DEPARTMENT OF ECONOMICS

UNIVERSITY OF MILAN - BICOCCA

WORKING PAPER SERIES

Influential Listeners: An Experiment on Persuasion Bias in Social Networks

Luca Corazzini, Filippo Pavesi, Beatrice Petrovich, Luca Stanca

No. 196 – August 2010

Dipartimento di Economia Politica
Università degli Studi di Milano - Bicocca
http://dipeco.economia.unimib.it

Influential Listeners: An Experiment on Persuasion Bias in Social Networks

Luca Corazzini*, Filippo Pavesi†, Beatrice Petrovich‡, Luca Stanca§

August 12, 2010

Abstract

This paper presents an experimental investigation of persuasion bias, a form of bounded rationality whereby agents communicating through a social network are unable to account for possible repetitions in the information they receive. The results indicate that network structure plays a significant role in determining social influence. However, the most influential agents are not those with more outgoing links, as predicted by the persuasion bias hypothesis, but those with more incoming links. We show that a boundedly rational updating rule that takes into account not only agents' outdegree, but also their indegree, provides a better explanation of the experimental data. In this framework, consensus beliefs tend to be swayed towards the opinions of influential listeners. We then present an effort-weighted updating model as a more general characterization of information aggregation in social networks.

*Department of Economics, University of Padua, Via del Santo, 33, 35123, Padova, Italy, and ISLA, Bocconi University, Via Roentgen, 1, 20136, Milan, Italy. Email: luca.corazzini@unipd.it

†Department of Economics, University of Milan - Bicocca. Piazza dell'Ateneo Nuovo 1, 20126 Milan, Italy, and Central European University (CEU), Department of Public Policy, Budapest, Hungary. E-mail: filippo.pavesi@unimib.it

‡Department of Economics, University of Milan - Bicocca. Piazza dell'Ateneo Nuovo 1, 20126 Milan, Italy. E-mail: beapetrovich@gmail.com

§Corresponding author. Department of Economics, University of Milan - Bicocca. Piazza dell'Ateneo Nuovo 1, 20126 Milan, Italy. E-mail: luca.stanca@unimib.it

1 Introduction

In many social and economic situations communication among individuals is determined by social networks (Jackson, 2007, Udry and Conley, 2001). Individuals learn by observing the behavior of those they are connected with in their local environment. In this context, an important issue concerns how dispersed information held by different individuals can be aggregated over time and, in particular, how the communication structure determined by a social network affects this process. This paper focuses on persuasion bias, a form of bounded rationality whereby individuals are incapable of accounting for possible repetitions in the information they receive (DeMarzo et al., 2003). We investigate experimentally the hypothesis that, under persuasion bias, the structure of the communication network affects social influence. More specifically, an individual's influence on group opinions may depend on the number of outgoing communication links associated with his position in the social network.

It is possible to imagine several situations in which different individuals have noisy signals regarding an underlying state of the world. This state of the world may represent the quality of a new product, the returns on a potential investment, the ability of a political candidate to carry out reforms or, more generally, any unknown condition or action that affects the payoff of all individuals in the same way. Consider, for example, a situation where the unknown state of the world is the level of crime in one's neighborhood. Each individual has some information on this issue due to his own personal experience. By observing whether other individuals with whom he is in direct contact are installing burglar alarms or purchasing more sophisticated locks or even carrying guns, an individual may draw inference on the information observed by his direct neighbors. Over time, because of lack of common knowledge about actions of all individuals in the community, an individual can try to infer his neighbors' knowledge of his neighbors' actions and the private information they reveal. It is apparent that the complexity of the learning problem increases over time.

Whether individuals are capable of rationally processing the information circulating in their social network is an empirical question. Moreover, different network structures may influence the information aggregation process by determining the nature and complexity of the inference problem faced by individuals. In this context, recent studies have attempted to model deviations from rationality. In particular, DeMarzo et al. (2003) model persuasion bias as the outcome of a mechanical updating process, in accordance with which individuals fail to account for repetitions of information when communicating within a network. The main implication of the persuasion bias hypothesis is that, after iterated communication, the members of a network converge to a consensus that is biased towards the private signals of the most influential agents, namely those with the highest number of outgoing links.

In this paper we test experimentally whether the evolution of beliefs of individuals communicating through a social network reflects the structure of the network itself, and in particular, whether these beliefs are consistent with the persuasion bias hypothesis. Our results indicate that network structure plays a significant role in determining convergence beliefs. However, contrary to the predictions of the persuasion bias hypothesis, we find that the most influential agents are not those with more outgoing links, but those with more incoming links: consensus beliefs tend to be swayed towards the opinions of *influential listeners*. In order to explain this finding, we propose a generalized updating rule to describe social learning that takes into account agents' indegree (the number of individuals they listen to) in addition to their outdegree (the number of individuals they talk to). We show that this alternative updating rule provides a much better characterization of the experimental data.

Finally, we present an effort-weighted updating model as a more general framework for understanding information aggregation in social networks. In this framework, based on Ballester et al. (2006), agents optimally choose the effort exerted in processing information. Intuitively, when aggregating information is costly, individuals will choose an optimal effort level on the basis of their position in the social network. Since individuals with higher indegree are in a better position to aggregate information, they generate a positive information externality for their neighbors. Individuals with lower indegree will therefore devote less effort to processing information. Thus, in equilibrium, less weight will be attributed to the beliefs of those who exert less effort, and consensus beliefs will be swayed towards the opinions of individuals with higher indegree.

The paper is organized as follows. Section 2 briefly reviews the related literature. Section 3 describes the experimental design and procedures. Section 4 presents the theoretical predictions and hypotheses to be tested. Section 5 discusses the experimental results. Section 6 presents a simple theoretical framework of information aggregation within a social network. Section 7 concludes.

2 Related Literature

This work relates to the extensive theoretical literature on social learning. This literature can be generally divided into two main strands: one that focuses on Bayesian learning and the other that deals with myopic or boundedly rational learning.

The literature on Bayesian learning originates from the contributions of Bikhchandani et al. (1992) and Banerjee (1992), who assume an exogenous sequential structure in which each agent, after observing all past actions, optimally updates her belief on an unknown pay-off relevant state of the world and makes a single irreversible choice accordingly. Subsequent papers by Smith and Sorensen (1998), Banerjee and Fudenberg (2004), Celen and Kariv (2004) and Acemoglu et al. (2010) consider situations where individuals observe only a subset of past actions. These studies differ from the present one in two aspects: first, agents act sequentially and each individual has only one decision node; second, they focus on characterizing the asymptotic properties of different social networks under Bayesian learning. In particular, they study whether in the limit, as the size of the social network becomes arbitrarily large, individuals converge to payoff-maximizing actions. In this context, as the action space is discrete while the signal space is continuous, optimality of convergence is non trivial under Bayesian learning.

The setup introduced by Acemoglu et al. (2010) is closer to our contribution, as it provides a representation of learning in social networks. In this framework, it is assumed that agents know the identity of the individuals whose information they observe. This is in contrast to Banerjee and Fudenberg (2004) and Smith and Sorensen (1998) where individuals observe a representative sample of the overall population without knowing the identity of those whose actions they observe. However, in Acemoglu et al. (2010) the network structure has a slightly different role compared with the present work, in that it simply determines the set of past actions observed once by the agents and therefore does not imply an ongoing interaction between network members. In this respect, our contribution is more closely related to Gale and Kariv (2003), who study the convergence and optimality of learning when the network structure is fixed and individuals repeatedly take simultaneous actions observing the past actions of those to which they are connected in the network.

Among the papers belonging to the non-Bayesian learning branch of the literature, the most closely related to ours are Bala and Goyal (1998, 2001), DeMarzo et al. (2003) and Golub and Jackson (2010). These papers study social learning in connected social networks. DeMarzo et al. (2003) and Golub and Jackson (2010) are particularly relevant for our work as they focus on the properties of consensus beliefs, in settings where individuals converge to the same opinions. These studies provide a characterization of social influence and analyse the likelihood that consensus beliefs will lead to optimal aggregation of information. In DeMarzo et al. (2003), where actions and signal spaces coincide, Bayesian updating would imply optimal information aggregation independently of the structure of the social network. If individuals are incapable of properly accounting for repetitions of information, and are therefore subject to persuasion bias, convergence beliefs will depend on the network structure. More specifically, under certain conditions, beliefs converge to a consensus given by a weighted average of agents' initial beliefs, where the weights represent individuals' social influence and depend on the number of outgoing links associated with the position of each agent. This setting provides an optimal testing ground to evaluate how social networks influence belief formation and is therefore the framework for our experiment.[1]

At the empirical level, starting from the seminal contribution of Bavelas (1950), several experiments have been conducted in social sciences to study how different network structures influence agents' learning and the information aggregation process. The economic experimental literature on learning in social networks is more recent and deals mainly with an experimental setup where three subjects who are connected by a given network are called upon to repeatedly choose between two binary actions.[2] In this paper, we adopt a continuous signals-continuous actions setup that differs from the ball-and-urn standard learning experiment, initially proposed by Anderson and Holt (1997) and applied by Choi et al. (2005, 2009) to the social network learning setting.[3] The motivation for our choice is that allowing actions to be continuous provides a richer context for analyzing how different social networks may affect the convergence of beliefs. When actions are binary, only switching behavior can be used to infer a change in a subject's beliefs, but it is impossible to observe whether individual beliefs tend to be swayed in the direction of more influential individuals. Continuous actions, instead, allow us to measure the direction and intensity of persuasion bias.

---

[1]The present work is also related to the mathematical sociology literature on social networks (DeGroot, 1974, Bonacich, 1987, Bonacich and Lloyd, 2001). Sociological studies on social networks have proposed several indexes of "power", "centrality", and "status" that are close to the notion of social influence discussed here, as they depend on the structure of the social network.

[2]Within this context, Choi et al. (2009) illustrate how the quantal response equilibrium (QRE) approach outperforms the pure Bayesian updating model. Choi et al. (2005) find that the experimental data exhibit significant differences in individual and group behavior among different network structures. Choi (2006) estimates a cognitive hierarchy model, assuming that individuals of heterogenous cognitive types make rational decisions considering the equilibrium cognitive types of the other individuals in the network.

[3]Celen and Kariv (2004) adopt a similar, although not identical, framework by implementing an experimental task with continuous signals, while still assuming binary actions.

3 The Experiment

The experiment is designed to test whether individuals communicating through a social network are subject to persuasion bias and, in particular, whether social influence ultimately reflects the structure of the social network. We consider a repeated learning problem, adapted from DeMarzo et al. (2003), where communication among individuals occurs within a social network. We implement two treatments by exogenously manipulating the structure of the network, in order to compare social learning in a balanced and an unbalanced network.[4] In this section, we describe the experimental task, treatments and procedures. The next section presents the theoretical predictions and hypotheses to be tested.

3.1 Baseline game

The experimental task involves four individuals interacting over 12 rounds). Each individual is assigned a letter (A, B, C and D) identifying his position in the social network. The position and identity of the group components remain unchanged and anonymous throughout the task. Detailed instructions distributed to the subjects are reported in Appendix 1.

At the beginning of the task, each subject is assigned an integer number (signal) randomly generated by the computer.[5] Subjects are only informed that the signals of the four group members are randomly drawn from a given (unknown) distribution. The aim of the task is to guess, in each of the 12 rounds, the mean of the signals received by the four components of the group at the beginning of the task.[6] Earnings depend on the accuracy of the guess on the basis of a triangular scheme: a subject's payoff in each round is equal to 20 euros minus the absolute value of the difference between his guess and the average of the four signals.[7] Actual earnings are determined on the basis

[4] A network is balanced if the total self-importance of those who listen to each agent is equal to one.

[5] Signals are generated as follows. Each group of four subjects is assigned an integer number drawn from a uniform distribution between 100 and 9999. Signals observed by group members are then drawn from a normal distribution with mean equal to the group-specific randomly selected number and standard deviation equal to 50.

[6] In DeMarzo et al. (2003) the signal is obtained by adding a normal random disturbance to an unknown parameter that has to be guessed. In our experiment, instead, each of the four group components receive a randomly drawn number, and the unknown parameter is the average of the four numbers. This substantially simplifies the experimental task, while leaving the theoretical properties unaffected.

[7] For example, if the mean of the four signals is 803.25 and the subject's choice in the selected round is 792, the absolute value of the difference is 11.25 and then the monetary reward is equal to 8.75 euro. If the difference is greater than 20, the payoff is 0; for

of one round randomly selected at the end of the experiment. There are no show-up fees.

From the second round onwards, each subject is informed by the computer of the choices made in previous rounds by the subjects he is connected to according to the network structure. In each round every subject thus observes the previous choices made by his "neighbors". The network structure is illustrated in the instructions and each subject's screen displays it during the task.

3.2 Treatments

The experiment is based on two treatments, T1 and T2, implemented in a between-subjects design. The two treatments differ with respect to the structure of the communication network, while keeping constant the set of signals received by the subjects.

In T1, the control treatment, the communication structure is determined by the circle network, represented in Figure 1. This is a strongly connected and balanced network, where each agent has one incoming link and one outgoing link.[8] It reproduces a situation in which each subject listens to one neighbor and talks to another neighbor.

Figure 1: Network structure in treatment 1

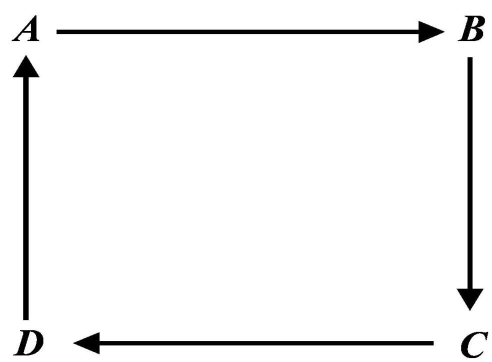

Treatment T2 is obtained by adding two links to the circle network, so that the choices made by subject A are observed by every other group member, while the choices of all other subjects are observed by only one subject (Figure 2). In this strongly connected and unbalanced network, subject A

instance, if the mean of the signals is 62.5 and the relevant guess is 30.5, the difference is 32 and the subject does not earn anything, as $20 - 32 < 0$.

[8]A network is connected if, for any two agents i and j, there is a sequence $i_1, \ldots, i_K$ such that $i_1 = i$, $i_K = j$, and i_k is connected to i_{k+1} for $k = 1, ..., K-1$.

has three outgoing links and one incoming link, B has one outgoing link and one incoming link, D and C have two incoming links and one outgoing link. This network reproduces a situation in which a *central* subject is listened to by everybody.

Figure 2: Network structure in treatment 2

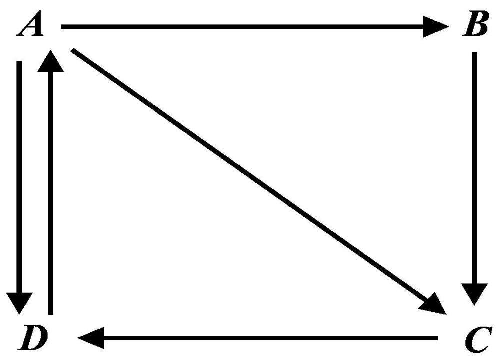

3.3 Procedures

We ran two sessions for T1 and four sessions for T2. In each session, 24 subjects were randomly divided into 6 groups of four subjects. The position and identity of the four group components remained unchanged throughout the session. Overall, the experiment involved 144 subjects, mainly undergraduate students of Economics recruited by email through an online system. The experiment was conducted at EELAB (University of Milan Bicocca) in November 2009 using z-Tree (Fischbacher, 2007). Each session lasted approximately 80 minutes, including instructions, control questions and payments. Average earnings were 14.5 euros.

Each session consisted of three phases of 12 rounds, for a total of 36 rounds in each session. In each phase, the experimental task was implemented with a new set of signals. The task was repeated three times in order to make it familiar to subjects, so that noise due to task misunderstanding was reduced. At the end of each phase, subjects were informed of the four signals, their mean, and the choices made by each subject. By choosing 12 communication rounds, we created a situation in which beliefs converge under both persuasion bias and rationality. If all agents were rational, four rounds would be sufficient for convergence. If all subjects followed a persuasion bias updating rule, after 12 rounds beliefs would be virtually identical.

The task was explained at the beginning of each session: instructions were read aloud and any questions about the game were answered individually.

Before the game started, every subject was asked written control questions in order to check if the task was fully understood. The instructions explicitly explained to participants that the best choice when knowing or being able to deduct a given number of signals is the average of these signals. The reason for this is that we wanted to ensure that subjects knew how to optimally aggregate information if all the signals were public information, thus being able to focus exclusively on the process of social learning.

In order to prevent possible mistakes caused by the fact that individuals might forget their guesses from previous rounds or other relevant pieces of information, each participant's monitor displayed the subject's past choices and the pieces of information received in previous rounds. In this way, we induced a game of perfect recall and controlled for memory effects on decision making. In order to minimize computation mistakes, we also provided subjects with a calculator on the computer screen.

4 Theoretical Predictions and Hypotheses

Let $y_{i,t}$ denote the guess of individual i in round t, and $\mathbf{y}_t$ the vector of guesses of all individuals within a group in round t. Let x_i denote individual signals, $\overline{x}$ the average of the four signals within a group and $\mathbf{x}$ the vector of signals. The structure of the network is represented as a directed graph, where $S_i \subset N$ denotes the set of agents who agent i listens to. The listening sets for T1 are $S_A = \{A, D\}, S_B = \{A, B\}, S_C = \{B, C\}, S_D = \{C, D\}$, while the listening sets for T2 are $S_A = \{A, D\}, S_B = \{A, B\}, S_C = \{A, B, C\}, S_D = \{A, C, D\}$. Denote with $q_{ij} \in \{0, 1\}$ an indicator function corresponding to S_i, where $q_{ij} = 1$, if agent i listens to agent j, and $q_{ij} = 0$ if there is no incoming communication link from j to i (note that $q_{ii} = 1$ since each agent listens to himself).

Communication occurs over repeated rounds. Since all agents have the same preferences, there is no scope for strategic communication and agents truthfully reveal their beliefs.[9] Although agents do not have information about the underlying distribution of the signals, a backward induction argument implies that, under both persuasion bias and Bayesian updating, agents have an incentive to truthfully report their signal in the first round: $\mathbf{y}_1 = \mathbf{x}$. Thereafter, in each round agents listen to the guesses of those in their listening set, update their beliefs and make a new guess. Whenever the communication process leads all individuals to converge to the same beliefs, these are defined consensus beliefs and denoted with $\mathbf{wx}$, where $\mathbf{w}$ represents the vector of weights attributed to the signal of each agent in the network.

[9] In what follows we therefore refer to beliefs, guesses, and actions as synonyms.

These weights can be interpreted as representing agents' social influence.

If all agents are rational, Bayesian updating allows each individual to extract all the private information in the network, so that consensus beliefs are efficient. In both T1 and T2, as signals are equally informative, all agents within a group converge to the same belief and in equilibrium each of the four group members is equally influential. Regardless of the network structure, consensus beliefs coincide with the arithmetic mean of the signals, $\mathbf{w}_e\mathbf{x}$, with $\mathbf{w}_e = [0.25, 0.25, 0.25, 0.25]$.

If individuals are subject to persuasion bias, as in DeMarzo et al. (2003), they treat the information received in each round as new and independent, without rationally discounting the fact that only a part of it is new, while the remaining part has already been communicated in previous rounds. The evolution of beliefs can be described by the updating rule

$$\mathbf{y}_t = L\mathbf{y}_{t-1}$$

where L is a matrix with elements $\ell_{ij} = (q_{ij}\pi_{ij}) / \left(\sum_j q_{ij}\pi_{ij}\right)$ and π_{ij} denotes the precision that agent i assigns to agent j's belief.[10] Assuming that each agent believes that all the agents he listens to have equal precision, the updating matrix in T1 is:

$$L = \begin{bmatrix} 1/2 & 1/2 & 0 & 0 \\ 0 & 1/2 & 1/2 & 0 \\ 0 & 0 & 1/2 & 1/2 \\ 1/2 & 0 & 0 & 1/2 \end{bmatrix}$$

while in T2 it is:

$$L = \begin{bmatrix} 1/2 & 0 & 0 & 1/2 \\ 1/2 & 1/2 & 0 & 0 \\ 1/3 & 1/3 & 1/3 & 0 \\ 1/3 & 0 & 1/3 & 1/3 \end{bmatrix}$$

[10] De Marzo et al (2003) allow for individuals to vary over time the weight they assign to their own beliefs relative to those of the others they are connected to. Thus we have

$$L_t = (1 - \lambda_t)I + \lambda_t L$$

where L is the updating matrix and $\lambda_t \in (0, 1)$. Values of λ_t closer to zero imply that individuals have more persistent opinions, while values closer to 1 imply that individuals place equal weights on the beliefs of all those they are connected to (including themselves). As long as agents are not too fixed in their beliefs and the network is strongly connected, under persuasion bias the beliefs of individuals should converge. Different values of λ_t can affect the speed of converge but not the convergence itself as well as the consensus beliefs, as long as $\sum_{t=1}^{\infty} \lambda_t = \infty$. Since we focus on convergence beliefs, we set $\lambda_t = 1$ in every period t, without loss of generality.

Beliefs in round t can therefore be written as $\mathbf{y}_t = L^{t-1}\mathbf{x}$. As we consider strongly connected networks, where no agent is isolated, beliefs will converge over rounds. Denoting with $\mathbf{w}_p$ the vector of weights characterizing consensus beliefs under persuasion bias, we obtain $\lim_{t\to\infty} y_t = \mathbf{w}_p\mathbf{x}$.

In T1, a balanced and strongly connected network, consensus beliefs under persuasion bias will be the same as under rationality. In T2, an unbalanced and strongly connected network, consensus beliefs under persuasion bias will instead differ from rational consensus beliefs. Given that agents fail to account for repetitions of information, more connected agents are more influential in equilibrium.[11] For the network structure in T2, consensus beliefs are given by $\mathbf{w}_p\mathbf{x}$ with $\mathbf{w}_p \simeq [0.42, 0.10, 0.16, 0.32]$. That is, consensus beliefs assign excessive weight to A's signal, as A is listened to by three subjects, whereas B, C and D communicate with one subject only. Moreover, D is also relatively more influent, as he communicates directly with A (indirect social influence). C and, to a greater extent, B should be less influential in T2, relative to T1, as they have only one outgoing link and do not communicate directly with A.

Summing up, the hypotheses to be tested can be stated as follows:

$$H_0 : w_A^{T1} = w_A^{T2} = 0.25 \quad \text{vs} \quad H_1\!: \ w_A^{T1} = 0.25 < w_A^{T2} = 0.42 \tag{1}$$

$$H_0 : w_B^{T1} = w_B^{T2} = 0.25 \quad \text{vs} \quad H_1\!: \ w_B^{T1} = 0.25 > w_B^{T2} = 0.10 \tag{2}$$

$$H_0 : w_C^{T1} = w_C^{T2} = 0.25 \quad \text{vs} \quad H_1\!: \ w_C^{T1} = 0.25 > w_C^{T2} = 0.16 \tag{3}$$

$$H_0 : w_D^{T1} = w_D^{T2} = 0.25 \quad \text{vs} \quad H_1\!: \ w_D^{T1} = 0.25 < w_D^{T2} = 0.32 \tag{4}$$

Under the null hypothesis of rationality, each agent should be equally influential in both T1 and T2. Under the alternative hypothesis of persuasion bias, social influence weights should be equal in T1 but different in T2. Subjects A and, to a lesser extent, D should be more influential in T2 than in T1, while subjects B and, to a lesser extent, C should be less influential in T2 than in T1.

5 Results

This section presents the results. We start with a descriptive analysis of the main features of the experimental data. We then present formal tests of the

[11] More precisely, an agent is more influential if he is listened to by many other agents (direct social influence) and if the agents who listen to him are themselves influential (indirect social influence).

hypothesis of persuasion bias by comparing across treatments the parameters that characterize social influence. Finally, we examine by simulation an alternative updating rule that takes into account both outdegree and indegree, showing that it provides a much better characterization of the experimental data.

5.1 Overview

In each of the 6 sessions, 6 groups of four subjects implement the experimental task three times, once for each phase, with 12 rounds for each phase. We therefore have a total of 5184 observations at individual level, with 432 observations for each round.

The behavior of individual beliefs over successive rounds indicates substantial heterogeneity at both subject- and group-level. Overall, in the first round, 92.4 per cent of the subjects report their own signal, while 96.3 per cent make a guess within 20 units from their own signal. Focusing on final beliefs, in the last round 28.7 per cent of the subjects correctly guess the average of the four signals within their group. Accounting for rounding errors, 37.7 per cent of the subjects make mistakes smaller than one unit. It should be observed that the share of correct beliefs in the final round is higher in T2 than in T1, reflecting the fact that, although the network structure in T2 is not symmetric, it provides more information due to the higher number of communication links.

The heterogeneity of opinions among subjects belonging to the same group, represented by the average variance of the beliefs held by the four group components, falls steadily over successive rounds in both treatments. While the average group-level variance of beliefs is initially higher in T1 than in T2, beliefs display a relatively larger variance in T2 in the final rounds.

5.2 Network Structure and Social Influence

We now turn to analyzing how the structure of the network affects social influence and, in particular, whether the differences in social influence between the two network structures are in the direction implied by the persuasion bias hypothesis. We study social influence by focusing on individual beliefs in the last round. We estimate the weights of individual signals in final beliefs, assuming that the observed final beliefs are a linear combination of private signals plus a random error term:

$$y_{i,T} = w_A x_i^A + w_B x_i^B + w_C x_i^C + w_D x_i^D + \varepsilon_{i,T} \tag{5}$$

where x_i^k is the private signal received by subject k in the group individual i is in, with $k \in (A, B, C, D)$; $y_{i,T}$ is the belief held by individual i in the last round (T) and $\varepsilon_{i,T}$ is an idiosyncratic error term.

The coefficient associated to a given signal measures the social influence of the corresponding subject. We control for the effect of outlying observations by eliminating from the sample the 10 per cent most extreme observations in either round 1 (misreported signals) or round 12 (divergent beliefs). Table 1 presents OLS estimation results. Confidence intervals at the 95% level are reported in square brackets in order to assess the statistical significance of the theoretical predictions within treatments.

Table 1: Estimates of social influence weights, by treatment

	T1	T2	T1 - T2
Signal A	0.24	0.25	0.01
	(9.90)	(14.29)	(0.27)
	[0.19,0.29]	[0.22,0.29]	[-0.05,0.06]
Signal B	0.26	0.20	-0.06
	(10.92)	(14.13)	(-2.29)
	[0.21,0.30]	[0.17,0.22]	[-0.11,-0.01]
Signal C	0.26	0.25	-0.02
	(12.45)	(21.24)	(-0.75)
	[0.22,0.31]	[0.22,0.27]	[-0.06,0.03]
Signal D	0.24	0.31	0.06
	(10.95)	(23.52)	(2.62)
	[0.20,0.29]	[0.28,0.33]	[0.02,0.11]
Number of observations	124	240	364

Note: the figures in the first two columns are the estimates of the social influence weights associated to the subject in the position indicated by the row heading, in T1 and T2, respectively. The third column reports differences in social influence weights across treatments. Dependent variable: individual subjects' guess at final round. t-statistics reported in round brackets, 95 per cent confidence intervals reported in square brackets.

In T1 (column 1), where the network structure is balanced, estimated coefficients for all network positions are similar and close to 0.25. Coefficient are jointly and individually not significantly different from 0.25. This indicates that, under the balanced network structure, final beliefs evenly reflect the private information held by the four group components. Estimated social influence weights are instead relatively different in T2 (column 2), where the network structure is unbalanced. As predicted by the persuasion bias hypothesis, B is the least influential subject, with a weight (0.20) that is significantly lower than 0.25, as indicated by the confidence interval. On the other hand, in contrast with the persuasion bias hypothesis, A is not the most influent subject. His estimated weight is indeed 0.25, virtually unchanged with respect to T1. The social influence weight of C is also 0.25, almost unchanged in T2 relatively to T1. Interestingly, the most influential subject in T2 is D, whose estimated social influence weight is 0.31, significantly higher than 0.25.

Between treatments, the null hypothesis of jointly equal coefficients in T1 and T2 can be strongly rejected (p-value 0.006). Focusing on individual coefficients, the difference between treatments, reported in column 3, is significant at the 5 per cent level for subject B and at the 1 per cent level for subject D. It should be observed that, because of repetition over rounds and phases, final round individual-level observations within the same group of subjects may not be independent. The four subjects in each group interact repeatedly over 12 rounds. In addition, the same group of four subjects performs the experimental task three times, with a different set of signals. In order to allow for the possible dependence of final-round beliefs within and across phases, we also considered test statistics and confidence intervals based on standard errors clustered at group and phase level, thus assuming 36 independent observations (6 groups for each of the 6 sessions). All the results reported above for the analysis within treatments are qualitatively unaffected. Between treatments, the change for subject B is not statistically significant, due to the larger standard errors. However, the null hypothesis of no change for the weight of D can be rejected at the 0.04 level when accounting for dependence across subjects within phases and at the 0.07 significance level when also accounting for dependence across phases.

In order to further assess the robustness of the results, we estimated social influence weights in T2 by network position, thus taking into account the possible non-convergence of beliefs. The results, reported in Table 2, indicate that the pattern described above is qualitatively robust. In particular, B and D are the least and the most socially influent subjects, respectively, for each of the four subjects in the different network positions. Subject A is not influent for any of the other three subjects in the social network.

Estimates of social influence are also qualitatively unchanged when considering each phase separately: weights range between 0.23 and 0.27 for A, 0.17 and 0.21 for B, 0.22 and 0.30 for C and between 0.25 and 0.33 for D, respectively. Finally, we also checked the robustness of the results with respect to the use of alternative identification criteria for outliers (eliminating either 5, 1 or 0 per cent most extreme observations). Although standard errors do vary across specifications, the size of the estimated coefficients is virtually unchanged. In particular, social influence weights range between 0.24 and 0.27 for subject A and between 0.29 and 0.32 for subject D.

Table 2: Social influence weights in T2, by netrole

	Subject A	Subject B	Subject C	Subject D
Signal A	0.30	0.24	0.22	0.23
	(7.77)	(6.93)	(7.13)	(6.71)
	[0.22,0.38]	[0.17,0.32]	[0.16,0.29]	[0.16,0.30]
Signal B	0.16	0.21	0.22	0.19
	(5.22)	(7.58)	(8.96)	(6.94)
	[0.10,0.22]	[0.16,0.27]	[0.17,0.27]	[0.14,0.25]
Signal C	0.24	0.21	0.28	0.26
	(9.18)	(8.93)	(13.54)	(11.36)
	[0.18,0.29]	[0.16,0.25]	[0.24,0.32]	[0.22,0.31]
Signal D	0.30	0.33	0.28	0.31
	(10.38)	(12.76)	(11.89)	(12.22)
	[0.24,0.36]	[0.28,0.38]	[0.23,0.32]	[0.26,0.37]
N. obs.	60	60	60	60

Note: the figures reported are the estimates of the social influence weights associated to the subject in the position indicated by the row heading for the subject in the position indicated by the row heading. Dependent variable: belief in final round of the subject indicated in the column heading. Standard errors reported in round brackets, 95 per cent confidence intervals reported in square brackets.

Summing up, the comparison of social influence weights across treatments indicates that, contrary to the predictions of the persuasion bias hypothesis, the social influence of A, the agent whose outdegree is exogenously increased in T2, is not higher in T2 relative to T1. Quite surprisingly, the social influence of D is instead significantly higher in T2. It should be noted that the latter result cannot be explained by indirect social influence, as A is not the most influent subject in T2. In the following subsection, we consider an alternative updating rule that may explain these findings.

5.3 A Generalized Updating Rule

The experimental analysis indicates that, relative to a balanced network structure, increasing the number of outgoing links of subject A does not lead to a higher social influence. On the other hand, subject D, who communicates to less agents but listens to more agents than A, becomes significantly more influential. One possible interpretation of this result is that, under a boundedly rational updating rule, social influence within a social network may depend not only on the number of subjects one talks to (outdegree), but also on the number of subjects one listens to (indegree). The updating rule proposed in the persuasion bias model does not take into account the fact that some individuals receive more information than others. Beliefs of different agents are given equal weights, regardless of the number of agents they, in turn, listen to. In real situations, instead, it is likely that agents take into account how informed their neighbors are when updating their own opinion.

We suggest that the persuasion bias updating rule should be considered as an extreme case, whereby only outdegree matters, of a more general updating rule that also takes into account agents' indegree. We thus propose an updating rule based on a more general weighted updating matrix, whose elements are defined as follows:

$$\ell_{ij}(\mathbf{d}, \boldsymbol{\rho}) = \frac{q_{ij} d_j^{\rho}}{\sum_j q_{ij} d_j^{\rho}} \tag{6}$$

where d_j is agent j's indegree, ρ is a parameter between 0 and ∞ and, as above, q_{ij} is equal to 1 if j belongs to i's listening set, and 0 otherwise.[12] We denote the weighted updating matrix with $L(\mathbf{d}, \boldsymbol{\rho}) = [\ell_{ij}(\mathbf{d}, \boldsymbol{\rho})]$, where $\mathbf{d}$ denotes the vector of agents' indegree, so that beliefs after t rounds of updating are given by $\mathbf{y}_t = L(\mathbf{d}, \boldsymbol{\rho})^{t-1}\mathbf{x}$.

It should be observed that this heuristic, as the one in DeMarzo et al. (2003), assumes that agents are incapable of properly accounting for repetitions of information. However, through the parameter ρ, it allows us to model in a continuous way the relative importance of indegree and outdegree. When ρ equals 0, we are in the case of persuasion bias: agents update their beliefs placing equal weights on the beliefs of those they listen to (including themselves), regardless of the number of their incoming links. When $\rho = 1$, agents update their beliefs aggregating the beliefs of those they listen to by using weights that are proportional to their indegree. When ρ tends to ∞, agents update their beliefs using only the beliefs of the most informed agent, i.e. the opinion(s) held by the agent(s) with the highest indegree.

In order to assess the explanatory power of this alternative heuristic, we simulate belief dynamics over rounds for the same sets of signals used in our experiment, for $\rho = 0$ (persuasion bias) and $\rho = 1$ (weights proportional to indegree). The results of the simulations, compared with the experimental results, are presented in Table 2 and displayed in Figure 3. The updating rule that proportionally takes indegree into account clearly outperforms the

[12] It is worth noticing that the number of incoming nodes of an agent's interlocutor are part of the agent's information set only from the second round onwards, since in the first round agents make their guesses on the basis of their own signal only. For the sake of simplicity our heuristic does not make this distinction, and assumes that agents take into account incoming nodes of those they are connected to in each communication round.

persuasion bias heuristic. When $\rho = 1$, not only is the pattern of social influence weights correctly predicted in the final round, but also the transition dynamics are remarkably similar to the ones observed in the experimental data.[13] Interestingly, while the restrictions implied by the persuasion bias hypothesis for the social influence weights of subjects A and D in T2 are rejected by an F-test (p-value 0.00), those implied by the alternative heuristic with weights proportional to indegree cannot be rejected (p-value 0.12).

Figure 3: Social influence in T2, estimated and simulated (ρ=1 and ρ =0)

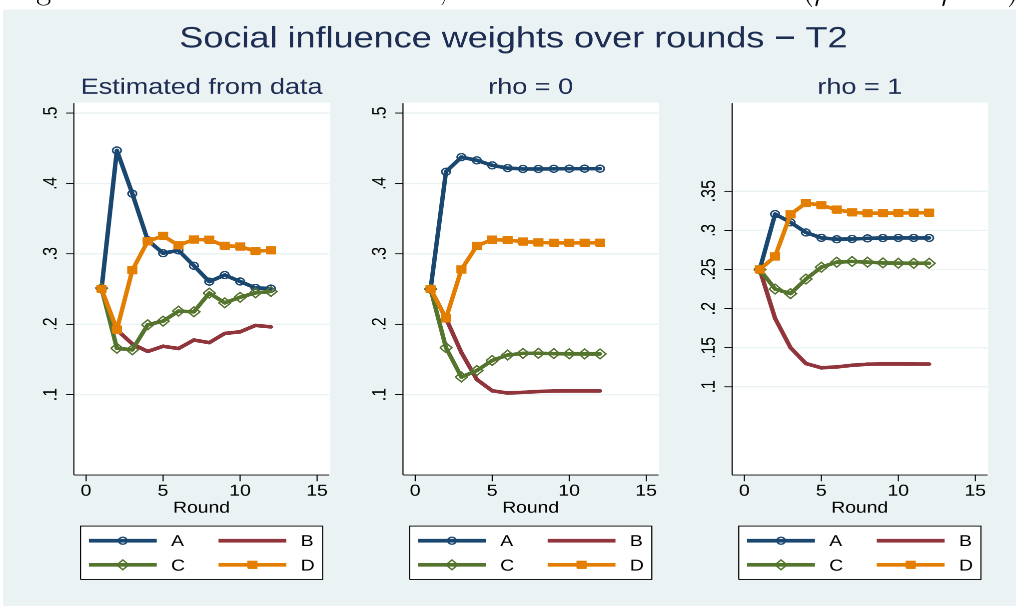

Next, rather than restricting the attention to the cases $\rho = 0$ and $\rho = 1$, we simulate the proposed updating rule for all possible values of ρ, in order to identify the value of ρ that maximizes explanatory power for the experimental data. We thus define the optimal ρ as the value that minimizes the sum of squared deviations, over all individuals, between observed (experimental) and

[13]Note that, as mentioned above, we simulate a simplified version of the persuasion bias model of DeMarzo et al (2003), where we assume that agents place equal weights on the beliefs of all those they listen to, which is equivalent to setting $\lambda_t = 1$ for every t in the original version of the model. While this assumption does not affect convergence beliefs (as long as $\sum_{t=1}^{\infty} \lambda_t = \infty$), it may affect the transition dynamics.

Table 3: Social influence weights over rounds in treatment 2

	Estimated				Simulated ($\rho = 0$)				Simulated ($\rho = 1$)			
Round	A	B	C	D	A	B	C	D	A	B	C	D
1	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25
2	0.45	0.19	0.17	0.19	0.42	0.21	0.17	0.21	0.32	0.19	0.23	0.27
3	0.39	0.17	0.16	0.28	0.44	0.16	0.13	0.28	0.31	0.15	0.22	0.32
4	0.32	0.16	0.20	0.32	0.43	0.12	0.13	0.31	0.30	0.13	0.24	0.34
5	0.30	0.17	0.20	0.33	0.43	0.11	0.15	0.32	0.29	0.12	0.25	0.33
6	0.30	0.17	0.22	0.31	0.42	0.10	0.16	0.32	0.29	0.13	0.26	0.33
7	0.28	0.18	0.22	0.32	0.42	0.10	0.16	0.32	0.29	0.13	0.26	0.32
8	0.26	0.17	0.24	0.32	0.42	0.10	0.16	0.32	0.29	0.13	0.26	0.32
9	0.27	0.19	0.23	0.31	0.42	0.11	0.16	0.32	0.29	0.13	0.26	0.32
10	0.26	0.19	0.24	0.31	0.42	0.11	0.16	0.32	0.29	0.13	0.26	0.32
11	0.25	0.20	0.24	0.30	0.42	0.11	0.16	0.32	0.29	0.13	0.26	0.32
12	0.25	0.20	0.25	0.31	0.42	0.11	0.16	0.32	0.29	0.13	0.26	0.32

Note: This table compares social influence weights for the subject indicated by the column heading in each round in three cases: estimation from experimental data (columns 1-4), simulation of the updating rule with $\rho = 0$ (persuasion bias, $\lambda_t = 1$), simulation of the updating rule with $\rho = 1$ (weights proportional to indegree).

simulated final beliefs:

$$\begin{aligned} \rho^* &= \arg\min(E) \\ E &= \sum_{g=1}^{72} \sum_{k=1}^{4} \left(y_{g,k,12} - \widehat{y}^{\rho}_{g,k,12}\right)^2 \end{aligned}$$

where $y_{g,k,12}$ are the observed last round beliefs of agent k in group g and $\widehat{y}^{\rho}_{g,k,12}$ are the corresponding simulated last-round beliefs using the updating rule in (6) for a given value of ρ. Figure 4 presents the results of the simulations for ρ between 0 and 2. As clearly shown, the value of ρ that maximizes explanatory power for the experimental data is very close to one (1.03). That is, agents' behavior in the experiment is consistent with an updating rule that takes into account both outdegree and indegree, and the latter is weighted proportionally.

Figure 4: Explanatory power for experimental data as a function of ρ

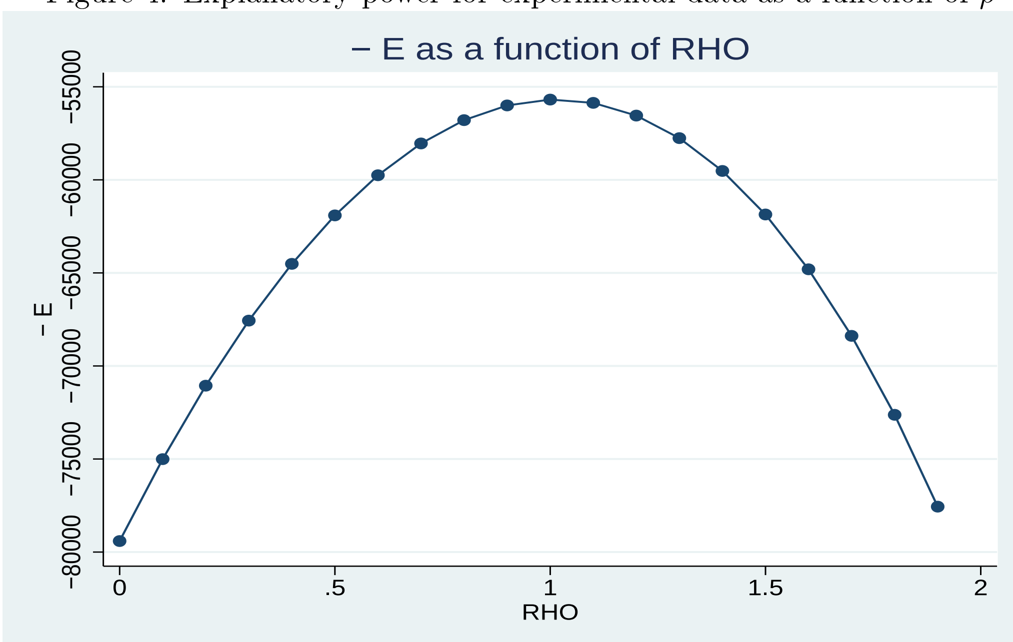

6 A Framework for Effort-Weighted Influence

In this section we propose a framework that is consistent with the generalized updating rule described above, and provides further insights on the individual behavior underlying our experimental results. The model we develop is an application of Ballester et al. (2006), that relates the choices of agents to the effort levels that the network members devote to processing information.

As shown in the previous section, when formulating their beliefs agents take into account how well informed are those they listen to. Agents may perceive a different relative importance of their private information based on their position in the network, and this may influence the effort they devote to combining their own beliefs on the true state of the world with the beliefs of those they listen to. More specifically, agents may consider their own effort in aggregating information and that of the others as strategic substitutes or complements. Since these cross effects (or externalities) of effort may differ based on the network structure and are unobservable, the model also enables us to identify the pattern of cross effects that better explain our experimental results. More formally, our claim is that the network structure may endogenously affect utility functions of agents by influencing the cross effects of effort. Altering the network structure therefore influences the way in which information is aggregated and, ultimately, consensus beliefs and social influence.

As an example, consider the unbalanced network structure of treatment 2. Our conjecture is that since A knows that D has more incoming links than him, he also knows that D is in a better position to aggregate information. Therefore, A's effort is a strategic substitute of D's. Moreover, since those who listen to A know that he listened to D, they will consider A's position for aggregating information to be more important than their own, and will therefore see their own effort as a strategic substitute for A's.

6.1 Model

As in DeMarzo et al. (2003), we assume that individuals are incapable of correctly discounting for repetitions of information, and adopt a Markovian updating rule. However, the weight assigned by subjects to the information received from neighbors depends on the level of effort that each agent devotes to processing information. In this case, effort in processing information is related to how agents aggregate the incoming stream of beliefs before making their own choice, and therefore communicating with their neighbors. The agents who devote more effort to processing information are those who are believed to more precisely aggregate the stream of incoming messages. Since agents are incapable of discounting repetitions by assumption, the best they can do is to correctly average the stream of beliefs they receive. Those who exert less effort, adopt a rule of thumb that is less informative on the true state of the world and, in equilibrium, a smaller weight will be attributed to their beliefs.

Effort levels chosen by individuals therefore determine the evolution of beliefs, but exerting effort involves a cost that can be interpreted as a cognitive cost. We assume that in each period, each agent maximizes his short-run utility, where the utility function of agent i in each period t can be written in the following way:

$$u_i(y_{i,t}, \theta, e_{i,t}, e_{-i,t}) = -[(y_{i,t} - \theta)^2] - c_i(e_{i,t}, e_{-i,t}) \tag{7}$$

where $c_i(e_{i,t}, e_{-i,t})$ represents the cost function of effort, that may differ between agents but is invariant over time and depends on own effort, $e_{i,t}$, and on the effort of the others, $e_{-i,t}$, in period t. In every communication round, each agent receives a stream of beliefs, one from each network interlocutor, and exerts a certain amount of effort to aggregate this information. Each belief is considered to be more or less informative on the signal received depending on the effort devoted to information processing by the agent that makes the guess in the previous round. We assume that effort levels may differ between agents and may also vary over time. Therefore, agents use the effort-weighted average stream of signals to assess the true value of θ:

$$E(\theta \mid \{y_{j,t-1}, e_{j,t-1}\}_{q_{ij}=1}) = E(\theta \mid y_{i,t}(\mathbf{e}_{t-1})), \text{where } y_{i,t}(\mathbf{e}_{t-1}) = \sum_j \frac{q_{ij} e_{j,t-1} y_{j,t-1}}{\sum_j q_{ij} e_{j,t-1}} \tag{8}$$

where $e_{j,t-1}$ represents the effort devoted to information processing by agent j in period $t-1$, and $\mathbf{e}_{t-1}$ is the corresponding vector of effort levels of all agents. We assume that processing information, by weighing incoming beliefs based on past effort levels, is costly. The effort devoted to processing this information determines the accuracy of each agent's belief in a given period. We denote P as the communication process that depends on each agent's choice of effort in a given period. The effective belief of agent i in round t, $y^P_{i,t}$ represents a signal on the effort-weighted stream of signals, for those who listen to agent i, according to the structure of the network:

$$y^P_{i,t} = y_{i,t}(\mathbf{e}_{t-1}) + \eta^P_{i,t}$$

We denote $\eta^P_{i,t}$ as the random deviation from the efficient belief that depends on the effort exerted by agent i in period t. The communication process depends on effort through $\eta^P_{i,t}$, normally distributed with zero mean and variance equal to a function of the effort level, so that $\sigma^2_{i,t} \equiv f(e_{i,t})$, where $f(e_{i,t})$ is a non-negative decreasing function of the agent's effort level in period t.

Ignoring the costs of effort, agent i's expected loss of utility in each period is therefore a function of effort:

$$-E[(y^P_{i,t} - \theta)^2 \mid \theta, e_{i,t}] = -f(e_{i,t}) \tag{9}$$

Notice that since the cost of effort is invariant over time, we can drop the time subscript from the expected utility, so that the expected utility of agent i including costs can be expressed as a function of effort:

$$E\left[u_i(e_{i,}e_{-i})\right] = -f(e_i) - c_i(e_i, e_{-i}) \tag{10}$$

From (9) it is apparent that own effort enters positively in the agent's utility because, independently of what the others do, passing on precise information improves information aggregation. Since all agents' expected payoffs are increasing in the accuracy of the aggregation process, everyone derives positive utility from exerting a positive amount of effort as this increases the quality of the information in the network.

In period 1, there are no costs of processing information since each agent observes his own signal only, and will therefore maximize utility by exerting the maximum amount of effort, implying that $\mathbf{y}_1 = \mathbf{x}$.

From the second round onwards, we assume that the cost of effort is represented by the following function:

$$c_i(e_i, e_{-i}) = \frac{1}{2}\phi_{ii}e_i^2 + \sum_{j \neq i} \phi_{ij}e_ie_j \tag{11}$$

the first term implies that each agent bears the same convex cost of effort where $\phi > 0$; e_i and e_j denote the effort devoted by individuals i and j respectively to information processing, and λ represents the weight of the interaction components. Each agent may display complementarities ($\phi_{ij} < 0$) or substitutabilities ($\phi_{ij} > 0$) with respect to the effort exerted by those to which he is connected. Whenever agent i listens to agent j (i.e. $q_{ij} = 1$), ϕ_{ij} is different from 0. When ϕ_{ij} is negative, effort levels are strategic complements. In other words, the greater is the effort that j devotes to processing information, the less costly it is for agent i to exert effort to process his neighbor's belief. When instead ϕ_{ij} is positive, this means that j's effort is a strategic substitute for i. This captures the behavioral assumption that agents may perceive different incentives of devoting effort to process the beliefs of those they listen to, based on their position in the network.

We assume that $\phi_{ii} = \phi$ for each i, and that $f(e_i)$ is linear and equal to $\delta - \alpha e_i$, where δ and α are constants such that $e_i \in [0, \delta/\alpha]$ for every i. Substituting (11) in (10) the expected utility of agent i in each period t becomes:

$$E[u_i(e_i, e_{-i})] = -\delta + \alpha e_i - \frac{1}{2}\phi e_i^2 - \sum_{j \neq i} \phi_{ij}e_ie_j \tag{12}$$

We assume that agent i's utility is concave in his own effort, so that in the absence of cross effects, all agents will exert a positive amount of effort implying that they always have an incentive to cooperate. The cross effects of effort depend both on the network structure and on the assumptions we make on how the utility of agents in different network positions, depends on the effort of those with whom they are connected. As mentioned previously we claim that the network structure endogenously determines the cross effects of effort and therefore the individual preferences. This affects the equilibrium levels of effort and therefore determines the way information is aggregated.

Given $\mathbf{e}$, the effort weighted listening links can be written in the following way:

$$\ell_{ij}(\mathbf{e}) = \frac{q_{ij}e_j}{\sum_j q_{ij}e_j} \tag{13}$$

where we denote the effort weighted listening matrix with $L(\mathbf{e}) = [\ell_{ij}(\mathbf{e})]$. Notice that the precision that agents assign to incoming messages, $\ell_{ij}(\mathbf{e})$ resembles that of (6), but in this case the weights are determined by the effort in aggregating information instead of the number of incoming links. As we will see, in equilibrium these effort levels reflect the network centrality of a given agent, which positively depends on the number of incoming links. When agents communicate repeatedly with their direct neighbors and make their guesses by carrying out an effort weighted average of their incoming stream of signals, passing it on to those that listen to them, the resulting beliefs after t completed communication rounds are:

$$\mathbf{y}_t = [L(\mathbf{e})]^{t-1}\mathbf{x}$$

The evolution of beliefs, $\mathbf{y}_t$, therefore depends exclusively on the levels of effort chosen by the other agents. Naturally, the weighted updating rule should not be seen as an exact algorithm that individuals will apply, as this seems very unrealistic. It should instead be seen as a general tendency to weigh the messages of others based on the network structure.

We denote $M = [-\phi_{ij}]$ as the square matrix that represents all cross effects of effort. We use M as a short-hand for the simultaneous move game with payoffs (12) and strategy spaces $\mathbb{R}_+$. We take a neutral stance in terms of the magnitude of the positive and negative cross effect, and simply distinguish between positive and negative effects. Following the procedure proposed by Ballester et al. (2006), we can characterize a Nash equilibrium of the game M, where the equilibrium strategies represented by effort levels are proportional to the measures of Bonacich centrality in the network of local complementarities derived from M.[14]

[14]See Appendix 2 for the derivation of equilibrium effort levels.

By substituting the equilibrium vector of effort $\mathbf{e}^*$ in the effort weighted updating matrix $L(\mathbf{e}^*)$, we obtain the consensus beliefs for a given network structure and matrix of effort complementarities. Notice that $L(\mathbf{e}^*)$ is related to $L(\mathbf{d}, \boldsymbol{\rho})$, since equilibrium effort levels reflect the network centrality of a given agent, which positively depends on the number of incoming links.[15]

Making different assumptions on cross effects may generate different predictions on consensus beliefs. Considering the unbalanced network of treatment 2 in our experiment, we therefore analyze different possible assumptions on the pair-wise complementarities versus substitutabilities of effort, in order to determine utility functions that are more consistent with our experimental results.

6.2 Application

The results of the experimental analysis are generally consistent with our initial conjecture on the heterogeneous role played by agents in aggregating information. More specifically, as shown in Table 4, numerical analysis suggests that the greater influence of agent D emerges only if preferences are such that both of the following conditions are satisfied:

1. A considers D's effort as a strategic substitute of his own.

2. C considers either A or B's efforts as strategic substitutes of his own.

These conditions have a rather intuitive interpretation that is consistent with our initial conjecture. Condition (1) implies that if D devotes a certain effort to processing information, A can avoid doing so and somehow conforms to D's belief without having to spend too much effort processing the information contained in D's belief. This assumption can for example be justified by the fact that since D has more incoming links, he potentially has greater scope for information aggregation.

The underlying rational of the second assumption is somewhat related to the first and relies on the fact that C's incoming links come from agents that received messages from a potentially more informed agent, namely D. In this case, therefore, either A or B (or both) are in a position to indirectly aggregate information, and C may consider either A or B's efforts as strategic substitutes of his own.

[15] Our game has a unique Nash Equilbrium where the equilibrium effort levels are proportional to the Bonacich network centrality measure. This measure was proposed in sociology by Bonacich (1987), and counts the number of all paths (not just shortest paths) that emanate from a given node and therefore is positively related to the number of incoming paths to a given node.

The model also highlights which conditions are necessary to generate consensus beliefs that go in the direction of persuasion bias, meaning that an agent with more outgoing links such as A will also have greater social influence. It turns out that this occurs in two possible cases. The first is if all efforts are complementary, in other words if from each agent's perspective, the efforts of those who communicate with him are strategic complements. The second case is characterized by efforts that are predominantly strategic substitutes, with only a few exceptions. An example of this can be seen in the last row of table (3), where only D considers C's effort as a strategic complement. In this last case, the social influence weights of A and D turn out to be very close to those of the Persuasion Bias model of DeMarzo et al. (2003).

Table 4: Social Influence weights under alternative parametrizations

Strategic Substitutes	w_A	w_B	w_C	w_D
None ($\forall \sigma_{ij} = 1$)	0.36	0.11	0.20	0.33
$\sigma_{AD} = -1$	0.32	0.16	0.20	0.32
$\sigma_{AD} = \sigma_{CA} = -1$	0.33	0.15	0.17	0.35
$\sigma_{AD} = \sigma_{CB} = -1$	0.33	0.15	0.17	0.35
All except $\sigma_{DC} = 1$	0.40	0.10	0.14	0.36

7 Conclusions

Humans learn most of what they know from others. Starting from this basic premise, this paper addressed a simple but important question: Are agents' opinions affected by the structure of the network through which they communicate? We presented an experimental investigation of the hypothesis of persuasion bias, whereby individuals communicating through a social network are unable to properly account for repetitions of information. Under persuasion bias, individuals' social influence ultimately reflects the structure of the network and, in particular, agents' outdegree.

The main conclusion of our analysis is that, as predicted by the persuasion bias hypothesis, agents fail to properly account for repetitions of information. As a consequence, the structure of the network plays a significant role in determining social influence, and opinions generally tend to converge towards those of the individuals who are better connected. However, contrary to the predictions of the persuasion bias model of DeMarzo et al. (2003), our experiment indicates that the most influential agents are not those who have more outgoing links, but rather those who have more incoming links.

We proposed a generalized boundedly rational updating rule that takes into account both agents' outdegree *and* their indegree, while nesting persuasion bias as a special case. Intuitively, the proposed heuristic is based on the idea that agents may take into account how informed their neighbors are when updating their own opinion. We showed that our generalized updating rule provides not only a more plausible characterization of aggregation of dispersed information, but also much higher explanatory power for the experimental data. We then presented a simple theoretical model, based on the structure of complementarity and substitutability among the efforts that agents devote to processing information, that provides a general framework for characterizing information aggregation in social networks.

Overall, our analysis indicates that most of what we know partially depends on the features of the social networks through which we communicate. In particular, due to a boundedly rational process in aggregating dispersed information, social influence depends not only on how much agents are listened to, but also on how much they listen to. As a result, in equilibrium, consensus beliefs tend to be swayed towards the opinions of *influential listeners.*

Appendix 1 - Instructions

[For T1 only, translated from Italian]

Welcome and thank you for taking part in this experiment. During the experiment talking or communicating with other participants is not allowed in any way. If you have a question at any time, raise your hand and one of the assistants will come to answer your question. By carefully following the instructions you can earn an amount of euro that will depend on the choices made by you and the other participants.

General Instructions

- 24 subjects will take part in this experiment
- The experiment takes place in 3 phases of 12 rounds each, for a total of 36 rounds.
- At the beginning of the experiment 6 groups of 4 subjects will be formed anonymously:
  - Each subject will interact exclusively within each group without knowing the identity of the other three subjects.
  - Each of the four subjects belonging to a group will be randomly and anonymously assigned one of four different roles: A, B, C and D.
  - The composition of each group and the roles assigned to the 4 components will remain unchanged throughout the experiment.

How earnings are determined

- In the first round of each of the 3 phases, the computer will randomly generate 4 integers that we will define *signals.* Each component of the group will be shown only one of the four signals. Signals will be denoted x_A, x_B, x_C, and x_D
- In each round of each of the 3 phases, each subject will be asked to guess the mean of the 4 signals extracted by the computer for that phase: $\overline{x} = \frac{x_A+x_B+x_C+x_D}{4}$
- The earnings of each subject will depend on the distance between his choice and $\overline{x}$:

- At the end of the experiment the computer will randomly select one of the 36 rounds
- Individual earnings will be equal to 20 euro minus the difference (in absolute value) between $\overline{x}$ and the choice made in the selected round.
- Both $\overline{x}$ and the choices made by each subject will contain at most 2 decimal points (i.e. 1412.00 or 21.50 or 516.33)
- If this difference turns out to be negative, the subject will earn 0 euro.

- Examples:
  - if $\overline{x} = 1424$ and the guess is 1424, the difference is 0 and earnings are 20 euro.
  - if $\overline{x} = 308$ and the guess is 311.5,the difference is 3.5 and earnings are 16.5 euro.
  - if $\overline{x} = 803.25$ and the guess is 792, the difference is11.25 and earnings are 8.75 euro.
  - if $\overline{x} = 62.5$ and the guess is 30.5, the difference is 32 and earnings are 0 euro, since $20 - 32 < 0$.

- In each round the choice that maximizes earnings depends on the information that each subject has on the signals:
  - if he knows only his own signal, the optimal choice is his own signal
  - if he knows or can deduce 2 signals, his optimal choice is the mean of the 2 signals
  - if he knows or can deduce 3 signals, his optimal choice is the mean of the 3 signals
  - if he knows or can deduce 4 signals, his optimal choice is the mean of the 4 signals

Information

- In each of the 3 phases
  - In the first round each subject knows his own signal

– From the second round onwards, before making his choice, each subject will be informed by the computer of the choices made in the previous rounds by some of the components of his group, based on the structure represented in the following figure:

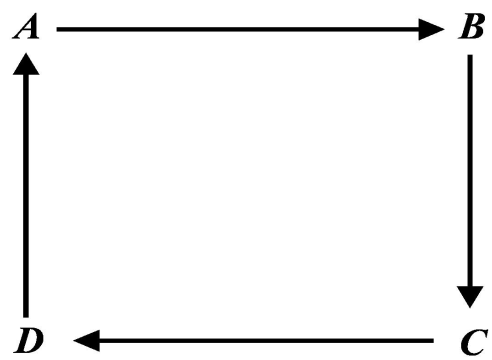

- Therefore, before making his choice
  - A will be informed of the choices made by D
  - B will be informed of the choices made by A
  - C will be informed of the choices made by B
  - D will be informed of the choices made by C

Feedback and Payments

- At the end of each phase the computer will show to each subject the 4 signals of his group, their mean, and the choices made.
- At the end of the experiment each subject will be shown the round the computer has selected to determine payments, the value of $\overline{x}$ for his group, the choice he made and the corresponding amount earned in euro.
- The experiment will terminate and the amount earned by each subject will be paid in cash.

Appendix 2 - Equilibrium Effort Levels

Following the procedure proposed by Ballester et al. (2006) we can decompose the matrix M in a concavity component, a global uniform substitutability component, and a local complementarity component. This decomposition gives us the non-negative matrix of local complementarities which we denote, $\mathbf{G} = [g_{ij}]$.This represents the adjacency matrix of a network $\mathbf{g}$ that reflects the pattern of existing payoff complementarities across all pairs of players.

In order to carry out the decomposition we define $\underline{\phi} = \min \{\phi_{ij} \mid i \neq j\}$ and $\overline{\phi} = \max \{\phi_{ij} \mid i \neq j\}$. We then let $\gamma = -\min \{\underline{\phi}, 0\} \geq 0$, $\lambda = \overline{\phi} + \gamma \geq 0$ and $g_{ij} = (\phi_{ij} + \gamma)/\lambda$ setting $g_{ii} = 0$, where by construction $0 \leq g_{ij} \leq 1$. The parameter g_{ij} measures the relative complementarity in efforts from i's perspective with respect to the benchmark value $-\gamma \leq 0$. This measure is expressed as a fraction of λ that is the highest possible relative complementarity for all pairs.

The decomposition of M allows us to rewrite the utility function of each agent i in the following way:

$$u_i(e_i, e_{-i}) = -\delta + \alpha e_i - \frac{1}{2}(\beta - \gamma) e_i^2 - \gamma \sum_{j=i}^{n} e_i e_j + \lambda \sum_{j=i}^{n} g_{ij} e_i e_j \quad \text{for all } i = 1, .., n$$

where γ corresponds to the weight assigned to the global substitutability component across all players represents, $\phi = (\beta - \gamma)$ is the second order derivative with respect to own effort (the concavity component), and $\lambda^* = \lambda/\beta$ denotes the weight of local interactions with respect to self-concavity.

The Nash equilibrium is unique and interior as long as λ^* is low enough, more specifically λ^* must be less than the inverse of the norm of the inverse of the largest eigenvalue of $\mathbf{G}$.

We therefore have that:

$$M = -\beta \mathbf{I} - \gamma \mathbf{U} + \lambda \mathbf{G}$$

where $\mathbf{I}$ is the n-square identity matrix and $\mathbf{U}$ is the n-square matrix of ones. Thus $\mathbf{G}$ captures all the heterogeneity in M. From Ballester et. al. (2006) we know that this game has a unique Nash Equilibrium $e^*(M)$ that is interior, where equilibrium effort levels, e_i^* can be expressed in relation to the total effort e^*:

$$e_i^* = \frac{b_i(\mathbf{g}, \lambda^*)}{b(\mathbf{g}, \lambda^*)} e^*$$

where $b_i(\mathbf{g}, \lambda^*)$ is the Bonacich centrality of agent i and $b(\mathbf{g}, \lambda^*)$ is the sum of the Bonacich centralities of all agents in the network. This allows us to write $\ell_{ij}(\mathbf{e})$ as:

$$\ell_{ij}(\mathbf{e}) = \frac{q_{ij} b_j(\mathbf{g}, \lambda^*)}{\sum_j q_{ij} b_j(\mathbf{g}, \lambda^*)}$$

Bonacich centrality of node i is defined as:

$$b_i(\mathbf{g}, \lambda^*) = \sum_{j=1}^{n} m_{ij}(\mathbf{g}, \lambda^*)$$

where $m_{ij}(\mathbf{g}, \lambda^*) = \sum_{k=0}^{\infty} (\lambda^*)^k (g_{ij})^k$ count the number of paths in $\mathbf{g}$ starting at i and ending at j, weighing paths of length k with $(\lambda^*)^k$.

We derive convergence beliefs for different assumptions on complementarities versus substitutabilities of efforts which generate different G matrices through the decomposition procedure. In order to ensure comparability of results we set the weight of local complementarities with respect to self-concavity, $\lambda^* = 1/2$. This value always satisfies the properties for the existence and uniqueness of a Nash Equilibrium mentioned above. In Table 4 we present convergence beliefs for five relevant cases.

References

[1] Acemoglu D., Dahleh M., Lobel I., and Ozdaglar A., "Bayesian Learning in Social Networks," Mimeo, MIT, 2010.

[2] Anderson, L. and Holt C., “Information Cascades in the Laboratory”, American Economic Review, 87 (5), pp. 847-62, 2007

[3] Bala V. and Goyal S., "Learning from Neighbours," The Review of Economic Studies, vol. 65, no. 3, pp. 595-621, 1998.

[4] Bala V. and Goyal S., "Conformism and Diversity under Social Learning," Economic Theory, vol. 17, pp. 101-120, 2001.

[5] Ballester C., Calvo-Armengol A., and Zenou Y., "Who is Who in Networks. Wanted: The Key Player.", Econometrica, 75(4), pp. 1403-1418, 2006.

[6] Banerjee A., "A Simple Model of Herd Behavior," The Quarterly Journal of Economics, vol. 107, pp. 797-817, 1992.

[7] Banerjee A. and Fudenberg D., "Word-of-mouth Learning," Games and Economic Behavior, vol. 46, pp. 1-22, 2004.

[8] Bavelas A., "Communication Patterns in Task-Oriented Groups." Journal of the Acoustical Society of America , vol. 22, pp. 725-30, 1950.

[9] Bikhchandani S., Hirshleifer D., and Welch I., "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," The Journal of Political Economy, vol. 100, pp. 992-1026, 1992

[10] Bonacich P., “Power and Centrality: A Family of Measures” The American Journal of Sociology, vol. 92, No. 5, pp. 1170-1182, 1987

[11] Bonacich P. and Lloyd P., “Eigenvector-like measures of centrality for asymmetric relations” Social Networks, 23, p.191-201, 2001

[12] Celen B. and Kariv S., "Observational Learning Under Imperfect Information," Games and Economic Behavior, vol. 47, no. 1, pp. 72-86, 2004.

[13] Celen B. and S. Kariv “Distinguishing Informational Cascades from Herd Behavior in the Laboratory.” American Economic Review, 94(3), pp. 484-497, 2004

[14] Choi S., "A cognitive Hierarchy Model of Learning in Networks,", ELSE Working Papers (238), 2006.

[15] Choi S., Gale D., and Kariv S., "Social Learning in Networks: A Quantal Response Equilibrium Analysis of Experimental Data," Berkeley Working paper, 2009

[16] Choi S., Gale D., and Kariv S., "Behavioral Aspects of Learning in Social Networks: An Experimental Study", Advances in Applied Microeconomics, vol. 13, Behavioral and Experimental Economics, edited by John Morgan, 2005

[17] DeGroot M., "Reaching a Consensus", Journal of the American Statistical Association, LXIX , pp.118-121, 1974

[18] DeMarzo P.M., Vayanos D., and Zwiebel J., "Persuasion Bias, Social Infuence, and Unidimensional Opinions," The Quarterly Journal of Economics, vol. 118, no. 3, pp. 909-968, 2003.

[19] Gale D. and Kariv S., "Bayesian Learning in Social Networks," Games and Economic Behavior, vol. 45, no. 2, pp. 329-346, 2003.

[20] Golub B. and Jackson M.O., "Naive Learning in Social Networks: Convergence, Influence, and the Wisdom of Crowds," American Economic Journal: Microeconomics, 2(1), pp. 112–49, 2010.

[21] Jackson M.O., "The Study of Social Networks In Economics," chapter in "The Missing Links: Formation and Decay of Economic Networks", edited by James E. Rauch; Russell Sage Foundation, 2007.

[22] Jackson M.O., *Social and Economic Networks*, Princeton University Press, Princeton, 2008

[23] Smith L. and Sorensen P., "Pathological Outcomes of Observational Learning," Econometrica, vol. 68, no. 2, pp. 371-398, 2000.

[24] Smith L. and Sorensen P., "Rational Social Learning with Random Sampling,' unpublished manuscript, 1998.

[25] Udry C. and Conley T., "Social Learning through Networks: the Adoption of New Agricultural Technologies in Ghana," American Journal of Agricultural Economics, vol. 83, pp. 668-673, 2001.

[26] Wasserman S. and Faust K., *Social Network Analysis: methods and applications*, Cambridge University Press, 1994.