



**İstanbul
Bilgi Üniversitesi**
LAUREATE INTERNATIONAL UNIVERSITIES



Working Paper Series

Working Paper No: 2014-03

An Experimental Study of Persuasion Bias and Social Influence in Networks

Jordi Brandts
Ayça Ebru Giritligil
Roberto A. Weber

October 2014

An Experimental Study of Persuasion Bias and Social Influence in Networks

Jordi Brandts*, Ayça Ebru Giritligil** and Roberto A. Weber***

October 9, 2014

Abstract

In many areas of social life individuals receive information about a particular issue of interest from multiple sources. When these sources are connected through a network then proper aggregation of this information by an individual involves taking into account the structure of this network. The inability to aggregate properly may lead to various types of distortions. In our experiment a number of agents all want to find out the value of a particular parameter unknown to all. Agents receive private signals about the parameter and agents can communicate their estimates of the parameter repeatedly through a network, the structure of which is known by all players. We present results from experiments with four different networks. We find that the information of agents who have more outgoing links in a network gets more weight in the information aggregation of the other agents than it optimally should. Our results are consistent with the model of “persuasion bias” of De Marzo et al. (2003) and at odds with an alternative heuristic according to which the most influential agents are those with more incoming links.

Keywords : Persuasion Bias, Experiments, Bounded Rationality

JEL Classification Codes : C92, D03, D83

Acknowledgements: Financial support by the Spanish *Ministerio de Ciencia e Innovación* (Grant: ECO2011-29847-C02-01), the *Generalitat de Catalunya* (Grant: 2009 SGR 820), the Barcelona GSE Research Support Program and *Consolider-Ingenio* is gratefully acknowledged. Aniol Llorente-Saguer and Julian Rode provided excellent research assistance.

* Instituto e Analisis Economico (CSIC) and Barcelona GSE. Address: Campus UAB, 08193 Bellaterra (Barcelona), Spain. e-mail: jordi.brandts@uab.es

** Murat Sertel Center for Advanced Economic Studies, İstanbul Bilgi University, Santral Campus L1-103, Kazım Karabekir Cad. 2/13, Eyüp 34060 İstanbul, Turkey. e-mail: ayca.ebru@bilgi.edu.tr

*** Department of Economics, University of Zurich, Blümlisalpstrasse 10, CH-8006 Zurich, Switzerland. email: roberto.weber@econ.uzh.ch

1. Introduction

In many social and economic situations individuals receive information about a particular issue from multiple sources and also send information to multiple others. Examples are the sharing of political opinions among voters or of information about prospective job candidates in an organization's hiring process. When people exchange opinions about such issues within a group some of the group members may have a stronger influence on the group's opinion than others due to the quality or accuracy of the information. However, often influence is also due to social factors like the resources some people have available to invest in spreading their opinions, how well considered their opinions are in society or how well connected they are to others.

One basic source of social influence occurs in communication through networks, where there exists the possibility of one person's information reaching a particular other person repeatedly. When individuals are connected through a network then they may receive information both directly and indirectly from the same source and send information directly and indirectly to a particular other source. In such situations perfectly rational aggregation of information by an individual involves taking into account the structure of the network and adjusting how one weighs the information one receives accordingly. However, boundedly rational agents may have difficulties with this process and aggregate in biased ways, by failing failing to adjust properly for repetitions of information.¹ In this paper we present results from laboratory experiments that shed light on how people aggregate information when they are connected in a network

¹ See, for instance, Gale and Kariv (2003), Golub and Jackson (2010) and Acemoglu et al. (2011).

DeMarzo et al. (2003) present a stylized version of an information aggregation situation in a network, together with a particular model of the bias to which boundedly rational agents will fall prey. They posit a situation in which a set of agents all want to find out the value of a numerical parameter. Each agent starts with some initial private information about the parameter, the aggregation of which is all the information available to the group of agents in the network. Agents then communicate their estimates about the true parameters to one another through the network. The network consists of a number of connections between the set of agents that specify who sends information to whom—or, alternatively, who listens to whom. There are multiple rounds of communication between the agents, a feature meant to represent a lengthy deliberation process. In each round each agent listens to the estimates of those who, following the network structure, send him information and sends his estimate to those who listen to him. After each round each agent can update his own estimates in order to approximate the true parameter based on any new information received from other agents.

There is, of course, a rational way to aggregate information in such a setting, which involves agents discounting information that reaches agents repeatedly through distinct channels in the network. But it is possible that people will not use the optimal process. A priori, there are many distinct ways in which information could be aggregated non-optimally. DeMarzo et al. (2003) propose a particular model of boundedly rational information aggregation, based on DeGroot (1974), which leads to “persuasion bias.” According to this model all agents treat all information they receive as new, ignoring the fact that an estimate received in a particular round uses information that has already been received previously—directly or indirectly—from another source. Agents treat the information they get each round as new and independent and do not adjust for the fact that over time the information of some

agents might contain more repetitions than that of others. The implications of this kind of biased information aggregation depend on the structure of the network. Some networks may cancel out the biased weighting of information, so that agents subject to the boundedly rational persuasion bias may nevertheless reach an unbiased estimate of the true parameter. However, other networks will not have this property and the consequence will be that the group arrives at a consensus estimate of the parameter that is biased.

The model of bias proposed by De Marzo et al. is a plausible one. Nevertheless, it is not a priori obvious that if people turn out to be biased, they will be so in that particular way. There are many ways to aggregate information and it is hardly possible to know what bias will occur without any empirical information. Therefore, we use laboratory experiments to explore the internal validity of the persuasion bias model. The laboratory provides an ideal environment to study the relation between network structures and the kind of information and communication processes necessary to test the persuasion bias model. The two main virtues of laboratory experiments are control and replicability. Causal knowledge requires controlled variation (Falk and Heckman 2009) and the laboratory allows for tight control over the environment in which interaction takes place. At the same time, it is possible to generate sufficient data in a simple way.

In this paper we present results from experiments with four different networks. The first two are directly inspired by the discussion in DeMarzo et al. (2003). We study behavior in their examples of both a “balanced” and an “unbalanced” network of four agents. In the balanced network four agents are located in a circle. They all receive one piece of information and all send their estimate of the parameter to and receive it from the two agents closest to them. In such a balanced network, DeMarzo et al. predict no bias in the consensus estimates of the group. In the unbalanced network the four agents all receive one piece of information, but some agents send information to a larger number of others and receive

information from a smaller number of agents than others. In such an unbalanced network, DeMarzo et al.'s model predicts a precise biased outcome to which the group's estimates should converge.

We find that observed behavior is consistent with persuasion bias, and with the precise predictions of DeMarzo et al.'s model. In the balanced network subjects converge to estimates in which the private information of all four players carries identical weights, as predicted by the theory. In contrast, in the unbalanced network subjects converge to biased consensus estimates that reflect greater weight on the private information of those with more outgoing communication channels and those connected to such agents, which is consistent with the theoretical predictions. Moreover, the weights of the different private signals correspond closely to those of the persuasion bias model. Hence, our data provide strong support for the theory.

In a set of experiments conducted independently Corazzini et al. (2012) find that the network structure plays a significant role in determining social influence, but that the most influential agents are not those with more outgoing links, as predicted by the persuasion bias hypothesis, but those with more incoming links. Their study presents data from a balanced and an unbalanced network that are different from those proposed by De Marzo et al. (2003). They present an explanation of their results in terms of a boundedly rational updating rule that takes into account not only agents' outgoing links but also their incoming links. This is an interesting alternative to the persuasion bias model in suggesting a different way in which bounded rationality might influence information updating in a network.

Given the qualitative discrepancy between the results of Corazzini et al. (2012) and our results we conducted new experimental sessions using their unbalanced network. Our results are again consistent with persuasion bias and not with the alternative model. In addition, we perform a simple counter-factual exercise by applying their model to our

unbalanced network. We find that the persuasion bias model outperforms their model, both in our new data using their network and in our original experiment using the networks proposed by DeMarzo et al. Hence, in contrast with their results, we provide much stronger support for the predictions of persuasion bias.

Our experiment also relates to a paper by Enke and Zimmermann (2014) on correlation neglect, which describes agents' tendency to overweight information, received from multiple sources, that is correlated due to its origin from the same source. They report a laboratory experiment in which subjects' exhibit correlation neglect and overweight information that they receive repeatedly through multiple channels. They also demonstrate, in a market experiment, how correlation neglect can lead to herding by market participants that results in bubbles and crashes. However, they do not test the precise predictions of a model of persuasion bias, as we do here.

2. Overview of Design and Theoretical Background

We now present the theoretical environment that we study experimentally, which is based on model of De Marzo et al. (2003). The aim of this section is to introduce some key concepts and the three theoretical results that motivate our experimental design.

In all four networks that we study there is a set of four agents $N = \{1, 2, 3, 4\}$ who all wish to find out the value of an unknown one-dimensional parameter, θ . There are sixteen rounds of interaction between the four agents. The reason why we chose this particular number of rounds is grounded in the persuasion bias model, which we explain in more detail below. Before the first round of interaction, each agent obtains some initial information about θ in the form of a noisy signal $x_{i0} = \theta + \varepsilon_i$, where ε_i is an error term with mean zero, independent across agents and normally distributed. Each agent i assigns an initial precision to the information of each other agent j , π_{ij0} .

Starting with the initial information, agents communicate according to a communication network that can be represented as a directed graph indicating whether agent i listens to agent j . We denote by $S(i)$ the set of agents to which i listens, and by q_{ij} the indicator function of $S(i)$, which takes the value of 0 if agent i does not listen to agent j and the value of 1 if he does; we consider that each agent listens to himself. We refer to the set $S(i)$ as the listening set of agent i and to the function S as the listening structure.

De Marzo et al. (2003) model communication and updating as follows. In the first communication round, agent i learns the signals of the agents in $S(i)$. Given normality and agents' fixed assessment of the precision of others' information, a sufficient statistic for these signals is their weighted average, with weight given by the precisions. They denote this statistic by x_{i1} and refer to it as agent i 's beliefs after one round of updating:

$$(1) \quad x_{i1} = \sum_j (q_{ij}\pi_{ij0}/\pi_{ii0}) x_{j0}$$

where $\pi_{ii0} = \sum_j q_{ij}\pi_{ij0}$ denotes the precision that agent i assigns to his own beliefs and x_{j0} represents j 's initial beliefs. Equation (1) is called agent i 's updating rule.

The updating rule (1) can be expressed in vector notation. Denote by x_t the matrix whose i th row is the vector x_{it} of agent i 's beliefs in communication round t . Denote also by T the listening matrix with elements:

$$(2) \quad T_{ij} = q_{ij}\pi_{ij0}/\pi_{ii1}.$$

Then the updating rule can be expressed as $x_1 = T x_0$.

A very important simplifying assumption is that all precisions are the same. This converts each element of the listening matrix into something very simple: the ratio between the indicator q_{ij} (0 or 1) and the total number of agents j that agent i listens to (1,2, 3 or 4). As a consequence of this simplification, the proposed updating consists simply of the averaging of others' estimates. In our experimental context this assumption is a very reasonable one, since (as will become clearer below) our experimental subjects will not have

much basis for attributing different precisions to distinct other agents. Nevertheless, it is an empirical question whether experimental subjects will behave in accordance with this assumption.

We now summarize the three key results of the De Marzo (2003) model pertaining to the convergence process of the updating rule that we will use in our analysis below. First, their Theorem 1 states that if the set of agents is strongly connected and another technical assumption holds then the vector of agents' social influences converges to the consensus given by the vector ω which is the unique solution $w = Tw$, where w is the "social influence" vector showing the weight of each agent's initial signal in the consensus. A set of agents being strongly connected simply means that every agent speaks to every other agent (directly or indirectly) and thus no agent is isolated from the others.

Second, their Theorem 2 states that, if agents follow the updating process introduced above, the beliefs to which agents converge are correct if and only if the network is balanced, i.e. if and only if for all i $\sum_j q_{ji} T_{ji} = 1$. In words, the sums of weights in the listening vectors are the same and equal to 1 for all agents. Third, their Theorem 3 states that if agents are fully rational (i.e. not subject to persuasion bias) the agents converge to the correct beliefs after at most N^2 rounds of updating. In our case $N^2 = 16$. Hence, we chose the number of rounds to be equal to N^2 , to allow for rational agents to converge to the true parameter.

The convergence to a consensus, the conditions under which the consensus is the mean of the signals—i.e., the rational Bayesian standard—and the time needed for rational agents to converge to the correct value will be studied with our data. We will also analyze the specific predictions of the persuasion bias model pertaining to the weights carried by the information of the different agents. To describe what this means more clearly, we now present the particular networks we use in Experiment 1.

3. Experiment 1

Figure 1 shows the “Balanced Network 1” used in our first experiment. The arrows indicate which agents send information to which other agents. For example, agent 1 sends information to agents 2 and 4 and also receives information from agents 2 and 4. This network is balanced, precisely because each agent talks and listens to the same number of other agents. Denoting by $S(i)$ the set of agents that agent i listens to we have that for the “Balanced Network 1” : $S(1) = \{1,2,4\}$; $S(2) = \{1,2,3\}$; $S(3) = \{2,3,4\}$; $S(4) = \{1,3,4\}$, including always the agent in question, i.e. each agent listens to himself.




Figure 1: Balanced Network 1

From here we obtain the corresponding listening matrix $T_{B1} = \begin{pmatrix} 1/3 & 1/3 & 0 & 1/3 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 1/3 & 1/3 & 1/3 \\ 1/3 & 0 & 1/3 & 1/3 \end{pmatrix}$. Solving

the system $wT=w$, we find that the weight of the four signals in the consensus value to which the updating process converges is given by the vector $w_{B1} = (1/4, 1/4, 1/4, 1/4)$, i.e. the social influence is the same for all four agents. In spite of the fact that the updating process does not

take into account repetition, the process converges to an unbiased set of weights, because the network is balanced.

Figure 2 shows the “Unbalanced Network 1,” where some agents speak and listen to a different number of other agents. In this case the sets of agents that agents listen to are the following $S(1) = \{1,4\}$; $S(2) = \{1,2,3\}$; $S(3) = \{1,2,3,4\}$; $S(4) = \{1,3,4\}$.




Figure 2: Unbalanced Network 1

From here one can construct the corresponding listening matrix $T_{UB1} = \begin{pmatrix} 1/2 & 0 & 0 & 1/2 \\ 1/3 & 1/3 & 1/3 & 0 \\ 1/4 & 1/4 & 1/4 & 1/4 \\ 1/3 & 0 & 1/3 & 1/3 \end{pmatrix}$.

and obtain the following vector of social influence weights: $w_{UB1} = (16/42, 3/42, 8/42, 15/42)$ for agents 1, 2, 3 and 4 respectively. In percentage terms, this vector is approximately equal to $w_{UB1} = (38\%, 7\%, 19\%, 36\%)$. Observe that the influence weights do not directly reflect the number of agents an agent speaks to. Agents 3 and 4 both speak to two other agents, but the corresponding influence weights are 19% and 36% due to the fact that agent 4 speaks to agent 1, who in turn talks to all three other agents.

In Experiment 1 we study behavior in the two networks that we have just introduced. The set-up was identical for all networks, with the exception of the links between subjects

corresponding to the different communication networks. Each person participated only in one session with one particular network. The shape of the relevant network was common information. A graph of the network was visible to the participants at all times.

Each participant was a member of a fixed group of four subjects.² Each group interacted consecutively in four blocks, each of which was based on a new target value. This was to allow for learning of the task. At the beginning of a block, participants did not know the target value. Instead, each of the four subjects initially received a private signal drawn from a normal distribution centered on the target value.³ The signals were drawn separately for the different target values.

Once each participant had received his or her signal, then the sixteen rounds of information exchange within the group for that block began. In each round each agent in a group sent an estimate of the target parameter to all agents in his or her listening set. After receiving estimates from others according to the network structure, the group moved to the next round. All 16 rounds proceeded according to the same rules.

The choice of sixteen rounds was guided by the result in De Marzo et al. (2003) that fully rational agents will converge to the true value in at most N^2 rounds of updating. Each subject was given a bundle of four record sheets, one for each round, to write down his or her own guess as well as the guesses of others in his/her listening set at each period. Each subject also received a sheet of paper depicting the communication network.

In every round, each participant's payoff was determined by the distance between the estimate provided by the participant in that round and the actual target value in that block. Specifically, we implemented the payoff function, $\pi = 400 - 0.5 * (\text{target number} -$

² The instructions can be found in Appendix A.

³ Participants did not receive any additional signals in subsequent rounds.

estimate)², where payoffs are denoted in Euro cents. In each block, one round was selected at random to count for payment. In principle, subjects payoffs could be negative, meaning that they could lose part of the 5 Euro participation payment, though this was unlikely. We did not allow subjects to lose more than the participation payment.

Each set of target values was randomly drawn from a uniform distribution over the range determined by taking the tenth highest and lowest numbers from one hundred draws from a normal with mean 0 and standard deviation 400.⁴ For robustness we investigated two complete information sets, i.e. two different sets of four target values with the corresponding set of signals for each of the four agents in the network. The two information sets were used in different sessions. Table 1 shows the distribution of groups per treatment for experiment 1. Since each group is composed of 4 subjects, the total number of subjects is 308.

We can now formulate specific hypotheses based on the persuasion bias model:

Hypotheses 1: In the sixteen rounds subjects will converge to a consensus in both networks.

Theorem 3 of De Marzo et al. (2003) proposes that perfectly rational players will converge in sixteen rounds. We therefore conjecture that sixteen rounds will be enough for convergence, regardless of whether participants behave perfectly rationally or not.⁵

Hypothesis 2: In the balanced network, subjects will converge to a consensus estimate in which the four initial signals will carry equal weight

Hypothesis 2 is based on the prediction from DeMarzo et al. for both perfectly rational agents and to those affected by persuasion bias. In either case, the balanced network should produce no bias in consensus estimates.

⁴ The range of the target values in the parameter set 1 and 2 are determined to be [-512.669 , 437.452] and [-501.17 , 418.394], respectively. The two parameter sets can be found in Appendix B.

⁵ Hypothesis 1 can be seen as a prerequisite for our analysis. Boundedly rational behavior distinct from persuasion bias may not lead to convergence to a consensus.

Hypothesis 3: In the unbalanced network, subjects will converge to a consensus estimate in which the four initial signals will carry unequal weights. These weights will be as predicted by the persuasion bias model.

Hypotheses 2 and 3 propose specific numerical parameters, either of equal weighting or precise unequal weights based on the analysis in Section 2, against which we can compare the data.

Table 1: Number of groups per treatment

Treatment (Parameter Set #)	Number of groups
Balanced 1 (1)	10
Balanced 1 (2)	10
Unbalanced 1 (1)	8
Unbalanced 1 (2)	9

4. Results for Experiment 1

We first look at our evidence pertaining to the convergence to a consensus. Figure 3 shows how the proportion of groups in which the variance of estimates is less than 4 evolves over time, separately for the Balanced and the Unbalanced Network 1. One can see that the proportion grows steadily for both networks. Figures 4 and 5 show the proportion of groups with variance less than 4, but now aggregated, separately, by parameter sets and by blocks. In Figure 4 one can see that there is no relevant difference between the parameter sets and in Figure 5 one can see how the convergence occurs in all four separate blocks, with the

convergence in block 4 appearing to reach higher levels. Overall, our data are consistent with Hypothesis 1 that behavior converges to a consensus.⁶

We now move to the examination of Hypotheses 2 and 3, pertaining to whether behavior converges to an unbiased estimate of the true target value, in which the individual private signals receive equal weight, or whether they converge to systematically biased consensus estimates that concord with the predictions of persuasion bias.

Figures 6 and 7 show, for parameter sets 1 and 2, respectively, the bias in average median estimate for the final four periods of each block. That is, for each group, we look at the median of the four estimates in the final four periods of a block. We then average these medians across all groups. For the balanced network the predicted values of bias are zero, according to Hypothesis 2. For the unbalanced network the persuasion bias predictions are provided by the dashed blue line. The two solid lines show the observed data. Figures 6 and 7 suggest that the estimates to which participants converge are in line with what the persuasion bias model suggests: no bias for the balanced network and a particular direction of bias for the unbalanced network. To put the analysis of the bias on a more solid basis we use regression analysis. Table 2 shows, in the first two columns, the results of absorbing regressions for the balanced network and the unbalanced network 1, where the mean estimates of groups are regressed on the value of the four signals and standard errors are clustered by group. The results for the Balanced Network show that the coefficients for all four signals are significant and close to the predicted and unbiased weights of 0.25. The point estimates are quite similar, varying between .22 and .28. The results of the F-tests show that the null hypothesis that all four coefficients are equal can not be rejected ($p < .75$). Hence, the

⁶ Figures which show the proportion of groups with variances less than 1 show increasing trends similar to Figures 3-5.

regression results are consistent with Hypothesis 2: in the Balanced Network participants' guesses converge to consensus estimates in which all four private signals carry roughly the same weights.

Table 2: Mean Estimates

Variable	Mean Estimate Balanced Network	Mean Estimate Unbalanced Network 1	Mean Estimate Unbalanced Network 2
Signal 1	.280*** (.070)	.384*** (.069)	.469*** (.054)
Signal 2	.235*** (.079)	.045 (.086)	.009 (.081)
Signal 3	.220*** (.027)	.199*** (.035)	.215*** (.049)
Signal 4	.250*** (.045)	.349*** (.049)	.299*** (.054)
Constant	2.616 (2.268)	-2.970 (2.268)	-.963 2.298

For the Unbalanced Network 1 the regression results reveal significant coefficients for signals 1, 3 and 4, but not for signal 2. One can also see that the point estimates are quite different and, indeed, the result of the F-test rejects the null hypothesis of equality of all four coefficient ($p < .0836$). The persuasion bias model predicts weights of 38%, 7%, 19% and 36% respectively for this network. The regressions results yield weights of 38%, 4%, 20% and 35%. The ordering of the estimated weights is the same as that of the predicted weights and even the closeness of the values is striking. We conclude that our data are consistent with Hypothesis 3; our results clearly support the precise point predictions of the persuasion bias model.

5. Experiment 2

After having conducted experiment 1 we became aware of the paper by Corazzini et al. (2012). This paper presents the results of experiments on persuasion bias using a balanced network and an unbalanced one, both different from the ones we use. They find results that are not in line with persuasion bias. Instead, in their data “social influence depends not only on being listened to by many others, but also on listening to many others” (p. 1276). In addition they present a model of a generalized boundedly rational updating rule and show that persuasion bias and the behavior observed in their experiment can both be seen as the outcomes of special cases of that updating rule.

The explanation given by Corazzini et al. (2012) is indeed an intriguing one. The question that arises is what the mechanism behind the rule that they use to explain their behavior is. The observation of the network structure may lead some of the participants to give more weights to the input of somebody who listens to many players guided by the notion that those players have a lot of information. In this sense the idea that influential listeners’ information should carry more weight can be interpreted as being more rational, i.e. as an attempt to look at the overall situation and not simply average the information obtained from others, as in the persuasion bias model. However, the results by Corazzini et al. seem to be at odds with our results in Experiment 1, using Unbalanced Network 1, which are largely consistent with the persuasion bias model in DeMarzo et al.

One possible source of the difference between their results and ours is that they employ different networks. Figure 3 shows “Balanced Network 2,” which is the balanced network studies by Corazzini et al. The listening sets for this network are $S(1) = \{1,4\}$; $S(2)$

$= \{1,2\}$; $S(3) = \{2,3\}$; $S(4) = \{3,4\}$. The listening matrix and the vector of influence weights

are given by $T_{B2} = \begin{matrix} & 1/2 & 0 & 0 & 1/2 \\ 1/2 & 1/2 & 0 & 0 \\ 0 & 1/2 & 1/2 & 0 \\ 0 & 0 & 1/2 & 1/2 \end{matrix}$ and $w_{B2} = (1/4, 1/4, 1/4, 1/4)$, respectively.




Figure 3: Balanced Network 2

The regression results reported in Corazzini et al. (2012)) for Balanced Network 2 find that the value of the estimated coefficients vary between .15 and .34, but that the null hypothesis of equality of equal weights for the four coefficients can not be rejected. This evidence adds to the support in our first experiment for Hypothesis 2, and suggests that behavior in balanced network is consistent both with the persuasion bias model and with rational information aggregation. Since the results of Corazzini et al using Balanced Network 2 are in line with our results, we did not repeat the experiment with their balanced network.

Figure 4 shows “Unbalanced Network 2”, which is the unbalanced network studied in Corazzini et al. (2012). Here the listening structure is given $S(1) = \{1,4\}$; $S(2) = \{1,2\}$; $S(3) = \{1,2,3\}$; $S(4) = \{1,3,4\}$ for which the listening matrix is given by $T_{UB2} =$

$\begin{matrix} 1/2 & 0 & 0 & 1/2 \\ 1/2 & 1/2 & 0 & 0 \\ 1/3 & 1/3 & 1/3 & 0 \\ 1/3 & 0 & 1/3 & 1/3 \end{matrix}$. The influence vector is $w_{UB2} = (42\%, 10\%, 16\%, 32\%)$. Agent 1 is the

most influential one, since he is the only one that speaks to all three others. Agent 2 is the least influential one, since he only speaks to agent 3 who only indirectly reaches agent 1. Agent 4 is the second most influential one; like agents 2 and 3 he only speaks to one other agent, but this agent turns out to be agent 1 who speaks to all three others.




Figure 4: Unbalanced Network 2

The weights that Corazzini et al. (2012) find for signals 1 to 4 are 25%, 20%, 25% and 30%. An F-test rejects the null hypothesis of equality of the weights. This configuration is certainly quite distinct from the ones predicted by the persuasion bias model of DeMarzo et al, of 42%, 32%, 16% and 10%. This discrepancy is intriguing and led Corazzini et al. to formulate an alternative model to the persuasion bias.

The question that arises is how their model and the persuasion bias model can be compared. Corazzini et al. (2012) present three results: two experimental results and the proposal of a new model. The results from their experiment with the unbalanced network is at odds with the persuasion bias model and the model they propose is an alternative to the persuasion bias model, which can explain the data from their unbalanced network experiment.

We employ two approaches to compare their work to ours. First, we derive the predictions of their model for our Unbalanced Network 1, from our Experiment 1, to see how

well their model does in explaining our independent data. For the unbalanced network from our Experiment 1, Unbalanced Network 1, Corazzini et al.’s model predicts weights, for signals 1 to 4, of 18%, 14%, 36% and 32%. This is quite distinct from the estimated weights we observe in the data, of 39%, 5%, 20% and 36%. Recall that these empirical weights are, instead, much more similar to those predicted by DeMarzo et al.’s persuasion bias model (38%, 7%, 19%, 36%). Hence, it is clear that in our data the persuasion bias model does better than the alternative model.

Second, we conduct new experimental sessions using the unbalanced network studied by Corazzini et al, Unbalanced Network 2, to see if their result replicates with our procedures, including the parameter values. In this second experiment, we used 10 groups with parameter set 1 and 9 groups with parameter set 2. Aside from using a different network and only studying an unbalanced network, procedures were identical to those in Experiment 1.

The third column in Table 1 shows the regressions results for our Experiment 2, indicating the weights placed on private signals in the consensus estimate arising in the final four periods of each block. The F-test strongly rejects the null hypothesis of equality of coefficients and indeed the point estimates are quite distinct from each other. Our estimated values for the vector of weights are 47%, 1%, 22% and 30%. These results are qualitatively in line with the persuasion bias model, which predicts weights of 42%, 10%, 16% and 32%. Hence, replicating our experiment, using the network studied by Corazzini et al. finds further support for the predictions of the persuasion bias model.

5. Concluding Remarks

We provide several experimental results that are consistent with the predictions of a model in which communication through networks yields biased information aggregation. Our

findings provide clear support for the prediction that people with more outgoing links and those connected to such people exert greater influence on group beliefs. Moreover, our results also yield estimated influence weights that are very close to those predicted by a model of persuasion bias.

The fact that the results from our balanced network 1 are consistent with these predictions is perhaps not surprising. However, the results from two different unbalanced networks are also consistent with persuasion bias. In addition we confront the data with an alternative model proposed by Corazzini et al. (2012) and do not find it to explain the data better than persuasion bias.

On the basis of our findings, we conclude that the persuasion bias model represents a good benchmark for the study of information aggregation. The model is based on a simple, intuitively defensible bias in information processing and we demonstrate that it generates predictions that are supported empirically in a careful test. To us this conclusion does not come as a surprise, since from the start the persuasion bias model seemed to us as a simple, perhaps very mechanical, but nevertheless natural way to behave. Of course, future research may produce results that may be at odds with persuasion bias, requiring the need for a refined or altogether new model. But, in at least a first test, the assumptions and predictions of persuasion bias appear vindicated.

REFERENCES

- Acemoglu, D., Dahle, M., Lobel, I., Ozdaglar, A., 2011. Bayesian learning in social networks. *Review of Economic Studies* 78(4), 1201-1236.
- Corazzini, L., Pavesi, F., Petrovich, B., Stanca, L., 2012. Influential listeners: An experiment on persuasion bias in social networks. *European Economic Review* 56, 1276-1288.
- Enke, B., Zimmermann, F., 2013. Correlation Neglect in Belief Formation.
- DeGroot, M., 1974. Reaching a consensus. *Journal of the American Statistical Association* 69, 118-121.
- Gale, D., Kariv, S. 2003. Bayesian learning in social networks. *Games and Economic Behavior* 45 (2), 329-346.
- Golub, B., Jackson, M.O., 2010. Naïve learning in social networks: convergence, influence, and the wisdom of crowds. *American Economic Journal: Microeconomics* 2 (1), 112-149.

APPENDIX A: Instructions

General Information

This is an experiment in decision-making. During the experiment, you will accumulate money. In addition to a participation fee of 5€, you will receive the amount you accumulate in cash after the experiment is finalized. The exact amount you receive will be determined during the experiment and will depend on your decisions and/or the decisions of other participants.

If you have any questions during the experiment, please raise your hand and wait for an experimenter to come to you.

Please do not try to communicate with other participants during the experiment.

At the bottom of this screen, you will see a letter (A, B, C, or D) and a number (1, 2, 3, or 4). This combination of letter and number is your participant number for the experiment. Your participant number will be the same for the entire experiment. This number is private and should not be shared with anyone. Please record this number at your stack of record sheets.

Your participation number

Periods, Rounds, Groups, and Roles:

During this experiment, you will be in a group with three other participants. You will be grouped with the three other people with the same letter in their participant number. That is, everyone with a letter “A” in his/her participation number will be in the same group, everyone with the letter “B” will be in the same group, and so on. Throughout the experiment, you will be grouped with the same three participants.

In the experiment, you will participate in 4 rounds of an activity. Each round will consist of 16 periods. That is, during the experiment you will participate in Rounds 1 through 4, and each of these rounds will consist of Periods 1 through 16.

For each round, we will randomly select one of the 16 periods for which to pay you. Since you will not know which period will be selected for payment until the end of the experiment, every period could influence your earnings for the experiment.

Round Instructions:

In each round, a number will be drawn at random for your group. This number will be a real number and can include three decimal places. The range from which we draw the number may change from round to round. You will not be given any other information about how the

number is determined. This number will be referred to as the “target number” for your group for that round.

At the beginning of each round, every member of your group will receive a “private estimate” of the target number. This private estimate will consist of the target number, plus or minus a random number. For each member of your group, the random number will be drawn from the normal distribution, with an average of zero. This means that your estimate will be related to the target number, but only imperfectly so. That is, your estimate will give you an idea of the target number, but will most likely not be the same as the target number. To be more precise, the private estimates will be normally distributed around the target number.

Each of you will receive a different private estimate. Each of the estimates will be drawn independently from the normal distribution around the target number. That is, for each of you we will start with the target number and then add a random number drawn from the normal distribution. Your private estimates will therefore be related to each other and to the target number, but only imperfectly so.

Are there any questions about the target number and estimates?

Period Instructions:

In each period, you will enter your guess of the target number. You will do so by entering a number (of up to 3 decimal places). The computer will prompt you for your guess and will then ask you to confirm this guess. If you want to enter a guess with a negative sign, make sure that you do not forget to put the negative sign.

In each period, your payoff will be determined by how close your guess is to the target number. That is, we will compare your guess with the target number. The closer your guess is to the target number, the higher your payoff will be for that period. Note that you will not be able to see either the target value or your payoff for each period until the end of the round.

IMPORTANT: YOU NEED TO USE THE DECIMAL POINT (.) TO ENTER THE DECIMALS.

Information Instructions:

At the end of each period, you will receive information on the guesses of other participants. Each of you will be able to observe the guess entered in that period by at least one other participant. Some of you may receive more information than others.

There are four participants in each group, with numbers from 1 to 4. The information received by each participant is presented in the table below. For each participant (row), the columns with an “x” indicate which guesses that participant is able to observe. Note that each participant is able to observe his or her own guess. In addition, every participant will be able to observe the guess of at least one other participant. For example, in addition to his/her

own guess, Participant 3 can observe the guesses of the rest of the participants whereas Participant 1 can observe only the guess of Participant 4.

	Observes the guess of participant:			
	Participant 1	Participant 2	Participant 3	Participant 4
Participant 1	X			X
Participant 2	X	X	X	
Participant 3	X	X	X	X
Participant 4	X		X	X

Therefore, the only information you will have on the target number is your private estimate as well as the guesses you receive from the other participants.

To summarize, in each period you will make a guess, this guess will determine your payoff for that period, and you will then observe the guesses of some of the other members of your group. Once this takes place, you will continue to the next period.

After 16 periods, the round will end, and you will be able to see the target number for that round, as well as your payoff in each of the 16 periods. At the end of the experiment, one of these 16 periods will be selected at random to determine your earnings for that round. You will be paid your earning across all 4 rounds, in addition to the 5 € participation fee, in cash.

Are there any questions?

Payoff Instructions:

Your payoff in every period will be determined by the following formula (in cents):

$$400 - 0.5 * (\text{target number} - \text{guess})^2.$$

To make sure everyone understands the payoff formula, let us go through a few examples:

- Example 1: If your guess is exactly the same as the target number, then your payoff in that period will be: $400 - 0.5 * (0) = 400$ cents (4 €).
- Example 2: If the target number is 100.0 and your guess is 80.0, then your payoff will be:

$$400 - 0.5 * (100 - 80)^2 = 400 - 0.5 * (400) = 400 - 200 = 200 \text{ cents (2 €)}.$$

- Example 3: If the target number is -60.0 and your guess is -20.0 , then your payoff will be:

$$400 - 0.5 * (-60.0 - (-20.0))^2 = 400 - 0.5 * (20.0 - 60.0)^2 = 400 - 0.5 * (1600) = 400 - 800 = -400 \text{ cents } (-4 \text{ €})$$

Please note three things about how your payoff will be determined.

First, a simple way to think about your payoff is that you start off with 400 cents, but lose money for the distance between your guess and the target number. Therefore, you want your guess to be as close as possible to the target number.

Second, remember that you will only receive payment for one of the 16 periods in a round. Since you will not know which period this will be until the end of the experiment, every period is equally likely to count towards determining your earnings.

Third, notice that it is possible to lose money in a period if your guess is far from the target number. If your payoff in a period is negative, and if that period is selected to count for that round, then the amount you lose will be subtracted from the amount that you accumulate in other rounds. This means that you may lose some of your participation fee if the sum of your payoffs for all 4 rounds is negative. However, we have designed the experiment so that this is very unlikely to be the case if you make decisions carefully. If you notice that you are accumulating negative payoffs regularly, please raise your hand so that we can make sure that you understand how payoffs are determined. Also, please double-check your guess before confirming it, since entering it incorrectly may cause you to lose a significant amount of money.

Are there any questions about your payoffs?

Quiz (Part 1)

Before proceeding to the experiment, we would like to ask you to answer a few questions to make sure that everyone understands the instructions. Please answer each of the following questions. Once you have completed the quiz, please wait for the experimenter to ask you to proceed.

1. The experiment consists of 4 rounds. Each round consists of 16 periods.
True False
2. At the beginning of each round, everyone in your group will be shown the target number.
True False
3. The target number in a round will be the same for everyone in your group.
True False
4. The private estimate you receive will be the same for everyone in your group.
True False

5. In each period, your payoff will be based on how close your guess is to the target number.
True False

6. At the end of every period, each of you will observe the guess of at least one other person in your group.
True False

Quiz (Part 2)

For the next three questions, remember that the formula for your payoff in a period is:
 $400 - 0.5 * (\text{target number} - \text{guess})^2$.

7. Suppose that in a period your guess is 50.0. If the target number is 50.0, then your payoff for that period is:

8. Suppose that in a period your guess is 200.00. If the target number is 190.0, then your payoff for that period is:

9. Suppose that in a period your guess is 500.00. If the target number is 470.0, then your payoff for that period is:

Final Instructions:

We are now ready to begin the experiment. We will proceed through 4 rounds, each of which will consist of 16 periods. At the end of each round, you will find out the target number as well as your payoff for each of the 16 periods. At the end of the experiment, one of these periods will be selected at random for each round to determine your earnings.

At the beginning of each round, please record your private estimate on your record sheet. Then, in each period, please also record your guess and the information you receive regarding the guesses of others.

If you have a question from this point on, please raise your hand and wait for the experimenter.

APPENDIX B: Target Values and Private Signals

Information set 1:

Target Value	Private Signals			
	Agent 1	Agent 2	Agent 3	Agent 4
81.405	63.884	85.355	93.682	78.521
-21.6	-3.518	-27.486	-51.65	-12.146
124.16	119.512	95.831	146.349	153.696
6.327	35.719	2.553	25.543	-11.991

Information Set 2:

Target Value	Private Signals			
	Agent 1	Agent 2	Agent 3	Agent 4
-62.072	-48.463	-92.322	-85.279	-42.989
-124.01	-129.756	-148.149	-123.711	-97.887
380.12	404.631	364.22	351.162	390.095
-294.56	-264.788	-319.406	-284.408	-290.207

FIGURES

Figure 3: Proportion of Groups with Variance Less than 4




Figure 4: Proportion of Groups with Variance Less than 4
(By InfoSet)




Figure 5: Proportion of Groups with Variance Less than 4
(By Block)




Figure 6: Bias in Average Median Choice
(Final 4 Periods of each Block, Information Set 1)




Figure 7: Bias in Average Median Choice
(Final 4 Periods of each Block, Information Set 2)

