MANAGEMENT SCIENCE
Vol. 58, No. 8, August 2012, pp. 1525–1548
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)

http://dx.doi.org/10.1287/mnsc.1110.1506

Social Learning Through Endogenous Information Acquisition: An Experiment

Boğaçhan Çelen
Graduate School of Business, Columbia University, New York, New York 10027,
bc2132@columbia.edu

Kyle Hyndman
Department of Economics, Maastricht University, Maastricht 6211 LM, The Netherlands,
k.hyndman@maastrichtuniversity.nl

This paper provides a test of a theory of social learning through endogenous information acquisition. A group of subjects face a decision problem under uncertainty. Subjects are endowed with private information about the fundamentals of the problem and make decisions sequentially. The key feature of the experiment is that subjects can observe the decisions of predecessors by forming links at a cost. The model predicts that the average welfare is enhanced in the presence of a small cost. Our experimental results support this prediction. When the informativeness of signals changes across treatments, behavior changes in accordance with the theory. However, within treatments, there are important deviations from rationality such as a tendency to conform and excessive link formation. Given these biases, our results indicate that subjects would, except when faced with a small cost, have been better off not forming any links.

Key words: social learning; information acquisition; link formation; herd behavior
History: Received August 8, 2010; accepted December 9, 2011, by Teck Ho, behavioral economics. Published online in *Articles in Advance* April 18, 2012.

1. Introduction

Evidence that the structure of social interactions is important in disseminating information has been found in many contexts. Kelly and Ó Gráda (2000) demonstrate the importance of social networks in market panics using data from two bank runs in 1850s New York. Foster and Rosenzweig (1995), Conley and Udry (2001), and Munshi (2004) show that technology diffusion is driven by an underlying social learning process. In particular, Conley and Udry (2001) report results of field work in rural areas of Ghana indicating that the main channel for the adoption of new technology in an agricultural sector is in fact social learning through the interaction of farmers. Also, it is well documented in the labor economics literature (see Ioannides and Loury 2004) that social networks are the main source of information about jobs. These are but a few examples exhibiting the significance of questions such as the following:

- How does interaction between individuals affect the dissemination of information?
- How does information aggregation affect the interaction of individuals?

To tackle these questions, we study a simple social learning model and test it experimentally. The model underlying our experiment extends the canonical social learning model to analyze the role of the information externality on how subjects interact with each other, as well as the impact of this interaction on the learning dynamics.

In the canonical social learning model, agents receive private signals regarding the state of the world and then make decisions sequentially, after having observed the action choices of all or some of their predecessors.[1] Bayesian agents infer valuable information from their observations, often leading to herd behavior or informational cascades. A critical assumption of these models is that the interaction between agents is exogenously determined. This paper relaxes this assumption by letting agents choose whom to observe. Put differently, to acquire more information, an agent can decide to form *links* to other agents to observe their decisions and make his decision thereafter.

Although the underlying model is quite stylized, it provides a rich structure to study certain real-world phenomena. One such example is the effect of social learning on voting behavior and political outcomes. Some early research on the topic include Lazarsfeld et al. (1944) and Katz and Lazarsfeld (1955). These studies demonstrate the impact of personal contact on information transmission and voting behavior. The literature is still quite active with contemporary research

[1] See, for example, Banerjee (1992), Bikhchandani et al. (1992, 1998), and Smith and Sørensen (2000), among others.

on the topic, including Huckfeldt and Sprague (1995), Beck et al. (2002), Huckfeldt et al. (2004). The paper that comes closest to ours is by Galeotti and Mattozzi (2011). They consider a theoretical framework of electoral competition where voters obtain information through word-of-mouth communication.

Connectivity and social networks in organizations are yet another example, where personal communication and information acquisition affect individual decisions. As Cross and Parker (2004) convincingly argue, connectivity is a key component to achieve performance and learning in a social network. One striking finding Cross and Parker (2004) report is that the formal organization chart of the company they analyze is radically different from the informal structure as revealed by their social network analysis. Moreover, they show that the informal structure is the key determinant of information flow. For instance, they report that a particular member can play a crucial role in overall flow of information because of his expertise. More importantly, in many cases the expertise of this member is associated with the network he formed. The fact that the informal structure is radically different from the formal one points out the importance of endogenous link formation. Indeed, in Cross and Parker's (2004) example, despite the existence of an exogenous communication structure (formal organization chart), it is the endogenous links that individuals form (informal structure) that seem to play the main role in determining information flow.

As all the examples that we mentioned indicate, information acquisition is one of the leading reasons why individuals form links to each other. This is the focus of the present paper; that is, we investigate how information acquisition motives affect individuals' decisions to form links. Therefore, the simplicity of the canonical social learning model, which is suitably extended to allow agents to decide whom they want to link to and observe, makes it attractive for our purposes.

Our experiment consists of a group of four subjects who sequentially make decisions on the same problem. Each subject is endowed with a piece of information regarding the fundamentals of the problem. In addition to his private information, each subject, before making a decision, is allowed to form links to his predecessors. Forming a link is costly, yet it allows a subject to observe the actions of those with whom he linked. Although the subjects cannot observe the actions without a link, they observe the link decisions of those who moved before them. Therefore, the observed link structure, in many cases, contains valuable information about one's predecessors, which should influence the subject's decision on whether or not to link and, if so, whom to link to. We have two different treatments: the high information (HI) treatment and the low information (LI) treatment. The two treatments differ such that in the former treatment, subjects always receive an informative signal, whereas in the latter, they receive a signal with probability less than one. This allows us to analyze how subjects react to changes in the quality of information in this environment.

In a nutshell, each subject first decides whether or not to form any links, and then, upon observing the actions through his links, he makes an action decision. This two-step procedure gives us the leverage to pose many interesting questions: How do subjects search for information? How do subjects respond to the cost of link formation? How much information is transmitted between subjects? How do subjects aggregate the information obtained from any link decisions?

Section 2 lays out the model that we implement in the laboratory and provides the predictions of the theory. We first describe what we call the information acquisition problem. Our model is parameterized by both the cost of forming links and by the probability that each agent receives an informative signal. Given these two parameters, we provide a characterization of the way agents collect information through what we call the Bayesian sequential link procedure (BSLP). In the process, we state the theoretical predictions of the model concerning the behavior of the Bayesian agents. One important theoretical prediction of our model is that when the cost is small enough, agents who come later in the decision line are actually better off (in expectation) relative to the case in which the cost is zero. This follows because, when the cost of link formation is positive, the observed link structure often contains information on whether or not one's predecessors received an informative signal and also, in some cases, whether the signals of one's predecessors were in agreement or not. This information that is inherent in the link structure is valuable to an agent. In contrast, when the cost of link formation is zero, because all possible links are formed, such valuable information is never revealed.

In §§3 and 4, we turn to a detailed examination of our experiment. The former section explains *what* happens in the laboratory as well as *how* it happens. We also explain some of our choices concerning the experimental design. Then, in §4, we provide a detailed analysis of the observed link and action decisions of the subjects in our experiment.

We break up our descriptive analyses of behavior in the experiment in two parts. First, we focus on the link decisions, and then we analyze action decisions separately. With respect to the link decisions, we show that subjects have a tendency to form too many links. For example, despite the fact that it is not optimal for an informed subject in the second position to form a link when the cost is strictly positive, we find that links are formed in more than 30% of such instances. Very often, this appears to be the catalyst for what we call "herding in link formation," where later decision makers become ever more likely to form links, even though such behavior is not justified by the information that they acquire by forming a link.

With respect to action decisions, we define a counting heuristic in which subjects basically "count" the information (private signals and observed actions) in favor of each of the possible states and choose the action that has the most support. This actually corresponds to the optimal decision rule if subjects were rational and is approximately optimal even in the presence of a small number of mistakes. Our results show that action decisions have a very high degree of consistency with the counting heuristic.

One interesting result concerns the action decisions taken when the counting heuristic suggests a tie between the two possible states. In such a situation, our results show that there are generally two types of subjects: conformists and nonconformists. A conformist is the one who almost always breaks the tie in favor of the action taken by one's predecessor, whereas a nonconformist is the one who breaks the tie in favor of his own private information. In both the HI and LI treatments, there is a nonnegligible fraction of conformists, but it is significantly more pronounced in the HI treatment.

Despite some apparent deviations from the Bayes-rational benchmark (more so for link decisions), we do find evidence that subjects in our experiments do in fact make more accurate predictions about the state when faced with a small, but positive, cost of link formation, as opposed to the case where the cost of link formation is zero. The reason for this result appears to follow the intuition outlined above. When the cost of link is zero, subjects almost always form a link, making the observed link structure uninformative. In contrast, when the cost of link is small, fewer links are formed (nonetheless still more than the theoretical predictions), making the observed link structure informative.

We end our discussion of the experimental results with a brief discussion of the results from a structural econometric model (contained in Appendix B) that provides insights into the reasons for certain deviations. This exercise highlights two points: First, subjects appear to overweight the information that confirms their prior beliefs. Second, when faced with contradictory information, *on average*, subjects have approximately Bayesian beliefs about the state. Combined, these two findings help explain why, for example, subjects in the second position link too frequently—they anticipate that the information they receive is more valuable than it actually is.

Finally, in §6 we offer some concluding remarks and directions for future research, and two appendices collect the proofs of our theoretical results and the instructions used in the experiments.

1.1. Related Literature

This paper is related to three strands of literature. The first strand of literature that the present paper is related to is known as social learning, which was introduced by Bikhchandani et al. (1992) and Banerjee (1992), and further generalized by Smith and Sørensen (2000). Roughly, these models assume that there is an exogenous sequence of Bayesian decision makers who make a binary decision when the outcome of their decision is uncertain. Each decision maker is endowed with a private information, and they are able to observe the entire history of decisions preceding them. The literature focuses on the dynamics of information aggregation and asymptotic properties of the sequences of decisions. A number of papers relaxed the assumption of perfect information and assumed that the decision makers can make a more limited observation from the past. Çelen and Kariv (2004), Banerjee and Fudenberg (2004), and Smith and Sørensen (1998) extend the literature in that direction. Çelen and Kariv (2004) assume that the decision makers can observe only the immediate predecessor's decision. Smith and Sørensen (1998) take a more general approach and explore the case in which everyone can see random unordered samples from the history of decisions. Banerjee and Fudenberg (2004) analyze a model with a continuum of decision makers and focus on the case where they can observe a representative sample from the past. The main distinction between our paper and this literature is that whereas the observation of past decisions is determined exogenously in the literature, we allow the decision makers to choose whose decisions they want to observe.

The second strand of literature focuses on subjects' preferences for noninstrumental information. As is well known, unless information is instrumental for Bayesian decision making it does not provide any value. Although this should be true for rational decision makers, recent studies by Kübler and Weizsäcker (2004), Eliaz and Schotter (2007, 2010), and Goeree and Yariv (2007) showed that behavior in the lab conflicts with this theoretical prediction. Specifically, in different settings, these papers found overwhelming evidence that subjects do appreciate noninstrumental information. This is in accordance with the behavior that we observe in our experiment: the majority of subjects pay for links that do not provide information in the equilibrium. So, our results provide additional support for preferences for noninstrumental information. Moreover, we observe the type of social interaction (herding in link formation) that may arise as a result of these preferences.

The third strand of literature that our paper is related to is the economics of social networks. The answers to these questions are sought in the realm of the economics of social networks and learning.[2] There are a number of studies concerning the question of how interaction between individuals affect the dissemination of information (see Bala and Goyal 1998, Gale and Kariv 2003). These papers address questions concerning how information propagates and how agents learn from each other in exogenously given structures. For instance, in the paper by Bala and Goyal (1998), players play a multiarmed bandit game. They can observe the actions and the outcomes of their neighbors and learn from these observations. Because the analysis of a full Bayesian model is intractable, Bala and Goyal (1998) take a simple approach, where the players take into account only the actions and outcomes of their neighbors and ignore possible information that they can deduce from the actions and outcomes that their neighbors might have observed. The main conclusion is that the actions of players eventually converge to the same action. In a recent paper, Acemoglu et al. (2011) considered a general model of learning in general social networks. The players, in addition to their private signal, observe the actions in their neighborhood (network). Acemoglu et al. (2011) assume that the neighborhood structure evolves in time with a given stochastic process and characterizes the conditions under which correct learning obtains. Our approach differs from these papers in that we allow players to choose their own social neighbors and observe the evolution of actions and links.

2. Theoretical Background

The design of our experiment is based on a parsimonious model of social learning in which agents can choose whose action to observe. In this section, we introduce the model and carefully discuss its theoretical predictions. These predictions offer the rational benchmark for the analysis of subjects' behavior in the laboratory.

2.1. Information Acquisition Problem

The basic structure of the problem—which we call the *information acquisition problem*—is a variant of that of Bikhchandani et al. (1992) and is described as follows. There are two equally likely states of the world $\theta \in \{-1, 1\}$. The problem consists of four agents who are randomly assigned to a position in a decision line indexed by $i = 1, 2, 3, 4$. Agents act sequentially in a predetermined order. The agents' problem involves correctly identifying the true state of the world. Precisely, each agent i is supposed to take an action $a_i \in \{-1, 1\}$, which we call the *action decision*. If an agent's action matches the true state, then he receives a payoff $m > 0$; otherwise, his payoff is zero. We assume that agent i's preferences are represented by the utility function

$$u_i(a_i; \theta) = \begin{cases} m & \text{if } a_i = \theta, \\ 0 & \text{otherwise.} \end{cases}$$

Before he makes a decision, agent i receives a private signal $\sigma_i \in \{-1, 0, 1\}$. We say that the agent is *informed* when the signal he receives is either -1 or 1. The signals $\sigma_i \in \{-1, 1\}$ are informative about the true state because, conditional on the true state, the probability that the signal matches the state is $p = 2/3$. In contrast, the signal $\sigma_i = 0$ is uninformative because, given $\sigma_i = 0$, the probabilities of state $\theta = -1$ and $\theta = 1$ are both 1/2. Hence, the agent cannot distinguish the states of the world based on his signal. Therefore, we say that agent i is *uninformed* if he receives the signal $\sigma_i = 0$. Finally, we assume that an agent is informed with probability $q \in (0, 1]$ and uninformed with probability $1 - q$.

The signals that agents receive are independently and identically distributed conditional on the true state. Table 1 summarizes the probabilities with which an agent receives each signal conditional on the state of the world.

After receiving a private signal but before making the action decision, an agent has the option to observe action decisions made by the preceding agents in the decision line. The decision whether to observe any of the preceding agents' actions—and, if so, whose actions to observe—is called the *link decision*. We make two assumptions on the process of information gathering through link formation:

(1) First, link decisions are assumed to be public information; that is, each agent observes all the link decisions made by the preceding agents but not their action decisions.

(2) Second, by forming a link to one of the preceding agents, not only does an agent observe the action decision of this agent, but also all the actions that this agent observed through his link decision(s). The cost of each link is assumed to be $c \geq 0$.

In fact, the first assumption is necessary to pose the questions about the importance of link structure on information transmission and vice versa. For instance,

Table 1 Information Structure

σ	$\theta = -1$	$\theta = 1$
1	$q/3$	$2q/3$
0	$1 - q$	$1 - q$
-1	$2q/3$	$q/3$

[2] For recent and comprehensive surveys of the literature, see Jackson (2003, 2006).

in the study by Cross and Parker (2004) that we discussed earlier, the company's informal social structure is public, and this structure plays a crucial role in the flow of information. Therefore, allowing link decisions to be public enables us to understand the effect of link structure both theoretically and behaviorally. The second assumption simplifies the analysis without changing the results qualitatively. If we relax the assumption and say that an agent can observe the action decisions of only those whom he is linked to, the equilibrium analysis does not change substantially. In fact, as we will discuss in the equilibrium analysis that follows, given the observed link structure, forming a link to a preceding agent already reveals the actions of the agents whom he observes. Therefore, aside from changes in the equilibrium cost thresholds, relaxing this assumption does not cause any substantial or qualitative deviation from our results. Underlying all of this is the assumption that decision makers observe the actions, and not the signals, of those with whom they link. On this point, the assumption is consistent with the existing literature (see, e.g., Bikhchandani et al. 1992, who also discuss some examples where this assumption is sensible).

2.2. Bayesian Sequential Link Procedure

The process of link formation for each agent is a sequential process. In other words, if an agent can form more than one link, he initially compares the cost and benefits of forming no links versus forming a link. After he forms the first link (if at all), he observes the relevant action choices and then, based on the information gained, weighs the costs and benefits of forming another link, and so on.

Let us be more specific and explain what we call the *Bayesian sequential link procedure* by use of an example, which is illustrated in Figure 1. It is the fourth agent's turn to move, and he observes the following: The second agent did not form a link to the first and took his action based on his private information. The third agent formed a link to the second, yet, after observing the action of the second, he did not form a link to the first agent. Therefore, he took his action based on the information deduced from the second agent's action and his private information. There are two links that the fourth agent can form: a link to the third agent, through which he can observe the actions of the third and the second agents, and a link to the first agent, through which he can observe the action of the first agent.[3]

According to the BSLP, the fourth agent evaluates the problem in the following way. Based on his own

Figure 1 Bayesian Sequential Link Procedure: An Example

information, he decides whether it is optimal to incur c and form a link to the third agent. But in doing so he keeps in mind that he could continue and form a link to the first agent by incurring the cost c again. More precisely, he considers all possible action profiles that he can observe through his link to the third agent. Also, conditional on his private information, he knows the probability with which he can observe each action profile. Furthermore, for each one of these contingencies, he considers the action profile he can observe by a second link and decides whether he will form the second link or not as if he is in that situation. Finally, with this continuation value in mind, he decides whether to form his first link or not. In what follows, we will explain this procedure more formally. All of the results that we will report in the following sections are based on the use of the formula we derive here.

As it is the case in the example, forming a link is equivalent to saying that the agent observes the outcome (the action profile) of an experiment (forming a link). For the purposes of the present paper, it is enough to look at the case where there are two random variables, X_1 and X_2, which are independent conditional on θ, but not necessarily identical. Let $\mathscr{X}_i$ be the set of all realizations of X_i. The realizations of the random variables are denoted by x_1 and x_2, respectively. Therefore, an agent facing X_1 and X_2 first decides whether to take an optimal action simply based on his private information or to experiment with X_1. If he decides to experiment with X_1, for any realization x_1, he specifies whether to take an optimal action or to further experiment with X_2. If he decides to experiment with X_2, for all realizations x_2, he specifies the optimal action he should take.

Let $s_0 = (\sigma)$, $s_1 = (\sigma, x_1)$, and $s_2 = (\sigma, x_1, x_2)$ denote the information nodes where the agent observes only his private information, his private information and the realization of X_1, and his private information and the realization X_1 and X_2, respectively. Note that the set of all realizations at node s_0 is $S_0 := \{-1, 0, 1\}$, at node s_1 is $S_1 := S_0 \times \mathscr{X}_1$, and at node s_2 is $S_2 := S_1 \times \mathscr{X}_2$. We denote the maximum expected utility an agent can get at a node $s \in \{s_0, s_1, s_2\}$ without further experimentation by

$$\underline{v}(s) := \max\{\Pr(\theta = 1 \mid s), \Pr(\theta = -1 \mid s)\}m. \quad (1)$$

Also, at node s_j, where $j \in \{0, 1\}$, the ex ante maximum expected utility from further experimentation is

$$\bar{v}(s_j) := \sum_{s_{j+1} \in S_{j+1}} \Pr(s_{j+1} \mid s_j) \max\{\underline{v}(s_{j+1}), \bar{v}(s_{j+1}) - c\}, \text{ and} \quad (2)$$

$$v(s_2) := \bar{v}(s_2),$$

[3] By Blackwell's (1951) celebrated theorem, it is straightforward to see that it is not optimal to form a link to the second, rather than the third agent. Similarly, it is not optimal to form a link to the first and then to the third agent.

because at each s_j an agent compares $\underline{v}(s_j)$ against $\bar{v}(s_j) - c$ to decide whether to experiment with X_{j+1} or to take the optimal action at s_j. To capture this, finally we define what we refer to as the *expected value of information* from further experimentation by

$$v(s_j) := \bar{v}(s_j) - \underline{v}(s_j) - c \quad \text{for } j \in \{0, 1\}. \tag{3}$$

Therefore, at s_j, an agent decides to further experiment if $v(s_j) > 0$; otherwise, he takes the optimal action at s_j without experimenting X_{j+1}.

By using the Equations (1)–(3) and backward induction, we can fully describe the optimal strategy of an agent. The following proposition formally states the complete characterization of the BSLP that we discussed.

Proposition 1. *The optimal policy of the Bayesian sequential link procedure for an agent is characterized by a pair* (τ, a^*) *such that*

$$\tau = \min\{j \in \{0, 1, 2\}: v(s_j) \le 0\}, \tag{4}$$

$$a^* \in \arg\max_a \left\{ \sum_\theta \Pr(\theta \mid s_\tau) u(a, \theta) \right\}. \tag{5}$$

Proposition 1 and the value of information, as defined by (3), provide us with the full characterization of the decision problem. In words, an agent stops at the first node at which the value of information is negative; otherwise, he keeps forming links. If the agent stops at the node s_τ, then he takes the action that maximizes his expected utility given his information. The equilibrium results that we present in the following section are derived by use of this characterization.

2.3. Theoretical Predictions

In this section we provide an intuitive discussion of the equilibrium behavior of agents $i \in \{1, 2, 3, 4\}$ and also enumerate the equilibrium link structures that can be observed in the information acquisition problem. In Appendix A we formally derive these results as corollaries to Proposition 1.

2.3.1. Optimal Behavior of Agents. *First Agent.* The decision problem of the first agent is easy because there is no preceding agent, and thus he takes an action based only on his private signal. If the first agent is informed, then he follows his signal, whereas if he is uninformed, he randomizes between the two possible actions. Therefore, unless $q = 1$, the second agent cannot determine the status of the first agent as either informed or uninformed. Figure 2 depicts this situation. We use the diamond to refer to an agent whose status cannot be determined.

Second Agent. The second agent faces a more interesting problem because he has the option to observe the first agent's action. Here, optimal behavior depends on whether or not the agent is informed.

Figure 2 After the First Agent

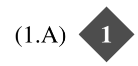

If the second agent is informed, by applying Proposition 1, we find that he does not form a link for any positive cost. We provide the formal statement (Corollary 1) in Appendix A. Intuitively, this is easy to see: if the second agent incurs the cost c and forms a link to the first, either he will observe an action that is the same as his own private information or he will observe the opposite action. For $q < 1$, he will favor the action that is in line with his signal, whereas when $q = 1$, he will become indifferent between the two actions. This suggests that forming a link does not provide any value to an informed second agent because he can always save the link cost and follow his own signal.

On the other hand, if he is uninformed, then it is optimal to form a link to the first agent if the cost is low enough, because there is a positive probability that the first is informed. If the cost is high, an uninformed second agent does not form a link; that is, there is a threshold cost of link formation, c^*, such that the uninformed second agent forms a link if and only if $c \le c^*$. It is also shown that the threshold value c^* increases linearly in the probability of being informed, q, and in the payoff from a correct action, m. Intuitively, when q is higher, the first agent is more likely to be informed, and thus, from the perspective of an uninformed second agent, it becomes more valuable to form a link to the first agent. In a similar vein, when the payoff from a correct decision is larger, the same amount of information increases the potential benefit from that information. Figure 3 depicts the link structures that can emerge in the equilibrium for $0 < c < c^*$, and $c \ge c^*$. The square indicates the event in which it can be deduced that the agent is uninformed, whereas the circle indicates that he is informed.

Third Agent. Note that when $c < c^*$, the third (and fourth) agent(s) can discern whether the second agent is uninformed ((2.A) in Figure 3) or informed ((2.B) in Figure 3) simply by observing his link decision. As such, the link structure itself contains valuable information and significantly affects the behavior of all agents coming after the second agent.

Figure 3 After the Second Agent

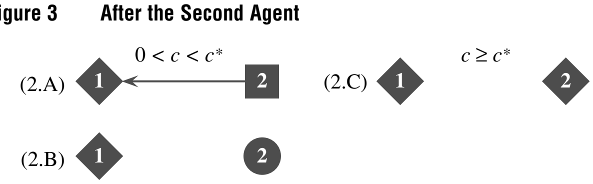

Figure 4 After the Third Agent

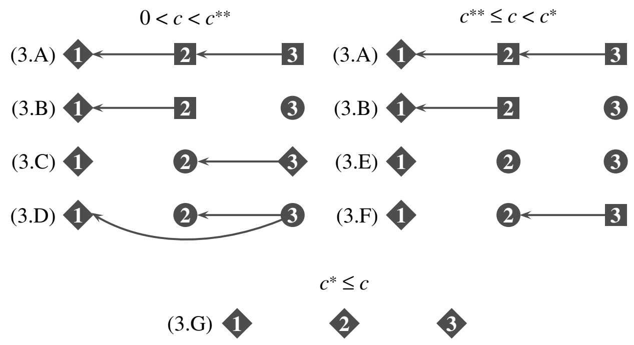

The decision problem for the third agent becomes interesting because the link structure becomes more significant in determining the optimal behavior. When the third agent observes a link between the second and the first agents, he rationally infers that the second agent is uninformed, and his action conveys no information. In this case, his problem is equivalent to the second agent's problem: if he is informed, he does not form a link for any positive cost, whereas if he is uninformed, he forms a link to the second agent for $c < c^*$.[4] The link structures (3.A) and (3.B) in Figure 4 exhibit these situations.

Suppose that $0 < c < c^*$ and the third agent faces the link structure (2.B) in Figure 3. When the third agent observes that there is no link between the first and the second agents, he knows that the second agent is informed. However, he knows that the first is informed only with probability $q \in (0, 1]$. Therefore, an uninformed third agent forms a link to the second and imitates his action ((3.C) and (3.F) in Figure 4). Note that after observing the second agent, the uninformed third agent is exactly the same as an informed second agent. Hence, a link to the first agent is worthless.

On the other hand, if the third agent is informed, we find that the value of information is positive only when the cost is low enough (i.e., $c < c^{**} = 2qm/39$). Suppose that $c < c^{**}$. An informed third agent starts forming a link to the second. If he finds out that the action of the second agent is the same as his signal, he imitates the second's action without forming a link to the first ((3.C) in Figure 4). Otherwise he proceeds with a link to the first agent ((3.D) in Figure 4). Because there is a chance an informed third agent does not form a link to the third, the fourth agent facing the link structure (3.C) in Figure 4 cannot tell whether the third is informed or not. On the other hand, when $c^{**} \leq c < c^*$, an informed third agent does not form a link to the second agent ((3.E) in Figure 4).

[4] Note that it is also optimal to form a link to the first agent, and the third observes only the first agent's action. To get around an unnecessary multiplicity of equilibria, we assume that an agent starts to form a link to the closest agent in a line.

Figure 5 The Fourth Agent's Problem: An Example

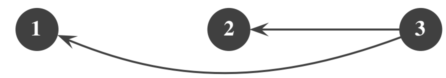

Finally, when the cost is high enough (i.e., $c > c^*$), the third agent observes that the second agent did not form a link. Because the third agent cannot distinguish whether the second agent is informed or not, he faces two opportunities of linking with equal amounts of information ((3.G) in Figure 4). However, because of the high cost, it is never optimal for him to form any link, regardless of whether he is informed or not.

We formally state this discussion in Appendix A (Corollary 2).

Fourth Agent. Because most of our attention will be focused on the first three agents in the analysis of the experimental data, we will not provide a full characterization of optimal behavior for agents in the fourth position. We can point out, however, that often the problem of the fourth agent will be strategically equivalent to that of either the second or the third agent. For example, consider the case where the cost of link formation is low enough, and the fourth agent observes the link structures displayed in Figure 5.

Given this link structure, the fourth agent can infer that both the second and third agents were informed. He can also infer that the signals of the second and third agents must have disagreed (because the third agent also formed a link to the first). Therefore, the signals of the second and third agents essentially cancel, leaving the only relevant signal that of the first agent—exactly as it is for the second agent. Therefore, if the fourth agent is informed, he should not link for any positive cost, whereas if he is uninformed, he should link provided that c is small enough. To be sure, there are situations in which the fourth agent's problem is not equivalent to that of one of his predecessors. However, because this is not our main focus, and the intuition is relatively straightforward, we will omit such details.

2.3.2. The Effect of Cost on Welfare. One noticeable feature of these predictions is the amount of information revealed at different cost levels. Indeed, starting from the third agent, the structure of previous links reveals valuable information. For an arbitrarily small cost, when the probability of being an uninformed agent is large enough, the ex ante welfare of the third agent is larger than what it is when the cost is zero. The intuition is simple. Recall that if the cost is positive, an informed second agent never forms a link to the first. This situation, in turn, reveals whether or not the second agent is informed. However, when the cost is zero, because there is always a link between agents, succeeding agents cannot make a similar deduction. When $1-q$ is high enough, the benefit of an uninformed third agent from this extra information dominates the loss of an informed second agent by not observing the first agent.

To demonstrate this claim, we compute the ex ante value of information $v(\cdot)$ conditional on the state of the world both when $c=0$ and $c>0$ for the third and the fourth agents; that is, we compute

$$v_i^c := \sum_{s\in S_i} v(s)\Pr(s \mid \theta=1),$$

where the subscript $i\in\{3,4\}$ denotes the agent, S_i is the set of all possible information nodes for agent i, and the superscript c is the cost.[5] We state this observation in what follows.[6]

Observation 1. There exists $\bar{q}<1$ such that for all $q\le\bar{q}$ and small enough $c>0$, we have $v_i^c > v_i^0$ for $i\in\{3,4\}$.

3. Experimental Design

The experiment was run at the Center for Experimental Social Sciences at New York University. The 144 subjects in this experiment were recruited from undergraduate classes at New York University and had no previous experience in social learning experiments. In each session, after subjects read the instructions they were also read aloud by an experimental administrator.[7] Each session lasted for about 1 hour and 15 minutes, and each subject participated in only one session. An \$8.00 participation fee and subsequent earnings, which averaged about \$12.26, were paid in private at the end of the session. Throughout the experiment, we ensured anonymity and effective isolation of subjects to minimize any interpersonal influences that could stimulate uniformity of behavior.[8]

In each session, subjects faced the information acquisition problem described in §2.1, for 40 independent rounds. At the beginning of each session, subjects' positions were randomly assigned as 1, 2, 3, or 4, and their positions were held fixed for the duration of the experiment. In each round, the cost of link formation, in experimental points, was a random draw from the set $\{0,2,\ldots,18,20\}$, with each cost equally likely. In each round, all members of a group faced the same cost of link formation, and this was common knowledge. In all sessions, the informativeness of the signal was held fixed at $p=2/3$.

Table 2 Summary of Experiments

	Number of subjects	c	q	p	Number of rounds
Session 1 (HI)	20	Random	1	2/3	40
Session 2 (HI)	16	Random	1	2/3	40
Session 3 (HI)	16	Random	1	2/3	40
Session 4 (HI)	8	Random	1	2/3	40
Session 5 (HI)	8	Random	1	2/3	40
Session 6 (LI)	20	Random	2/3	2/3	40
Session 7 (LI)	16	Random	2/3	2/3	40
Session 8 (LI)	16	Random	2/3	2/3	40
Session 9 (LI)	24	Random	2/3	2/3	40

We conducted two treatments, which we call (i) the high information treatment and (ii) the low information treatment. In the HI treatment, all the subjects received an informative signal (i.e., $q=1$), whereas in the LI treatment, we set $q=2/3$—hence, with probability 1/3 a subject received an uninformative signal. Table 2 summarizes the details of our experiment.

In each round, subjects earned $m=100$ points for correctly guessing the state and 0 points otherwise. Their net point total was determined by subtracting the appropriate number of points for each link that a subject made from the points they collected in that round for guessing the state. At the end of the experiment, the computer randomly selected three rounds for which subjects would be paid. The total number of points was then converted back to dollars at the rate of \$1.00 = 15 points.

3.1. Some Remarks on the Design

3.1.1. Subjects' Positions. In the experiment, subjects engage in a two-step decision process. Our decision to fix the subjects' positions throughout the session aims to allow them to develop a *strategy* and play accordingly. Therefore, for fixed p and q, subjects decide whom to link to at different levels of cost. Note that in each session there were either 16 or 20 subjects, and the groups (of four subjects) were reshuffled for each round. Therefore, a subject knew that in each round he was a member of a (possibly) different group.

3.1.2. Random Cost. As §2.3 discussed in detail, the cost of link formation is one of the critical parameters that affects rational behavior according to the BSLP. Therefore, having a subject play at different cost levels, randomly drawn for each of 40 rounds,

[5] We assume that when the cost is zero, all agents form a link.

[6] The computations are simulated by Matlab. The source file is available upon request.

[7] At the end of the first round, subjects were asked whether there were any misunderstandings. No subject reported any problems with understanding the procedures or using the computer program.

[8] Participants' workstations were isolated by cubicles, making it impossible for participants to observe others' screens or to communicate. We also made sure that all remained silent throughout the session. At the end of a session, participants were paid in private according to the number on their workstation.

allows us to observe how his behavior responds to the cost of link formation. The costs were chosen so that subjects in each (nontrivial) position would experience link costs such that the optimal action would be to form a link (if c is low) or not form a link (if c is high). In each decision round, the cost of link formation was the same for all members of a group, and the cost was known at the beginning of each round. This design feature was common knowledge to all subjects.

3.1.3. HI and LI Treatments. Similar to the level of cost, different values of q generate different behaviors according to the BSLP.[9] The two treatments, HI and LI, give us the opportunity to compare, across treatments, the effect of varying the probability with which subjects receive a private signal and determine whether the predicted comparative statics hold true.

3.1.4. Payoffs. We randomly chose three rounds to determine the payoffs of the subjects. Cubitt et al. (1998) demonstrates that this incentive scheme generates reliable data. Such a design helps to ensure homogenous behavior during the experiment by diminishing the salience of wealth effects.

4. Descriptive Analysis

We organize our descriptive analyses into two categories that focus on the linking and action decisions of subjects. We first report the linking decisions and then focus on the action decisions as a function of the observations that are made via the linking decisions.

Throughout this paper, we take the subjects as the unit of observation in our analysis. Therefore, when conducting tests, for each subject, we average over all relevant decisions and then proceed to conduct the appropriate test based on subject averages.[10]

4.1. Understanding Link Decisions

4.1.1. The Second Subject. Table 3 reports the decomposition of the second subjects' link decisions in the HI and LI treatments. In particular, it shows the frequency with which a link is formed when $c = 0$ and $c > 0$ as well as across all cost values. The shaded cells indicate when a link should form according to theory.

In the HI treatment, the theory predicts that a subject in the second position should not form a link unless the cost is zero. However, we see that subjects form a link to the first agent 33% of the time when the cost of link formation is strictly positive. We observe a similar pattern in the LI treatment: Although, theoretically, an informed subject should never form a link to the first subject when the cost is positive, we see that this event happens 30.6% of the time.[11] Even though the subjects tend to form links slightly less frequently (when $c > 0$) in the LI treatment (30.6%) than in the HI treatment (33.0%), the difference is not statistically significant ($t_{34} = 0.24$, $p = 0.81$).[12] These findings suggest that subjects exhibit a strong bias in favor of link formation, even when it is not optimal to do so.

Table 3 **Linking Behavior of the Second Subject in the HI and LI Treatments**

	Frequency of link formation		
	Overall	$c = 0$	$c > 0$
HI	0.362 (0.315)	**0.742** (0.392)	0.330 (0.313)
	Overall	$c = 0$	$c > 0$
LI (informed)	0.359 (0.280)	**1.000** (0.000)	0.306 (0.294)
	Overall	$c < 11$	$c > 11$
LI (uninformed)	0.653 (0.311)	**0.800** (0.309)	**0.480** (0.398)

Notes. Standard deviations, based on subject averages, are in parentheses. For informed subjects, it is optimal to form a link only when the cost is 0, whereas for uninformed subjects, it is optimal to form a link provided that the cost is below 11.

From the discussion of §2.3 (and in particular Figure 3) we know that the optimal behavior of uninformed second subjects depends on the level of cost. Given the parameters of the experiment, an uninformed subject should always form a link when the cost is less than 11, and should avoid forming a link when it is above 11. In fact, we observe a significant difference in the behavior when the cost is above and below 11. As Table 3 shows, whereas subjects link 48% of the time when $c > 11$, they link 80% of the time when $c < 11$. A matched-pairs test rejects the null hypothesis that these frequencies are equal ($t_{18} = 4.23$, $p < 0.01$).

Observation 2. For the subjects in the second position, linking behavior is not substantially different

[9] The exact behavioral differences are evidenced in the threshold costs of link formation. In general, the higher the probability that subjects receive signals, the higher the threshold cost.

[10] When called for, we condition on a particular event such as whether or not the subject is informed or whether the cost of link formation is in a particular range.

[11] Although it is not a mistake to link to the first subject when the cost is zero, because of the chance that the first subject did not receive a signal, the second subject should always ignore the action taken by the first.

[12] Specifically, for each subject, we calculate the average frequency of link formation, conditional on $c < 11$, and the average frequency of link formation, conditional on $c > 11$, and then conduct a matched-pairs t-test (using subject averages as the unit of observation). We come to a similar conclusion using the nonparametric signed rank test.

across treatments. In both treatments, contrary to the prediction of the theory—that it is suboptimal to form a link—informed subjects link approximately 30% of the time when the cost of link formation is positive.

4.1.2. The Third Subject. The subjects in the third position face two possible link structures: either there is a link between the first and the second subjects or not. Table 4, panels (A) and (B), summarizes subjects' link decisions in the HI and LI treatments, respectively, for these two link structures. The data are organized according to the critical cost threshold, the treatment, and which link structure the subjects observe. For example, in the HI treatment, when the third subject does not observe a link between the first and the second subjects, he should form a link to the second if and only if the cost of link formation is less than 5. Similarly, a cost of 3 is the critical threshold for an informed subject in the LI treatment, and a cost of 11 is the critical threshold for an uninformed subject in the LI treatment. The shaded cells indicate when a link should form according to theory.

We first observe that in all cases the subjects link more frequently when the cost is below the threshold than when it is above the threshold. Moreover, in four (bold in Table 4) of all six instances, a matched-pairs test rejects the null hypothesis that the frequencies above and below the thresholds are the same at the 5% level or better.

Second, in all the cases, subjects are more likely to form a link when there is a link between the first and second subjects. If we focus on the overall frequency of link formation, then in the HI treatment, the frequency of link formation is 69.2% when there is a link between the first and second subjects, but only 32.1% when there is no link. A matched-pairs test rejects the null hypothesis with $p < 0.01$. In the LI treatment, for informed subjects the frequencies are 59.7% and 26.9%, respectively, whereas for uninformed subjects the frequencies are 81.2% and 49.5%. In both cases, a matched-pairs test rejects the null hypothesis that these frequencies are equal with $p < 0.01$.

The comparison of the linking behavior of the second and third subjects (compare Tables 3 and 4) provides us with the first evidence of what we call *herding in link formation*. Before we focus on the comparison, let us explain what we mean by this term. Consider two link structures, l_1 and l_2, and assume that (i) there are more links in l_1 than l_2 and that (ii) l_1 is not more informative than l_2. Then, we say that there is herding in link formation if there is a higher propensity to form a link to l_1 than to l_2. In other words, higher number of preexisting links makes subjects more likely to link, even though doing so does not necessarily provide more information.

Now, we compare the linking behavior of the second and the third subjects when there is a link between the first and the second subjects. From our earlier discussions we know that in the LI treatment these two observations are equally informative simply because an informed subject never forms a link to the first subject. When we compare Tables 3 and 4 we see that in the LI treatment, informed subjects in the second position formed a link to the first subject 35.9% of the time. When there is a link between the first and second subjects, the subjects in the third position link 59.7% of the time—an increase of 23.8 percentage points. Also, for uninformed subjects in the LI treatment, we find that the same difference is 15.9 percentage points. So, the subjects are more prone to form a link when they observe links between previous subjects, although it is actually less informative.

Our findings, so far, suggest that subjects' linking behavior is substantially and systematically different from the predictions of the theory. To further investigate this issue, let us compare linking behavior of informed subjects in the LI and HI treatments. We make two critical observations. First, when there is no link between the first and the second subjects, and the cost is below the critical threshold, the frequency of link formation by the third subject is actually lower in the LI treatment than in the HI treatment (although the difference is not significant ($t_{22} = 1.41$, $p = 0.17$)).

Table 4 Linking Behavior of the Third Subject in the HI and LI Treatments

(A) HI treatment

Observed link structure	Frequency of link formation: Overall	$c < 5$	$c > 5$
1 2	0.321 (0.280)	**0.700** (0.397)	**0.236** (0.275)
1 ← 2	0.692 (0.312)	**0.850**[a] (0.336)	**0.575**[a] (0.361)

(B) LI treatment

Observed link structure	Informed subjects: Overall	$c < 3$	$c > 3$	Uninformed subjects: Overall	$c < 11$	$c > 11$
1 2	0.269 (0.330)	0.438 (0.496)	0.260 (0.323)	0.495 (0.392)	**0.647** (0.464)	**0.401** (0.418)
1 ← 2	0.597 (0.344)	**0.804** (0.334)	**0.486** (0.393)	0.812 (0.304)	0.840 (0.301)	0.622 (0.481)

Notes. Standard deviations, based on subject averages, are in parentheses. Bold cells indicate that the difference of linking frequency between "high" and "low" cost is significant at the 5% level or better. The costs in the table represent the critical thresholds, below which (according to the theory) it is optimal to form a link.

[a]In these cases, the subjects observe an out-of-equilibrium link structure (i.e., there is a link between the first and the second subjects).

This is remarkable because, even though (i) the cost is, on average, lower in the LI treatment, and (ii) the lack of link between the first and second theoretically means that the second subject is actually informed, subjects find linking more attractive in the HI treatment. Second, when there is a link between the first and second subjects, there is virtually no difference between the frequency of link formation between the HI and LI treatments ($t_{34} = 0.87$, $p = 0.39$). This is also interesting because, theoretically, in the LI treatment, a link between the first and the second subjects reveals that the second subject in uninformed. In such cases, the informed third subject should never form a link. Yet, we observe that they do form a link as much as in the HI treatment.

We summarize our discussion of the third agent as follows:

Observation 3. Link decisions of the subjects in the third position systematically deviate from the predictions of the theory: We observe that subjects in both treatments are significantly more likely to form a link when they observe a link between the first and second subjects even though it is not necessarily optimal to do so (especially in the LI treatment).

4.1.3. The Fourth Subject. Table 5 shows the linking behavior of the subjects in the fourth position in the HI and LI treatments for each one of the observed link structures. First observe that, across all link structures, subjects never link significantly more frequently in the LI treatment than in the HI treatment, although in two cases (highlighted in Table 5) the frequency of link formation is significantly higher in the HI treatment.

We observe herding in link formation also for the fourth subjects. Specifically, subjects are substantially more likely to form a link to the third subject if, by doing so, they will observe all of the actions of their predecessors. To be more specific, consider the LI treatment and compare the third row (where the only link is between the second and third subject) and the fourth row (where the third subject linked to both the first and the second subject) of Table 5. In the former case, according to the theory, one can infer that the second subject was informed and that the third subject was either informed (and the signals of the second and third agreed) or uninformed. In the latter case, not only can we infer that both the second and third subjects were informed, but also that they received different signals. This makes the decision of the first subject—who may or may not be informed—the tiebreaker. Thus, the former link structure is actually more informative than the latter. However, the linking frequency was only 28.2% in the more informative structure, and 60.8% in the less informative structure.[13]

We can make a similar argument for the link structures in the third and the fifth rows. In a similar vein, when the fourth subject observes a link between the first and second subjects while the third subject has no link, the fourth subject infers that the second subject is uninformed, and the third subject was informed. Thus, if the fourth subject forms a link, he should do so to the third subject. However, almost 80% of the time, when a subject forms a link, it is actually to the second subject.

Table 5 **Linking Behavior of the Fourth Subject in the HI and LI Treatments**

	Frequency of link formation		
Observed link structure	HI	LI (informed)	LI (uninformed)
1 2 3	0.241	0.236	0.419
	(0.289)	(0.315)	(0.434)
$1 \leftarrow 2$ 3	0460	0.474	0.525
	(0.381)	(0.369)	(0.465)
1 $2 \leftarrow 3$	**0.529**	**0.282**	0.458
	(0.371)	(0.381)	(0.500)
$1 \leftarrow 3$, $2 \leftarrow 3$	0.824	0.608	0.611
	(0.295)	(0.439)	(0.486)
$1 \leftarrow 2 \leftarrow 3$	**0.873**	**0.681**	0.801
	(0.203)	(0.270)	(0.330)

Notes. Standard deviations, based on subject averages, are in parentheses. Bold cells indicate that the difference between treatments is significant at the 10% level or better.

4.2. Understanding Action Decisions

In this section, we aim to understand how subjects, after having possibly gathered information through their link decisions, make their action choices. We organize our discussion around a variable that we call `count`. Specifically, let $a_i \in \{-1, 1\}$ denote the action decision of subject i, and $\sigma_i \in \{-1, 0, 1\}$ by subject i's signal. Finally, let $s_i^j = 1$ if subject i observes the action of subject $j < i$, and let $s_i^j = 0$ otherwise. Then, we define

$$\texttt{count}_i = \sigma_i + \sum_{j<i} s_i^j a_j.$$

In fact, the variable `count` characterizes the optimal action decision on the equilibrium path of the game: A subject i chooses $a_i = 1$ if $\texttt{count}_i > 0$ and $a_i = -1$ if $\texttt{count}_i < 0$. In the HI treatment, for $\texttt{count}_i = 0$, subject i is indifferent between the two actions. In the LI treatment, for $\texttt{count}_i = 0$, an informed subject i chooses the action that matches his own signal because some of his predecessors may have been uninformed. In what

[13] A logit regression where the dependent variable takes a value of 1 if there is a link to the third subject and 0 otherwise, with independent variables for the link cost and dummies for the two networks under consideration, easily allows us to reject that the linking frequencies are the same ($p = 0.014$).

follows, for each of the second, third, and fourth subjects, we analyze the frequency of decisions that are consistent with count. Note also that a decision maker who makes his action decision according to the counting heuristic should not change his linking behavior. For example, an informed second subject should never form a link for any positive cost because either $\texttt{count} = 0$ or $|\texttt{count}| = 2$. In both cases, the subject will make the same decision as she would without forming any link; therefore, it is better to save the cost and not to link. Thus, we feel that the counting heuristic represents a sensible and cognitively simple rule of thumb that will be close to optimal provided that the frequency of mistakes is not too large.

4.2.1. Action Decisions When $|\texttt{count}| > 0$. Table 6 summarizes the action decisions of subjects based on whether or not they are consistent with the variable when $|\texttt{count}| > 0$. The decision rule does a very good job of organizing the data. In fact, on average, *at least* 88.9% of decisions and often 100% of decisions are consistent with count. A couple slight anomalies are that informed subjects in the third position, when faced with $|\texttt{count}| = 1$, and informed subjects in the fourth position, when faced with $|\texttt{count}| = 2$, are significantly less consistent than their uninformed counterparts in the same position.

Furthermore, a random effects Tobit model confirms that the frequency of decisions that are consistent with count is increasing in $|\texttt{count}|$; that is, as the evidence in favor of a particular action increases, the probability that this action is taken increases. To be precise, let $c_{i,j}$ denote the frequency that subject i makes the correct action decision when $|\texttt{count}| = j$. Then our regression model is given by

$$c_{i,j} = \begin{cases} c^*_{i,j} & \text{if } c^*_{i,j} \in (0,1), \\ 0 & \text{if } c^*_{i,j} \le 0, \\ 1 & \text{if } c^*_{i,j} \ge 1, \end{cases}$$

where $c^*_{i,j} = \beta_0 + \beta \cdot j + \mathbf{x}_\mathbf{i}'\mathbf{\Gamma} + \epsilon_{i,j} + \nu_i$ is the latent variable. Note that $j \in \{1, 2, 3, 4\}$ and captures the value of $|\texttt{count}|$, whereas $\mathbf{x}$ is a vector of control variables, Γ is the coefficient vector for the control variables, and $\epsilon_{i,j}$ and ν_i are error terms. Estimation results indicate that the marginal effect (at the mean of the independent variables) on j is 0.048 ($p < 0.01$), which means that for every unit increase in $|\texttt{count}|$, the frequency of consistent decisions increases by about 4.8 percentage points.

4.2.2. Action Decisions When $\texttt{count} = 0$. The case of $\texttt{count} = 0$ is particularly interesting because it allows us to determine subjects' tiebreaking rule when their signal goes against the majority opinion among those whose actions the subject observed. We say that a subject's action is "conformist" if his action follows the majority opinion, which is inconsistent with his own signal. Table 7 reports the frequency of conformist actions. We observe that for all subjects, the frequency of conformist behavior is higher in the HI treatment than in the LI treatment, but the difference is statistically significant only for the second subject.[14] Such a result is to be expected because in the LI treatment, a subject must entertain the possibility that the subject whose action he chose to observe was actually uninformed. Therefore, as noted above, even when $\texttt{count} = 0$, an informed subject should still follow his

Table 6 Frequency of Choices Consistent with count

|count|	Frequency		
	HI	LI (informed)	LI (uninformed)
	(A) The second subject		
1	0.898	0.922	0.910
	(0.159)	(0.153)	(0.170)
2	0.962	1.000	
	(0.010)	(0.000)	
	(B) The third subject		
1	0.900	**0.921**	**0.981**
	(0.136)	(0.085)	(0.069)
2	0.962	0.888	0.919
	(0.139)	(0.183)	(0.166)
3	1.000	0.989	
	(0.000)	(0.043)	
	(C) The fourth subject		
1	0.912	0.903	0.901
	(0.151)	(0.145)	(0.200)
2	**1.000**	**0.910**	**1.000**
	(0.000)	(0.144)	(0.000)
3	1.000	0.909	1.000
	(0.000)	(0.302)	(0.000)
4	0.990	1.000	
	(0.040)	(0.000)	

Notes. Standard deviations are in parentheses. Bold cells indicate that the adjacent columns are significantly different at the 5% level or better.

Table 7 Frequency of Conformist Actions

Subject	Frequency of conformists	
	HI	LI (informed)
DM 2	**0.528**	**0.240**
	(0.397)	(0.352)
DM 3	0.524	0.333
	(0.466)	(0.471)
DM 4	0.463	0.391
	(0.373)	(0.408)

Notes. Standard deviations are in parentheses. Bold cells indicate that the adjacent columns are significantly different at the 5% level or better.

[14] Pooling across all subjects, we can also reject the null hypothesis that subjects conform with equal frequency.

Figure 6 Frequency of Conformity by Informed Subjects When `count` = 0

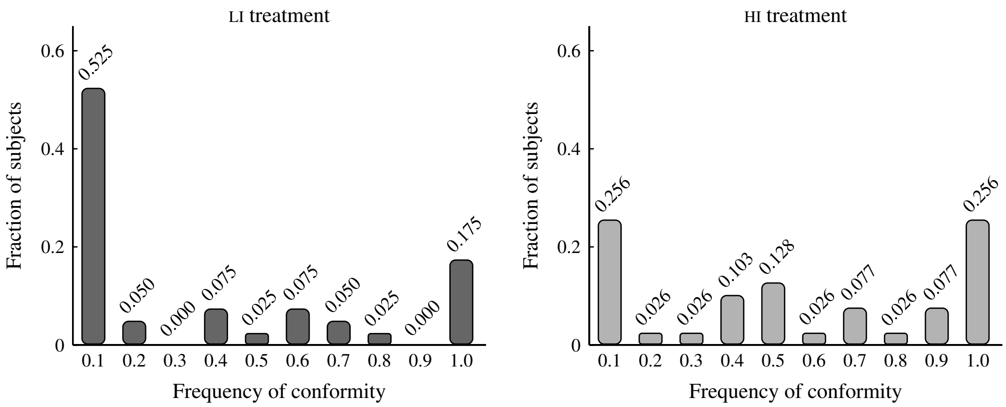

own signal. Despite the fact that conformist behavior is less frequent in the LI treatment, we still observe a nonnegligible amount of such behavior. To better understand conformity behavior, we look at the data at the individual level.

Figure 6 plots two histograms showing the average frequency of conformist behavior for subjects in LI and HI treatments. Clearly, there is a lot of heterogeneity in conformist behavior. Although the distribution is skewed leftward (i.e., less frequent conformist behavior) in the LI treatment, we see that in both treatments, the distribution is bimodal, with many subjects almost never conforming and many subjects almost always conforming, and there is also a group of subjects who conform some of the time.

We summarize the results of our analysis of subjects' action decisions as follows:

Observation 4. The action decisions of subjects are well described by `count`. In both treatments, some subjects have a tendency to conform to the action(s) of one's predecessor(s) when `count` = 0. At least for the second subject, this tendency is significantly more pronounced in the HI treatment.

5. Further Data Analysis

This section aims to achieve three goals: (1) We further investigate the "herding in link formation" phenomenon that we came across in the descriptive analyses. (2) Also, we want to investigate the effect of cost on the behavior of subjects. (3) Finally, we briefly report on subject-specific heterogeneity and summarize the results of a more structural econometric approach that is undertaken in Appendix B.

Table 8 reports a random effects logit procedure used to see how linking behavior depends on cost and the number of links formed by one's predecessor. We pool the data across the second, third, and fourth subjects, but we include dummy variables for the third and fourth subjects. `LL` is the log-likelihood, `N` is the total number of observations, and `subjects` is the number of subjects in the regression. Finally, ρ captures the fraction of the variance that is due to subject-specific heterogeneity (i.e., the random effects). For the coefficients, we report the marginal effects for each variable. The variable `#links` ≥ 1 is a dummy variable that takes a value of 1 if the subject observed one or more links by her predecessors, and the variable `#links` = 2 is a dummy variable that takes a value of 1 if the subject observed exactly two links by her predecessors. Because of the way we define these variables, the coefficient on the former captures the effect of going from zero to one observed link, whereas the coefficient on the latter captures the effect of observing a second link. In particular, a positive

Table 8 Random Effects Logit Regressions: Linking Behavior by the Second, Third, and Fourth Subjects

	HI	LI (informed)	LI (uninformed)
`cost`	−0.067***	−0.060***	−0.045***
	(0.007)	(0.009)	(0.014)
`#links` ≥ 1	0.354***	0.407***	0.207***
	(0.058)	(0.066)	(0.079)
`#links` = 2	0.371***	0.157**	0.096*
	(0.085)	(0.072)	(0.055)
`DM 3`	0.076	−0.175	−0.081
	(0.221)	(0.222)	(0.152)
`DM 4`	−0.004	−0.234	−0.314
	(0.218)	(0.211)	(0.233)
ρ	0.673***	0.712***	0.793***
	(0.050)	(0.050)	(0.048)
`LL`	−674.2	−572.1	−275.9
`N`	2,040	1,544	736
`subjects`	51	57	57

Notes. Standard errors are in parentheses. The table reports the marginal effects assuming that the random effect is zero, calculated at mean of independent variables except as noted. Marginal effects for `#links` ≥ 1 assumes `#links` = 2 is equal to 0. Marginal effects for `#links` = 2 assumes `#links` ≥ 1 is equal to 1. Marginal effects for `DM 3` assumes `DM 4` is equal to 0. Marginal effects for `DM 4` assumes `DM 3` is equal to 0.

*Significant at the 10% level; **significant at the 5% level; ***significant at the 1% level.

coefficient indicates a tendency toward herding. The variable `cost` is simply the cost of link formation.

In the next two subsections, we will discuss the results of Table 8, first focusing on responsiveness to cost, and then on herding.

5.1. Responsiveness to Cost

Whereas theoretically it is a mistake for an informed second subject to form a link to the first subject, it is even a bigger mistake to do so when the cost is higher.[15] Therefore, one would expect a negative relationship between the frequency of forming links and the cost of doing so. Not surprisingly, this is exactly what we observe in the first row of Table 8. In all cases, the coefficient is negative and significant at the 1% level. Therefore, the higher the cost is, the less likely subjects are to form links. We state this in the following observation:

Observation 5. The frequency with which subjects form links is decreasing with the cost of link formation. In particular, for each unit increase in the cost of link formation, the probability that a link is formed declines by between 4.5 and 7.4 percentage points (assuming that the random effect is 0).

5.2. Herding in Link Formation

Table 8 helps us to understand the issue of herding in link formation in more detail. Specifically, we see that the coefficient on `#links` $=2$ is always positive. Therefore, subjects are substantially more likely to form a link when they observe two existing links than when they merely observe one. If we compare the magnitude and significance of this marginal effect across treatments, we see that it is smaller in magnitude and only significant at the 5% level among informed subjects in the LI treatment, whereas the effect is large and significant at the 1% level in the HI treatment. Thus, we observe herding in link formation in all cases, but the fact that the signals are more informative in the HI treatment makes people herd even more.

We summarize these results as follows:

Observation 6. There is a strong tendency toward herding in link formation: With each successive link, the next subject is ever more likely to form a link (assuming that the random effect is 0). This phenomenon is present in both the HI and LI treatments, but is more pronounced in the former.

5.3. Welfare and Linking Behavior

One of our theoretical predictions (see Observation 1) was that the ex ante expected utility would be higher for a small cost of link formation than what it would be for zero cost. The reason behind this prediction is that a second informed agent does not form a link when the cost is positive and, as a result, reveals that his action is based on a more informative signal. Therefore, such a valuable source of information helps increase the ex ante welfare of succeeding agents.

Our earlier results clearly indicate that the subjects are not as strategic as the theory depicts them. Nevertheless, the subjects are still responsive to the cost, and as a result the behavior is significantly different across treatments. Therefore, even though linking decisions are less strategic vis-à-vis theory, they are still influenced by the informational content of the environment.

To test our hypothesis, let us look at Table 9, which exhibits the average earnings and prediction accuracy of subjects—which is the fraction of decisions that are correct—as a function of the cost of link formation. To compute this, for each subject and each cost level, we compute the average earnings and average prediction accuracy. Then, at each cost level, we average over all subjects. If we look at the columns labeled "Overall," we can see that both the earnings and the prediction accuracy are higher at $c=2$ and $c=4$ than those at $c=0$, before dropping off quite substantially at $c=6$ and leveling off thereafter. The results are similar in other columns, though in the HI treatment, the peak appears to occur at $c=2$.

The two left-hand columns of Table 10 seek to test the significance of this apparent finding through a random effects regression. We include dummy variables for $c=2$, $c=4$, and $c=6$ and thereafter assume that cost enters linearly via the term $c\cdot(c>6)$; that is, we multiply the cost by the dummy variable $(c>6)$, which takes a value of 1 for $c>6$ and a value of 0 otherwise. The baseline, therefore, is $c=0$, and the significant positive coefficients on $c=2$ and $c=4$ indicate that both the earnings and the prediction accuracy are significantly higher when subjects face a small cost of link formation.

In theory, by forming links, subjects should gain an informational advantage and predict the state more accurately. Of course, in practice, this need not actually be true because of deviations from the rational benchmark. We briefly investigate whether or not, even ignoring the cost of forming links, linking increases subjects' predictive power. Our results are summarized in the two right-hand columns of Table 10, where we regress the earnings and the prediction accuracy on a dummy variable for whether or not the subject formed at least one link (`link`), as well as other explanatory variables.[16] As can be seen,

[15] Similar logic applies for subjects in the third and fourth positions.

[16] The dummy variable `informed` indicates whether or not the subject is informed, and `DM 3` and `DM 4` are dummy variables that indicate the position of the subjects.

Table 9 Analysis of Prediction Accuracy and Earnings by Link Cost Pooling Over the Second, Third, and the Fourth Subjects

	Average earnings				Prediction accuracy			
cost	Overall	HI	LI ($\sigma \neq 0$)	LI ($\sigma = 0$)	Overall	HI	LI ($\sigma \neq 0$)	LI ($\sigma = 0$)
0	58.50	60.72	61.47	52.03	0.585	0.607	0.615	0.520
	(38.7)	(32.9)	(37.5)	(46.3)	(0.387)	(0.329)	(0.375)	(0.463)
2	65.50	69.94	66.56	59.05	0.671	0.716	0.681	0.608
	(34.4)	(27.2)	(32.7)	(42.7)	(0.342)	(0.270)	(0.327)	(0.425)
4	67.54	59.78	72.30	70.95	0.703	0.627	0.746	0.740
	(34.4)	(32.7)	(34.4)	(35.7)	(0.344)	(0.323)	(0.343)	(0.361)
6	47.68	48.54	48.24	45.66	0.517	0.528	0.515	0.505
	(34.9)	(30.7)	(33.0)	(43.0)	(0.344)	(0.307)	(0.320)	(0.426)
8	57.27	59.15	65.70	43.14	0.619	0.637	0.695	0.489
	(34.7)	(27.4)	(32.0)	(42.5)	(0.350)	(0.276)	(0.314)	(0.444)
10	58.36	61.38	67.88	44.36	0.636	0.664	0.720	0.510
	(34.9)	(28.9)	(25.0)	(45.1)	(0.348)	(0.295)	(0.243)	(0.455)
12	59.48	58.28	65.94	52.24	0.642	0.625	0.697	0.590
	(35.3)	(28.1)	(36.6)	(40.4)	(0.359)	(0.291)	(0.356)	(0.420)
14	54.93	61.40	53.22	49.57	0.597	0.665	0.561	0.563
	(37.9)	(31.7)	(38.0)	(43.9)	(0.376)	(0.325)	(0.376)	(0.426)
16	56.46	63.20	58.45	44.02	0.620	0.669	0.627	0.542
	(34.6)	(25.7)	(35.0)	(41.9)	(0.342)	(0.249)	(0.347)	(0.432)
18	61.82	63.79	66.14	53.73	0.666	0.675	0.703	0.606
	(35.2)	(29.3)	(34.5)	(42.0)	(0.352)	(0.288)	(0.343)	(0.432)
20	59.25	67.09	61.28	45.43	0.638	0.711	0.654	0.514
	(35.8)	(28.4)	(38.4)	(39.3)	(0.358)	(0.261)	(0.370)	(0.430)

Notes. Standard deviations are in parentheses. For each subject, we average over all those situations in which the subject faced each given link cost. The table then reports the average over all subjects.

Table 10 Random Effects Regressions of Prediction Accuracy and Earnings Pooling Over the Second, Third, and the Fourth Subjects

	By link cost		By link decision	
	Earnings	Prediction accuracy	Earnings	Prediction accuracy
$c = 2$	7.929*	0.0788*		
	(4.049)	(0.040)		
$c = 4$	9.843**	0.109***		
	(3.887)	(0.039)		
$c = 6$	−10.32***	−0.0796**		
	(3.445)	(0.034)		
$c \cdot (c > 6)$	0.043	0.00258		
	(0.181)	(0.002)		
link			−5.150**	0.013
			(2.305)	(0.023)
informed	10.91***	0.0893***	11.09***	0.101***
	(2.192)	(0.022)	(2.853)	(0.029)
DM 3	−3.386	−0.0215	−1.365	−0.005
	(2.149)	(0.020)	(2.887)	(0.029)
DM 4	−2.804	−0.0157	0.041	0.010
	(2.253)	(0.024)	(2.790)	(0.029)
constant	51.95***	0.542***	52.99***	0.531***
	(3.102)	(0.031)	(3.714)	(0.037)
ρ	0.003	0.005	0	0
R^2	0.0381	0.0288	0.0777	0.0519
N	1,566	1,566	308	308
groups	108	108	108	108

Note. Robust standard errors are in parentheses (clustering at the subject level).

*Significant at the 10% level; **significant at the 5% level; ***significant at the 1% level.

forming a link has virtually no effect on the prediction accuracy of subjects, even though it leads to significantly lower earnings. Thus, Table 10 indicates that, except when the cost of link formation is small but positive, subjects tend to pay too much for the information gained from links, and may have been better off not forming any links at all. Finally, observe that subject-specific heterogeneity plays almost no role in the results presented in Table 10 because the fraction of variance due to the random effect, ρ, is never more than 5%.

5.4. Subject-Specific Heterogeneity

We conclude our descriptive analyses with a brief discussion of subject-specific heterogeneity. Obviously, there are numerous ways in which subjects could be heterogeneous. In the interest of space we will consider two:[17] the frequency of optimal link decisions by subjects in the second position and the frequency with which subjects' action choices are consistent with the counting heuristic described above. The results are on display in Figure 7.

First, consider the linking behavior of the second subject. As can be seen in panel (A), although a significant fraction of subjects made the optimal link decision in 85% or more of the rounds, there is also a nonnegligible fraction of subjects who made the optimal link decision in fewer than 50% of the rounds.

[17] We have already shown in Figure 6 significant amounts of heterogeneity with regard to one's tendency to conform.

Figure 7 **A Look at Subject-Specific Heterogeneity**

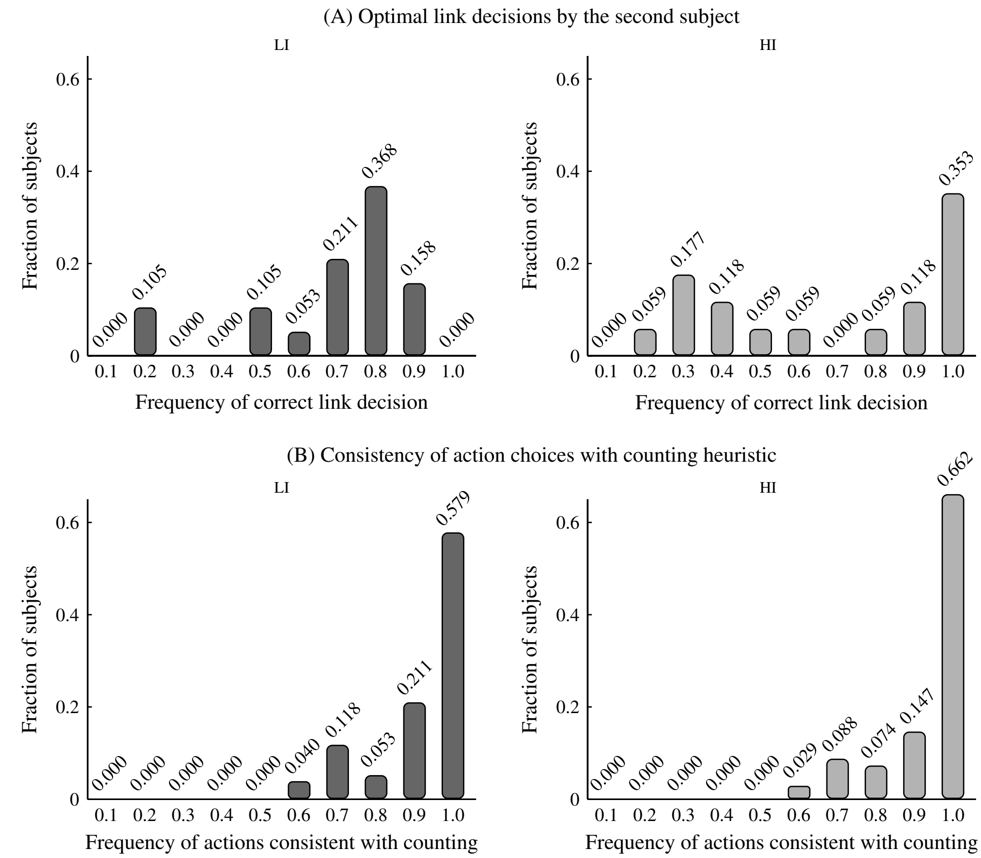

Next, consider the action decisions of players and, in particular, whether or not these decisions are consistent with the counting heuristic. As can be seen in panel (B), there is still heterogeneity in terms of action choices, but it seems to be less with no subjects being consistent with counting in fewer than 60% of the rounds.

5.5. Quantal Response

One can go beyond the descriptive analyses that we have conducted and build a more structural model of behavior. As a first step, we assume that subjects are prone to make mistakes and only stochastically choose the best response. The probability of taking an action increases in the expected payoff from that action. In this decision problem, the key parameters are the conditional beliefs that subjects form about the likelihood of the state following the formation of links. Our theory is based on the assumption that decision makers update using Bayes' rule. However, a structural model can directly estimate these beliefs to gain insights into the decision biases that subjects may have.

In Appendix B, we take a step in this direction by examining the behavior of the second and third subjects. The details are rather technical, but we briefly summarize the important results. First, the results indicate that subjects appear to overweight the information that confirms their prior beliefs. Second, when faced with contradictory information, our results suggest that, *on average*, subjects have approximately Bayesian beliefs about the state. Combined, these two findings go a long way to explain why, for example, subjects in the second position link too frequently—they anticipate that the information they receive is more valuable than it actually is.

Third, when there is a link between the first and the second subjects, and their actions are different, as one might expect, the third subject's signal becomes decisive. Moreover, whereas in the HI treatment the actions of the first and second subjects essentially cancel each other out, in the LI treatment, the third subject appears to put more weight on the first subject's action. Thus, it appears that subjects do realize that a link reveals the second subject to be uninformed.

The punchline of this exercise is as follows: First, subjects overreact to information that confirms their prior. Second, their reactions are (on average) consistent with Bayesian theory when the information conflicts with their own prior. These two together create a strong incentive to form links as we observe in the data. Therefore, we can attribute the systematic deviation from the theory to the bias of observing confirming information. However, when this bias is taken into account, the action decisions are in line with rational behavior.

Remark 1. A word of caution is necessary regarding the second result. Taken at face value, it suggests that any tendency to conform to one's predecessor(s) is not due to any inherent bias, but instead to the fact that subjects only stochastically best respond. However, in the structural framework, because of data constraints, we are unable to obtain parameter estimates based on individual-level data, and instead we pool over all subjects. Because of this, the estimates ignore subject heterogeneity. As we saw in Figure 6, some subjects conform nearly all of the time, and others almost never conform. It is not plausible that such strong tendencies are due to mistakes. Rather, we argue that this is one instance where a more nuanced individual-level study is needed to fully understand the data.

6. Concluding Remarks

In this paper we have presented the results of an experiment on endogenous information acquisition and social learning. One of the key theoretical insights is that, for some parameterizations of the information acquisition problem, social learning is actually enhanced by the presence of a small positive cost of link formation. When links come for free, all possible links will be formed, rendering the observed link structure uninformative. In contrast, with a small positive cost, the observed link structure actually contains a great deal of useful information (in particular, it often indicates which subjects are in fact informed), which leads to increased predictive power and higher expected payoffs. As we reported, this important result is largely borne out in our experiment.

More generally, in terms of behavior, our experiment has identified some deviations from the Bayes-rational benchmark. Specifically, we observe that subjects in the second position tend to form too many links, which then leads to what we called "herding in link formation" by the third and fourth decision makers. In some cases, especially at higher costs of link formation, this actually leads subjects to make worse decisions—so much so, that forming no links at all may have been a better decision.

We also obtained some insights into the motivation for these apparent deviations. In particular, as the results of our structural econometric model highlighted, subjects place a great deal of weight on information that confirms their priors. Because of this, they form more links than what is optimal for them.

With respect to action decisions, our results suggest that subjects' behaviors are highly consistent with a counting heuristic; that is, subjects' action choices are consistent with the majority opinion, which includes their own private information and the actions observed through their linking decisions. Such a rule is quite natural, and in fact, it is optimal assuming that subjects never made mistakes, and is approximately optimal even in the presence of a small number of mistakes. Therefore, it is comforting that most subjects appear to follow this rule.

One other interesting finding is the tendency to conform to one's predecessors when the observed action conflicts with one's own private information. Particularly for the second subject, there is a nonnegligible amount of such behavior (and it is more pronounced in the HI treatment). Interestingly, this does not appear to be a trait uniformly held by all subjects. Instead, there appears to be some "conformists" at one extreme, who almost always conform, as well as some "nonconformists," who almost never conform to their predecessors.

We also note that despite the apparent deviations from Bayes rationality (more so with respect to link formation), a number of comparative statics are consistent with the theory. First, as noted, the tendency of agents in the second position to conform in the HI treatment is less pronounced in the LI treatment. Second, agents are also less likely to join the herd of link formation in the LI treatment than what was observed in the HI treatment. Third, there is also a fairly strong negative relationship between the frequency of link formation and the cost of link formation. Therefore, despite the difficult nature of the task, there is some reason to be optimistic about the subject's behavior. Indeed, the key "obstacle" to overcome would be to teach subjects not to place so much emphasis on information that confirms their prior beliefs.

Acknowledgments

Syngjoo Choi was an early contributor to this project but later withdrew; the authors gratefully acknowledge his contributions. They also acknowledge Guillaume Fréchette, Andrew Schotter, seminar participants at Texas A&M University and the University of Pittsburgh, as well as participants at various conferences for their valuable comments. Finally, the authors thank department editor Teck Ho, an associate editor, and three anonymous referees for their thoughtful comments and suggestions, which greatly improved this paper. This research was supported by the Center for Experimental Social Sciences and the C. V. Starr Center at New York University.

Appendix A. Omitted Proofs

Corollary 1. *The optimal decision rule of the second agent is characterized as follows*:

1. *Let* $\sigma_2 \in \{-1, 1\}$, *and let* $q \in (0, 1]$. *For any* $c > 0$, *the second agent does not form a link and takes action* $a_2 = \sigma_2$.

2. *Let* $\sigma_2 = 0$, *and let* $q \in (0, 1]$. *There exists a threshold* $c^* = (qm)/6$ *such that for any* $c < c^*$, *the second agent links to the first and takes action* $a_2 = a_1$; *if* $c \geq c^*$, *the second agent does not form a link and randomizes between the two actions.*

Proof. 1. Let $q \in (0, 1]$, and without loss of generality, suppose that $\sigma_2 = 1$. Then, the value of information from linking to the first agent when the unit cost of a link is $c > 0$ is

$$\begin{aligned} v(\sigma_2) &= \sum_{a_1} \Pr(a_1 \mid \sigma_2) \max_{a_2} \left\{ \sum_{\theta} \Pr(\theta \mid \sigma_2, a_1) u(a_2, \theta) \right\} \\ &\quad - \max_{a_2} \left\{ \sum_{\theta} \Pr(\theta \mid \sigma_2) u(a_2, \theta) \right\} - c \\ &= \sum_{a_1} \Pr(a_1 \mid \sigma_2) \sum_{\theta} \Pr(\theta \mid \sigma_2, a_1) u(1, \theta) \\ &\quad - \sum_{\theta} \Pr(\theta \mid \sigma_2) u(1, \theta) - c < 0. \end{aligned}$$

The second equality comes from the fact that choosing $a_2 = 1$ is always optimal, regardless of a_1 and hence the first two terms are canceled out. Therefore, it is never optimal to form a link to the first agent for any positive unit cost $c > 0$. The case of $\sigma_2 = -1$ is similar.

2. Let $q \in (0, 1]$, and suppose that $\sigma_2 = 0$. Then a simple computation leads the value of information from linking to the first agent to be

$$v(\sigma_2) = \frac{qm}{6} - c.$$

Therefore, when $c < (qm)/6$, the second agent forms a link to the first. However, when $c \geq (qm)/6$ the second agent does not form a link. □

Corollary 2. *The optimal decision rule of the third agent is characterized as follows*:

1. *Suppose there is a link between the first and the second agents.*
   (a) *Let* $\sigma_3 \in \{-1, 1\}$, *and let* $q \in (0, 1]$. *Then, for any* $c > 0$ *the third agent does not form a link and takes action* $a_3 = \sigma_3$.
   (b) *Let* $\sigma_3 = 0$, *and let* $q \in (0, 1]$. *Then, for any* $c < c^*$ *the third agent links to the second and takes action* $a_3 = a_2$; *if* $c \geq c^*$, *the third agent does not form a link and randomizes between the two actions.*
2. *Suppose there is no link between the first and the second agents.*
   (a) *Let* $\sigma_3 \in \{-1, 1\}$, *and let* $q \in (0, 1]$.
      (i) *There exists a threshold* $c^{**} = (2qm)/39$ *such that for any* $c < c^{**}$, *the third agent links to the second agent. If* $a_2 = \sigma_3$, *then he does not form a link to the first and takes action* $a_3 = a_2$. *If* $a_2 \neq \sigma_3$, *then he links to the first and takes action* $a_3 = a_1$.
      (ii) *For any* $c \geq c^{**}(q, m)$, *the third agent does not form a link and takes action* $a_3 = \sigma_3$.
   (b) *Let* $\sigma_3 = 0$, *and let* $q \in (0, 1]$.
      (i) *For any* $c < c^*$, *the third agent links to the second and takes action* $a_3 = a_2$.
      (ii) *For any* $c \geq c^*$, *the third agent does not form a link and randomizes between the two actions.*

Proof. Suppose there is a link between the first and the second agents. When there is a link between the first and the second agents, the decision problem of the third agent is equivalent to that of the second agent (see the proof of Corollary 1).

1. Suppose that there is no link between the first and the second agents.
   (a) Let $\sigma_3 \in \{-1, 1\}$, and let $q \in (0, 1]$.
      (i) The value of information is

$$\begin{aligned} v(\sigma_3) &= -c + \sum_{a_2} \Pr(a_2 \mid \sigma_3) \max\{v(\sigma_3, a_2), 0\} \\ &= -c + \frac{4}{9}\left(\frac{qm}{6} - c\right) = \frac{2qm}{27} - \frac{13}{9}c. \end{aligned}$$

Therefore, when $c < (2qm)/39$, we have $v(\sigma_3) > 0$; hence, the third agent links to the second agent. If $a_2 = \sigma_3$, the value of information is $v(\sigma_2, \sigma_3) = v(\sigma_2)$; hence, the third does not form a link. If $a_2 \neq \sigma_3$, then the value of information is equivalent to $v(\sigma_3 = 0)$. We already know that in this case the third forms a link.

      (ii) If $c \geq (2qm)/39$, we have $v(\sigma_3) < 0$; hence, the third agent does not form a link to the second agent.
   (b) Let $\sigma_3 = 0$, and let $q \in (0, 1]$.
      (i) The value of information is

$$\begin{aligned} v(\sigma_2) &= \sum_{a_2} \Pr(a_2 \mid \sigma_3) \max_{a_3} \left\{ \sum_{\theta} \Pr(\theta \mid \sigma_3, a_2) u(a_3, \theta) \right\} \\ &\quad - \max_{a_3} \left\{ \sum_{\theta} \Pr(\theta \mid \sigma_3) u(a_3, \theta) \right\} \\ &\quad - c + \sum_{a_2} \Pr(a_2 \mid \sigma_3) \max\{v(\sigma_3, a_2), 0\} = \frac{m}{6} - c, \end{aligned}$$

where $v(\sigma_3, a_2) < 0$ for any (σ_3, a_2). Thus, for $0 < c < (qm)/6$, it is optimal to form a link to the second agent. After forming a link, for any action he observes, the value of information for the third agent is identical to that of an informed second agent. Therefore, the third agent does not form a link to the first and takes action $a_3 = a_2$ (see the proof of Corollary 1) and chooses the same action as the second agent.

      (ii) When $c \geq (qm)/6$, the value is $v(\sigma_3) < 0$. Therefore, the third does not form a link to the second agent. □

Appendix B. Partial Quantal Response Analysis

In our analysis we have described behaviors of the second, third, and fourth subjects in various situations. The analysis showed that subjects often made systematic deviations from the rational theory as espoused by the BSLP. Our goal in this section is to uncover the origins of these deviations. In particular, we will build an econometric model, which allows us to estimate the subjects' beliefs about the state of the world as a function of their information (i.e., signals and link decisions). We show that many of the deviations can be explained by *overoptimism* about the information obtained by forming links.

The Second Agent

In the standard model of stochastic best response, agents experience a random shock to each of their possible decisions. For instance, suppose that the second agent receives signal $\sigma_2 = 1$. Let $\ell_{1,2} = 1\,(0)$ denote the existence (lack) of a link between the first and second agents. The expected utility of forming a link is given by

$$\begin{aligned} \varphi(\sigma_2 = 1, \ell_{1,2} = 1) &:= r \max\{sm, (1-s)m\} \\ &\quad + (1-r)\max\{tm, (1-t)m\} - c + \epsilon_1, \end{aligned} \tag{B1}$$

where $r = \Pr(a_1 = 1 \mid \sigma_2 = 1)$, $s = \Pr(\theta = 1 \mid \sigma_2 = a_1 = 1)$, and $t = \Pr(\theta = 1 \mid \sigma_2 = 1, a_1 = -1)$, whereas the expected utility of not forming a link is given by

$$\varphi(\sigma_2 = 1, \ell_{1,2} = 0) = \max\{pm, (1-p)m\} + \epsilon_0, \quad \text{(B2)}$$

where $p = \Pr(\theta = 1 \mid \sigma_2 = 1)$. Under standard assumptions on ϵ_1 and ϵ_0, the probability that a link is formed is given by

$$\Pr(\ell_{1,2} = 1 \mid \sigma_2, c) = [1 + \exp\{\lambda^l(\varphi(\sigma_2, \ell_{1,2} = 0) - \varphi(\sigma_2, \ell_{1,2} = 1))\}]^{-1}, \quad \text{(B3)}$$

where λ^l is a parameter to estimate. It captures the subject's ability to best respond in his link decision.

Because subjects jointly made both a link and an action decision, we can also consider these two decisions jointly. For example, if the second subject forms a link to the first and $\sigma_2 = a_1$, then

$$\Pr(a_2 = 1 \mid \sigma_2 = 1, a_1 = 1) = [1 + \exp\{\lambda^a(1 - 2s)m\}]^{-1}, \quad \text{(B4)}$$

where λ^a is a parameter to be estimated. It captures the subject's ability to best respond in his action decision. By taking both the link and action decisions into account, we can write the *full* likelihood function for the subject's decision problem and obtain maximum likelihood estimates of both λ^l and λ^a.

Before proceeding to the results, note that in the experiment subjects were told the value of p but not the values of r, s, s^{ui}, or t. One of two approaches can be taken here. First, we could estimate these parameters directly alongside our estimation of the rationality parameters. Second, with enough algebra, it is possible to derive all of these parameters as functions of exogenous variables and the rationality parameter of the first subject. For example, one can show that

$$r = \frac{\sum_{\theta, \sigma_1} \Pr[a_1 = 1 \mid \sigma_1, \theta] \Pr[\sigma_1 \mid \theta] \Pr[\sigma_2 = 1 \mid \theta] \Pr[\theta]}{\sum_\theta \Pr[\sigma_2 = 1 \mid \theta] \Pr[\theta]},$$

where all terms except $\Pr[a_1 = 1 \mid \sigma_1, \theta]$ are exogenously given, and $\Pr[a_1 = 1 \mid \sigma_1, \theta]$ can be estimated from the data. Similar computations can be made for s, t, and s^{ui}.

Although the second approach is desirable in that one need only estimate the rationality parameters, given the results from the descriptive analysis, our belief is that the fit will be significantly worse. The reason is as follows. According to the BSLP, an informed second agent should *never* form a link to the first at a positive cost.[18] However, as reported above, we observed a significant number of links being formed as well as agents conforming in the face of contradictory information. Now, if we assume that subjects have rational expectations of their predecessors, then, because subjects in the first position made mistakes more than 10% of the time, this will make forming a link *less valuable* and impose a real penalty to conforming to one's predecessor when $\sigma_2 \neq a_1$.

[18] Recall that if he forms a link he will either observe $a_1 = \sigma_2$ (in which case his belief increases to $s > 2/3$), or he will observe $a_1 \neq \sigma_2$ (in which case his belief decreases to $1 - t \leq 1/2$). In both cases, the subject would have done just as well to simply not form a link and choose $a_2 = \sigma_2$.

Table B.1 Second Subject: Uncovering Beliefs

	Link only		Link and actions	
	HI	LI	HI	LI
(A) Beliefs and rationality parameters				
λ^l	3.174***	3.511***	3.087***	3.480***
λ^a	NA	NA	1.143***	1.717***
s	0.898***	0.868***	0.880***	0.844***
t	0.50	0.56	0.5165	0.5875*
s^{ui}	NA	0.641***	NA	0.645***
r	0.556	0.537	0.556	0.537
LL	−380.547	−417.745	−636.830	−638.022
N	680	760	680	760
(B) Rationality parameters under rational expectations				
λ^l	1.775***	1.901***	1.775***	1.901***
λ^a	NA	NA	1.195***	1.930***
s	0.764	0.738	0.764	0.738
t	0.553	0.587	0.553	0.587
s^{ui}	NA	0.585	NA	0.585
r	0.539	0.528	0.539	0.528
LL	−397.197	−453.824	−653.587	−684.093
N	680	760	680	760

Notes. In panel (A), shaded cells indicate that the variable was fixed at the theoretical value, whereas in panel (B), shaded cells indicate that the variable was fixed at its "rational expectation" value (i.e., given the actual behavior of the first subject). For λ we computed likelihood ratio (LR) tests that the coefficient is equal to zero, whereas for the other variables, we conducted LR tests that the coefficient is equal to the theoretical value.

*Significant at the 10% level; **significant at the 5% level; ***significant at the 1% level.

Table B.1 reports both sets of results, with panel (A) reporting estimates of the underlying belief parameters and panel (B) deriving the belief parameters according to rational expectations on the first subject's behavior. In both cases, we report results for the link decision separately and a joint estimation of the link and action decision. In panel (A), for the link-only estimation, we could not separately identify r, s, and t. Therefore, we chose to estimate s and fix the other two parameters at their theoretical values (shaded in the table). For the full decision problem, all parameters are, in principle, identified, though practical identification of r appeared to be rather poor; therefore, we omitted it from our estimates and fixed it at the theoretical value. In panel (B) the rational expectations values of r, s, and t are reported (shaded in the table). Asterisks denote the level of significance of the parameter according to a likelihood ratio test. For the two λ terms, the null hypothesis is that $\lambda = 0$, whereas for the other terms, the null hypothesis is that the parameter equals the theoretical value.

Consider first panel (A), and notice that both λ^l and λ^a are highly significantly different from zero. Regarding the other parameters, in all cases, we see that $s > s^{\text{BSLP}}$ ($s^{\text{BSLP}}_{\text{HI}} = 0.8$; $s^{\text{BSLP}}_{\text{LI}} = 0.7586$); that is, subjects overweight the informational gain from forming links, which causes them to form more links than it is optimal. Looking at t, we see that subjects appear to (almost) optimally weight contradictory information. Accordingly, it appears that any tendency to conform to the actions of one's predecessor is due to the fact that $\lambda^a < \infty$. However, as we saw in Figure 6, the frequency of conformist behavior is very bimodal, with some subjects conforming nearly all the time and others almost never conforming. Therefore, although on average it appears that there is no bias, our descriptive analyses caution us from drawing overly strong conclusions from the pooled estimates of our structural model; that is, a nonnegligible portion of our subject pool displays a strong tendency to conform.

Briefly, consider panel (B). As expected, when we fix r, s, and t at their rational expectations values, the overall fit goes down a lot. Furthermore, our estimate of λ^l is much lower. Again, this is to be expected because linking is now a bigger mistake. Thus, it appears as though subjects are less rational. The fit on λ^a actually appears to increase slightly. This makes sense given that s is lower and the payoff difference from following one's signal when $\sigma_2 = a_1$ is $m(2s-1)$. Therefore, to match the action data, λ^a must increase.

The Third Agent

The empirical analyses of the third agent is in the same spirit as for the second agent. However, matters are substantially more complicated. In particular, the third agent may face one of two different networks and has the option of forming zero, one, or two links to his predecessors. Instead of giving a full derivation of the empirical choice model, we simply note the key steps required. Given the above discussion, we also focus on the case in which underlying beliefs are estimated directly, rather than obtained via rational expectations.

In addition to the required calculations from the previous section, we need more tedious calculations to be able to solve for the value of information. The task is even more complicated because we assume that the third agent anticipates the errors made by the first and the second agents. For example, consider the case in which the third agent observes that the second linked to the first. For ease of exposition, let us introduce the notation in Table B.2.

Then, the ex ante expected utility of the third who observes a link between the first and the second is

$$\varphi_j = a_{1j}\max\{c_{1j}m, (1-c_{1j})m\} + a_{2j}\max\{c_{2j}m, (1-c_{2j})m\} \\ + a_{3j}\max\{c_{3j}m, (1-c_{3j})m\} + a_{4j}\max\{c_{4j}m, (1-c_{4j})m\},$$

Table B.2 Notation for Parameters Needed for Bayesian Updating

	σ_3		
(a_1, a_2)	-1	0	1
	(A) $\Pr(a_1, a_2 \mid \sigma_3, \ell_{1,2}=1)$		
(1, 1)	a_{11}	a_{12}	a_{13}
(1, −1)	a_{21}	a_{22}	a_{23}
(−1, 1)	a_{31}	a_{32}	a_{33}
(−1, −1)	a_{41}	a_{42}	a_{43}
	(B) $\Pr(\theta=1 \mid a_1, a_2, \sigma_3, \ell_{1,2}=1)$		
(1, 1)	c_{11}	c_{12}	c_{13}
(1, −1)	c_{21}	c_{22}	c_{23}
(−1, 1)	c_{31}	c_{32}	c_{33}
(−1, −1)	c_{41}	c_{42}	c_{43}

where $j=1$ for $\sigma_3=-1$, $j=2$ for $\sigma_3=0$, and $j=3$ for $\sigma_3=1$, and an informed third agent will form a link if and only if

$$\max\{pm, (1-p)m\} + \epsilon_1 < \varphi_j - c + \epsilon_0.$$

If no link is observed, the decision is even more complicated because the agent must decide whether to link to the first agent, link to the second agent, or not link at all. Then, if he decides to link, he must decide whether or not to form another link, and only then to make his action decision.

First, consider the aforementioned case in which the third agent observes a link between the second and the first agents. To reduce the number of parameters to estimate, we make the following restrictions:

$$\begin{aligned} c_{32} &= 1-c_{22}, & c_{42} &= 1-c_{12}, \\ c_{13} &= 1-c_{41}, & c_{23} &= 1-c_{31}, \\ c_{33} &= 1-c_{21}, & c_{43} &= 1-c_{11}. \end{aligned} \tag{B5}$$

In the HI treatment, when we consider the full decision problem; this means that there are four c_{ij} parameters to estimate as well as λ^l and λ^a, a total of six parameters. In the LI treatment, there are two more parameters to estimate (c_{12} and c_{22}), for a total of eight parameters. So, if we consider only the link decisions, the identification problem becomes even more severe. Therefore, we immediately consider the full decision problem.

For the HI treatment, the results are presented in panel (A) in Table B.3. We see that, when the actions of the first and second agents agree, the third agent becomes *overconfident* that the state is consistent with the observed actions and that this is true whether or not the observed actions are consistent with his signal. When the observed actions are different from each other, as we saw in the descriptive analyses, the third agent's signal becomes decisive, and the estimated beliefs match very closely to their theoretical values. With regard to the estimated λ values, we simply note that they are fairly similar to those obtained in our estimation of the second agent's problem. This suggests that herding in link formation is rather due to biases in underlying beliefs, and not due to irrationality (as captured by λ^l).

We repeat the analysis in panel (B) in Table B.3 for the LI treatment. The results are not too different. As was the case with the HI treatment, the third agent is generally overconfident (relative to the theoretical prediction) whenever $a_1 = a_2$. This overconfidence is even greater when also $\sigma_3 = a_1 = a_2$. Looking at the cases in which $a_1 \neq a_2$, we see that in all cases the subject's own signal is decisive; however, it appears that the subject places more weight on the action decision of the first subject. In particular, when $\sigma_3 = a_1 \neq a_2$, the third subject has very strong (overconfident) beliefs about the state, whereas when $\sigma_3 = a_2 \neq a_1$, the third subject is actually slightly underconfident about the true state (relative to the theoretical prediction).

The results above provide insight only for the case in which the third agent observes a link between the first and the second agents. Ideally, we would like to repeat the analysis for the more general setting in which the third agent can observe either of the two possible networks and make his decisions accordingly. However, as is often the case in

Table B.3 Third Subject: Uncovering Beliefs

	σ_3		
(a_1, a_2)	-1	0	1
(A) HI treatment: $\Pr(\theta = 1 \mid a_1, a_2, \sigma_3, \ell_{1,2} = 1)$			
(1, 1)	0.661 [0.523]	NA	1.000 [0.814]
(1, −1)	0.332 [0.345]	NA	0.679 [0.678]
(−1, 1)	0.321 [0.322]	NA	0.668 [0.655]
(−1, −1)	0.000 [0.186]	NA	0.339 [0.477]
λ^I	2.804	LL	−194.59
λ^a	1.202	N	245
(B) LI treatment: $\Pr(\theta = 1 \mid a_1, a_2, \sigma_3, \ell_{1,2} = 1)$			
(1, 1)	0.587 [0.446]	0.717 [0.617]	0.917 [0.763]
(1, −1)	0.285 [0.257]	0.469 [0.409]	1.000 [0.580]
(−1, 1)	0.000 [0.420]	0.531 [0.592]	0.715 [0.743]
(−1, −1)	0.083 [0.237]	0.283 [0.383]	0.413 [0.554]
λ^I	2.708	LL	−288.750
λ^a	1.279	N	337

Note. For each c_{ij}, the theoretical value is in brackets.

structural models of this sort, doing so would introduce many more parameters, making it difficult to obtain precise estimates. As an alternative, we now discuss a more reduced form approach, which yields the same conclusions in a less computationally demanding way.

It is fairly obvious that the more times the third agent observes $a_i = 1$ (or $\sigma_3 = 1$), the more confident he should be that the true state of the world is actually $\theta = 1$.[19] Therefore, consider the following latent variable model:

$$y_i = \beta_0 + \beta_1\sigma_3 + \beta_2 a_1 + \beta_3 a_2 + \beta_4 a_2 \cdot (\ell_{1,2} = 1) + u_i,$$

where $y_i^* = \mathbb{I}[y_i \geq 0]$ are what we observe, and $(\ell_{1,2} = 1)$ is a dummy variable that takes a value of 1 if there is a link between the first and second subjects. Next, $a_i = 0$ if the third subject does not observe the action choice of subject i, and $a_i \in \{-1, 1\}$ is the action choice of subject i if the third subject does observe it.

The results of this exercise are reported in Table B.4. We find that the marginal effects on one's signal and the observed actions are all positive and significant at the 1% level. Thus, as our descriptive analyses showed, subjects respond quite strongly to the information that they observe. In the HI treatment, we cannot reject that $\sigma_3 = a_i$ for $i = 1, 2$ or $a_1 = a_2$ (in all cases, $p > 0.1$). However, in the LI treatment, consistent with the fact that subjects may be uninformed, we can reject that $\sigma_3 = a_1$ at the 1% level, and we can reject $\sigma_3 = a_2$, though only at the 6.6% level. However, when we condition on a link being present between the first and second subjects (which should signal that the second subject is uninformed), we are able to reject the null hypothesis at the 1% level. Thus, consistent with the results of our structural estimation, there is a sense in which subjects discount the information learned from the second subject when the third observes that he formed a link to the first. Next, observe that in neither treatment are we able to reject that $a_1 = a_2$. Finally, we see that the estimate of ρ, which captures the importance of subject specific heterogeneity is, in both treatments, small and not significant. Therefore, it appears that there is little heterogeneity in terms of action decisions.

[19] Of course, by how much each observation influences the third agent's beliefs depends on the parameters of the underlying structural model, but the direction of influence is clear in the reduced form framework on which we focus.

Table B.4 Third Subject: Decomposing Action Decisions

Variable	HI	LI
σ_3	0.528*** [0.037]	0.623*** [0.047]
a_1	0.573*** [0.064]	0.430*** [0.052]
a_2	0.429*** [0.084]	0.485*** [0.072]
$a_2 \cdot (\ell_{1,2} = 1)$	0.057 [0.113]	−0.089 [0.091]
ρ	2.1e−7 [9.5e−6]	0.034 [0.034]
LL	−191.58	−259.98
N	680	760
subjects	17	19

Notes. Standard errors are in brackets. For coefficients, we report the marginal effects at mean of independent variables and assuming the random effect is zero.
***Significant at the 1% level.

Appendix C. Instructions

General Instructions

This is an experiment in the economics of decision making. Your earnings will depend partly on your decisions and partly on chance. By following the instructions and making careful decisions you will earn varying amounts of money, which will be paid at the end of the experiment. Details of how you will make decisions and earn money will be provided below.

In this experiment, you will participate in 40 independent rounds, each of which contains four decision positions in a decision queue. In each round you will be asked which of two urns has been randomly chosen (called *action decision*); however, before making your action decision, some subjects will be able to observe the actions of those who have gone before them (called *link decision*(*s*)) by paying a cost that will be determined by the computer at the beginning of each round.

Before the first round, you will be randomly assigned to a position in the decision queue labeled **1**, **2**, **3**, or **4**. One-fourth of the participants will be randomly assigned to each of the four positions. Your position depends solely on chance and will remain constant in all rounds throughout the experiment. When you are called to make decisions, in the center of the computer screen you will be informed of your position and any link decisions made by those in preceding decision positions; however, you will not observe their action decisions.

A Decision Round

Each round starts by having the computer randomly form groups of four participants by selecting one participant from each of the four positions. The groups formed in each round depend only on chance and are independent of the groups formed in any of the other rounds.

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

To help you determine which urn has been selected, you will be allowed to observe one ball, drawn at random, from the urn *at no cost*. In addition, if you are in position **2**, **3**, or **4**, you will be given a chance to see action decisions in preceding positions at a cost determined by the experimental software.

Your private draw in each round is independent of the draw received by any other participant. The result of your draw will be your private information and should not be shared with any of the other participants. You will see your private draw in the middle portion of the computer screen.

After each draw, the ball will be returned to the urn before making a private draw for the next participant. This is done by the experimental software.

Participants assigned to position **1** may see the following screen on your computer screen:

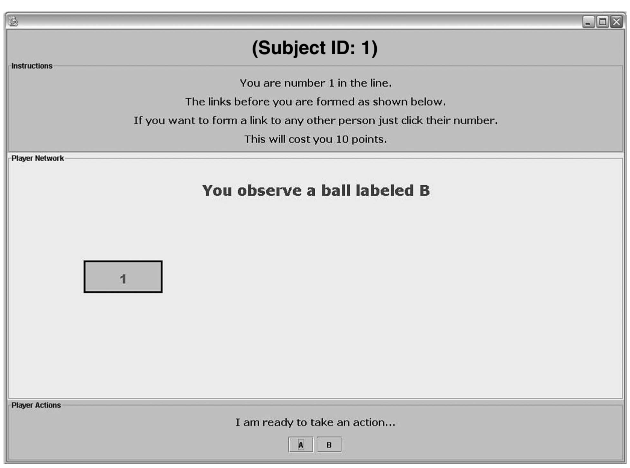

In this case, since you are in the first position in a decision queue, all you need to do is make your action decision based on your private information. This is done at the bottom of the screen by simply clicking on either **A** or **B**.

For participants assigned to positions **2**, **3**, and **4**, there will be other participants in the same group who have already made their action and link decisions. In addition to your private draw, you will have the opportunity to observe the action decisions of those in preceding positions at a cost determined by the experimental software at the beginning of each round. When it is your turn to move, you will see a graphical representation of all the link decisions made by those who precede you in the decision queue.

For example, suppose that you are assigned to position **3** in the queue. You may see the following screen:

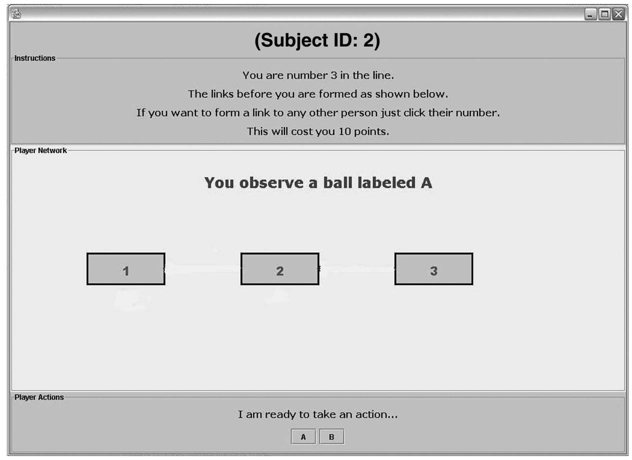

In this example, your private draw was a ball labeled **A**. In addition, you observe that the participant in the second position chose *not* to form a link to the first position; that is, the participant in the second position predicted which urn was more likely to be chosen, while having chosen not to observe the action decision by the participant in the first position.

Continue with the example above and suppose that you wish to form a link to the second position. To do this, simply click on the box labeled **2**. Then you will observe the action decision made by the participant in position **2** while incurring the cost of forming a link. This is depicted below:

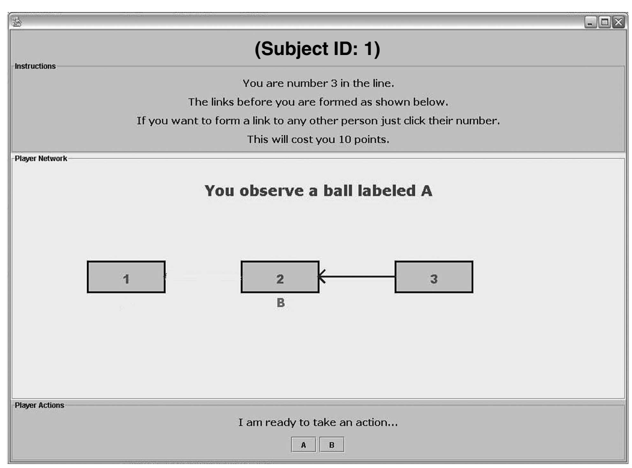

Note that once you form a link to one preceding participant, you see not only his/her action choice but also *the action choices of all those with whom that person linked.*

For example, suppose that the second person *had* actually formed a link to the first person in the queue. In this situation, by forming a link to the second position in the queue you would see the action decisions of *both* participants in the first and the second positions in the queue.

In this example, if you wish to observe more information, you may form a link to the first position, *at an additional cost*, and observe his or her action decision. If not, and you are ready to make your decision, simply click on the box labeled **A** or **B** at the bottom of your screen, corresponding to which urn you think was more likely to have been chosen.

Once you have made your decision for that round, you will be informed which urn was actually used and what your potential payoffs are for that round. By clicking on the *OK* button you will be taken to a waiting screen and then the next person in the line will be able to make his or her decisions.

This concludes one decision round. All of the participants will then be randomly placed into a new group of four people. In total, you will repeat 40 independent rounds with various levels of costs.

Remember: In each round, the same urn applies to all members of a group. That is, the experimental software picks *one* urn for each group in each round.

Cost of Forming Links

Now, we will describe in detail how the cost of forming a link will be determined in each of the 40, independent, rounds. In all rounds throughout the experiment, the cost of forming a link can be any *even* number between 0 and 20, inclusive; that is, the cost will be one of the following numbers, $0, 2, 4, \ldots, 16, 18, 20$.

In each round, the computer will randomly assign a cost to each group of four. The chance that the computer selects any even number between 0 and 20 points is exactly the same. That is, the chance that a cost of 2 is selected is the same as the chance that a cost of 14 is selected and so on. Moreover, the cost assigned in one decision round is independent of the cost in any other decision round.

Remember: The cost of forming a link in each round is the same for all members of a group. Moreover, the cost for each link is the same (e.g., if you form one link at a cost of 10 points, you are free to form another link by paying an additional cost of 10 points).

Payoffs

Your potential earnings for each round are determined as follows. If you made the correct action decision regarding which urn was used, you will be awarded 100 points for that round; otherwise, you will be awarded nothing. From this amount, either 100 or 0, we will subtract the appropriate cost for *each* link decision that you made. For example, if, in round 10, the cost of link formation was 18 points, then in determining your potential earnings for round 10, 18 points will be subtracted, from either 100 or 0, for every link decision that was made. For example, if, after having made one link decision, you correctly guessed which urn was chosen, your potential earnings would be $100-18=82$ points.

At the end of the 40 rounds, the experimental software will randomly select three rounds from which you will be paid. The total number of points earned will be summed up for each of these three rounds—100 points for each correct decision, from which we will subtract the appropriate number of points for each link decision. This will be converted to a dollar amount according to the rule:

$$\$1 = 15 \text{ points.}$$

This amount will then be added to the $8.00 participation fee to give your payment for this experiment. Payments will be made in private via petty cash vouchers at the conclusion of the session.

Rules

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

Your participation in the experiment and any information about your earnings will be kept strictly confidential. Your receipt of payment and consent form are the only places on which your name will appear. This information will be kept confidential in the manner described in the consent form.

If you have any questions please ask them now. If not, we will proceed to the experiment.

References

Acemoglu, D., M. A. Dahleh, I. Lobel, A. Ozdaglar. 2011. Bayesian learning in social networks. *Rev. Econom. Stud.* **78**(4) 1201–1236.

Bala, V., S. Goyal. 1998. Learning from neighbours. *Rev. Econom. Stud.* **65**(3) 595–621.

Banerjee, A. 1992. A simple model of herd behavior. *Quart. J. Econom.* **107**(3) 797–817.

Banerjee, A., D. Fudenberg. 2004. Word-of-mouth learning. *Games Econom. Behav.* **46**(1) 1–22.

Beck, P. A., R. J. Dalton, S. Greene, R. Huckfeldt. 2002. The social calculus of voting: Interpersonal, media, and organizational influences on presidential choices. *Amer. Political Sci. Rev.* **96**(01) 57–73.

Bikhchandani, S., D. Hirshleifer, I. Welch. 1992. A theory of fads, fashion, custom and cultural change as informational cascade. *J. Political Econom.* **100**(5) 992–1026.

Bikhchandani, S., D. Hirshleifer, I. Welch. 1998. Learning from the behavior of others: Conformity, fads and informational cascades. *J. Econom. Perspect.* **12**(3) 151–170.

Blackwell, D. 1951. Comparison of experiments. J. Neyman, ed. *Proc. Second Berkeley Sympos. Math. Statist. Probab.*, University of California Press, Berkeley, 93–102.

Çelen, B., S. Kariv. 2004. Observational learning under imperfect information. *Games Econom. Behav.* **47**(1) 72–86.

Conley, T., C. Udry. 2001. Social learning through networks: The adoption of new agricultural technologies in ghana. *Amer. J. Agricultural Econom.* **83**(3) 668–673.

Cross, R. L., A. Parker. 2004. *The Hidden Power of Social Networks: Understanding How Work Really Gets Done in Organizations*. Harvard Business Press, Cambridge, MA.

Cubitt, R., C. Starmer, R. Sugden. 1998. On the validity of the random lottery incentive scheme. *Experiment. Econom.* **1**(2) 115–131.

Eliaz, K., A. Schotter. 2007. Experimental testing of intrinsic preferences for noninstrumental information. *Amer. Econom. Rev.* **97**(2) 166–169.

Eliaz, K., A. Schotter. 2010. Paying for confidence: An experimental study of the demand for non-instrumental information. *Games Econom. Behav.* **70**(2) 304–324.

Foster, A. D., M. R. Rosenzweig. 1995. Learning by doing and learning from others: Human capital and technical change in agriculture. *J. Political Econom.* **103**(6) 1176–1209.

Gale, D., S. Kariv. 2003. Bayesian learning in social networks. *Games Econom. Behav.* **45**(2) 329–346.

Galeotti, A., A. Mattozzi. 2011. Personal influence: Social context and political competition. *Amer. Econom. J.: Microeconomics* **3**(1) 307–327.

Goeree, J. K., L. Yariv. 2007. Conformity in the lab. Working paper, California Institute of Technology, Pasadena.

Huckfeldt, R. R., J. Sprague. 1995. *Citizens, Politics, and Social Communication: Information and Influence in an Election Campaign.* Cambridge University Press, Cambridge, UK.

Huckfeldt, R. R., P. E. Johnson, J. D. Sprague. 2004. *Political Disagreement: The Survival of Diverse Opinions Within Communication Networks*. Cambridge University Press, Cambridge, UK.

Ioannides, Y., L. D. Loury. 2004. Job information networks, neighborhood effects, and inequality. *J. Econom. Literature* **42**(4) 1056–1093.

Jackson, M. O. 2003. A survey of models of network formation: Stability and efficiency. G. Demange, M. Wooders, eds. *Group Formation in Economics: Networks, Clubs and Coalitions*. Cambridge University Press, Cambridge, UK.

Jackson, M. O. 2006. The economics of social networks. R. Blundell, W. Newey, T. Persson, eds. *Advances in Economics and Econometrics (Proceedings of the Ninth World Congress of the Econometric Society)*. Cambridge University Press, Cambridge, UK.

Katz, E., P. F. Lazarsfeld. 1955. *Personal Influence: The Part Played by People in the Flow of Mass Communication*. Free Press, New York. [Reprinted with a new introduction by Elihu Katz and a forward by Elmo Roper. Transaction Publishers, New Brunswick, NJ, 2006.]

Kelly, M., C. Ó Gráda. 2000. Market contagion: Evidence from the panics of 1854 and 1857. *Amer. Econom. Rev.* **90**(5) 1110–1124.

Kübler, D., G. Weizsäcker. 2004. Limited depth of reasoning and failure of cascade formation in the laboratory. *Rev. Econom. Stud.* **71**(2) 425–441.

Lazarsfeld, P. F., B. Berelson, H. Gaudet. 1944. *The People's Choice. How the Voter Makes Up His Mind in a Presidential Campaign*. Duell, Sloan & Pearce, New York.

Munshi, K. 2004. Social learning in a heterogeneous population: Technology diffusion in the Indian Green Revolution. *J. Development Econom.* **73**(1) 185–213.

Smith, L., P. Sørensen. 1998. Rational social learning with random sampling. Working paper, University of Michigan, Ann Arbor.

Smith, L., P. Sørensen. 2000. Pathological outcomes of observational learning. *Econometrica* **68**(2) 371–398.