Treasure Hunt*

Markus M. Möbius
Harvard University and NBER

Tuan Phan
Harvard Business School

Adam Szeidl
UC Berkeley and NBER

June 23, 2010

Abstract

We seed a large real-world social network with binary information and analyze subsequent social learning. A unique feature of our field experiment is that we measure both the pre-existing social networks *and* the actual conversation network. Our experiment allows us to test how rational agents behave when processing information that originates within their social network. We find that information decays quickly with social distance and that agents mainly incorporate information within social distance 2. Conversations trough common friends do not increase the weight that a subject places on signals from direct friends but linearly increases the weight on signals from indirect friends. This suggests that agents are able to avoid double-counting information from indirect friends. We propose a simple "streams model" of social learning that is consistent with the evidence from our experiment.

JEL Classification: C91, C93, D83
Keywords: information processing, de Groot model, streams model

*Preliminary and incomplete. Please do not cite. We are grateful to Abhijit Banerjee, Ben Golub and Matthew Jackson for helpful discussions. We would also like to thank many seminar participants for their feedback. We thank the National Science Foundation for financial support.

1 Introduction

Social networks are an important source of information. Lazarsfeld, Berelson and Gaudet (1944) found that US voters in the 1940 Presidential election were more influenced by friends than mass media. Katz and Lazarsfeld (1955) identified opinion leaders along various dimensions such as local politics, fashion and movies. Development economists have documented the importance of networks for the spread of new farming technologies (Conley and Udry 2010, Duflo, Kremer and Robinson 2006), deworming drugs (Kremer and Miguel 2007) and labor markets (Munshi 2003). Online social networks such as Facebook and MySpace have experimented with viral marketing techniques to allow advertisers to exploit word-of-mouth when introducing new products such as movie releases and electronic gadgets.

In this paper, we use a field experiment to analyze in detail how agents aggregate information in a real-world social network. We are particularly interested in how information *decays* as it travels through a social network and whether agents correctly *tag* the source of information when aggregating information originating within their social network. For example, consider the following communication network:

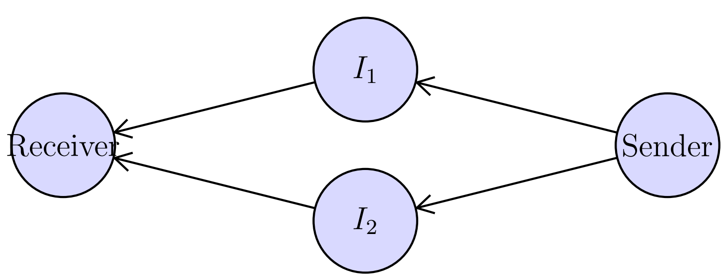

In this example, the sender has a signal which she communicates to two intermediaries who subsequently both talk to the receiver. If information is "untagged" during this transmission process, the receiver might "double-count" the sender's signal as two signals from independent senders. On the other hand, if information is "tagged" the receiver will count the sender's signal at most once. A theory of social learning has to fall somewhere within this spectrum: the most popular social learning model, the de Groot model, implicitly assumes that information is untagged (DeMarzo, Vayanos and Zwiebel 2003, Golub and Jackson 2010b). We propose an alternative "streams" model at the other extreme end of spectrum where instead all signals are perfectly tagged.[1]

The extent to which agents in a network can tag information has important empirical implications: (1) without tagging, agents tend to pay too much attention to information that originates within their social island and (2) agents tend to overweigh information signals from their social network relative to their prior. The intuition for the latter effect is that an untagged signal from a single source can reach the receiver through multiple paths and is erroneously interpreted as coming from several senders. This effect can provide a rational for gossip and rumors in tightly knit social networks.

Our experiment was conducted with 843 students at a large private university. We had already measured the social network of students in the previous academic year for other network experiments (Leider, Mobius, Rosenblat and Do 2009, Leider, Mobius, Rosenblat and Do 2010). We *seeded* the network with information on three distinct question, each of which had two possible answers. We suggested to every student one answer for each of the three questions and we told them that the majority of students would receive the correct suggestion. We encouraged them to talk to each other and they could update their guess online as often as they liked. The game ended randomly

[1] Acemoglu, Bimpikis and Ozdaglar (2010) have independently proposed this specification.

after 4 days and participants who got the answers to all three questions right were eligible to receive two movie tickets.

Our first main finding is that information decays quickly with social distance: we find that agents incorporate information up to social distance 2. Second, we find tentative evidence that agents in our experiment tag the signals that they received from direct and indirect friends. This provides tentative evidence for the streams model.

Our experimental design has several features which mitigate alternative explanations for our experimental results. Most importantly, we measure *both* social networks *and* actual conversations conducted. This allows us to distinguish actual indirect conversations from direct conversations that we mistakenly classify as indirect because of a mis-measured social network. Moreover, we precisely control the strength of each individual's information and we make participants in our study aware of this fact.[2] Hence, there are no a priori reasons to expect that agents should attach a lower weight to the information of socially distant neighbors.

Our experiment builds on a rapidly growing theoretical literature on social networks and a smaller experimental literature. The focus of the theory literature has been to (a) identify conditions on the social network that guarantee convergence to correct beliefs after many rounds of communication (DeMarzo et al. 2003, Golub and Jackson 2010b, Acemoglu et al. 2010); (b) characterize the speed of convergence (Golub and Jackson 2010a) and (c) identify the influencers in a social network (Bala

[2]This distinguishes us from more naturalistic field experiment such as Kremer and Miguel (2007). These authors also measure social learning in a field experiment with deworming medication. However, they do not seed the network directly with information but instead give some random subset of the population early access to medication.

and Goyal 1998). Many of these models rely on random graphs without information decay. Experimentalists have usually conducted lab experiments with artificially induced networks of small size: such a setting is less suitable for testing tagged information transmission as real-world social networks where information can arrive from a very large number of sources.

The balance of the paper is organized as follows. Section 2 discussed our theoretical framework. In section 3 we explain the design of our experiment and in section 4 we summarize the data. Section 5 describes the main results on information decay and signal tagging which we discuss in section 6. Section 7 concludes.

2 Theoretical Framework

In this section, we flesh out two application which illustrate the importance of tagged vs. untagged information processing.

2.1 Social Learning and Gossip

Consider the social network from the introduction and assume that the sender observes a binary symmetric signal $s \in \{0, 1\}$ whether some agent, which we call "Bob", is a criminal. This signal is correct with probability $q > \frac{1}{2}$. The receiver has a prior belief μ that Bob is a criminal. We denote the receiver's posterior belief with $\hat{\mu}$.

In the de Groot model the receiver will listen to the sender's signal and her posterior belief will satisfy:

$$\log(\frac{\hat{\mu}}{1-\hat{\mu}}) = \log(\frac{\mu}{1-\mu}) + 2 \times (2s-1) \times \log(\frac{q}{1-q}) \tag{1}$$

In the streams model the receiver will realize that she double-counted the receiver's signal and her posterior belief will satisfy:

$$\log(\frac{\hat{\mu}}{1-\hat{\mu}}) = \log(\frac{\mu}{1-\mu}) + (2s-1) \times \log(\frac{q}{1-q}) \tag{2}$$

Note, that the de Groot updater overweights the information that she receives from the social network relative to her prior. We call this the "gossip effect".[3] This model of gossip might be particularly appropriate to explain phenomena such as "urban legends" which are transmitted infrequently and where the original source is therefore hard to ascertain.

2.2 Echo-chamber Effect

To illustrate the echo-chamber effect consider the following example:

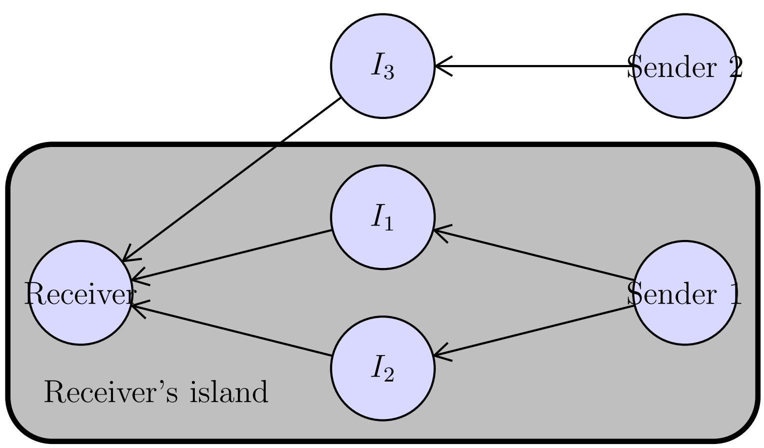

Receivers are more likely to have *strong links* with more than one common friend to senders in their own social island, such as their school, university or workplace,

[3]Interestingly, an informal rule of journalism specifically requires journalists to verify a story with three *independent* sources in order to prevent such double-counting.

compared to senders outside their social island (Granovetter 1973). In the de Groot model this type of social structure implies over-weighting of information that originates within their social island which we refer to as the echo-chamber effect.[4]

3 Experimental Design

Social Network Elicitation. For this experiment, we used social network data that was originally collected for two unrelated experiments on directed altruism (Leider et al. 2009, Leider et al. 2010). To focus on relationships with a high frequency of interaction, we developed the trivia task technique to elicit the social network for the second wave. Each subject listed 10 friends about whom they would answer questions. Over several weeks, several of the listed friends were randomly selected and each were sent an e-mail asking him to answer a multiple choice question about himself (e.g. What time do you get up in the morning?). The subject then received an e-mail directing her to a web page where she had 15 seconds to answer the same question about her friend. If the subject and her friend submitted identical answers, they both won a prize. The trivia task provides subjects with incentives to list friends that the subject spends time with frequently (and thus is more likely to know the friends habits).

Seeding Information. Subjects who had participated in the social network elicitation were invited to a new online experiment by email. The email contained a link to a "treasure hunt" website. The goal of the treasure hunt was to find to correct an-

[4]This effect is different from Cass Sunstein's notion of the echo-chamber effect that is based on homophily. Sunstein postulates that we like to listen to people who share our opinion which leads us to listen too little to outside opinions.

swer to three questions (see figure 1). Subjects were told that each questions had two possible answers (which we told to subjects). Hence, subjects had to find the correct set of answers out of eight possible combination. Every participant who correctly guessed the answer to all three questions received two movie ticket vouchers.

Subjects were given a suggested answer for each of the three questions (see figure 2). They were told that the *majority* of participants would see the correct suggestion (the exact share was 57.5 percent).

Submitting a guess. Subjects could submit a guess for each of the three questions (see figure 3). After a submission they received an email with a link that allowed them to login and update their submission. We encouraged subjects to log into the experiment as often as they liked. Only the last submission would count for determining winners. We left the end date of the experiment vague since we wanted to encourage subjects to update their guess frequently rather than only once shortly before the end of the experiment.

Eliciting conversations. After submitting a guess, we asked subjects about the conversations they had with other participants since their last submission (or since the start of the experiment if submitted for the first time). We programmed this page in such a way that we first asked subjects to *how many* people they talked to in the meantime. After selecting a number, an input form appeared where we asked subjects to list conversations partners by name. We took great care to make this part of the webpage as user-friendly as possible: whenever a participant started to type a name, the site would suggest a drop-down menu of student names for auto-completion. The universe of names consisted of all undergraduates at the university rather than just participants in our treasure hunt – this was done on purpose in order to distinguish truthful submissions from "random" submissions. As an added incentive we paid subjects 25 cents for each submitted conversation partner who would reciprocate the submission when updating her own action.

4 Data Description

Information on social networks was collected in December 2004 through an online trivia task advertised on the popular student social website `facebook.com`.[5] 2,360 students signed up, generating were 12,782 links between participants out of 23,600 total links[6] with 6,880 symmetric links. In total, 5,576 out of the 6,389 undergraduates either participated in the trivia task or were named by a participant. Upperclassmen had higher participation rates, with only 34 percent of freshman responding, but with 45, 52, and 53 percent of sophomores, juniors, and seniors participating, respectively. The social OR-network of 5,576 individuals contains a single component (meaning all individuals are connected) with a mean path length of 4.2 between participants.

We conducted the "treasure hunt" experiment in May 2006. Since the network data was by now more than one year old and seniors from the previous had already left the college we only included current juniors and seniors in the experiment.[7] 843 out of 1392 eligible subjects participated in the treasure hunt (about 25 percent of all juniors and seniors).

On average, subjects reported at each update to have talked to 2.32 people. They named 1.74 people by name (75 percent). We define a named conversation partner

[5]More than 90 percent of undergraduates were already members of facebook.com at that time.

[6]Subjects could also list non-participants in our experiment as friends.

[7]We excluded current sophomores since the trivia game participation rate had been low among freshmen in the previous academic year.

as *informative* if (a) the partner was eligible for our experiment and (b) had logged in *before* the current subject submitted its updated guess. If subjects named random partners we would have expected only 25 percent of links to be informative. Instead, 83 percent of links were informative.[8]

5 Results

We first use only social network data rather than conversation data for estimating how agents process information from their social network. The limitations of this analysis highlight the importance of using conversation data.

5.1 Inferring Social Learning from Network Data

For each agent and question we consider the agent's final submission. We then utilize a simple logit approach to estimate variants of the following equation:

$$\begin{aligned} \text{Guess} &= \alpha + \beta \cdot \text{own signal} + \gamma \cdot \text{sum of signals of direct friends} + \\ &= \delta \cdot \text{sum of signals of indirect friends} + \epsilon \end{aligned} \tag{3}$$

The left-hand side variable is either 0 or 1 depending on whether the agent's guess is correct. On the right-hand side, we code an agent's own signal as -1 if her signal was incorrect and $+1$ if it was correct. An implicit assumption in this coding is that the agent's own signal has symmetric but oppositive effects when being correct or

[8] Submitting an uninformative link is not necessarily evidence of subject error or a random submission: students could have acted as conduits of information even if they were not direct participants in the treasure hunt.

incorrect. We separately test this hypothesis by estimating two separate coefficient and cannot reject equality.

The signals of *direct friends* in the social network (who also participated in the treasure hunt) are also coded using the same $+1/-1$ scheme. We then assign each direct friend to one of several *bins* and sum up the signals of all agents in the same bin. Formally, the 10 bin contains all direct friends with whom the agent is connected through 0 other common friends while the 11 bin contains all direct friends with whom the agent is connected through 1 common friend. We analogously define the 12 and 13 bins except that the latter bin contains *all* direct friends with 3 or more common friends. Note, that by construction each direct friend is contained in exactly one bin.

The signals of *indirect friends* who participated in the treasure hunt are similarly aggregated: the 20 bin contains all friends at distance 2 with whom the agent is connected through 0 other common friends. Analogously, we define the bins 21, 22 and 23.

While we also calculated bins for friends at distance 3 and higher, these coefficient are statistically insignificant at the 5 percent level and the magnitudes are also economically insignificant.

We treat all agent/question responses as independent observations – our results do not change qualitatively when we estimate each question separately except for somewhat larger standard errors. The logit estimates are reported in table 1. Two patterns stand out: (1) the estimated weights on indirect friends' information are substantially smaller than the estimated weights on direct friends' information. This indicates substantial *decay of information*. (2) Common friends do not increase the weight on direct friends' information but increase, almost linearly, the weight on indirect friends' information.

It is difficult to interpret the latter result: measurement error in our network elicitation might imply that the number of common friends is a proxy for having a direct link to a social neighbor. To exclude this possibility we have to utilize conversation data.

5.2 Inferring Social Learning from Conversation Data

Conversation data allows us to approximate the actual communication exchange within the social network much more precisely. Moreover, conversation data is *timed* which allows us to only count communication paths where signals are transmitted in a chronologically correct order. Two such sample networks are shown in figure 5.

We estimate the same logit regression as we did for the social network data except that we define all right-hand side variables using only the conversation network. The result are reported in table 2. The main difference between both tables is that the estimated coefficients on direct and indirect friends' information are scaled up by a factor of about 5. Own information is still weighed more heavily than direct friends' information – however, both coefficients are now of the same order of magnitude. This scale effect is expected since the social network provides a more noisy proxy for the actual communication structure than the conversation network.

However, the two main patterns we observed in figure 1 survive. There is substantial information decay with the weights on direct information 2 to 3 times larger than the weights on indirect information. Moreover, additional intermediaries increase the weight of indirect information but not the weight on direct information.

6 Discussion

It might be tempting to interpret the data as supporting the de Groot model since additional paths between receiver and sender increase the weight on indirect senders' information. However, neither the simple de Groot nor the simple streams model can handle information decay. Moreover, the fact that the weight on direct friends' information does not increase with more common friends is inconsistent with the predictions of the simple de Groot model.

To incorporate information decay in both models we assume that intermediaries always transmit direct information but "forget" to transmit indirect information with probability $1-p$. In the sample network of the introduction with two intermediaries, the receiver attaches a weight of $2p$ on the sender's information in the de Groot model. In the streams model, she attaches a weight of $2p-p^2$ where the second term capture the double-counting correction. Note, that for small p the two models become difficult to distinguish empirically. In both models, additional paths between sender and receiver increase the attention that the receiver pays to the sender's information.

However, only the streams model predicts that the weight on direct friends' information does not increase with additional paths connecting sender and receiver. The de Groot model predicts the same increasing pattern as for indirect information.

7 Conclusion

We conduct a unique experiment to study social learning in a real-world social network. Our field experiment gives us complete control over how much information each agent in the network possesses. In particular, we can equalize the precision of initial information across agents which makes it much easier to estimate a social learning model.

We find strong evidence for information decay in the social network. This is an important finding because standard models of social learning have focused on convergence of long-run beliefs and have, for the most part, ignored information decay. We also find tentative evidence that information is "tagged" and that agents might be smarter in how they update information than the assumptions of the simple de Groot model allow for. The streams model provides an elegant alternative to the de Groot model which is also analytically very tractable (in contrast to models of observational learning in networks or models of Bayesian learning from posteriors).

References

Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar, "Communication Dynamics in Endogenous Social Networks," working paper, MIT 2010.

Bala, Venkatesh and Sanjeev Goyal, "Learning from Neighbors," *The Review of Economic Studies*, 1998, *65* (3), 595 – 621.

Conley, Timothy G. and Christopher R. Udry, "Learning About a New Technology: Pineapple in Ghana," *American Economic Review*, March 2010, *100* (2), 1–28.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel, "Persuasion Bias, Social Influence and Unidimensional Opinions," *Quarterly Journal of Economics*, August 2003, *118* (3), 909–968.

Duflo, Esther, Michael Kremer, and Jonathan Robinson, "Understanding Technology Adoption: Fertilizer in Western Kenya, Evidence from Field Experiments," working paper, MIT 2006.

Golub, Benjamin and Matthew Jackson, "How Homophily affects Diffusion and Learning in Networks," Technical Report January 2010.

___ and ___ , "Naive Learning in Social Networks: Convergence, Influence and the Wisdom of Crowds," *American Economic Journal: Microecnomics (forthcoming)*, 2010.

Granovetter, Mark, "The Strength of Weak Ties," *American Journal of Sociology*, 1973, *78*, 1360–1380.

Katz, Elihu and Paul F. Lazarsfeld, *Personal influence: The part played by people in the flow of mass communications*, Glencoe, Illinois: Free Press, 1955.

Kremer, Michael and Ted Miguel, "The Illusion of Sustainability," *Quarterly Journal of Economics*, August 2007, *122* (3), 10071065.

Lazarsfeld, Paul, Bernard Berelson, and Hazel Gaudet, *People's Choice: How the Voter Makes Up His Mind in a Presidential Campaign*, Columbia University Press (3rd edition), 1944.

Leider, Stephen, Markus M. Mobius, Tanya Rosenblat, and Quoc-Anh Do, "Directed Altruism and Enforced Reciprocity in Social Networks," *Quarterly Journal of Economics*, November 2009, *124* (4), 18151851.

___ , ___ , ___ , and ___ , "What Do We Expect from Our Friends?," *Journal of the European Economic Association*, March 2010, *8 (forthcoming)* (1).

Munshi, Kaivan, "Networks in the Modern Economy: Mexican Migrants in the US Labor Market," *Quarterly Journal of Economics*, May 2003, *118* (2), 549–599.

Figure 1: Screen shot with questions and list of possible answers for "Treasure Hunt" experiment

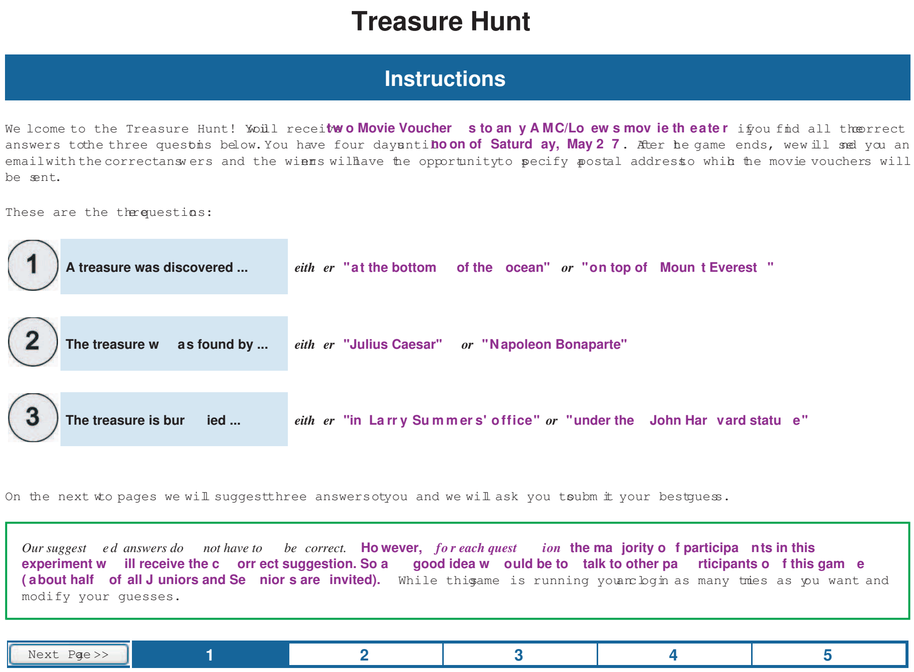

Figure 2: Screen shot with suggested answers for "Treasure Hunt" experiment

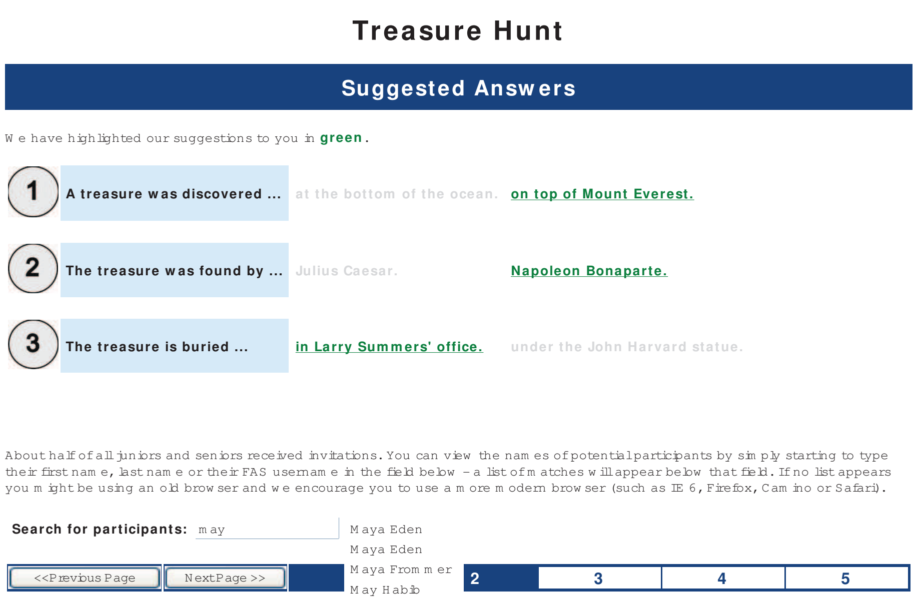

Figure 3: Screen shot with answer submission screen

Figure 4: Screen shot showing elicitation of conversations for "Treasure Hunt" experiment

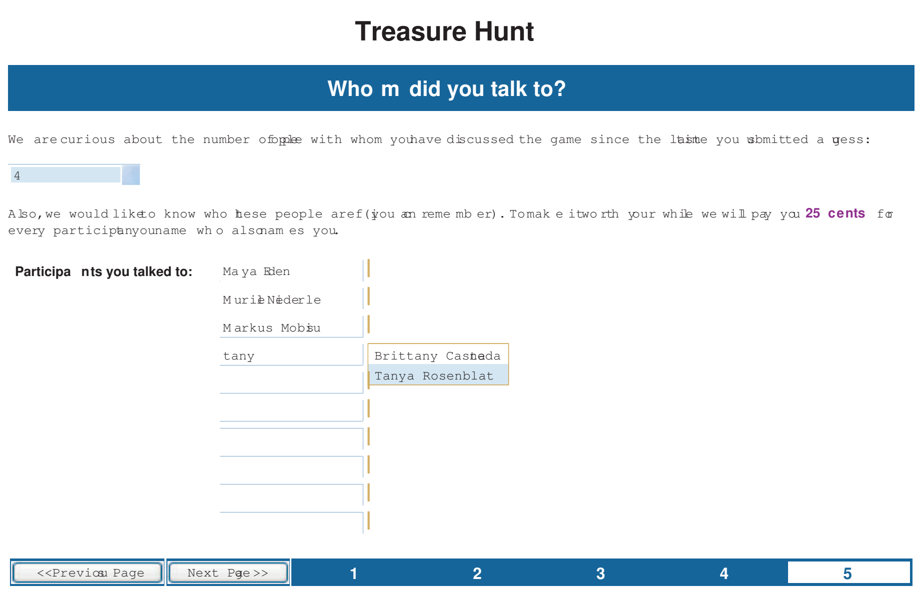

Figure 5: Two sample conversations

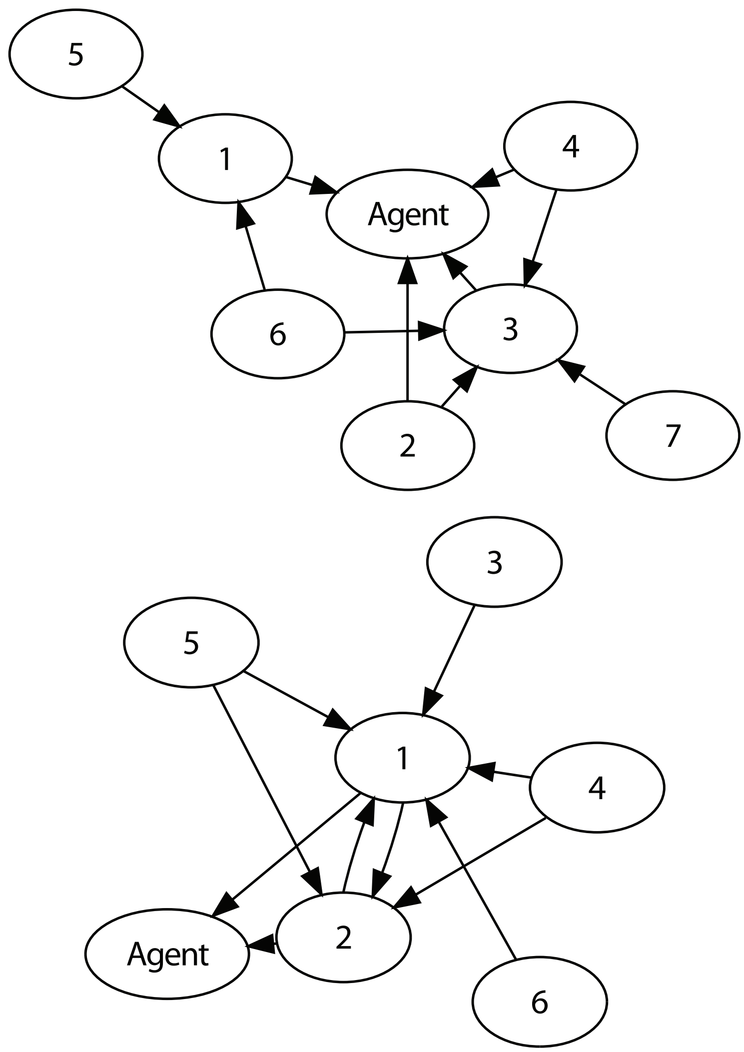

Table 1: Estimating social learning from network data only

Variable	
Intercept	0.383**
	(0.053)
CORRECTSIGNAL	0.972**
	(0.071)
INFO_D1_0	0.084**
	(0.032)
INFO_D1_1	0.104*
	(0.042)
INFO_D1_2	0.091**
	(0.037)
INFO_D1_3	0.095*
	(0.045)
INFO_D2_0	0.023†
	(0.012)
INFO_D2_1	0.050*
	(0.025)
INFO_D2_2	0.184*
	(0.071)
INFO_D2_3	0.155
	(0.120)
N	2,367

Table 2: Estimating social learning from conversation data

Variable	
Intercept	0.134^{**}
	(0.042)
CORRECTSIGNAL	1.173^{**}
	(0.042)
INFO_D1_0	0.590^{**}
	(0.052)
INFO_D1_1	0.490^{**}
	(0.063)
INFO_D1_2	0.630^{**}
	(0.078)
INFO_D1_3	0.490^{**}
	(0.064)
INFO_D2_0	0.115^{**}
	(0.023)
INFO_D2_1	0.204^{**}
	(0.053)
INFO_D2_2	0.338^{**}
	(0.083)
INFO_D2_3	0.318^{**}
	(0.080)
N	2,367