvey and Palfrey (1995, 1998) with the cognitive hierarchy model. We call it *CH-QRE* throughout the section.

For tractability, we assume that all types receive idiosyncratic preference shocks to the expected payoff from an action in each period and that the shocks follow the type I extreme value distribution. We also assume that the preference shocks are independently and identically distributed across subjects, decision periods, and types. Then the likelihood function can be

⁴We found that some subjects appear to use different decision rules in similar decision situations. Such behavioral instability across games might be caused by learning effects or experimentation motives. But there is no evidence of learning effects when we split the data between the first-half and second-half of a single session.

⁵We tried to estimate the cognitive hierarchy model under the alternative assumption of constant types. The estimation results under this specification usually yielded overestimation of $T0$ and had worse performance in terms of the goodness of fit.

further written as follows:

$$\begin{aligned} & \mathcal{L} \left(\{\beta_t\}_{t=1}^T, \{\pi_k\}_{k \in \{T0, T1, SB\}} \mid \left\{ \{a_{nt}, I_{nt}\}_{t=1}^T \right\}_{n=1}^N \right) \\ &= \prod_{n=1}^N \left\{ \sum_{k \in \{T0, T1, SB\}} \pi_k \left\{ \prod_{t=1}^T \Pr \left(a_{nt} \mid I_{nt}, k, \{\beta_s\}_{s=1}^t, \{\pi_k\}_{k \in \{T0, T1, SB\}} \right) \right\} \right\}, \end{aligned}$$

where $\Pr \left(a_{nt} \mid I_{nt}, k, \{\beta_s\}_{s=1}^t, \{\pi_k\}_{k \in \{T0, T1, SB\}} \right)$ denotes the probability that action a_{nt} is chosen conditional on the information set I_{nt} when the type is k .

The equilibrium choice rule for type $k \in \{T0, T1, SB\}$ in the CH-QRE model can be summarized by the following logistic choice probability function: for any period t with information set I_t

$$\Pr(a_t = 1 \mid I_t, k) = \frac{1}{1 + \exp(-\beta_t \Delta_t(I_t, k))},$$

where a_t is the action of the subject at time period t , I_t is the subject's information set at period t , β_t is a coefficient parameter determining the sensitivity of the choice function to expected payoff differences, and $\Delta_t(I_t, k)$ is the difference between the expected payoffs for type k from actions $a_t = 1$ and $a_t = -1$. Since $T0$ does not use any information, we adopt the convention that the expected payoff difference for $T0$ is zero in all time periods and thus the choice probability for $T0$ is always $1/2$.

The specific cognitive types for SB vary across networks and information treatments. In the complete network and the center of the star network under full information, SB is specified as cognitive type 3, whereas SB is specified as type 4 in the periphery of the star network under full information.⁶ In all other treatments, SB is specified as cognitive type 5.

The estimation proceeds treatment-by-treatment. We assume that one beta coefficient in each period applies to all types and all network positions. In the complete and circle network, we estimate a single type distribution applying to all positions because the network is informationally symmetric. Thus, we have 8 parameters to estimate in the complete network under each information treatments: a series of beta coefficients, $\{\beta_t\}_{t=1}^6$, and two type probabilities, (π_{T0}, π_{T1}) . In the circle network under each information treatment, there are 7 parameters to be estimated: a series of beta coefficients, $\{\beta_t\}_{t=1}^6$, and one type probability, π_{T0} . On the other hand, in the star network, due to the asymmetry of the network, we separately estimate the type distributions for the center and the periphery. Therefore, we have 10 parameters to estimate: a series of beta coefficients, $\{\beta_t\}_{t=1}^6$, two type probabilities in the center, (π_{T0}^A, π_{T1}^A) , and two type probabilities in the periphery, (π_{T0}^B, π_{T1}^B) .

The likelihood function is constructed recursively using the beta coefficients at previous periods and also the type distribution. The method can be illustrated with reference to the

⁶We tried alternative specifications for SB up to the fifth turn at each treatment. The results are essentially unchanged.

complete network and SB type. In order to compute the choice probability in the first decision period, the difference in expected payoffs between two actions, $\Delta_1(I_1)$, is calculated using the information (a private signal) from the first period. (Indeed, all types except the 0-type make this computation in the first period). Then, the choice probability is determined by the first-period beta coefficient, β_1 , and the payoff difference $\Delta_1(I_1)$. In the second decision period, processing information of their neighbors' actions in the first period requires agents to infer whether their neighbors are $T0$ and whether they “trembled” in choosing their actions. Because of these inferences, the difference in expected payoffs in the second period depends on π_{T0} and β_1 , and so is denoted by $\Delta_2(I_2; \pi_{T0}, \beta_1)$. ($T1$ conducts this step of information processing if he is uninformed but otherwise his second-period payoff difference is the same as $\Delta_1(I_1)$). Then, the choice probability in the second period has the logistic functional form and depends on β_2 and $\Delta_2(I_2; \pi_{T0}, \beta_1)$. In the third decision period, the information processing gets more complicated. Since agents need to consider how their neighbors made inferences and how such inferences were reflected in their second period actions, SB players need to make an inference about whether their neighbors are $T0$, $T1$, or SB , and also whether they trembled in the second period. Thus, the difference in expected payoffs at the third period is a function of π_{T0} , π_{T1} , and the previous betas, denoted by $\Delta_3(I_3; \pi_{T0}, \pi_{T1}, \beta_1, \beta_2)$. This payoff difference, together with β_3 , determines the choice probability in the third period. Continuing in this way, the entire sequence of choice probabilities for each type can be constructed and these probabilities are then used to construct the conditional likelihood functions on types. The procedure in other networks is analogous, but must take into account the process of updating beliefs that is specific to that network. The details are relegated to Appendix I.

We used the method of maximum likelihood estimation separately for each of the nine treatments. In order to get around the problem of local maxima, we used the simplex algorithm of Nelder and Mead to maximize the likelihood function with a large step size and a variety of starting values to be confident that a global maximum was achieved.⁷

3.5 Estimation results

Table 4 presents the maximum likelihood estimation results for the CH-QRE model. The standard errors are reported in parentheses.

[Insert Table 4]

All the beta coefficients and type probability estimates are significantly positive and some beta coefficients are quite high in the star network under low information.

⁷The Fortran 90 program is used to do the ML estimation. The program codes are available on request to the author.

One of the surprising findings is that SB is dominant in all networks and information treatments. The estimated probabilities of SB are at least 50% in all networks and information treatments except the star network under low information. This implies that the majority of subjects in the experiment behave as if Bayes-rational. On the other hand, the estimated distribution varies significantly across networks and information treatments. In fact, the comparative static properties of the estimated distribution are largely consistent with the predictions of the cognitive hierarchy model, as is clearly seen below.

First, we summarize the comparative-static comparisons of estimated probabilities of $T0$ and SB across information treatments in a given network.

Result 1 (Information) *(i) In all network locations except the center of the star network, the estimated probability of $T0$ is lowest under full information and highest under low information. In the center of the star network, the estimated probability of $T0$ is lowest under high information and highest under low information. (ii) In the complete network and the star network, the estimated probability of SB is highest under high information and lowest under low information. In the circle network, the estimated probability of SB is highest under full information and lowest both under high and low information.*

Figure 5 graphically shows the changes of estimated probabilities of $T0$ and SB across information treatments.

[Insert Figure 5]

Part (i) of Result 1 strongly confirms Conjecture I. Part (ii) of Result 1, however, is less clear cut but the probability of SB is noticeably lower under the low information treatment as required by Conjecture I. We conclude that the predictions of the cognitive hierarchy model, associated with the change of q , are overall confirmed by the experimental data.

Next, we describe the comparative-static comparisons of estimated probabilities of SB across networks in an information treatment.

Result 2 (Network) *(i) Under all information, the estimated probability of SB is higher in the center (location A) of the star network than in the periphery (location B and C) of the star network. (ii) Under high and low information, the estimated probability of SB is higher in the complete network than in the periphery of the star network. Under full information, the estimated probability of SB is higher in the periphery of the star network than in the complete network.*

Part (i) of Result 2 strongly supports Conjecture II. However, part (ii) of Result 2 is also consistent with it, since the estimated probability of SB is lower in the periphery of the star network than in the complete network under high and low information. The comparison of estimated probabilities of SB between the periphery of the star network and the complete

network under full information is less significant since the estimates are within one standard deviation of each other. We conclude that the experimental data also confirms the predictions related to the change of network architecture from the cognitive hierarchy model.

3.6 How differently did individuals behave?

The estimation results of the CH-QRE model also lead us to investigate individual-level behavior across network and information treatments. Given a sequence of actions and the associated evolution of the information sets, $\{a_t, I_t\}_{t=1}^6$, the individual type distribution is computed as follows, using the estimated values of the parameters: for any type $\tau = T0, T1$, and SB ,

$$\Pr\left(\tau \mid \{a_t, I_t\}_{t=1}^6\right) = \frac{\prod_{t=1}^6 \Pr(a_t \mid I_t, \tau) \Pr(\tau)}{\sum_{\tau'} \left\{ \prod_{t=1}^6 \Pr(a_t \mid I_t, \tau') \Pr(\tau') \right\}}$$

For each subject in all treatments, we have 15 sequences of actions and associated evolution of information sets. The individual-level estimated type distribution often yields sharp predictions on what type the behavior of a subject is close to. The analysis confirms that there are significant patterns of behavioral heterogeneity in the subject pool at each network and information treatment.

Figure 6 presents a collection of scatter diagrams for estimated type distributions across 15 games for representative subjects across networks and information treatments. An estimated distribution is shown as a point in the two dimensional simplex in each scatter diagram. Each vertex of the simplex stands for the degenerate distribution where one type has probability 1 and the other two types have probability 0.

[Insert Figure 6]

Interestingly, some subjects are consistently close to SB or $T1$ in all plays of the game in the laboratory. Other subjects often switched among several decision rules. In other words, they sometimes played the game as if they were SB but other times they seemed to use less sophisticated rules such as $T0$ and $T1$.

In summary, there are significant levels of behavioral heterogeneity across subjects. Sometimes, there is even behavioral heterogeneity across games for a single subject. The findings of the individual-level analysis strongly suggest that behavioral heterogeneity in learning is an important ingredient of any structural model of this data.

3.7 Is sophistication rewarded?

Each decision rule in the cognitive hierarchy model is associated with a tradeoff between the complexity of decision problems and the accuracy of the resulting decision. On the one hand, the

complexity of information processing increases in later periods. On the other, more periods of information processing leads to higher accuracy. It is interesting to see whether actual individual performance in the lab is related to the degree of sophistication in a decision rule, as the theory would predict. We compute the average probability of being SB for each subject over all treatments and compare these probabilities with the average actual earnings they would have obtained if they were paid in each period of a given game.

Figure 7 presents the scatter diagram between average probabilities of being SB and actual earnings for subjects over all treatments. The bottom horizontal line in the figure represents expected earnings when randomization is used. The top horizontal line stands for expected earnings for a hypothetical agent who has access to the complete signal distribution in a network under full information.

[Insert Figure 7]

The simple linear regression (red line) reveals the relationship between the two variables:

$$\widehat{Earning} = \hat{\alpha} + \hat{\beta} * \Pr(SB),$$

where $\hat{\alpha} = 13.78$ (0.84) and $\hat{\beta} = 7.07$ (1.22). Interestingly, the relationship between the actual earnings and average probabilities of SB is significantly positive. Note that the top horizontal line is an upper bound on the theoretical expected payoff whereas actual payoffs are random because they depend on the signal distribution.

3.8 Model comparison and goodness of fit

In this section we measure the goodness of fit of the CH-QRE model and also compare it with that of the Bayesian QRE model. First, we reproduce the maximum likelihood estimation results of the Bayesian QRE model with the game-by-game estimation⁸. Table 5 presents the results of game-by-game maximum likelihood estimation of the Bayesian QRE model. Standard errors are reported in parentheses.

[Insert Table 5 here]

We briefly summarize the estimation results of the Bayesian QRE model. First, all beta coefficients across all networks and information treatments are significantly positive. Secondly, the magnitude of beta estimates is between 1.5 and 4 in most cases. Finally, the estimated beta series in the complete network are stable across decision periods. However, in the star and circle networks, the beta series decrease slightly across decision periods. For more details, see CGK.

⁸In CGK06, we pooled homogeneous data in each decision period to estimate a single beta coefficient. For example, we used the data from all treatments to estimate one beta coefficient in the first period. On the other hand, we estimated separately the betas for the subjects in the center and in the periphery of the star network because of information asymmetry. Overall, the estimation results are consistent with each other.

Interestingly, the estimated series of beta coefficients from the CH-QRE model are uniformly higher than those from the Bayesian QRE model. The following result summarizes the difference between them.

Result 3 (Trembling) *Over all treatments, the estimated beta series of the CH-QRE model are uniformly higher than those of the Bayesian QRE model. That is, the rates of trembling of the CH-QRE model are much lower than those of the Bayesian QRE model.*

The intuition for this result is quite easy to understand. The Bayesian QRE model treats any deviation from the best response action as trembling, given a single Bayesian decision rule. However, in the CH-QRE model, noise in the data is decomposed into two parts: one resulting from heterogeneous decision rules and the other from pure trembling given each decision rule. Therefore, the CH-QRE model in principle explains the behavior in the laboratory with greater precision within a specified set of the decision rules.

Another notable comparison of the betas between the two models is that the beta series of the CH-QRE model is not always stable across decision turns. For example, in the complete network, there are some big jumps in the beta series at later decision turns. This may suggest that many of the deviations at later turns in the complete network are explained by heterogeneity.

Now we turn attention to each model’s goodness of fit. First, we compare the changes of the log likelihood values at the ML estimates between the CH-QRE model and the Bayesian QRE model. Over all treatments, the log likelihood values of the CH-QRE model are significantly higher than those of the Bayesian QRE model. Under the current specification of *SB* type in the cognitive hierarchy model, it is not necessarily true that the cognitive hierarchy model encompasses the Bayesian model as a degenerate case. But the results suggest that it is nearly true.

The most noticeable increase in the log likelihood values occurs in the complete network under full information, the star network under full information, and the circle network under high information. In the complete network under full information, the log likelihood value increased from -441.40 to -237.64 , the highest value among all information treatments. Similarly, in the star network under full information, the log likelihood value increased from -552.78 to -343.10 . And in the circle network under high information it increased from -719.41 to -643.87 .

In order to test the goodness of fit of each model, we compare the simulated aggregate behavior of each model with the aggregate behavior of the experimental data. This is mainly done by comparing the average herd behavior from the experimental data with the simulated herd behavior from the two models. Herd behavior is said to arise in the laboratory when, from some decision period on, all subjects in a group take the same action. Table 6 presents the average level of herd behavior from the experimental data, the simulated data from the

estimated Bayesian QRE model and the simulated data from the estimated CH-QRE model.

[Insert Table 6]

For the purpose of measuring the goodness of fit, we compare the frequencies of herd behavior from the two models and the experimental data in each network and information treatment.

Result 4 (Herd behavior) *(i) The CH-QRE model generates almost the same frequency of herd behavior across decision periods as the experimental data in the complete and star networks under full information and in the circle network under high information. Over other treatments, the frequency of herd behavior in the CH-QRE model is lower than in the experimental data. (ii) Over all treatments, the Bayesian QRE model generates a frequency of herd behavior that is uniformly lower than the experimental data.*

Support for these results is presented in Figure 8 which shows, period by period, the relative frequencies of herd behavior in the two models and in the experimental data, in each network under all information treatments.

[Insert Figure 8]

One caveat should be noted regarding the comparison of herd behavior from the simulated data for the two models and the experimental data: the measure of herd behavior is highly discontinuous because uniformity of actions in earlier decision periods may not be counted as a herd if there is a single deviation in the last period. Nonetheless, the CH-QRE model fits the experimental data well in regard to herd behavior. In particular, the goodness of fit in the complete and star networks under full information is nearly perfect. However, there are other treatments in which the CH-QRE model produces lower frequencies of herd behavior than the experimental data.

Another feature of the CH-QRE model is its success in replicating the patterns of learning. In the complete network under full information, the frequency of herds in the experimental data increases dramatically from the first period to the second period and increases slowly later on. This implies that learning in the laboratory mainly occurs in the second decision period and subjects rapidly converge to an uniform action from that period onwards. In the star network under full information, learning mainly continues until the third period and, as a result, the frequency of herds jumps up between the second period and the third period.

The apparent overall failure of the Bayesian QRE model to fit the empirical frequency of herd behavior calls our attention to the earlier caveat that the measure of herd behavior is highly discontinuous. The Bayesian QRE model usually generates more trembling in simulations because of the lower values of the beta estimates. The only way that the Bayesian QRE model can capture noise in the data from heterogeneous decision rules is by having the lower beta estimates.

The cognitive hierarchy model’s goodness of fit highlights the performance of individual-level heterogeneity in replicating and interpreting patterns of aggregate behavior. One conclusion that can be drawn from the CH-QRE model’s superior performance is that individual-level heterogeneity is too important to be ignored.

4 Discussion

There exist various hierarchical models of bounded rationality designed to analyze heterogeneous strategic behavior. Most of them focus on static normal form games, especially dominance-solvable games such as the beauty-contest game or sequential games with one-shot decision. The literature includes Nagel (1995), Stahl and Wilson (1995), Ho, Camerer and Weigelt (1998), Costa-Gomes, Crawford and Broseta (2001), Costa-Gomes and Crawford (2004), Camerer, Ho and Chong (2004) and Kübler and Weiszsäcker (2004). One of the main motivations of the literature is the apparent lack of empirical evidence confirming Nash equilibrium even in simple dominance-solvable games, which might reflect limited ability of real players to do inductive thinking. The models are quite successful in explaining the behavior of players in the laboratory and often perform better than Nash equilibrium.

In particular, Camerer, Ho and Chong (2004) provide a synthetic approach where an agent mistakenly believes that his strategy is the most sophisticated with respect to the number of steps of iterative thinking. This behavioral assumption is often justified by psychological evidence of persistent overconfidence (e.g., Camerer and Lovallo (1999)) and the limitations of the brain’s capacity, such as working memory.

Their approach could generate a hierarchy of cognitive types in our setting but it turns out to be less clear. Suppose we start with 0-step players who simply randomize without using any information. Then 1-step players believe that all others are 0-step and thus it is optimal for them to use only the information they have in the first period, that is, their private signals. The 2-step players believe that their neighbors are either 0-step or 1-step and so would reveal at most their own private signals. However, this does not guarantee that 2-step players should necessarily stop processing information after the first period, because they might think that by continuing to observe others’ behavior they may learn whether they are 0-step or 1-step. In order to make this approach fit the dynamic structure of our games, one may need to impose some *ad hoc* restrictions such as that k -step thinkers stop using information after k -th period.

Our model of cognitive hierarchy is different from those found in the literature in several respects. First, we focus on dynamic games rather than static games. The dynamic structure in our games naturally defines steps of thinking, whereas the literature imposes a “pseudo-sequential” structure of thinking in order to analyze the data from normal-form games. Secondly, we use an equilibrium approach in which all agents have the same expectations regarding the potential benefit of information processing, but each of them chooses a different cognitive type

because of their different levels of cognitive costs. The literature, by contrast, relaxes mutual consistency of beliefs and adopts a disequilibrium approach. Finally, the literature does not provide any rationale for how the type distribution emerges in different task environments. In contrast, our equilibrium approach provides a good sense of how the type distributions relate to different task environments.

In addition, the current approach in the paper advances the literature by linking the type distributions of cognitive hierarchy with the difficulty of task environments. The model produces intuitive comparative-static predictions of how type distributions change with respect to the parameters determining the difficulty of task environments. Overall, such predictions are confirmed by the data. The methodological tool and the findings deliver a very important message of studying the bounded rationality in line with Simon’s scissors metaphor.

One main empirical conclusion in the literature is that on average people seem to conduct very few steps of thinking in playing the normal-form games. For example, Camerer et al. find that an average of 1.5 steps of thinking fits data from many games well. This seems at variance with one main finding in the present paper: the dominant type is *SB* across all networks and information treatments. We hesitate to make precise comparisons between the findings in the literature and in this paper for the following reasons. First of all, the nature of the games we consider is different from those in the literature: dynamic versus static. We hypothesize that the dynamic structure helps people conduct longer chains of reasoning. Secondly, the *SB* type is the same as lower cognitive types in some treatments. For example, given the estimated type distribution, *SB* in the complete network under full information behaves no differently from type 2. The question of how differently real people behave in dynamic games compared to static games will be better understood when there have been more empirical studies on this important topic.

5 Conclusion

This paper proposed a cognitive hierarchy model in games of learning in networks. The model is flexible enough to capture a diversity of cognitive types whose behavior ranges from random to substantively rational. More importantly, the model predicts different distributions of cognitive types as the task environment changes with respect to the architecture of networks and structure of information. This allows a set of conjectures regarding comparative static comparisons of type distributions, which can be tested using empirical data.

Using experimental data produced in the laboratory by Choi, Gale and Kariv (2005, 2006), we estimated a cognitive hierarchy model. Quite surprisingly, we found that the dominant cognitive type across all treatments is closely related to Bayes-rational behavior. Despite the presence of multiple cognitive types in the subject pool, this finding tells us that the classical Bayesian paradigm has considerable explanatory power as far as the majority of subjects is

concerned. This seems at variance with the findings in the literature on hierarchical models of bounded rationality, *e.g.*, Camerer, Ho and Chong (2004). They found that an average of 1.5 steps of thinking fits the data from many static games such as the beauty-contest game, as experimentally studied by Nagel (1995). We resist any direct comparisons between the findings in the literature and in this paper for the following reasons. First of all, the nature of the games we consider is different from those in the literature: dynamic versus static. We hypothesize that the dynamic structure helps people conduct longer chains of reasoning. Secondly, the behavior of types resembling Bayesian rationality is the same as that of lower cognitive types in some treatments. For example, given the estimated type distribution, the sophisticated Bayesian type in the complete network under full information behaves no differently from type 2. The question of how differently real people behave in dynamic games compared to static games is an important future research topic.

We also tested a set of conjectures from the theory. One of the findings is that, holding constant the network architecture, the probability of a type whose behavior is random increases as the value of q (the probability of receiving a private signal) decreases. At the same time, we also found that the probability of the sophisticated Bayesian type decreases as the value of q decreases. Regarding changes of network structures, the theory has a weaker prediction on the type distribution. Nonetheless, somewhat consistently with the theory, we found that, holding constant the value of q , the probability of the sophisticated Bayesian type is significantly higher in the center of the star network than in the periphery of the star network. One important research topic for the future is to extend the theoretical framework and experimental design to larger and more complex network architectures. Especially, the design of experiments with larger networks should be done in a way that allows clear-cut theoretical predictions of changes on equilibrium type distributions across networks.

Finally, we performed a goodness-of-fit test for the cognitive hierarchy model. This has been mainly done by comparing the simulated average herd behavior with the average herd behavior from the data. Surprisingly, the cognitive hierarchy model closely replicates the empirical patterns of herd behavior across networks and information structures at the aggregate level. This highlights the potential importance of structural approaches allowing individual heterogeneity in explaining patterns in aggregate data.

The model and techniques we have developed in this paper provide a foundation for future theoretical and empirical research for various dynamic games. One of the many interesting questions is how to apply the notion of a cognitive hierarchy to dynamic games in which the standard equilibrium concept requires backward induction. One natural starting point may be applying a cognitive hierarchy with various levels of limited foresight to such dynamic games.

References

- [1] Camerer, C.F. (2003) *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
- [2] Camerer, C. F., T. Ho, and J. Chong (2004), "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119(3), 861-898.
- [3] Camerer, C.F. and D. Lovo (1999), "Overconfidence and Excess Entry: An Experimental Approach," *American Economic Review*, 89, 306 - 318.
- [4] Choi, S., D. Gale and S. Kariv (2005), "Behavioral Aspects of Learning in Social Networks: An Experimental Study," *Advances in Behavioral and Experimental Economics* (in the Advances in Applied Microeconomics series), edited by John Morgan, JAI Press.
- [5] Choi, S., D. Gale and S. Kariv (2006), "A Quantal Response Equilibrium Analysis of Social Learning in Networks," working paper, UCL.
- [6] Costa-Gomes, M. and V. P. Crawford (2004), "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," working paper, UC San Diego.
- [7] Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001), "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, 69(5), 1193-1235.
- [8] Gale, D. and S. Kariv (2003), "Bayesian Learning in Social Networks," *Games and Economic Behavior*, 45(2), 329-346.
- [9] Ho, T., C. Camerer and K. Weigelt (1998), "Iterated Dominance and Iterated Best-response in p -Beauty Contests," *American Economic Review*, 88, 947-969.
- [10] Keynes, J. M.(1936) *The General Theory of Interest, Employment and Money*. London. MacMillan.
- [11] Kübler, D. and G. Weizsäcker (2004), "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory," *Review of Economic Studies*, 71, 425-441.
- [12] McKelvey, R.D. and T.R. Palfrey (1995), "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6-38.
- [13] McKelvey, R.D. and T.R. Palfrey (1998), "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1, 9-41.
- [14] Nagel, R. (1995), "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85, 1313-1326.

- [15] Nelder, J.A. and R. Mead (1965), "A Simplex Method for Function Minimization," *Computer Journal*, 7(4), 308-313.
- [16] Simon, H.A. (1990), "Invariants of Human Behavior," *Annual Review of Psychology*, 41, 1-19.
- [17] Stahl, D. O. and P. Wilson (1995), "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10, 218-254.

6 Appendix

This section presents the construction of likelihood functions for the CH-QRE model. The CH-QRE model assumes that each cognitive type receives idiosyncratic preference shocks. Formally, for an agent i of type $\tau \in \{T0, T1, SB\}$, the random utility function from choosing action $a_{it} = 1$ rather than $a_{it} = -1$ is given by

$$U_{it}(I_{it}; k) = \beta_t \Delta(I_{it}, k) + \varepsilon_{it},$$

where $\Delta(I_{it}, k)$ is the difference between the expected payoffs for type k from actions $a_{it} = 1$ to $a_{it} = -1$, conditional on information set in period t , I_{it} , and coefficient β_t parametrizes the sensitivity of action choices to payoff differences. $\Delta(I_t, k)$ is defined as

$$\Delta(I_{it}, k) = \begin{cases} M [2 \Pr(\omega = 1 | I_{it}) - 1] & \text{if } t \leq k \\ M [2 \Pr(\omega = 1 | I_{ik}) - 1] & \text{if } t > k \end{cases}.$$

Random variable ε_{it} represents the agent's preference shock, which is assumed to be privately observed by that agent. For tractability, we have the following assumptions on the error structure. First, ε_{it} is assumed to be independently and identically distributed according to the logistic distribution with cumulative distribution $F(\varepsilon) = 1 / (1 + \exp(-\varepsilon))$, for each agent and each period $t = 1, 2, \dots, T$. Secondly, ε_{it} is independent of the type of each agent. Then the stochastic choice function has the following form of logistic distribution:

$$\Pr(a_{it} = 1 | I_{it}, k) = \frac{1}{1 + \exp(-\beta_t \Delta(I_{it}, k))}.$$

The posterior probability of a unknown state conditional on an information set is determined by Bayes' rule. The Bayesian belief updating contains essentially two inference problems: *inferring other agents' private signals as well as other agents' types*. Such inference problems vary the degree of complexity, depending on networks and information treatments.

Let $\bar{N}_i = \{i\} \cup N_i$ denote a set of agent i 's neighbors and himself and let $I_{it} = \{\sigma_i, (a_{js})_{s=1}^{t-1} \mid j \in \bar{N}_i\}$ denote agent i 's information set in period t . The posterior belief that the true state is $\omega = 1$ conditional on I_{it} is given by

$$\Pr(\omega = 1 | I_{it}, k) = \begin{cases} \sum_{\{\tau_j\}_{j \in N_i}} \Pr(\omega = 1 | I_{it}, \{\tau_j\}_{j \in N_i}) \Pr(\{\tau_j\}_{j \in N_i} | I_{it}) & \text{if } t \leq k \\ \sum_{\{\tau_j\}_{j \in N_i}} \Pr(\omega = 1 | I_{ik}, \{\tau_j\}_{j \in N_i}) \Pr(\{\tau_j\}_{j \in N_i} | I_{ik}) & \text{if } t > k \end{cases}.$$

The first probability in each summand on the right side, $\Pr(\omega = 1 | I_{it} \text{ (or } I_{ik}), \{\tau_j\}_{j \in N_i})$, is the posterior belief on the state $\omega = 1$ conditional on the assumption that his neighbor j 's type is τ_j . The second probability in each summand, $\Pr(\{\tau_j\}_{j \in N_i} | I_{it} \text{ (or } I_{ik}))$, represents agent i 's inference on types for each of his neighbors. For expositional convenience, suppose throughout the section that agent i is of k -type, for $k \geq t$.

The probability of the true state being $\omega = 1$ conditional on I_{it} and $\{\tau_j\}_{j \in N_i}$ can be written further by, via Bayes' rule,

$$\Pr(\omega = 1 | I_{it}, \{\tau_j\}_{j \in N_i}) = \frac{\Pr(\sigma_i | \omega = 1) \prod_{s=1}^{t-1} \prod_{j \in N_i} \Pr(a_{js} | \omega = 1, I_{is}, \tau_j)}{\sum_{\omega} \Pr(\sigma_i | \omega) \prod_{s=1}^{t-1} \prod_{j \in N_i} \Pr(a_{js} | \omega, I_{is}, \tau_j)}.$$

The posterior belief of his neighbors' types, conditional on information set at turn t , I_{it} , can be also further written into, via Bayes' rule,

$$\Pr(\{\tau_j\}_{j \in N_i} | I_{it}) = \frac{\Pr(\{\tau_j\}_{j \in N_i}) \prod_{s=1}^{t-1} \Pr((a_{js})_{j \in N_i} | I_{is}, \{\tau_j\}_{j \in N_i})}{\sum_{\{\tau'_j\}_{j \in N_i}} \Pr(\{\tau'_j\}_{j \in N_i}) \prod_{s=1}^{t-1} \Pr((a_{js})_{j \in N_i} | I_{is}, \{\tau'_j\}_{j \in N_i})},$$

where

$$\Pr((a_{js})_{j \in N_i} | I_{is}, \{\tau_j\}_{j \in N_i}) = \sum_{\omega} \Pr(\omega | I_{is}, \{\tau_j\}_{j \in N_i}) \prod_{j \in N_i} \Pr(a_{js} | \omega, I_{is}, \tau_j).$$

The precise form of $\Pr(a_{js} | \omega, I_{is}, \tau_j)$ depends on the observational structure of networks. Appendix I in Choi, Gale and Kariv (2006) illustrates the details of this probability with regard to the network structure, when agents are Bayesian. The derivation of the probability, $\Pr(a_{js} | \omega, I_{is}, \tau_j)$, is straightforward while taking into account that agents are different with respect to the number of periods in which they process new information. Thus, for more details, we refer to Appendix I in CGK.

Table 1. Simulated marginal values of information across time periods in the Bayesian model across networks and information

Information	Network	Position	Time period					
			1	2	3	4	5	6
$q = 1$	Complete	All	1.98	0.75	0.00	0.00	0.00	0.00
		Circle	1.98	0.00	0.32	0.21	0.00	0.00
		Star	1.98	0.75	0.00	0.00	0.00	0.00
	B&C	A	1.98	0.00	0.60	0.00	0.00	0.00
		B&C	1.98	0.00	0.00	0.00	0.00	0.00
		All	1.32	0.55	0.12	0.00	0.00	0.00
$q = 2/3$	Complete	All	1.32	0.35	0.12	0.00	0.04	0.02
		Circle	1.32	0.55	0.12	0.00	0.00	0.00
		Star	1.32	0.35	0.20	0.06	0.00	0.00
	B&C	A	1.32	0.40	0.16	0.03	0.00	0.00
		B&C	1.32	0.40	0.16	0.00	0.00	0.01
		All	0.66	0.40	0.20	0.00	0.00	0.00
$q = 1/3$	Complete	All	0.66	0.40	0.08	0.09	0.00	0.00
		Circle	0.66	0.40	0.08	0.09	0.00	0.00
		Star	0.66	0.40	0.08	0.09	0.00	0.00
	B&C	A	0.66	0.40	0.08	0.09	0.00	0.00
		B&C	0.66	0.40	0.08	0.09	0.00	0.00
		All	0.66	0.40	0.08	0.09	0.00	0.00

* Payoff is $M = 2$ and randomization is the tie-breaking rule. The length of the time period is equal to $T=6$.

Table 2. A comparison of average herd behavior from the data and the Bayesian model
 (the percentage of games in which players follow a herd from a given period on from the experimental data (2A) and from the simulation (2B))

		<u>Herd behavior from the experimental data</u>					
Herds		1	2	3	4	5	6
Full	Complete	33.3	58.9	64.5	66.7	66.7	71.1
	Star	28.9	40.0	60.0	62.2	65.6	71.1
	Circle	20.0	31.1	35.6	45.6	51.1	67.8
High	Complete	25.3	46.7	49.3	58.7	62.7	69.3
	Star	12.2	24.4	32.2	35.6	41.1	52.2
	Circle	13.3	21.1	27.8	32.2	34.4	48.9
Low	Complete	7.8	14.4	20.0	34.4	43.3	63.3
	Star	6.7	14.4	21.1	22.2	31.1	51.1
	Circle	10.7	18.7	25.3	33.3	38.7	48.0

Table 2B Herd behavior simulated under the Bayesian model

Herds		1	2	3	4	5	6
Full	Complete	33.3	100.0	100.0	100.0	100.0	100.0
	Star	33.3	55.6	100.0	100.0	100.0	100.0
	Circle	33.3	53.0	68.0	100.0	100.0	100.0
High	Complete	28.8	78.4	85.4	88.1	88.2	88.6
	Star	28.8	54.2	79.6	85.1	85.2	85.5
	Circle	28.7	47.0	58.9	67.9	74.7	87.3
Low	Complete	25.8	48.7	71.7	74.2	74.4	76.2
	Star	26.0	54.3	66.8	83.0	83.2	85.6
	Circle	25.9	47.9	67.0	68.9	70.8	77.6

Table 3. Simulated marginal values of information in the CH model with various type distributions

(Pr(T0), Pr(T1), Pr(SB)) = (0.2, 0.2, 0.6)

Information	Network	Position	Time period				
			1	2	3	4	5
Full (q = 1)	Complete	All	1.98	0.45	0.00	0.03	0.00
	Circle	All	1.98	0.00	0.00	0.00	0.00
	Star	A	1.98	0.45	0.00	0.00	0.02
		B&C	1.98	0.00	0.16	0.06	0.02
High (q = 2/3)	Complete	All	1.32	0.30	0.16	0.03	0.00
	Circle	All	1.32	0.30	0.04	0.03	0.00
	Star	A	1.32	0.30	0.20	0.00	0.02
		B&C	1.32	0.30	0.04	0.06	0.02
Low (q = 1/3)	Complete	All	0.66	0.30	0.16	0.00	0.00
	Circle	All	0.66	0.30	0.04	0.03	0.00
	Star	A	0.66	0.30	0.16	0.00	0.00
		B&C	0.66	0.30	0.04	0.06	0.00

(Pr(T0), Pr(T1), Pr(SB)) = (0.4, 0.0, 0.6)

Information	Network	Position	Time period				
			1	2	3	4	5
Full (q = 1)	Complete	All	1.98	0.15	0.08	0.00	0.00
	Circle	All	1.98	0.00	0.00	0.00	0.00
	Star	A	1.98	0.10	0.04	0.03	0.00
		B&C	1.98	0.00	0.00	0.03	0.04
High (q = 2/3)	Complete	All	1.32	0.20	0.12	0.03	0.00
	Circle	All	1.32	0.25	0.00	0.00	0.02
	Star	A	1.32	0.25	0.08	0.03	0.02
		B&C	1.32	0.25	0.00	0.03	0.00
Low (q = 1/3)	Complete	All	0.66	0.20	0.12	0.03	0.00
	Circle	All	0.66	0.20	0.08	0.00	0.00
	Star	A	0.66	0.20	0.12	0.00	0.02
		B&C	0.66	0.20	0.04	0.03	0.00

(Pr(T0), Pr(T1), Pr(SB)) = (0.0, 0.4, 0.6)

Information	Network	Position	Time period				
			1	2	3	4	5
Full (q = 1)	Complete	All	1.98	0.75	0.00	0.00	0.00
	Circle	All	1.98	0.00	0.20	0.15	0.00
	Star	A	1.98	0.75	0.00	0.00	0.00
		B&C	1.98	0.00	0.40	0.00	0.00
High (q = 2/3)	Complete	All	1.32	0.55	0.08	0.03	0.00
	Circle	All	1.32	0.40	0.04	0.03	0.00
	Star	A	1.32	0.55	0.16	0.00	0.00
		B&C	1.32	0.40	0.08	0.09	0.00
Low (q = 1/3)	Complete	All	0.66	0.35	0.20	0.03	0.00
	Circle	All	0.66	0.35	0.04	0.06	0.02
	Star	A	0.66	0.35	0.24	0.00	0.00
		B&C	0.66	0.35	0.12	0.03	0.04

* Payoff is M = 2 and randomization is the tie-breaking rule.

Table 4. Maximum likelihood estimates of the cognitive hierarchy QRE model

Complete	Turn	Full	High	Low
	1	8.67 (2.45)	5.65 (0.93)	5.15 (1.00)
	2	5.54 (0.69)	6.33 (1.07)	6.74 (1.46)
	3	6.08 (0.98)	5.73 (0.97)	7.36 (1.74)
	4	16.23 (5.58)	7.84 (2.20)	5.11 (0.73)
	5	8.51 (3.13)	6.37 (1.30)	5.49 (0.90)
	6	19.89 (15.98)	9.94 (3.28)	20.69 (9.03)
	Type			
	T0	0.08 (0.02)	0.14 (0.02)	0.26 (0.04)
	T1	0.33	0.14	0.24
	SB	0.59 (0.07)	0.72 (0.04)	0.50 (0.06)
	logL	-237.64	-450.98	-806.39

Star	Turn	Full		High		Low	
	1	5.82 (0.71)		5.75 (0.98)		5.24 (1.10)	
	2	4.19 (0.43)		3.19 (0.36)		9.70 (1.91)	
	3	9.99 (3.66)		4.11 (0.59)		7.49 (0.72)	
	4	6.09 (1.51)		3.11 (0.43)		38.04 (24.14)	
	5	12.08 (7.80)		3.77 (0.54)		10.50 (2.07)	
	6	5.72 (0.69)		2.52 (0.34)		30.48 (24.64)	
	Type	A	B&C	A	B&C	A	B&C
	T0	0.12 (0.02)	0.03 (0.02)	0.04 (0.03)	0.08 (0.03)	0.33 (0.06)	0.38 (0.04)
	T1	0.10	0.32	0.14	0.26	0.20	0.22
	SB	0.78 (0.06)	0.65 (0.07)	0.81 (0.07)	0.65 (0.06)	0.47 (0.07)	0.39 (0.05)
	logL	-343.10		-651.68		-838.30	

Circle	Turn	Full	High	Low
	1	5.72 (0.76)	5.83 (0.93)	10.62 (1.43)
	2	4.63 (1.70)	6.87 (0.81)	8.55 (1.13)
	3	2.64 (0.32)	6.06 (0.77)	6.67 (0.85)
	4	3.11 (0.35)	5.29 (0.90)	6.42 (1.19)
	5	3.24 (0.53)	4.60 (0.76)	9.65 (2.74)
	6	3.38 (0.49)	5.29 (1.20)	5.64 (0.84)
	Type			
	T0	0.08 (0.02)	0.18 (0.02)	0.17 (0.02)
	SB	0.92	0.82	0.83
	logL	-599.75	-643.87	-592.03

Standard errors are given in parentheses.

Table 5. Maximum likelihood estimates of the Bayesian QRE model

Complete	Turn	Full	High	Low
	1	3.25 (0.31)	2.79 (0.31)	2.67 (0.42)
	2	2.58 (0.33)	2.46 (0.33)	2.34 (0.38)
	3	2.82 (0.40)	2.59 (0.42)	2.51 (0.40)
	4	2.54 (0.48)	2.52 (0.40)	2.20 (0.36)
	5	2.60 (0.37)	2.66 (0.40)	2.08 (0.37)
	6	2.69 (0.58)	2.32 (0.35)	2.85 (0.43)
	log L	-441.40	-556.18	-891.30
Star	Turn	Full	High	Low
	1	3.38 (0.32)	4.10 (0.41)	2.60 (0.40)
	2	2.47 (0.32)	2.25 (0.25)	2.43 (0.36)
	3	2.26 (0.47)	2.72 (0.31)	2.33 (0.47)
	4	2.60 (0.42)	1.99 (0.27)	2.70 (0.49)
	5	2.08 (0.60)	2.44 (0.30)	1.98 (0.38)
	6	2.17 (0.42)	1.60 (0.23)	2.45 (0.44)
	log L	-552.78	-682.91	-921.96
Circle	Turn	Full	High	Low
	1	4.04 (0.41)	3.99 (0.52)	6.48 (1.53)
	2	2.41 (0.32)	2.83 (0.34)	3.29 (0.52)
	3	2.11 (0.23)	2.51 (0.34)	3.62 (0.58)
	4	2.05 (0.27)	2.13 (0.36)	3.09 (0.61)
	5	2.20 (0.28)	2.43 (0.33)	3.02 (0.57)
	6	2.19 (0.32)	1.56 (0.30)	2.34 (0.39)
	log L	-630.27	-719.41	-637.90

Standard errors are given in parentheses.

Table 6. A comparison of simulated herd behavior between Bayesian QRE and CH QRE
 (the percentage of games in which players follow a herd from a given period on from the simulations of the estimated Bayesian QRE (6A)
 and the estimated CH QRE (6B))


Table 6A **Herd behavior simulated under the estimated Bayesian QRE model**

	Herds						
	1	2	3	4	5	6	
Full	Complete	20.0	28.6	33.7	41.7	55.4	71.7
	Star	10.2	13.0	17.7	27.2	37.1	57.4
	Circle	13.9	16.2	21.1	28.4	42.0	58.8
High	Complete	8.2	12.6	16.8	23.5	34.2	49.5
	Star	6.0	7.4	12.0	15.9	27.9	43.2
	Circle	5.9	8.0	11.1	15.9	24.9	40.9
Low	Complete	1.3	1.9	4.2	8.1	16.7	37.3
	Star	0.6	1.4	3.6	6.7	11.6	35.0
	Circle	2.4	5.2	8.6	12.7	21.3	39.8

Table 6B **Herd behavior simulated under the estimated CH QRE model**

	Herds						
	1	2	3	4	5	6	
Full	Complete	28.2	56.7	60.3	64.0	65.5	70.5
	Star	29.3	41.5	58.5	61.3	65.4	69.3
	Circle	21.0	23.7	27.0	35.5	45.3	61.7
High	Complete	19.6	34.2	41.0	46.9	51.9	65.1
	Star	12.0	15.9	20.8	24.5	32.4	47.4
	Circle	13.0	20.1	24.4	28.5	34.9	52.0
Low	Complete	5.1	8.3	11.8	15.6	23.5	44.4
	Star	3.3	6.4	10.1	15.7	21.4	40.3
	Circle	7.3	14.7	20.6	26.1	32.3	47.2

Figure 1. Three-person connected networks



* A line segment between two locations indicates that they are connected and the arrowhead points to the agent whose action can be observed.

Figure 2. Marginal values of information across networks and information parameter

Figure 2A. Marginal values of information in the complete network in the Bayesian model




Figure 2B. Marginal values of information in the circle network in the Bayesian model




Figure 2C. Marginal values of information in the center (A) of the star network in the Bayesian model




Figure 2D. Marginal values of information in the periphery (B&C) of the star network in the Bayesian model




Figure 2E. Marginal values of information when $q = 1$ in the Bayesian model




Figure 2F. Marginal values of information when $q = 2/3$ in the Bayesian model




Figure 2G. Marginal values of information when $q = 1/3$ in the Bayesian model




Figure 3. Relative frequencies of decision rules when one's neighbor continues to choose actions different from one's own signal




Figure 4. A comparison of average herd behavior from the experimental data and the Bayesian model under full information
 (the percentage of games in which players follow a herd from a given period on from the data and the Bayesian model)

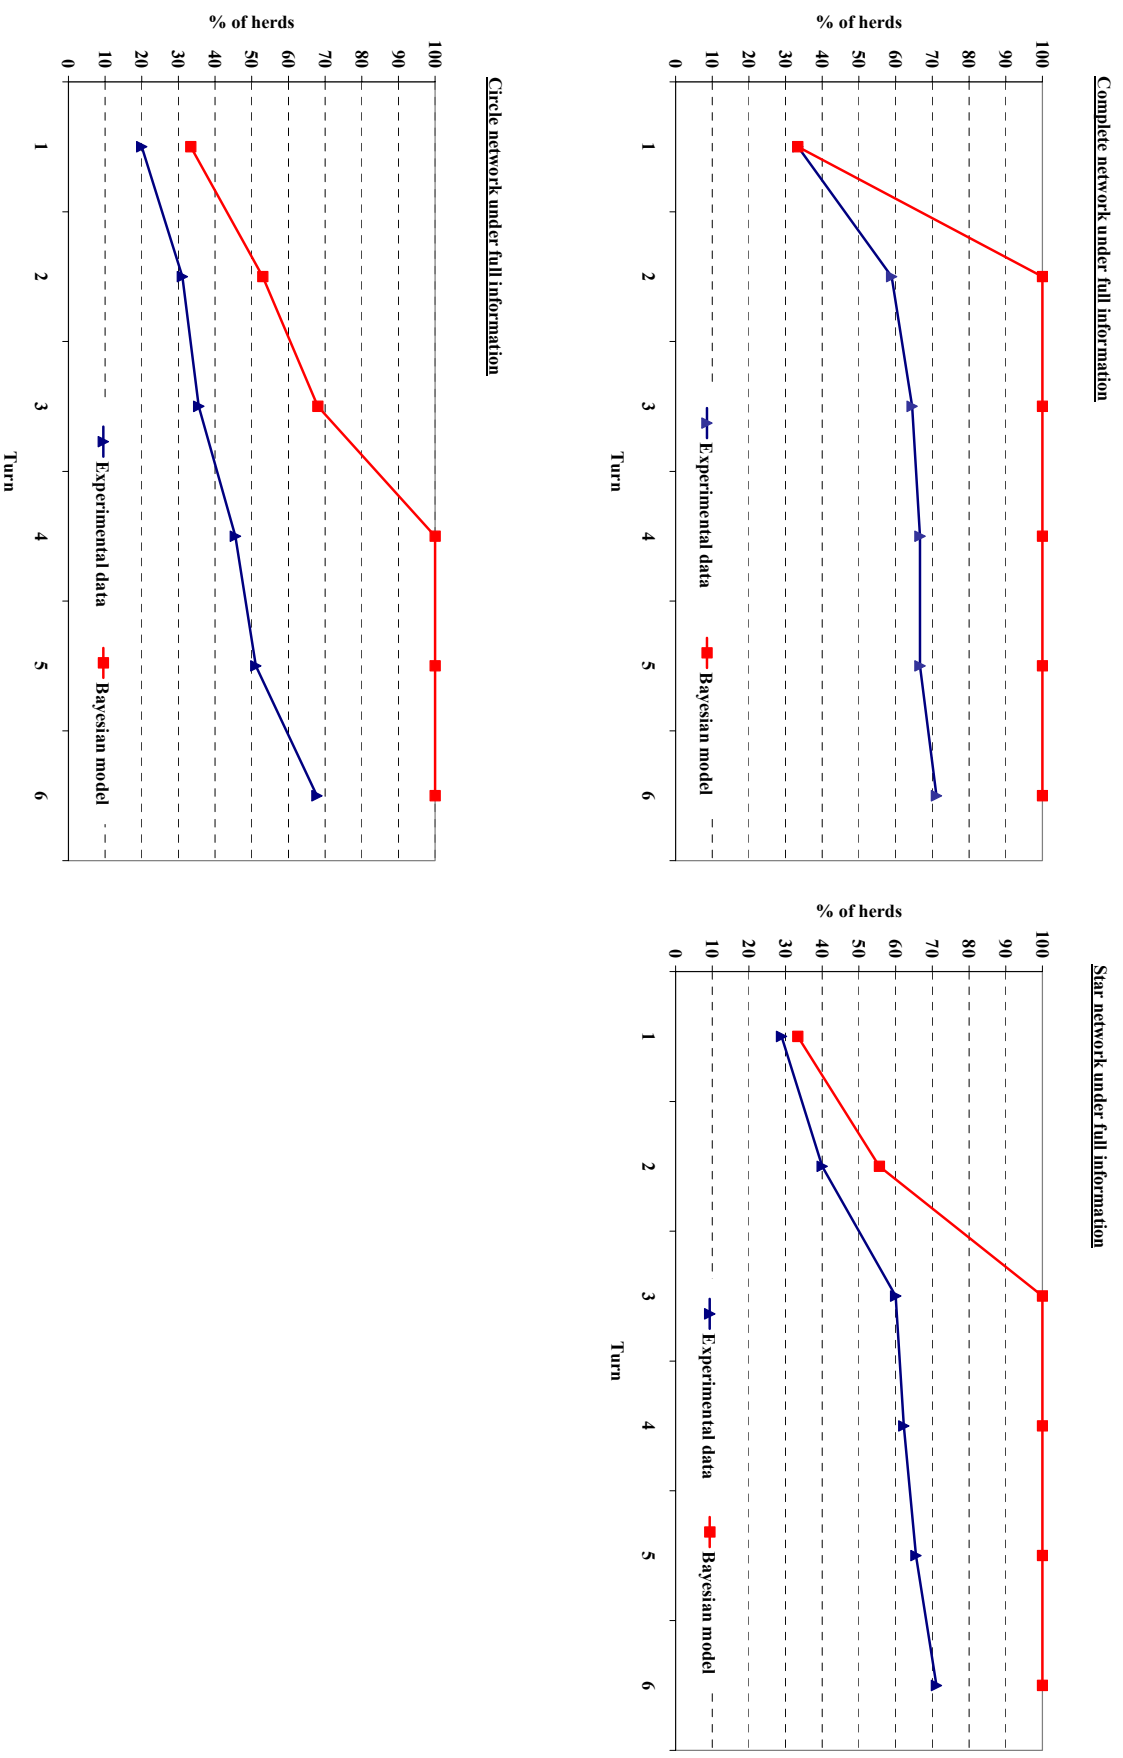


Figure 5. CH-QRE Estimated probabilities of T0 and SB across networks and information treatments

Figure 5A. Estimated Prob(T0)




Figure 5B. Estimated prob(SB)




Figure 6. Selected individual-level estimated type distributions across treatments





Figure 7. The scatter diagram between average probabilities of SB and average actual earnings



* The Red line is obtained from a simple linear regression of average actual earnings on the probabilities of SB with a constant term.

* The bottom horizontal line represents average expected actual earnings when randomization is used.


* The top horizontal line stands for average expected actual earnings for a hypothetical agent who has an access to the complete signal distribution in a network under full information.

Figure 8. A comparison of average herd behavior from the experimental data and the simulations of the estimated Bayesian QRE and CH QRE models


(the percentage of games in which players follow a herd from a given period on from the data and the simulations from each estimated model)

Complete Network


Herd behavior in the complete network under full information



Herd behavior in the complete network under high information




Herd behavior in the complete network under low information




Star network


Herd behavior in the star network under full information



Herd behavior in the star network under high information




Herd behavior in the star network under low information




Circle network

Herd behavior in the circle network under full information



Herd behavior in the circle network under high information



Herd behavior in the circle network under low information

