

Treasure Hunt: Social Learning in the Field*

Markus Mobius

Microsoft Research, University of Michigan and NBER

Tuan Phan

National University of Singapore

Adam Szeidl

Central European University, CEPR

March 1, 2015

Abstract

We seed noisy information to members of a real-world social network to study how information diffusion and information aggregation jointly shape social learning. Our environment features substantial social learning. We show that learning occurs via diffusion which is highly imperfect: signals travel only up to two steps in the conversation network and indirect signals are transmitted noisily. We then compare two theories of information aggregation: a naive model in which people double-count signals that reach them through multiple paths, and a sophisticated model in which people avoid double-counting by tagging the source of information. We show that to distinguish between these models of aggregation, it is critical to explicitly account for imperfect diffusion. When we do so, we find that our data are most consistent with the sophisticated tagged model.

JEL Classification: C91, C93, D83

*András Kiss and Jenő Pál provided outstanding research assistance. We are grateful to Abhijit Banerjee, Ben Golub, Bo Honore, Matthew Jackson and seminar participants for helpful comments and suggestions. Mobius and Szeidl thank for support the National Science Foundation (award #0752835). Phan thanks for support the National University of Singapore under grant R-253-000-088-133. Szeidl thanks for support the Alfred P. Sloan Foundation and the European Research Council under the European Unions Seventh Framework Program (FP7/2007-2013) ERC grant agreement number 283484.

1 Introduction

Social learning in networks is central to many areas in economics. Farmers talk about fertilizers with their neighbors; voters discuss politics with their friends; and fund managers pass information about stocks to their colleagues.¹ The mechanism of social learning has two key components: diffusion, which governs how far information travels; and aggregation, which describes how people combine information to form their opinions. Both limited diffusion and biased aggregation have been proposed as explanations for aggregate inefficiencies, such as the slow adoption of superior technologies or persistent disagreements about objective facts.² Existing research has mostly studied limited diffusion and biased aggregation separately. We argue that, to understand the exact mechanism of social learning and its contribution to aggregate inefficiencies, we need to study these components in combination.

In this paper we use a field experiment with college students on social learning to analyze both diffusion and aggregation. We find substantial social learning, much of which is realized through conversations between people who are not close friends. Thus precisely measuring the path of diffusion is important for fully capturing learning. But even in the network of conversations, diffusion is limited: information only travels two steps and is passed on imperfectly. We next show that—because it affects which signals are aggregated into opinions—explicitly accounting for imperfect diffusion is critical when trying to distinguish between models of information aggregation. When we do this, we find that people in our setting, probably by tagging the source of information, avoid the “double-counting” aggregation bias (DeGroot 1974) that has been emphasized as a key potential factor driving incorrect beliefs and persistent disagreements (DeMarzo, Vayanos and Zwiebel 2003, Golub and Jackson 2012). In combination with laboratory studies in which tagging is not permitted but diffusion is perfect, and which do find support for the DeGroot model (Chandrasekhar, Larreguy and Xandri 2012, Grimm and Mengel 2014), our results suggest that the DeGroot bias is more likely to emerge when the lack of tagging and the volume of information make updating difficult. Thus, ironically, biases may be stronger for precisely those topics that people discuss more.

In Section 2 we describe the experiment, which took place in May 2006 with 789 junior and senior students at a private US university. The experiment was framed as an online game of “Treasure Hunt”. Subjects had to find the correct answer in three binary choices about a hypothetical treasure (e.g., “the treasure was found by either Julius Ceasar or Napoleon Bonaparte”), and those who did received movie tickets. Each subject received a signal for

¹For fertilizers see Conley and Udry (2010), for political opinions see Lazarsfeld, Berelson and Gaudet (1944), and for investments see Hong, Kubik and Stein (2005).

²For example, Banerjee, Chandrasekhar, Duflo and Jackson (2013) emphasize limited diffusion as a factor behind the slow adoption of microfinance, and Golub and Jackson (2012) show how biases information aggregation can lead to persistent disagreements. We review below in more detail the literatures on limited diffusion, biased aggregation, and their overlap.

each question in the form of a suggested answer. Subjects were informed that their candidate answer need not be correct. They were also told that the majority of participants in the game had the correct answer, and that therefore by talking to others they could improve their information and guesses. During a four-day period subjects could log in as often as they liked and update their guesses. Each time they logged in, they were also asked to report with whom they talked about the game, resulting in direct data on conversations. In addition, from an earlier, unrelated project (Leider, Mobius, Rosenblat and Do 2009) we have independent data on subjects' friendship networks.

We can think about this experiment as being an analogue of social learning about whether or not to switch to a new technology, such as using fertilizer. Each person gets her own imperfect private signal about whether switching is beneficial, and can talk to others to learn about their private signals. In this environment, the data on conversations allow us to trace the path of diffusion; and the data on signals and guesses allow us to make inference about information aggregation.

In Section 3 we document a set of stylized facts. We start by asking how much people learn. In our full sample, the share of correct signals is 55% and the share of correct decisions is 61.3%, thus social learning improved guesses by 6.3 percentage points. In the subsample of subjects for whom at least one conversation is reported, the corresponding improvement is 10.3 percentage points; and in the subsample with no conversations it is 2.7 percentage points. Thus there is substantial learning, and conversations correlate with the extent of learning.

We next turn to explore diffusion. We begin with comparing the network of conversations and the network of friendships. We find that the majority of conversations took place between people who are neither direct nor indirect friends.³ Moreover, the structure of friendship and conversation networks is different: conversations are more concentrated by geography, as measured with the dorms of the subjects. These results suggest that by focusing only on friendship data we might miss a substantial part of social learning.

We then ask from whom people learn. To answer this, we regress a person's decision on her own signal and on the signals of others at different social distances. When we do this in the network of *friendships*, we find that direct and indirect friends' signals have very small effects on a person's decision. In contrast, when we do this in the network of *conversations*, we find that direct partners' signals have a weight of about 40%, and indirect partners' signals have a weight of about 10% as high as the weight on the person's own signal. Signals of people at higher social distance have insignificant and small effects. Thus even in the conversation network there are substantial diffusion frictions. We can also use this regression to decompose total learning into learning from different sources. In our full sample, over half of the 6.3 percentage point improvement in correctness can be attributed to learning from direct and

³By an indirect connection we refer to a pair of agents who are at social distance 2 in the network.

indirect conversation partners. In the sample of people with at least one conversation the result is even stronger: over 85 percent of the 10.3 percentage point improvement can be attributed to direct and indirect conversation partners. These results confirm that diffusion is limited to small distances.

We then turn to explore information aggregation. As DeMarzo et al. (2003) emphasize, a plausible and important mistake in aggregation, formalized by the DeGroot model, is persuasion bias: that a person—because she does not realize that they come from the same source—may put a larger weight on signals that reach her through multiple paths. To explore whether our data feature such double-counting, we again regress a person’s decision on other subjects’ signals, but now distinguish between these other subjects as a function of how many paths connect them to the person. Absent other forces, double-counting implies that subjects connected through more paths should have a larger influence on the person’s decision.

Our results seem initially puzzling. Consistent with the lack of bias, we do not find that influence weights increase in the number of paths when the source is a *direct* conversation partner. But, consistent with the bias, we do find that influence weights increase when the source is an *indirect* conversation partner. One way to reconcile these observations is to recall that we have documented significant frictions in information transmission. Such frictions may provide an alternative explanation for the increasing weights for indirect partners. If intermediaries transmit indirect information with less than full probability, then having more connecting paths can lead to increasing influence weights simply because there is a higher probability that the signal reaches the decision maker. This logic suggests that explicitly accounting for limited diffusion is important to identify the mechanism of information aggregation.

To explore these issues formally, in Section 4 we develop two theoretical models of learning. Our first model, which we call the “streams” model, is based on Acemoglu, Bimpikis and Ozdaglar (2014). This model assumes that information is “tagged”, i.e., that the name of the source is always transmitted together with the signal. By using tags, people avoid double-counting signals. Our second model is a “naive” model which is analogous to the DeGroot model in that it features double-counting, but to match our setting, is formulated not with continuous but with binary signals. We also develop a nested model parameterized by $\lambda \in [0, 1]$ which collapses into the streams model when $\lambda = 0$ and the naive model when $\lambda = 1$. We add diffusion frictions to this framework in the form of probabilistic transmission and underweighting of others’ signals. The resulting model generates a structural equation which provides microfoundations for the reduced-from regressions we used to explore double-counting.

In Section 5 we structurally estimate this model. Our first approach is a minimum distance procedure in which we look for parameters to match the reduced-from regression coefficients. Here we estimate λ very close to zero, i.e., we find support for the streams model. The intuition is straightforward. The streams model can explain that influence weights do not increase for

direct partners because direct transmission always takes place. It can also explain that weights do increase for indirect partners, because each indirect transmission takes place with less than full probability. In contrast, the naive model predicts increasing coefficients for both direct and indirect partners. Standard errors for λ are small enough that the naive model can be statistically rejected. We also estimate our model using a more powerful method of simulated moments, in which we look for parameters to match all decisions of all agents. Here we find an even more precisely estimated λ of approximately zero, suggesting that the additional information exploited in this estimation is also more consistent with the streams model.

It is useful to compare our results to the findings of Chandrasekhar et al. (2012) and Grimm and Mengel (2014), who use lab experiments to test the DeGroot model against the rational Bayesian alternative. These papers find support for the DeGroot model. Two salient differences between the environments in these papers and in ours are that (1) these have restrictive communication protocols in which only guesses—not signals or tags—can be transmitted; and (2) these feature perfect diffusion. Both of these differences can plausibly contribute to the differences in results. Having the opportunity to transmit tagged signals naturally makes it easier to avoid double-counting. And with limited diffusion a person hears fewer messages, reducing the burden of keeping track of their correlations. When the results of these papers and ours are taken together, they suggest that strong diffusion—which makes tagging difficult because signals travel far, and makes aggregation difficult because of many reports—may be necessary for the DeGroot bias to have large effects in the field. This observation suggests that people will exhibit stronger biases about precisely those topics which they discuss more. Whether many real-world topics feature sufficiently intensive diffusion for the DeGroot effect to be relevant is an open question.⁴

We build on and contribute to a body of work that studies how communication in a network leads to social learning.⁵ Models of limited diffusion go back to at least Bass (1969). Several papers explore diffusion of technologies using data on social connections, including Duflo and Saez (2003), Kremer and Miguel (2007), and in particular Conley and Udry (2010) who also have data on conversations about farming. In more recent work Banerjee et al. (2013) and Banerjee, Chandrasekhar, Duflo and Jackson (2014) combine theory and evidence to study how the precise mechanism of diffusion interacts with network structure. These papers do not focus on information aggregation. The foundational model of naive information aggregation in

⁴Also closely related are the lab experiments on observational learning by Choi, Gale and Kariv (2012) who test for the implications of a Bayesian learning model in three-player networks; by Mueller-Frank and Neri (2014) who test for general properties of the rules of thumb people use in updating; and by Corazzini, Pavesi, Petrovich and Stanca (2012) and Brandts, Giritligil and Weber (2014) who test between different variants of boundedly rational models. These studies too feature restricted communication and perfect diffusion.

⁵A closely related literature, emanating from Banerjee (1992) and Bikchandani, Hirshleifer and Welch (1992), builds models of observational social learning. Important studies that focus on networks include Bala and Goyal (1998), Gale and Kariv (2003), Eyster and Rabin (2010), Acemoglu, Dahleh, Lobel and Ozdaglar (2011) and Mueller-Frank (2013).

networks is the DeGroot (1974) model. DeMarzo et al. (2003) and Golub and Jackson (2010) explore the conditions under which this model can generate systematic biases in opinions, and Golub and Jackson (2012) show that in societies with homophily it predicts persistent disagreements. These papers do not focus on limited information diffusion.

Among the few papers that study both diffusion and aggregation when learning occurs via communication is the work of Ellison and Fudenberg (1995), which shows that when agents use rules of thumb for social learning, the strength of diffusion can affect information aggregation. Although their mechanism is different, they also predict that more diffusion can make learning worse. Acemoglu et al. (2014) explore the interaction of diffusion and aggregation in a tagged learning model in which agents who know enough can choose to stop participating in conversations. We do not focus on this mechanism in the current paper, but their framework forms the basis of our streams model.

2 Experimental Design and Data Description

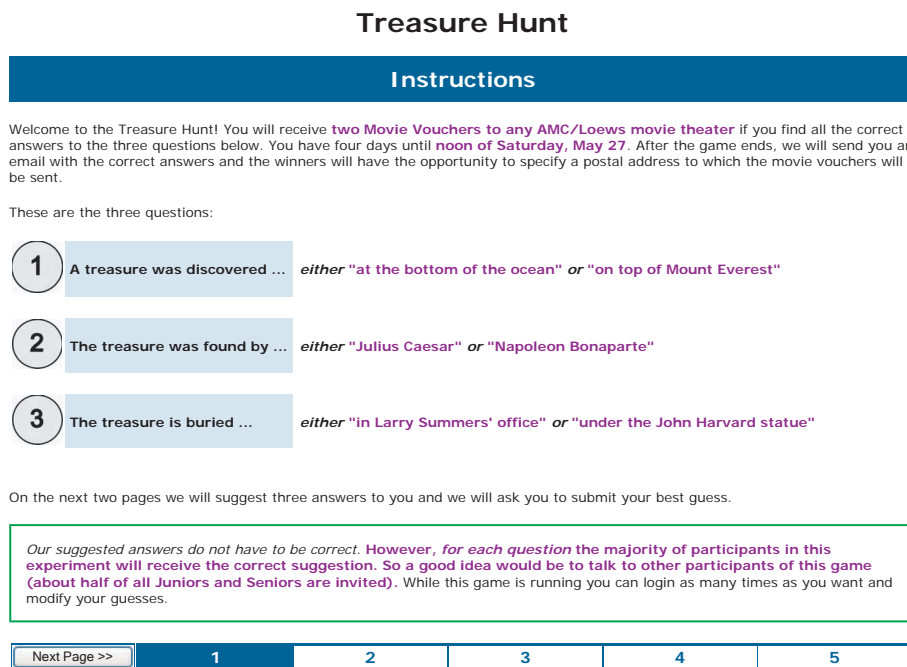
2.1 Design

We conducted our experiment in May 2006 with junior and senior students at a private US university. In both the experimental design and the analysis we use data, collected in December 2004 for two unrelated experiments, on the friendship network of these students. We first describe the friendship network data and then the design of the Treasure Hunt experiment.

Friendship network elicitation. Markus Mobius, Paul Niehaus and Tanya Rosenblat collected data on the friendship networks of undergraduates at the university in December 2004 for two unrelated experiments on directed altruism (Leider et al. 2009, Leider, Mobius, Rosenblat and Do 2010). Participants were recruited through a special signup campaign on [facebook.com](https://www.facebook.com).⁶ During the signup phase the experimenters collected the list of Facebook friends of each participant. Because the average number of Facebook friends exceeded 100, the experimenters also used a *trivia game technique* to elicit participants’ “close” Facebook friends. Each subject was asked to select a subset of 10 Facebook friends about whom they would be later asked trivia questions. Participants then received random multiple-choice questions by email about themselves (such as “What time do you get up in the morning?”). Afterwards, subjects who had listed these people as close friends received an email directing them to a web page where they had 15 seconds to answer the same multiple choice question about their friend. If the answers matched, both parties won a prize. This game provided incentives to list friends with whom subjects spent a sufficient amount of time to know their habits. We use all close friends listed by participants to construct the friendship network.

⁶More than 90% of undergraduates were already Facebook members at the time.

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



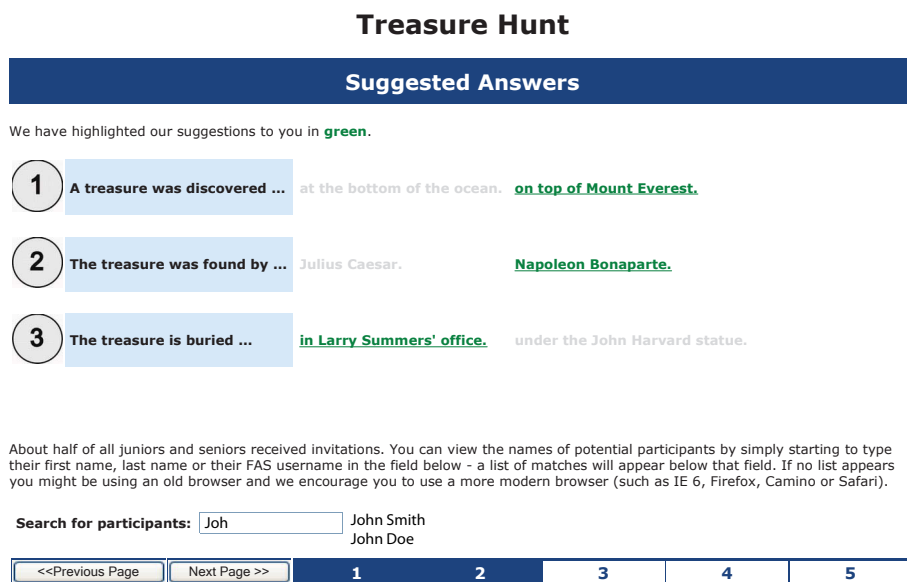
Participation rates in the friendship network elicitation were high, particularly for upper class students: 34, 45, 52, and 53 percent of freshmen, sophomores, juniors, and seniors participated. In total 2,360 out of 6,389 undergraduates participated, generating 12,782 links between participants and 10,818 links between participants and non-participants, for a total of 23,600 links, of which 6,880 were symmetric links.

Treasure Hunt. In May 2006 we invited by email subjects who had participated in the social network elicitation to a new online experiment. Because students who were freshmen in December 2004 had a lower participation rate in the friendship elicitation, and because students who were seniors in December 2004 had already graduated by 2006, we only invited students who had been sophomores or juniors in December 2004. These students were juniors or seniors in May 2006. The invitation email went out to 1392 eligible subjects (out of about 3,200 juniors and seniors in total). 789 students—about 25 percent of all juniors and seniors—participated in the Treasure Hunt.

The goal of the Treasure Hunt was to find the correct answer to three binary-choice questions. Every participant who correctly guessed the answer to all three questions received two movie ticket vouchers. Importantly, the Treasure Hunt game was non-rival: whether or not another student won movie tickets had no direct effect on a student’s winning.

The invitation email contained a link to the Treasure Hunt website. An image of the

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



“Instructions” screen, containing the three questions and the possible answers, is shown in Figure 1. On the same screen subjects were also informed that on the next screen they will receive suggested answers—which may or may not be correct—to each of these three questions. The screen encouraged subjects to talk to other participants in the game to learn more about the correct answers.

The next screen, shown in Figure 2, gave the subject’s suggested answers—in effect, signals—to the three questions. Subjects were told that the *majority* of participants would see the correct suggestion. Signals were independently drawn across subjects and questions and were correct with 57.5 percent probability.⁷ This screen also allowed subjects to search for the names of other participants in the game.

As shown in Figure 3, subjects could use the same web interface to submit a guess for each of the three questions. After submission they received an email with a link that allowed them to log back in and update their submission any time they wished. We encouraged subjects to repeatedly log in to the experiment and update their guesses. The last submission determined whether they won. In our emails we left the end date of the experiment vague, because we wanted subjects to update their guesses frequently rather than only once shortly before the end of the experiment.

Eliciting conversations. After the subject submitted her guesses, she was asked about the conversations she had with other participants since her last submission, or—if submitting for the first time—since the start of the experiment. On this page, shown in Figure 4, we first

⁷The share of correct signals in our data is slightly different due to sampling.

Figure 3: Screen shot with answer submission screen

Treasure Hunt

You can return to this page as often as you like during the next four days and update your choices if you receive new information or change your mind. If all your choices are correct, you will receive your movie tickets.

You submitted your last guess on **Friday 26th of May 04:50:37 AM** which is shown below. Please modify your guess if you changed your mind on a guess. You can use the invitation email to login again as often as you like!

1 A treasure was discovered ... at the bottom of the ocean. on top of Mount Everest.

2 The treasure was found by ... Julius Caesar. Napoleon Bonaparte.

3 The treasure is buried ... in Larry Summers' office. under the John Harvard statue.

Please don't forget to move to the next page so that your new guess gets saved!

<<Previous Page Next Page >> 1 2 3 4 5

asked subjects with *how many* people they talked in the meantime. After selecting a number, an input form appeared which asked subjects to list conversation partners by name. We took great care to make this part of the webpage user-friendly: whenever a participant started to type a name, the site suggested a drop-down menu of student names for auto-completion. To help distinguish true submissions from random submissions, the universe of names consisted not only of participants in Treasure Hunt, but of all undergraduates at the university. As an added incentive we paid subjects 25 cents for each submitted conversation partner who would reciprocate the submission when updating her own action.

To the best of our knowledge, our study is the first field experiment on social learning that directly records conversations.

2.2 Conversation Network

On average, our 789 subjects named 3.2 people as conversation partners. We define a named conversation partner to be *informative* if (a) the partner was eligible for our experiment and (b) had logged in *before* the current subject submitted her or his updated guess. These conditions ensure that the partner had information to share. About 83 percent of conversation links are informative, and in the analysis below we only use these conversations.⁸

⁸Submitting an uninformative partner is not necessarily evidence of error or random submission: students could have acted as conduits of information even if they were not participants in the Treasure Hunt. Repeating the analysis of the paper with uninformative conversations also included has very small effects on our results.

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

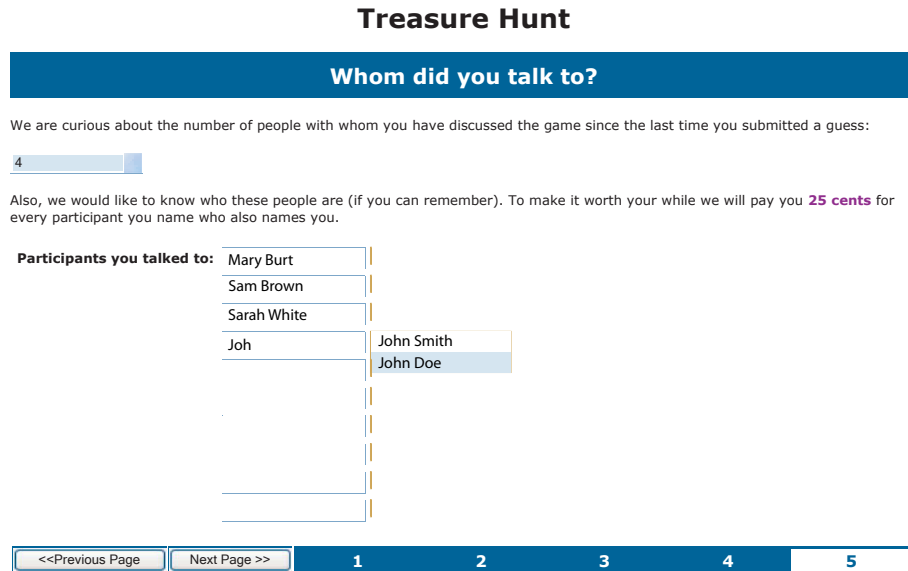



Figure 5: Sample conversation network from Treasure Hunt experiment



We next introduce the concept of a timed *conversation network*. Formally, we denote the set of conversations happening at any point in time t by g_t , such $(i, j) \in g_t$ implies that i talks to j at t (hence, j listens to i at t). The conversation network is then the set $\cup_t g_t$ of all conversations. We can draw a conversation network as a directed network such as in Figure 5, in which the direction captures the flow of information and each edge is labeled with the time at which conversation over that edge took place. Repeat conversations can be shown as parallel edges with distinct time stamps.

To construct an empirical conversation network from our data on conversations, we assume that in each reported conversation, information flows both ways. We also assume that the time at which a conversation is reported by a subject is a reasonable proxy for the time at which the conversation actually took place. By design, a reported conversation must have occurred before the time of the report. The assumption that it did not occur much before is supported by the feature of the design that the experiment’s end date was left vague, which encouraged subjects to log in multiple times.

For conversation links which are reported only once, the above assumptions uniquely generate a time stamp. For conversation links which are reported exactly once by both parties, we use the earlier of the two reporting times as a proxy for the conversation date. Finally, when conversations are reported at least twice by one of the two parties—approximately 6 percent of all conversation links fall into this category—we assume that multiple conversations occurred, and use the following procedure to estimate the timing. For each such pair A and B, we order the times at which A reports a conversation with B, and also the times at which B reports a conversation with A. We assign the first conversation time to the the smaller of the first elements of these two lists. We then assign the second conversation time to be the smaller of the second elements on these lists, and so on, until both lists are finished. For example, suppose that agent A reported conversations with B at times 1 and 5 while agent B reported a conversation with A only at time 3. Then our lists are (1,5) respectively (3); and in our conversation network there will be two conversations between A and B, taking place at times 1 and 5. Intuitively, we assume that the report of B at time 3 refers to the conversation A had already reported at time 1.⁹

Importantly, information in a conversation network can only flow in chronologically correct order. We call paths which respect this order *feasible paths*: for example, in the Figure the path $C \xrightarrow{t=1} B \xrightarrow{t=2} A$ is feasible while $B \xrightarrow{t=1} E \xrightarrow{t=1} A$ is not feasible.

Finally, we introduce the notion of the distance- d conversation network of a person. This is simply the union of all feasible paths of length at most d that end in that person. In most of the analysis we will focus on distance-2 conversation networks; For example, Figure 5 shows

⁹We also experimented with other protocols. The exact choice of timing for repeated conversations—probably because they constitute only about 6 percent of all conversations—has essentially no effect on our results.

Table 1: Summary statistics

	Full sample	Agents with conv. networks
Number of friendship links	4.85	5.63
Number of conversation links	3.19	6.81
Size of conversation networks ($SD \leq 2$)	12.60	25.74
Share of female decision makers	0.58	0.62
Year of college (either 3 or 4)	3.50	3.45
Number of subjects	789	370
Number of questions	2,367	1,110

that network of a person A in our data.¹⁰

2.3 Sample Definition and Summary Statistics

In our analysis we work with two samples. Our full sample contains all decision makers who submitted a guess. To more directly focus on people who engaged in conversations about Treasure Hunt, we also introduce the restricted sample of people with conversation links. These are subjects in the full sample who have at least one connection in the conversation network: either they named a conversation partner, or another subject named them. There are 789 people in our full sample, and 370 people in the restricted sample.

Table 1 reports summary statistics of key variables in these two samples. Counting also friends who are not in the sample, people in the full sample had on average 4.8 friendship links, while people in the restricted sample had 5.6 friendship links. Because the friendship data predate the realization of conversations, this difference is not mechanical, and suggests that people who have a conversation network are more social. These people also have more conversation links and more people in their distance-2 conversation networks, but these differences are mechanical. The two samples have about the same share of female students (58% and 62%) and are both distributed about evenly between juniors and seniors.

3 Reduced-form Analysis of Social Learning

In this section we identify some key facts in our data. In Section 4 we use these facts to develop theoretical models of social learning, and in Section 5 we structurally estimate these models.

¹⁰For readability, in the figure we replaced the exact time-stamps with integers without changing the set of feasible paths ending at person A.

Table 2: The extent of social learning

	Full sample (1)	Agents with conv. networks (2)	Agents without conv. networks (3)
Share of correct signals	55.0 %	56.7 %	53.5 %
Share of correct decisions	61.3 %	67.0 %	56.2 %
Improvement due to learning	6.3 pp	10.3 pp	2.7 pp

3.1 The Extent of Learning

Table 2 reports summary statistics, for three groups of subjects, on the extent of learning. Column 1 shows that in our full sample, the share of correct signals across all questions is 55%, while the share of correct guesses across all questions is 61.3%. Thus word-of-mouth learning improves the share of correct guesses by 6.3 percentage points. By column 2, in the sample of decision makers with conversation networks, the corresponding improvement due to learning is 10.3 percentage points. And by column 3, among people who do not have a conversation network, the improvement due to learning is much smaller, only 2.7 percentage points. These results show that the experiment succeeded in generating social learning, and confirm that the conversations data contain information on learning.

Fact 1. *There is substantial social learning.*

3.2 Friendships and Conversations

We now turn to explore information diffusion. We first ask through what links people pass information. The empirical literature on social learning often proxies conversations by measuring a range of social interactions, including friendships.¹¹ But how similar are conversation and friendship networks in practice?

To answer this question, we begin by looking at the overlap between friendships and conversations in our data. Out of the 1,655 conversation links we observe, only 471 (28%) are between direct friends; and another 322 (19%) are between indirect friends. Therefore, we miss over half of the conversations in our sample if we only focus on direct and indirect friendship links.

A natural explanation for this finding is that many conversations take place between people who are *not* close friends. A possible alternative explanation, however, is missing data on some

¹¹For example, Kremer and Miguel (2007) use data on friends, relatives, and additional social contacts with whom people tend to talk frequently; and Banerjee et al. (2013) use data measuring a range of social interactions, including kin, visits, borrowing, lending, advice, and others. These connections measure typical interactions, and hence may be distinct from the conversations people have on any given topic.

Table 3: Homophily by geographic and demographic type in the friendship and conversation network

	Same cluster of dorms	Dorms within 300 meters	Same dorm	Same gender	Same major	Same year
Friendship network	0.47	0.45	0.43	0.12	0.08	0.77
Conversation network	0.66	0.57	0.40	0.20	0.04	0.49

Notes: Coleman-style homophily measure. Higher numbers mean greater homophily. Zero means no homophily, one means that all links are between people who have the same type. There are two clusters of students dorms on campus which are about 1 mile apart.

friendships. For example, because the friendship network was elicited one year prior to the Treasure Hunt experiment, it is possible that some conversations took place between people who became close friends in the interim period.

To address this possibility, we next compare the *structure* of friendship and conversation networks. It is known that social networks exhibit homophily: similar people are more likely to be connected (Currarini, Jackson and Pin 2009). In particular, people tend to be connected with others who are geographically close (Marmaros and Sacerdote 2006) or have similar demographics. To measure homophily, we compute a version of the Coleman (1959) homophily index.¹² Fix a network, and consider a partition of its members by some geographic or demographic type (e.g., by major). Denote by H the share of all links in the network that are between people who have the same type, and by w the share of all pairs of people who have same type. Thus w measures the share of links between people of the same type in the complete network, or in a random network. Our “Coleman-style” homophily measure is then defined as:

$$\text{Homophily} = \frac{H - w}{1 - w}. \tag{1}$$

Intuitively, if there is no homophily then the share of same-type network links should be w . In that case our measure is equal to 0. Higher values mean that a larger share of links connect people who have the same type, and hence reflect higher homophily. If every observed link is between people who have the same type then our homophily measure is equal to 1.

We report this measure of homophily for various partitions in Table 3. Consistent with the literature, both networks exhibit strong homophily along both geographic types (measured with dorms) and demographic types (gender, major and year). However, conversation partners are far more likely than friends to be in neighboring dorms. At the same time, conversation partners are less likely than friends to be in the same major or year. This pattern suggests

¹²Note, the original Coleman index is computed for a group, whereas our Coleman-style index is computed for a partition.

that the incomplete overlap between friendship and conversation networks is not only due to measurement error: the networks are also structurally different. In fact, the observed pattern is consistent with the model of Rosenblat and Mobius (2004) in which agents prefer connections who have the same type, but are constrained by the scarce supply of such people in their neighborhood. If having a same-type partner is more important for a friendship than for a conversation, then conversation links will be relatively more concentrated by geography.

Fact 2. *Many conversations take place between non-friends. Conversations are more concentrated than friendships by geography.*

Fact 2 suggests that to study social learning we need to pay particular attention to conversation networks.

3.3 Decomposing Social Learning

From whom do people learn? A natural way to approach this question is to “bin” the information accessible to the decision maker according to social distance. Thus bin 0 contains the decision-maker’s own signal, bin 1 contains the signals of direct conversation partners, bin 2 the signals of indirect conversation partners and so on. Using this classification, we estimate the following reduced-form logit model:

$$\text{Decision}_{iq} = \text{sign} \left[\alpha + \beta_0 \cdot \text{own signal}_{iq} + \sum_{d=1}^m \beta_d \cdot \text{Info}_{iq}(d) + \epsilon_{iq} \right]. \quad (2)$$

Each person i decides on three questions q , hence each observation is a (person, question) pair. The left-hand side category variable is coded 1 or -1 depending on whether the person’s final answer for that particular question is correct. The right hand side includes a constant, the signals of various agents in the network, and an error term. Each signal is also coded as 1 or -1 depending in whether it is correct. By coding signals in this way we assume that agents react equally strongly to a correct signal and to an incorrect signal. The “own signal” variable is the signal of the decision maker for this particular question. The $\text{Info}_{iq}(d)$ variables equal the sum of signals of others at social distance d from the agent in the network (which is either the friendship network or the conversation network). Because signals are randomly assigned, the β_d coefficients measure the *causal* effects of others’ signals on the person’s decision. The constant α generates a shift in the share of correct guesses which is due to other factors such as information coming from socially distant agents. Finally, ϵ is an extreme value error term which can emerge from forgetting, misreporting, or updating mistakes.

To help interpret the coefficient estimates, imagine a world in which agents fully incorporate signals within some social distance d , gain no information from subjects further than d , and make guesses which are the correct Bayesian posterior perturbed by an extreme value error

Table 4: The effect of signals on decisions as a function of social distance

Dependent variable: decision is correct	Friendship network (1)	Conversation network (2) (3)	
Own signal	1.000	1.000	1.000
Info(1)	0.095*** (0.021)	0.405*** (0.041)	0.403*** (0.041)
Info(2)	0.024** (0.011)	0.108*** (0.019)	0.112*** (0.019)
Info(3)	0.010* (0.006)	0.015 (0.015)	
Constant	0.310*** (0.064)	0.248*** (0.050)	0.261*** (0.049)
Observations	2,367	2,367	2,367

Notes: Normalized logit regressions. Signals are coded as +1 if correct and -1 if incorrect. Info(d) is the sum of signals of agents at social distance d in the conversation network. Standard errors in parenthesis.

term. Because the signals are interchangeable, in such an environment we would observe $\beta_0 = \beta_1 = \dots = \beta_d > 0$ and $\beta_{d+1} = \dots = \beta_m = 0$. In turn, the magnitude of the β_0, \dots, β_d coefficients is pinned down by the variance of the error term. This logic suggests that the economically relevant quantities are not the levels but the ratios of the estimated coefficients. We work with the normalization $\hat{\beta}_d = \beta_d / \beta_0$; these normalized coefficients measure the weight the decision maker puts on the signal of a person at social distance d relative to the weight she puts on her own signal.

Table 4 reports the normalized coefficients from estimating (2) in various specifications.¹³ Note, the normalized coefficient on own signal is by construction equal to 1. Column 1 shows estimates in the friendship network in which friends up to social distance three are included. We estimate significant but small $\hat{\beta}_1 = 0.095$: the signal of a direct friend counts for less than 10% of the signal of the decision maker. The coefficients associated with higher social distances are even smaller. Consistent with Fact 2, the friendship network does not seem to capture well the flow of information.

The second column, which estimates the same specification for the conversations network, shows stronger results. The relative weight on a direct conversation partner is $\hat{\beta}_1 = 0.405$, and the relative weight on a partner at distance two is also significant at $\hat{\beta}_2 = 0.108$. The

¹³We estimate this equation in a pool, treating all question responses of all agents as independent observations. Appendix B explains how we compute standard errors for the normalized coefficients. Conducting robust inference (Cameron, Gelbach and Miller 2011) in which we allow for clustering either at the level of the decision maker or at the level of algorithmically detected communities (Blondel, Guillaume, Lambiotte and Lefebvre 2008) has only marginal effects on our standard errors, hence we chose to report simple standard errors.

signals of people at social distance three have an insignificant and small effect on decisions, and as column 3 shows, leaving them out of the regression has a negligible effect on the other coefficients.

Fact 3. *There is stronger information diffusion in the conversation network than in the friendship network.*

Fact 4. *Influence weights decline by social distance. People are not influenced by others at distance 3 or higher.*

Fact 4 may reflect limited information diffusion, forgetting, lower trust in other people’s signals, or other mechanisms. Given this result, in the rest of the analysis we focus only on the signals of agents who are within distance two in the conversation network. We also note that the point estimate for the normalized constant $\hat{\alpha} = \alpha/\beta_0$ is significantly positive in all specifications, indicating the presence of other sources of learning, for example, conversations which are missing from our data.

Decomposing learning. How much of learning is captured by the conversations data? To answer this question, it is useful to interpret regression (2) as a behavioral rule. This rule takes as inputs the signals of the decision maker and her direct and indirect conversation partners, and adds a constant and an error term to produce a decision. We can then ask how much the various terms contribute to total learning.

Table 5 reports the results. In Panel A we focus on the sample of all decision makers, and use the coefficients from column 3 of Table 4 for the behavioral rule. As we have seen, in this sample, relative to the share of correct signals (55 percent) social learning improves the share of correct guesses by 6.3 percentage points. Panel A shows that, if people in our data formed decisions according to this rule but using only their own signals and that of their direct conversation partners, then relative to the baseline of 55 percent we would observe an 1.8 percentage point increase in correctness. Thus direct conversation partners are responsible for about $(1.8pp/6.5pp =)$ 29% of social learning. When the signals of indirect conversation partners are also included, correctness increases by an additional 1.6 percentage points, or by 25 percent of the total learning effect. Thus in this sample direct and indirect conversation partners are responsible for more than half of social learning. Other sources and mistakes—captured by the constant and the error term—explain the rest.

Panel B repeats this calculation for our restricted sample of people who have a conversation link.¹⁴ Because in this sample we have better data on conversations, we expect to be able to account for a higher share of learning. This is indeed what we find. Out of 10.3 percentage points of total learning, 4.9 percentage points (about 47 percent) are due to direct friends, and

¹⁴We also re-estimate the behavioral rule for this sample.

Table 5: Decomposing social learning in the conversation network

<i>Panel A: Full sample</i>			
Share of correct signals	55.0 %		
Total learning	6.3 pp		
Learning from social distance 1	1.8 pp	Friends	0.9 pp
		Other conv. partners	0.9 pp
Learning from social distance 2	1.6 pp	Friends	0.2 pp
		Other conv. partners	1.4 pp
Other effects	2.9 pp		
<i>Panel B: Agents with conversation networks</i>			
Share of correct signals	56.7 %		
Total learning	10.3 pp		
Learning from social distance 1	4.9 pp	Friends	2.4 pp
		Other conv. partners	2.5 pp
Learning from social distance 2	4.0 pp	Friends	0.6 pp
		Other conv. partners	3.4 pp
Other effects	1.4 pp		

4 percentage points (about 39 percent) are due to indirect friends. Other sources and mistakes explain only 1.4 percentage points (about 14 percent).

Fact 5. *Direct and indirect conversation partners generate most of social learning.*

The second column in the table further decomposes learning from direct and indirect conversation partners as a function of how close they are in the person’s *friendship* network. In both the full and the restricted sample, less than half of the incremental learning from conversation partners can be attributed to direct and indirect friends. This result confirms that in our setting the conversation network—not the friendship network—captures best the paths of social learning.

3.4 Double-counting in Social Learning

The reduced form regression equation (2) treats agents at the same social distance symmetrically. In particular, the regression ignores that sometimes subjects listen to the same person multiple times, or access the signal of a person through multiple paths. AsDeMarzo et al. (2003) emphasize, a key distinguishing prediction of naive learning models is persuasion bias: that people should overweight information which they receive multiple times. We now turn to look for such effects in our data.

Repeated conversations. We begin by looking at the effect on learning of the same pair of people talking multiple times, a situation we label “repeated conversations”. We include in regression (2) a variable which equals the sum of signals of partners with whom—according to the conversations data—the decision maker talked more than once. In this regression (results not reported) we find that repetition increases the weight on direct partners’ signals from 0.391 to 0.636, or by about 60% – however, the effect is only marginally significant (p-value of 0.054). Importantly, this effect may be an artifact of our experimental design: while we had asked subjects to only list conversation partners to whom they had talked since their last submission, they might have instead simply listed conversation partners whose signals they remembered. Such recall bias could explain the increased weight placed on repeated conversations. Given the marginal significance and the plausible alternative explanation, we cannot conclusively determine whether subjects overweight signals from repeated conversations. However, since only about 6.4% of conversation links feature repetition, their effect on overall learning in our experiment is likely to be small. For the rest of the analysis we therefore ignore repeat conversations.

Multiple paths. People may also overweight—effectively double-count—signals they receive through different paths. To explore this effect, we consider the following refinement of our earlier reduced-form regression:

$$\text{Decision}_{iq} = \text{sign} \left[\alpha + \beta_0 \cdot \text{own signal}_{iq} + \sum_{l=0}^h \gamma_l \cdot \text{Info}_{iq}(1l) + \sum_{l=0}^h \delta_l \cdot \text{Info}_{iq}(2l) + \epsilon_{iq} \right]. \quad (3)$$

On the right-hand side we now decompose the signals of agents at different social distances into various “sub-bins” depending on the number of paths connecting that agent and the decision maker. More concretely, the $\text{Info}(dl)$ variables represent the sum of signals of others at social distance d to whom there are l additional length-two connecting paths. Thus for example $\text{Info}(10)$ is the sum of signals of direct connections to whom there are no other connecting paths of length two. And $\text{Info}(11)$ is the sum of signals of direct connections to whom there is exactly one connecting path of length two, i.e., direct partners with whom the agent has one common partner. Analogously, $\text{Info}(20)$ is the sum of signals of others at social distance two to whom there is no other path of length two besides the path which defines them as being at social distance two. For example, in the network of Figure 5 the signal of conversation partner E would be in $\text{Info}(11)$ of decision-maker A. Finally, we put the signals of all agents with h or more paths in the highest bin $\text{Info}(dh)$. We only focus on the signals of partners within social distance two; by construction, all such agents are contained in exactly one bin.

Our main interest is in the normalized coefficients $\hat{\gamma}_l = \gamma_l/\beta_0$ and $\hat{\delta}_l = \delta_l/\beta_0$. Table 6 reports estimates of (3) in the conversations network when $h = 3$ and when $h = 4$. The results in these two specifications are similar. Beyond the basic patterns familiar from Table 4—significant

Table 6: The effect of signals on decisions as a function of social distance and number of paths in the conversation network

Dependent variable: decision is correct	up to 3 paths (1)	up to 4 paths (2)
Own signal	1.000	1.000
Info(10)	0.416*** (0.060)	0.416*** (0.060)
Info(11)	0.447*** (0.076)	0.444*** (0.076)
Info(12)	0.568*** (0.097)	0.571*** (0.097)
Info(13)	0.384*** (0.066)	0.301*** (0.106)
Info(14)		0.428*** (0.081)
Info(20)	0.077*** (0.023)	0.076*** (0.023)
Info(21)	0.146*** (0.048)	0.139*** (0.049)
Info(22)	0.178** (0.088)	0.178** (0.089)
Info(23)	0.339*** (0.085)	0.286** (0.125)
Info(24)		0.412*** (0.116)
Constant	0.257*** (0.049)	0.258*** (0.049)
Observations	2,367	2,367

Notes: Normalized logit regressions. Signals are coded as +1 if correct and -1 if incorrect. Info(dl) is the sum of signals of agents at social distance d to whom there are l connecting 2-paths. The category with the highest l value also includes the signals of partners with more than l connecting paths. Standard errors in parenthesis.

learning, and coefficients declining with social distance—new patterns also emerge.

Fact 6. *For direct partners, weights do not increase in the number of paths. For indirect partners, weights do increase in the number of paths.*

These results might seem puzzling. The increasing weight on indirect partners seems consistent with persuasion bias under naive learning. But it is puzzling that we do not find the same pattern for direct partners. One possible explanation is that persuasion bias is active, but because our estimates are noisy the pattern is weak for direct connections. Another possible explanation is based on Fact 4 which suggests the presence of frictions in information transmission. If intermediaries transmit information they heard from other people probabilistically, then having more connecting paths can lead to increasing coefficients simply because there is a higher probability that the signal reaches the decision maker.

To be able to statistically distinguish between these explanations we next develop a formal

framework that allows for both naive and sophisticated learning, and can incorporate frictions in information transmission.

4 Models of Social Learning

We begin with models of sophisticated and naive information aggregation under perfect diffusion, and then add frictions in information transmission.

4.1 Environment

We consider a finite set of agents who each observe a signal about the state of the world. By talking to each other, agents can pool their information and improve their knowledge. Formally, Nature draws a state of the world $\theta \in \{-1, 1\}$ such that both realizations have equal prior probability. Agents do not observe Nature’s draw; but they each observe a conditionally i.i.d. binary signal $s_i \in \{-1, 1\}$ which is correct with probability $q > 0.5$.

There is a finite set of time periods $t = 1, 2, \dots, T$ during which agents can talk to each other. In each period t , a set of conversations, g_t , is realized: $(i, j) \in g_t$ implies that i talks to j (hence, j listens to i) at time t . The union of all conversations creates the conversation network. We do not model the process that generates the conversation network but instead focus on information aggregation *conditional* on its particular realization. While in our empirical analysis we assume that conversations are always bidirectional, in the model we also allow conversations which are directed.

After the conversation phase, each person i has to take a decision $a_i \in \{-1, 1\}$. The person derives utility 1 if her action is equal to the state of the world ($a_i = \theta$) and 0 otherwise. Therefore, she will always choose the action corresponding to the state of the world that she deems the most likely.

Recall from Section 2 that a *feasible path* from j to i by time t is a timed sequence of conversations along which information from j can reach i by time t . Formally, it is a sequence $j_0 \xrightarrow{t_1} j_1 \xrightarrow{t_2} \dots j_n \xrightarrow{t_{n+1}} j_{n+1}$ with $j_0 = j$ and $j_{n+1} = i$ such that $(j_{k-1}, j_k) \in g_{t_k}$, $t_{k-1} < t_k$ for all $1 \leq k \leq n + 1$, and $t_n \leq t$. Clearly, only information that reaches the decision maker on a feasible path by time t can influence her opinion at t . We thus define the *influencer set* F_i^t of agent i to include herself as well as all partners from whom there is a feasible conversation path reaching i by time t .

4.2 Sophisticated Learning: the “Streams” Model

In this model we follow Acemoglu et al. (2014) and assume that agents have unlimited memory and can exchange arbitrarily complex messages. This allows agents to communicate the full

stream of direct and indirect information—including, for each signal, a tag for the identity of the original source—they have accumulated up to that point.

To define this “streams” model formally, we define the set L_i^t of *last* conversations that an agent had with any conversation partner j up to time t . Thus $(j, \tau) \in L_i^t$ if j talked to i the last time at time $\tau < t$. We assume that every agent j who talks to agent i at time t sends a report R_{ji}^t which contains both the indirect and direct knowledge of j . Specifically, R_{ji}^t is the union of the *last* reports that agent j has herself received from her own conversation partners *other than* i (we exclude i ’s report because it would not be helpful to i) and of j ’s own tagged signal. Thus

$$R_{ji}^t = \bigcup_{(k, \tau) \in L_j^t, k \neq i} R_{kj}^\tau \cup \{(j, s_j)\}. \quad (4)$$

In particular, if j has not received any information from others then her report is simply her own (tagged) signal $R_{ji}^t = \{(j, s_j)\}$.

When it is time to make her decision, person i computes the union of the last reports she received from all her conversation partners and her own (tagged) signal. Taking the union gets rid of duplicate signals from the same source. Person i then sums all these independent signals to compute her *net knowledge* S_i^t . This is just the net number of +1 signals to which she has had access.

Given the binary information environment, person i ’s rational Bayesian posterior belief that $\theta = 1$, denoted μ_i^t , is simply determined by

$$\log \left[\frac{\mu_i^t}{1 - \mu_i^t} \right] = S_i^t \cdot \log \left[\frac{q}{1 - q} \right] \quad (5)$$

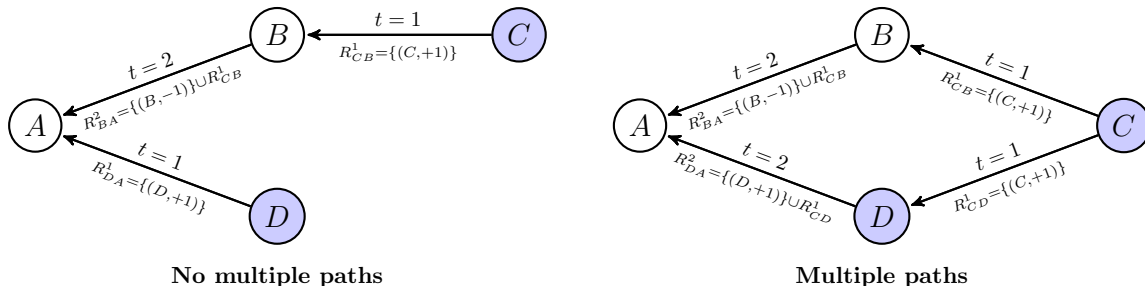
and hence her optimal decision is

$$\text{Decision}_i^t = \text{sign} \left[S_i^t \cdot \log \left[\frac{q}{1 - q} \right] \right]. \quad (6)$$

It is helpful to illustrate the mechanics of the streams model with two simple examples shown in Figure 6. Shaded nodes receive +1, white nodes receive -1 signals. In both networks, person A has access to all signals by $t = 2$; because the number of +1 and -1 signals is the same, her correct Bayesian posterior agrees with the prior. In the network on the left, there is no possibility of double counting, and hence it is easy to see that A correctly aggregates all information. In the network on the right, A aggregates the message $\{(B, -1), (C, +1)\}$ from the path in the top and $\{(D, +1), (C, +1)\}$ from the path in the bottom. Taking the union of these sets with $\{(A, -1)\}$ yields $\{(A, -1), (B, -1), (C, +1), (D, +1)\}$ and by summing the signals i arrives at the correct net knowledge $S_A^2 = 0$ and forms the correct posterior $\mu_A^2 = 0.5$.

It is straightforward to generalize the logic of these examples.

Figure 6: Information aggregation in streams model



Proposition 1. *In the streams model, the net knowledge of person i at time t is*

$$S_i^t = \sum_{k \in F_i^t} s_k. \quad (7)$$


In particular, i 's belief at time t is the Bayesian posterior of all signals received by agents in her influencer set F_i^t .

4.3 Naive Learning: A DeGroot-style model

The natural starting point for naive learning is the DeGroot (1974) model commonly used to explore wrong beliefs and disagreements (DeMarzo et al. 2003, Golub and Jackson 2012). This model makes two important simplifying assumptions: (1) instead of reporting all information they have received, agents only report a summary statistic of their opinion; (2) agents treat the reports they receive from their neighbors as independent. As a result of these two assumptions, a person will double-count information that comes from the same source but reaches her through multiple paths.

In the DeGroot model the state space is continuous, and the summary statistics people report are their mean beliefs. Here we develop a similar naive model in which the state space is binary, and the summary statistics people report are net sums of signals. Formally, we assume that when agent j talks to agent i at time t , she sends a summary report r_{ji}^t which is simply the *sum* of j 's signal and the *last summary* reports that agent j has herself received from any of her conversation partners *other than* i . Thus the basic structure exactly parallels the streams model, but instead of taking the union of complex messages, here we take the sum

Figure 7: Information aggregation in naive model



of simple summary statistics:

$$r_{ji}^t = \left(\sum_{(k,\tau) \in L_j^t, k \neq i} r_{kj}^\tau \right) + s_j. \quad (8)$$


In particular, if j has not received any information from others then her summary report is simply her signal $r_{ji}^t = s_j$. When it is time to make a decision, agent i sums the last reports of all her neighbors and adds her own signal to calculate her *net knowledge* \tilde{S}_i^t . She acts on this knowledge the same way as a fully rational agent:

$$\text{Decision}_i^t = \text{sign} \left[\tilde{S}_i^t \cdot \log \left[\frac{q}{1-q} \right] \right]. \quad (9)$$

To illustrate the mechanics of the naive model, consider the two simple examples shown in Figure 7. In the network on the left there is no possibility of double-counting, hence A correctly aggregates all information in the naive model as well. But in the network on the right, A now sums the message $(-1 + 1)$ from the path in the top, $(1 + 1)$ from the path in the bottom, and her own signal -1 to obtain a summary net knowledge of $+1$. She now fails to account for the fact that the signal of C reaches her on two paths, and hence double-counts that signal.

In order to precisely characterize information aggregation in the naive model, we introduce some concepts. First, we call feasible paths which only differ in the time stamps *congruent* feasible paths. These paths only exist in the presence of repeat conversations: for example, in the network in the left panel of Figure 8, $C \xrightarrow{1} B \xrightarrow{2} A$ and $C \xrightarrow{1} B \xrightarrow{3} A$ are two feasible congruent paths. Second, we say that a feasible path has *no reflection* if it does not contain any path segment where the same agent appears two steps apart. For example, in the right panel of Figure 8 the feasible path $C \xrightarrow{1} B \xrightarrow{2} D \xrightarrow{3} B \xrightarrow{4} A$ has reflection because B appears twice in

Figure 8: Conversation networks with congruent paths (left) and with reflection (right)



the segment $B \xrightarrow{2} D \xrightarrow{3} B$.¹⁵ Third, we call a feasible path without reflection that connects a person to herself a *loop*. Let P_{ki}^t denote the number of different non-reflective non-congruent feasible paths from j to i by time t .¹⁶

Proposition 2. *In the naive model, the net knowledge of person i at time t is*

$$\hat{S}_i^t = s_i + \sum_{k \in F_i^t} P_{ki}^t \cdot s_k. \quad (10)$$

In particular, in a given conversation network the naive model induces the same beliefs as the streams model for all signal realizations if and only if (1) for all i and j , all feasible paths without reflection from j to i are congruent, and (2) there are no loops.

The intuition is straightforward. Because people use only the last report from a conversation partner, an additional congruent path does not induce double-counting. Similarly, because a person does not echo back a partner’s report to her, an additional reflective path does not induce double-counting either. However, given the naivete in updating, all other occurrences of multiple feasible paths—including loops—do induce double counting. As an illustration of the Proposition, go back to the two conversation networks in Figure 7. In the left panel, there is a single feasible path from C to A , and hence A counts C ’s signal only once. In the right panel, there are two non-congruent feasible paths from C to A , and hence A counts C ’s signal twice.

4.4 Discussion

Our naive DeGroot-style model requires an agent to keep track of the last report received from *every* conversation partner. This specification is slightly more demanding for the agent than

¹⁵Note that there is also a non-reflective feasible path $C \xrightarrow{1} B \xrightarrow{4} A$ through which C can influence A .

¹⁶Formally, this is the highest number of different feasible paths from j to i by t such that no path is reflective and no two paths are congruent.

the standard DeGroot model in which people only keep track of a *single* number, their current belief. We could obtain a similarly naive model in our setting by having agents just keep track of the net sum of all reports and simply add every new report to that running tally.

The main reason why we do not use that “very naive” model is that it has implausible implications when multiple conversations take place between the same parties. Recall the two conversation networks in Figure 8. In the network on the left, in the very naive model A will double-count the signal of B simply because of listening to her twice. And in the network on the right, B will double-count her own signal because of reflection: D reflects back that B’s own signal in her report in period 3. As we have seen in Proposition 2, our preferred naive model deals with both of these situations correctly. We do note, however, that given the small number of repeated conversations these two naive models are essentially equivalent in our data.

4.5 Empirical Model

We now present an empirical model which nests the sophisticated and the naive learning models and adds information frictions. We then show that the resulting framework provides microfoundations for the reduced-form regressions we estimated in Section 3.4.

Information frictions. We assume that agents always incorporate their own signal in their report, but they only include indirect information with probability r . This assumption is consistent with the declining influence weights documented in Fact 4, and can be viewed as a model for the intermediary forgetting to transmit another person’s signal.¹⁷ For both the sophisticated and naive models, information frictions can be applied to feasible paths of arbitrary length. However, given that in the data information only seems to travel up to social distance two, for simplicity we only consider feasible paths of length at most two.¹⁸

We also allow agents to underweight everyone else’s signal relative to their own signal by a factor $o \leq 1$. This second type of information friction is also motivated by Fact 4, in particular, by the observation that people put a lower weight on the signals of direct conversation partners than on their own signals. Because these direct conversations are explicitly reported by one of the two parties, forgetting seems to be an unlikely explanation; a more plausible reason may be that people trust others’ signals less than they trust their own signal. We model this effect in a reduced-form fashion using the parameter o .

Behavioral implications. Because direct transmission always takes place, and because for indirect transmissions and aggregation people use the *last reports* from their partners, even with frictions repeated conversations do not affect the probability with which a signal reaches a decision maker.¹⁹ Frictions do alter the prediction of Proposition 1 about aggregation in the

¹⁷See Banerjee et al. (2013) for a related structural model of probabilistic diffusion.

¹⁸Our estimated transmission probability r is low enough that considering longer paths would have very small effects on our estimates.

¹⁹One can imagine alternative models—such as the very naive model introduced in section 4.4—in which

sophisticated model: while there is still no double-counting, now only signals that do not get “lost in transmission” affect opinions. And frictions also alter the prediction of Proposition 2 about aggregation in the naive model. As we just noted congruent paths still do not matter; but now the restriction on path length also rules out loops. Thus double-counting is governed simply by the number of non-congruent paths of length at most two through which a signal actually reaches the decision maker.

To derive the behavioral implications of these observations, let J_q^{jki} be an indicator for whether indirect transmission from j via k to i on question q was successful. This indicator is only defined when a $j \rightarrow k \rightarrow i$ feasible path exists. The J_q^{jki} are i.i.d., independent of the signals, and equal one with probability r . Then—in both the streams and the naive model—the number of times the signal of a *direct* partner j reaches i on question q equals

$$M_q^1(i, j) = 1 + \sum_{j \rightarrow k \rightarrow i \text{ feasible}} J_q^{jki}. \quad (11)$$

Here the 1 represents direct transmission that succeeds with certainty, and the sum represents indirect transmissions that each succeed with probability r . Similarly, the number of times the signal of an *indirect* partner j reaches i on question q equals

$$M_q^2(i, j) = \sum_{j \rightarrow k \rightarrow i \text{ feasible}} J_q^{jki}. \quad (12)$$

Here we only have indirect transmissions that each succeed with probability r .

In the naive learning model, a signal that reaches a person m times is counted m times. In the sophisticated learning model, a signal that reaches a person m times is counted $\min(1, m)$ times. For estimation purposes, we now introduce a λ -naive model which nests these two cases. For $\lambda \in [0, 1]$, we assume that a signal heard m times is counted $\lambda(m) = \lambda m + (1 - \lambda) \min(1, m)$ times. When $\lambda = 1$ this framework simplifies to the naive model, and when $\lambda = 0$ it simplifies to the streams model.

Combining these assumptions yields the main behavioral equation of our structural model. Person i makes a decision on question q based on the sign of

$$\text{opinion}_{iq} = s_{iq} + o \cdot \sum_{j \in N_1(i)} \lambda(M_q^1(i, j)) \cdot s_{jq} + o \cdot \sum_{j \in N_2(i)} \lambda(M_q^2(i, j)) \cdot s_{jq} + a + \eta_{iq}. \quad (13)$$

Here s_{iq} just corresponds to i taking into account her own signal. The two sums correspond to signals from direct and indirect friends, double-counted to an extent governed by λ , and

repeated conversations do affect the probability of transmission. Because only about 6 percent of conversations are repeated, it is unlikely that we can distinguish between models that only differ in the role of repeated conversations.

underweighted by o . We also allow for a constant shifter, the term a , which represents other sources of information that are not captured by our conversation network. And we model mistakes in updating and decision making with η_{iq} , an i.i.d. extreme value error term.

Equation (13) has a more intuitive form which uses the $\text{Info}(dl)$ variables introduced earlier and provides microfoundations for the reduced form regressions in Section 3.

Proposition 3. *In the structural empirical model (13) the opinion of person i on question q , the sign of which determines her decision, can be written as*

$$\text{opinion}_{iq} = \tilde{\alpha} + \text{own signal} + \sum_{l=0}^{\infty} \tilde{\gamma}_l \cdot \text{Info}_{iq}(1l) + \sum_{l=0}^{\infty} \tilde{\delta}_l \cdot \text{Info}_{iq}(2l) + \epsilon_{iq} \quad (14)$$

where $\tilde{\alpha} = a$, $\tilde{\gamma}_l = o(\lambda(1 + rl) + 1 - \lambda)$ and $\tilde{\delta}_l = o(\lambda r(1 + l) + (1 - \lambda)(1 - (1 - r)^{1+l}))$ and ϵ is mean-independent of all variables on the right hand side.

The form of equation (14) differs from that of the reduced-form regression (3) in Section 3.4 in only one respect: while in that regression the error term has an i.i.d. extreme value distribution, here we only know that ϵ is mean-independent of the right hand side. The intuition underlying Proposition 3 is that in equation (13) we can replace the random coefficients $\lambda(M_q^1(i, j))$ and $\lambda(M_q^2(i, j))$ that measure stochastic transmission with their expected values, and move the residual to the error term. Because the transmission probability depends only on social distance (denoted d) and the number of paths (denoted l), the expected coefficients of all signals in a given (d, l) bin are the same, allowing those signals to be grouped together into the $\text{Info}(dl)$ variables. And because J_q^{jki} are independent of the signals, the residuals that go into the error term are mean-independent of the $\text{Info}(dl)$ variables.

The only remaining step is to understand why $\tilde{\gamma}_l$ and $\tilde{\delta}_l$ have the form given in the proposition. Consider first $\lambda = 0$, i.e., the streams model. Because a signal of a direct connection is always heard, in the absence of double-counting that signal is counted exactly once, hence $\tilde{\gamma}_l = o$. These influence weights do not depend on the number of paths. In contrast, the signal of an indirect partner is heard if at least one of the l common partners pass it on, and in that event is counted exactly once. The probability of that happening gives $\tilde{\delta}_l = o(1 - (1 - r)^{1+l})$. Here, because indirect transmission is imperfect, influence weights do increase with the number of paths.

Consider next the naive learning model. Now the signal of a direct connection (with whom there are l extra indirect paths) is heard and counted, in expectation, $\tilde{\gamma}_l = o(1 + rl)$ times. Here, due to double-counting, even for a direct partner influence weights increase with the number of paths. A similar logic gives $\tilde{\delta}_l = or(1 + l)$, i.e., for an indirect partner influence weights also increase. Finally, when λ takes an intermediate value, we must compute the convex combination of these extremes.

5 Structural Estimation

We denote the parameter vector of our structural model by $\theta = (r, o, \lambda, a, \sigma_\eta)$. We now turn to estimate this model using two techniques: minimum distance estimation and a method of simulated moments.

5.1 Minimum distance estimation

Here we look for structural parameters to match the normalized logit coefficient estimates reported in Table 6. We denote the vector of the estimated normalized logit coefficients by $\hat{v} = (\hat{\alpha}, \hat{\beta}_1, \dots, \hat{\beta}_h, \hat{\gamma}_1, \dots, \hat{\gamma}_h)$. Our goal is to find parameters θ such that the normalized logit coefficients $v(\theta)$ implied by our structural model are as close as possible to the \hat{v} observed in the data. Specifically, denoting the estimated covariance matrix of \hat{v} by $\hat{\Sigma}$, we solve

$$\min_{\theta} (\hat{v} - v(\theta))' \hat{\Sigma} (\hat{v} - v(\theta)) \quad (15)$$

and compute standard errors for these estimates using standard minimum distance procedure (Cameron and Trivedi 2005).

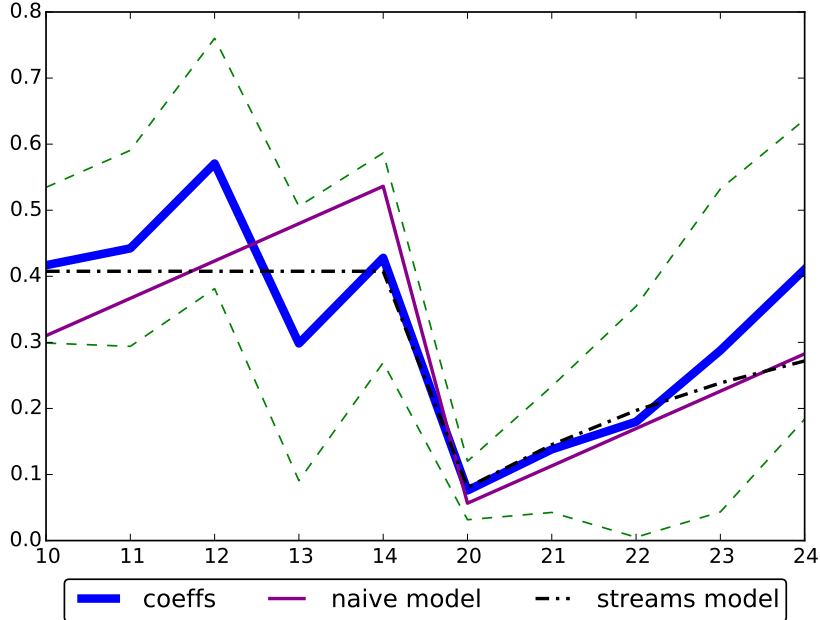
We implement this minimum distance estimation in two ways: first using analytically computable approximate formulas, and second using exact numerically computed values for $v(\theta)$. In the analytic approach, we approximate the model-implied logit coefficients $v(\theta)$ with the coefficients in equation (14) of Proposition 3. We denote the vector of these coefficients by $\tilde{v}(\theta) = (\tilde{\alpha}, \tilde{\beta}_1, \dots, \tilde{\beta}_h, \tilde{\gamma}_1, \dots, \tilde{\gamma}_h)$. The vector $\tilde{v}(\theta)$ is only an approximation for $v(\theta)$ because logit assumes that the error term has i.i.d. extreme value distribution, whereas in equation (14) we only know that the error term ϵ is mean-independent of the right-hand side variables. As we show below, this distinction has very small effects on our parameter estimates. The benefit of using $\tilde{v}(\theta)$ is that—as Proposition 3 shows—its components are simple functions of θ .

In the numerical approach, we estimate logit regressions in Monte-Carlo simulations of the structural model for a grid of θ values to infer the model-implied $v(\theta)$ coefficient vector. Here, the fact that $\theta = (r, o, \lambda, a, \sigma_\eta)$ takes values in a five-dimensional space imposes a computational burden. Because the approximate analytical formulas $\tilde{v}(\theta)$ do not depend on the last two parameters a and σ_η , it seems plausible that the exact numerical values $v(\theta)$ will also only be marginally affected by these parameters. We therefore simply set a and σ_η at the values implied by the logit regression estimates in the data.²⁰

We then take values of (r, o, λ) from a grid, and for each parameter vector simulate the

²⁰Recalling our notation that α stands for the constant and β_0 stands for the coefficient on own signal in our logit regression, we set $a = \alpha/\beta_0$ and $\sigma_\eta = \pi/(\sqrt{3}\beta_0)$ using the estimates in Table 6. We have experimented with other choices for these values, and these changes had negligible effects on our estimates.

Figure 9: Estimated and model-implied normalized coefficients



Notes: Horizontal axis shows (d, l) bins, where d is social distance and l is the number of paths between sender and receiver. Vertical axis shows normalized logit coefficients.

structural model.²¹ Specifically, holding fixed the conversations network and the signals of all agents, we take 150 independent draws from the J_q^{jki} variables that govern information transmission and the η_{jq} variables which represent mistakes in decision making and updating. For each realization we estimate the normalized logit coefficients in these simulated data, and average them across realizations to obtain $v(\theta)$. We also verify that using 300 draws instead results in small changes in these values.

Results. We begin with Figure 9 which summarizes the intuition behind our identification. The Figure is constructed from the estimates in column (2) in Table 6, which includes up to four additional paths for each informer. The horizontal axis corresponds to the (d, l) —distance and number of paths—bins of different informers: thus the left half of the figure refers to direct, the right half to indirect friends. The heavy line plots the empirical normalized logit coefficients corresponding to these bins—these are the values reported in Table 6. The dashed lines show the standard errors of the normalized logit coefficients. The optimization problem (15) corresponds to finding the parameters for which the model-implied coefficients get closest to the heavy line, trying to match in particular coefficients with small standard errors.

The figure also plots the implied normalized coefficients under the best-fit naive learning

²¹We restrict $(r, \alpha, \lambda) \in [0.1, 0.3] \times [0.35, 0.65] \times [0, 0.2]$ and use a $40 \times 40 \times 40$ grid. Widening the intervals within the unit cube affects neither point estimates nor standard errors.

Table 7: Minimum distance estimates of structural model

	Analytical		Monte Carlo		
	up to 3 paths	up to 4 paths	up to 4 paths unconstr. λ	up to 3 paths	up to 4 paths
	(1)	(2)	(3)	(4)	(5)
r	0.196*** (0.040)	0.196*** (0.041)	0.192*** (0.040)	0.218*** (0.048)	0.218*** (0.049)
o	0.414*** (0.054)	0.414*** (0.053)	0.425*** (0.053)	0.458*** (0.056)	0.450*** (0.054)
λ	0 (0.325)	0 (0.268)	-0.092 (0.265)	0.005 (0.181)	0.026 (0.171)

(thin line) and the best-fit streams (dash-dot line) specifications, obtained using our analytical approximate formulas by solving (15) under the constraint of $\lambda = 1$ respectively $\lambda = 0$. The figure reveals a better fit for the streams than for the naive model. Intuitively, the naive model predicts increasing coefficients for both direct and indirect informers: in the former case because of the double-counting logic, and in the latter case both because of double-counting and increasing transmission probability. In contrast, the streams model predicts flat coefficients for direct informers due to the absence of double-counting, but increasing coefficients for indirect friends due to increasing transmission probability. The data are more in line with the latter pattern.

Table 7 reports our point estimates and standard errors, allowing inference on the statistical confidence of this result. The first two columns report estimates based on the approximate analytical approach, in which we match the coefficients of the first respectively the second logit specification in Table 6. The point estimates for λ are zero in both specifications, showing that the pure streams model fits the data best. Because $\lambda = 0$ is at the boundary of the feasible set, asymptotic standard errors may not be valid. Therefore in column (3) we estimate the statistical model of column (2) with unconstrained λ . Here we find λ to be marginally negative, with the standard error essentially unchanged. In each of columns (1)-(3) standard errors are small enough that the naive learning model can be rejected.

Columns (4) and (5) of Table 7 report the results from the numerical (Monte-Carlo) implementation of the estimation. These coefficients are very similar: λ is estimated to be very close to zero and standard errors are again small enough that the naive learning model can be rejected. Estimating specification (5) with unconstrained λ (not reported) yields results which are identical to those reported in column (5). Taken together, the minimum distance estimates thus provide support for the streams model.

Table 8: Method of simulated moments estimates of structural model

	constrained λ	unconstrained λ
r	0.213*** (0.043)	0.213*** (0.041)
o	0.558*** (0.069)	0.558*** (0.067)
λ	0.041 (0.061)	0.041 (0.121)

5.2 Method of simulated moments estimation

A weakness of the minimum distance procedure is that—by only matching reduced-form regression coefficients—it does not use all information in the data. To address this problem, we next estimate our structural model using a method of simulated moments in which we try to match as closely as possible every decision of every agent. Because of the computational burden of searching over a five-dimensional space of θ values, similarly to the approach we took in the minimum distance Monte Carlo estimates, we set the parameters a and σ_η at the values implied by the logit regression in column (1) of Table 6.²² Then, for each parameter vector (r, o, λ) in a grid, every agent i and question q , we compute the model-implied probability that the agent chooses the right answer on that question, denoted $p_{iq}(r, o, \lambda)$. We do this by using the same 150 draws from the structural model used in the minimum-distance procedure, and computing the share of times that i made the right choice on question q .

Denoting the actual decision of agent i on question q by D_{iq} , coded as one or zero depending on whether it is correct, we then solve the following program:

$$\min_{r,o,\lambda} \sum_{i,q} (D_{iq} - p_{iq}(r, o, \lambda))^2. \tag{16}$$

We compute standard errors using a bootstrap in which we draw from the set of (person, question) pairs, with replacement, a sample as large as our actual data. Then for each (person, question) in this bootstrap sample, we take her conversation network, and use the resulting collection of conversation networks, signals and decisions as “data” in which we re-estimate (16) using the above simulated method of moments. We do this for 1000 bootstrap draws, and use the resulting distribution of parameter estimates to compute standard errors.

Table 8 reports the results. They align well with the minimum-distance estimates: λ is estimated at 0.041 and the standard errors are low enough to reject the naive learning

²²Here too, using other values for these parameters have negligible effects on our estimates.

specification. When we do not constrain λ to be in the unit interval, the results are similar. Interestingly, in these specifications the standard error of λ is lower than in the minimum-distance estimates. One possible reason is that the MSM approach exploits data in which an informer is connected to a decision maker through many paths, which is ignored in our minimum distance specifications and may contain information that helps distinguish the two models. We conclude that the data favor the streams model over the naive model.

5.3 Discussion

It is useful to compare our results to the findings of Chandrasekhar et al. (2012) and Grimm and Mengel (2014) who use lab experiments to test between naive and Bayesian models of social learning. In their environments, like in ours, the state space is binary and all subjects receive a noisy binary signal. However, the communication protocol is different: in these experiments agents observe in each round the guesses of all their exogenously assigned network neighbors. In this environment, the papers compare the DeGroot model with a Bayesian alternative, and find support for naive learning.

There are two salient differences between the environments in these papers and in ours: (1) these have communication protocols in which only guesses—not signals or tags—can be transmitted. (2) These feature perfect diffusion. One plausible reason for the differences in results is that these features make sophisticated learning more challenging, motivating subjects to instead rely on a rule of thumb. Indeed, when transmitting tagged signals is not allowed, netting out correlations for a given conversation network is a much less intuitive and much more complicated problem. And with unlimited diffusion the size and complexity of the conversation network is also larger, further increasing the burden of sophisticated learning. In contrast, with unstructured communication and limited diffusion people can both pass tagged signals and need only to keep track of a small number of signals which originate from sources they likely know. Thus both remembering tags and performing the aggregation is simpler.

The idea that psychological biases may be stronger in more difficult—though not necessarily more computationally complex—problems is not new. One prominent fact that can be explained using this logic is choice overload (Iyengar and Lepper 2000). When faced with a small set of jams, choosing is straightforward, and people make the utility-maximizing decision. But when faced with a large set of jams, choosing is difficult, which makes the default option of not purchasing more attractive. Thus the psychological bias of the “default effect” is more prominent in the more difficult decision.

In the social learning context this logic suggests that in real-world conversations naive learning should be more prominent when diffusion is stronger. Because real-world conversations are unstructured, tagging is in principle possible. But when—due to strong diffusion—signals come from many people including some who are far away in the network, keeping track is more

difficult, allowing the bias to emerge. This argument suggests that, ironically, wrong beliefs and persistent disagreements will emerge precisely in those topics that people tend to discuss more. Of course, even if one accepts this hypothesis, it is still an open question whether many real-world topics feature sufficient diffusion for naive learning to be an important source of bias.

6 Conclusion

We used a field experiment to learn about the micro mechanism of social learning. We found substantial social learning, and showed that most of it occurs through conversations between people who are not close friends. We also found that even in these conversation networks information transmission is highly imperfect. We then explored the mechanism of information aggregation by comparing a naive learning model that features double-counting and sophisticated model in which people tag the sources of information. We showed that to distinguish between these models, it is critical to explicitly incorporate imperfect information diffusion; and after doing so, we found that our evidence is most consistent with the sophisticated “streams” model. Finally, we argued that the hypothesis that naive learning is more likely to emerge in more difficult learning environments helps explain both other evidence from the lab and our evidence from the field.

Our findings suggest at least two new research directions. First, it would be interesting to explore whether variation in the intensity of diffusion affects the extent to which people make systematic learning mistakes. Second, it would be useful to better understand the choice of what people talk about. The fact that diffusion is limited suggests that people make choices about what information to pass. These choices are likely to be shaped by the environment. Changing incentives about topic choice can affect both the extent of diffusion and the extent of bias. We need more experiments and perhaps new theories to investigate these issues.

References

- Acemoglu, Daron, Kostas Bimpikis, and Asuman Ozdaglar**, “Dynamics of information exchange in endogenous social networks,” *Theoretical Economics*, 2014, 9 (1), 41–97.
- , **Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar**, “Bayesian Learning in Social Networks,” *Review of Economic Studies*, 2011, 78, 1201–1236.
- Bala, Venkatesh and Sanjeev Goyal**, “Learning from Neighbors,” *The Review of Economic Studies*, 1998, 65 (3), 595 – 621.
- Banerjee, Abhijit**, “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 1992, 107, 797–817.

- , **Arun G. Chandrasekhar**, **Esther Dufo**, and **Matthew O. Jackson**, “The Diffusion of Microfinance,” *Science*, 2013, *341* (6144).
- , — , — , and — , “Gossip: Identifying Central Individuals in a Social Network,” Working Paper 2014.
- Bass, Frank M.**, “A New Product Growth for Model Consumer Durables,” *Management Science*, 1969, *15* (5), pp. 215–227.
- Bikchandani, Sushil**, **David Hirshleifer**, and **Ivo Welch**, “A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades,” *Journal of Political Economy*, 1992, *100*, 992–1026.
- Blondel, Vincent D**, **Jean-Loup Guillaume**, **Renaud Lambiotte**, and **Etienne Lefebvre**, “Fast unfolding of communities in large networks,” *Journal of Statistical Mechanics: Theory and Experiment*, 2008, *2008* (10), P10008.
- Brandts, Jordi**, **Aya Ebru Giritligil**, and **Roberto A. Weber**, “An Experimental Study of Persuasion Bias and Social Influence in Networks,” 2014.
- Cameron, A. Colin**, **Jonah B. Gelbach**, and **Douglas L. Miller**, “Robust Inference With Multiway Clustering,” *Journal of Business and Economic Statistics*, 2011, *29* (2), 238–249.
- Cameron, Colin A.** and **Pravin K. Trivedi**, *Microeconometrics. Methods and Applications*, Cambridge, UK: Cambridge University Press, 2005.
- Chandrasekhar, Arun G.**, **Horacio Larreguy**, and **Juan Pablo Xandri**, “Testing Models of Social Learning on Networks: Evidence from a Framed Field Experiment,” 2012.
- Choi, Syngjoo**, **Douglas Gale**, and **Shachar Kariv**, “Social Learning in Networks: A Quantal Response Equilibrium Analysis of Experimental Data,” *Review of Economic Design*, 2012, *16*, 135–157.
- Coleman, James S.**, “Relational Analysis: The Study of Social Organizations with Survey Methods,” *Human Organizations*, 1959, *17* (4), 28–36.
- Conley, Timothy G.** and **Christopher R. Udry**, “Learning About a New Technology: Pineapple in Ghana,” *American Economic Review*, March 2010, *100* (2), 1–28.
- Corazzini, Luca**, **Filippo Pavesi**, **Beatrice Petrovich**, and **Luca Stanca**, “Influential listeners: An experiment on persuasion bias in social networks,” *European Economic Review*, 2012, *56* (6), 1276–1288.
- Currarini, Sergio**, **Matthew Jackson**, and **Paolo Pin**, “An Economic Model of Friendship: Homophily, Minorities and Segregation,” *Econometrica*, July 2009, *77* (4), 10031045.
- DeGroot, M. H.**, “Reaching a Consensus,” *Journal of the American Statistical Association*, 1974, *69*, 118–121.
- 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 and Emanuel Saez**, “The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment,” *Quarterly Journal of Economics*, August 2003, 118 (3), 815–842.
- Ellison, Glenn and Drew Fudenberg**, “Word-of-mouth Communication and Social Learning,” *Quarterly Journal of Economics*, 1995, 110 (1), 93–125.
- Eyster, Erik and Matthew Rabin**, “Nave Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2010, 2, 221–243.
- Gale, Douglas and Shachar Kariv**, “Bayesian Learning in Social Networks,” *Games and Economic Behavior*, 2003, 45, 329–346.
- Golub, Benjamin and Matthew O. Jackson**, “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2010, 2 (1), 112–49.
- and — , “How Homophily Affects the Speed of Learning and Best-Response Dynamics,” *Quarterly Journal of Economics*, August 2012, 127 (3), 1287–1338.
- Grimm, Veronika and Friederike Mengel**, “Experiments on Belief Formation in Networks,” 2014.
- Hong, Harrison, Jeffrey D Kubik, and Jeremy Stein**, “Thy Neighbor’s Portfolio: Word-of-Mouth Effects in the Holdings and Trades of Money Managers,” *Journal of Finance*, 2005, 60 (6), 2801–2824.
- Iyengar, Sheena S. and Mark R. Lepper**, “When Choice is Demotivating: Can One Desire Too Much of a Good Thing?,” *Journal of Personality and Social Psychology*, 2000, 79 (6), 995–1006.
- Kremer, Michael and Ted Miguel**, “The Illusion of Sustainability,” *Quarterly Journal of Economics*, August 2007, 122 (3), 1007–1065.
- 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), 1815–1851.
- , — , — , and — , “What Do We Expect from Our Friends?,” *Journal of the European Economic Association*, March 2010, 8 (1).
- Marmaros, David and Bruce Sacerdote**, “How Do Friendships Form?,” *Quarterly Journal of Economics*, 2006, 121 (1), 79–119.
- Mueller-Frank, Manuel**, “A General Framework for Rational Learning in Social Networks,” *Theoretical Economics*, 2013, 8, 1–40.
- and **Claudia Neri**, “Social Learning in Networks: Theory and Experiments,” 2014.

Rosenblat, Tanya S. and Markus M. Mobius, “Getting Closer or Drifting Apart,” *Quarterly Journal of Economics*, August 2004, 119 (3), 971–1009.

A Proofs

A.1 Proof of Proposition 1.

The signal of every influencer will reach decision maker i through a sequence of conversations. Every signal is counted at most once due to tagging. Therefore, agent i will form correct beliefs based on the signals of all of her influencers.

A.2 Proof of Proposition 2.

From equation (8) it follows that the signal s_j of person j adds to the net knowledge \hat{S}_i^t of person i for each instance of a feasible path from j to i by t , denoted $j = j_0 \xrightarrow{t_1} j_1 \xrightarrow{t_2} \dots j_n \xrightarrow{t_{n+1}} j_{n+1} = i$ along which (1) for all $u = 1, \dots, n$, $j_{u-1} \neq j_{u+1}$; and (2) for all $u = 1, \dots, n + 1$, the report j_{u-1} sends to j_u at t_u is the *last report* before t_{u+1} (here we let $t_{n+2} = t + 1$). Condition (1) follows because in (8) a person never reports back to her conversation partner. Condition (2) follows because in (8) only last reports are being reported.

It is clear that condition (1) is equivalent to the path having no reflection. We next show that condition (2) is equivalent to counting each feasible congruent non-reflective path exactly once. These two observations imply the statement of the Lemma.

Partition the set of all feasible non-reflective paths by the equivalence relation of congruence. We claim that in each partition set exactly one path will satisfy condition (2): namely the one for which $\sum_{u=0}^{n+1} t_u$ is the highest among all paths in the partition set. Denote this “maximal” path by π .

We first show that path π does satisfy condition (2). Suppose not. Then for some u , the report j_{u-1} sends to j_u at t_u is *not* the last report before t_{u+1} . Thus another $j_{u-1} \rightarrow j_u$ conversation must take place at some t'_u such that $t_u < t'_u < t_{u+1}$. But then by replacing t_u with t'_u in π we get another congruent path which has a higher sum of time stamps, contradicting our initial assumption.

We next show that no other path congruent with π satisfies condition (2). Towards a contradiction, suppose that path π' congruent with π also satisfies the condition. Let $j_{u-1} \rightarrow j_u$ be the last communication in which these two paths differ: $t_u \neq t'_u$, but for all $u' > u$ the time stamps on the two paths agree. If $t_u < t'_u$ then by replacing t_u for t'_u in P we increase the sum of time stamps, a contradiction. Conversely, if $t_u > t'_u$, then in path π' the report at t'_u is not the last report to j_u before t_{u+1} , contradicting the assumption that π' satisfies condition (2).

A.3 Proof of Proposition 3.

Begin with the streams model. Net knowledge is the sum of direct and indirect signals that reach the decision maker. Direct signals are always passed on and hence have a coefficient of o in net knowledge. Each indirect signal is potentially passed on as many times as there are feasible paths from the source to the decision maker. Each of these passings occur with probability r . Thus the contribution to net knowledge of the signal of an indirect friend j who is connected by $l + 1$ feasible paths is $B(l + 1)s_{jq}$ where $B(l + 1)$ is the maximum of $l + 1$ independent binary random variables, each of which equals one with probability r . We can write this term as $EB(l + 1) \cdot s_{jq} + (B(L + 1) - EB(L + 1)) \cdot s_{jq}$. The second term here, which we denote ϵ_{iq}^j , is mean-independent of s_{jq} and, because signals and transmission events are independent, also independent of all other signals. Computing the expected value $EB(l + 1)$, the contribution of s_{jq} to net knowledge can be written as $(1 - (1 - r)^{l+1}) \cdot s_{jq} + \epsilon_{iq}^j$. Summing across all j indirect partners and adding the independent extreme-value error term gives the result.

Now consider the naive learning model. The basic logic is the same, but now we need to compute the expected number of times i hears a direct respectively indirect partner's signal under the naive model. Because there are no repeat conversations, for a direct partner person i counts the signal once through the direct conversation. Person i counts this signal with probability r for each indirect path, and because there are no repeated conversations, each such path must go through a different conversation partner. This logic explaining $\hat{\gamma}_l = (1 + rl)$. For indirect partners, person i counts the signal of j with probability r for each distinct indirect path, explaining $\hat{\delta}_l = r(1 + l)$.

B Estimating the Normalized Logit²³

Setup

The agent makes a 0-1 decision based on the sign of the latent variable ($y_i = 1$ when $y_i^* \geq 0$):

$$y_i^* = h(x_i, \beta) + \varepsilon_i$$

where $\varepsilon_i \sim \text{Logistic}(0, 1)$ (standard logistic distribution). Notation: i indexes the observation, p indexes the parameters ($p = 1, \dots, P$), y_i and ε_i are scalars, x_i is a $1 \times P$ row vector, and β is a $P \times 1$ column vector. The difference from the usual logistic regression is that the latent variable model is not linear in parameters, but is defined by the function

$$h(x_i, \beta) = \beta_n \cdot (\beta_1 x_i^1 + \dots + \beta_{n-1} x_i^{n-1} + x_i^n + \beta_{n+1} x_i^{n+1} + \dots + \beta_P x_i^P)$$

²³We thank Andras Kiss for helping us develop this methodology.

instead. In this setup, $\beta_{p \neq n}$ are the $(P - 1)$ normalized parameters and β_n is the parameter that normalizes all the others (relative to the linear case).

Log-likelihood function

Probability of observing the data (y_i, x_i) with a given set of parameters:

$$f(y_i, x_i | \beta) = [G(h(x_i, \beta))]^{y_i} \cdot [1 - G(h(x_i, \beta))]^{1-y_i}$$

where $G(\cdot)$ is the standard logistic CDF.

The log-likelihood of observation i :

$$l_i(\beta) \equiv \log f(y_i, x_i | \beta) = y_i \log [G(h(x_i, \beta))] + (1 - y_i) \log [1 - G(h(x_i, \beta))]$$

The maximum likelihood estimator of β chooses $\hat{\beta}$ to maximize $\sum_{i=1}^N l_i(\hat{\beta})$.

Score function

The score is the $P \times 1$ vector of partial derivatives of $l_i(\beta)$. The p th element of the score:

$$\begin{aligned} s_i^p(\beta) &\equiv \frac{\partial l_i(\beta)}{\partial \beta_p} = y_i \frac{g(h(x_i, \beta)) \cdot \frac{\partial h(x_i, \beta)}{\partial \beta_p}}{G(h(x_i, \beta))} - (1 - y_i) \frac{g(h(x_i, \beta)) \cdot \frac{\partial h(x_i, \beta)}{\partial \beta_p}}{1 - G(h(x_i, \beta))} \\ &= \frac{g(h(x_i, \beta)) \cdot \frac{\partial h(x_i, \beta)}{\partial \beta_p} \cdot [y_i - G(h(x_i, \beta))]}{G(h(x_i, \beta)) \cdot [1 - G(h(x_i, \beta))]} \end{aligned}$$

In vector-form:

$$s_i(\beta) = \frac{g(h(x_i, \beta)) \cdot \nabla_{\beta} h(x_i, \beta) \cdot [y_i - G(h(x_i, \beta))]}{G(h(x_i, \beta)) \cdot [1 - G(h(x_i, \beta))]}$$

where:

$$\nabla_{\beta} h(x_i, \beta) = \begin{bmatrix} \beta_n x_i^1 \\ \vdots \\ \beta_n x_i^{n-1} \\ \frac{h(x_i, \beta)}{\beta_n} \\ \beta_n x_i^{n+1} \\ \vdots \\ \beta_n x_i^P \end{bmatrix}$$

is the Jacobian-vector of function $h(x_i, \beta)$.

Since the distribution of the error term is standard logistic with the following CDF:

$$G(z) \equiv \Pr(\varepsilon_i < z) = \frac{e^z}{e^z + 1}$$

and PDF:

$$g(z) \equiv G'(z) = \frac{e^z}{(e^z + 1)^2},$$

the score function actually simplifies a lot:

$$s_i(\beta) = \frac{\frac{e^{h(x_i, \beta)}}{(e^{h(x_i, \beta)} + 1)^2} \cdot \nabla h \cdot [y_i - G(h(x_i, \beta))]}{\frac{e^{h(x_i, \beta)}}{e^{h(x_i, \beta)} + 1} \cdot \left[1 - \frac{e^{h(x_i, \beta)}}{e^{h(x_i, \beta)} + 1}\right]} = \nabla_{\beta} h(x_i, \beta) \cdot [y_i - G(h(x_i, \beta))].$$

The score for the usual logistic model is $x'_i \cdot [y_i - G(x_i \beta)]$, which is exactly what we get by substituting $x_i \beta$ for $h(x_i, \beta)$.

Hessian

To derive the Hessian-matrix $H_i(\beta)$ (the Jacobian of the score function), we need to take partial derivatives again. Proceeding step-by-step:

$$\frac{\partial}{\partial \beta_p} s_i(\beta) = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ x_i^p \\ 0 \\ \vdots \\ 0 \end{bmatrix} [y_i - G(h(x_i, \beta))] - \nabla_{\beta} h(x_i, \beta) \left[g(h(x_i, \beta)) \cdot \frac{\partial h(x_i, \beta)}{\partial \beta_p} \right] \quad \text{for } p \neq n$$

and

$$\frac{\partial}{\partial \beta_n} s_i(\beta) = \begin{bmatrix} x_i^1 \\ \vdots \\ x_i^{n-1} \\ 0 \\ x_i^{n+1} \\ \vdots \\ x_i^P \end{bmatrix} [y_i - G(h(x_i, \beta))] - \nabla_{\beta} h(x_i, \beta) \left[g(h(x_i, \beta)) \cdot \frac{\partial h(x_i, \beta)}{\partial \beta_n} \right].$$

In matrix-form:

$$H_i(\beta) = \nabla_{\beta}^2 h(x_i, \beta) \cdot [y_i - G(h(x_i, \beta))] - \nabla_{\beta} h(x_i, \beta) \cdot [\nabla_{\beta} h(x_i, \beta)]' \cdot g(h(x_i, \beta))$$

where

$$\nabla_{\beta}^2 h(x_i, \beta) = \begin{bmatrix} 0 & \cdots & 0 & x_i^1 & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & x_i^{n-1} & 0 & \cdots & 0 \\ x_i^1 & \cdots & x_i^{n-1} & 0 & x_i^{n+1} & \cdots & x_i^P \\ 0 & \cdots & 0 & x_i^{n+1} & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & x_i^P & 0 & \cdots & 0 \end{bmatrix}$$

is the matrix of second derivatives of function $h(x_i, \beta)$.

In the usual linear-in-parameters logistic setup, $\nabla_{\beta}^2 h(x_i, \beta) = 0$ and the $P \times P$ Hessian is expressed as $H_i(\beta) = -x_i' \cdot x_i \cdot g(x_i\beta)$, which is also what we get by substituting $x_i\beta$ for $h(x_i, \beta)$.

Point estimates

Because the normalized model is statistically equivalent to the standard logit model, estimating the normalized model by maximum likelihood yields identical results to those obtained by first estimating standard logit and then normalizing the coefficients. Hence we follow this latter approach to obtain our point estimates.

Standard errors

For the asymptotic variance matrix, we implemented the estimator:

$$\left[-\sum_{i=1}^N H_i(\hat{\beta}) \right]^{-1}$$

which is the standard one for ML. Alternatively, we could use the outer product of the score:

$$\left[\sum_{i=1}^N s_i(\hat{\beta}) \cdot s_i(\hat{\beta})' \right]^{-1}$$

which is asymptotically equivalent, but—supposedly—behaves poorly in small to moderate samples.

We get standard errors for the estimates by taking the square root of the diagonal elements of the estimated variance matrix. These are similar to those yielded by the delta method following the original logit procedure (sometimes smaller, sometimes larger, but of the same magnitude).