Trust and Manipulation in Social Networks

Manuel Förster[*]   Ana Mauleon[†]

Vincent J. Vannetelbosch[‡]

July 17, 2014

Abstract

We investigate the role of manipulation in a model of opinion formation. Agents repeatedly communicate with their neighbors in the social network, can exert effort to manipulate the trust of others, and update their opinions about some common issue by taking weighted averages of neighbors' opinions. The incentives to manipulate are given by the agents' preferences. We show that manipulation can modify the trust structure and lead to a connected society. Manipulation fosters opinion leadership, but the manipulated agent may even gain influence on the long-run opinions. Finally, we investigate the tension between information aggregation and spread of misinformation.

Keywords: Social networks; Trust; Manipulation; Opinion leadership; Consensus; Wisdom of crowds.

JEL classification: D83; D85; Z13.

---

[*]CEREC, Saint-Louis University – Brussels, Belgium. E-mail: manuel.forster@usaintlouis.be.

[†]CEREC, Saint-Louis University – Brussels; CORE, University of Louvain, Louvain-la-Neuve, Belgium. E-mail: ana.mauleon@usaintlouis.be.

[‡]CORE, University of Louvain, Louvain-la-Neuve; CEREC, Saint-Louis University – Brussels, Belgium. E-mail: vincent.vannetelbosch@uclouvain.be.

1   Introduction

Individuals often rely on social connections (friends, neighbors and coworkers as well as political actors and news sources) to form beliefs or opinions on various economic, political or social issues. Every day individuals make decisions on the basis of these beliefs. For instance, when an individual goes to the polls, her choice to vote for one of the candidates is influenced by her friends and peers, her distant and close family members, and some leaders that she listens to and respects. At the same time, the support of others is crucial to enforce interests in society. In politics, majorities are needed to pass laws and in companies, decisions might be taken by a hierarchical superior. It is therefore advantageous for individuals to increase their influence on others and to *manipulate* the way others form their beliefs. This behavior is often referred to as lobbying and widely observed in society, especially in politics.[1] Hence, it is important to understand how beliefs and behaviors evolve over time when individuals can manipulate the trust of others. Can manipulation enable a segregated society to reach a consensus about some issue of broad interest? How long does it take for beliefs to reach consensus when agents can manipulate others? Can manipulation lead a society of agents who communicate and update naïvely to more efficient information aggregation?

We consider a model of opinion formation where agents repeatedly communicate with their neighbors in the social network, can exert some effort to manipulate the trust of others, and update their opinions taking weighted averages of neighbors' opinions. At each period, first one agent is selected randomly and can exert effort to manipulate the social trust of an agent of her choice. If she decides to provide some costly effort to manipulate another agent, then the manipulated agent weights relatively more the belief of the agent who manipulated her when updating her beliefs. Second, all agents communicate with their neighbors and update their beliefs using the DeGroot updating rule, see DeGroot (1974). This updating process is simple: using her (possibly manipulated) weights, an agent's new belief is the weighted average of her neighbors' beliefs (and possibly her own belief)

---

[1]See Gullberg (2008) for lobbying on climate policy in the European Union, and Austen-Smith and Wright (1994) for lobbying on US Supreme Court nominations.

from the previous period. When agents have no incentives to manipulate each other, the model coincides with the classical DeGroot model of opinion formation.

The DeGroot updating rule assumes that agents are boundedly rational, failing to adjust correctly for repetitions and dependencies in information that they hear multiple times. Since social networks are often fairly complex, it seems reasonable to use an approach where agents fail to update beliefs correctly.[2] Chandrasekhar et al. (2012) provide evidence from a framed field experiment that DeGroot "rule of thumb" models best describe features of empirical social learning. They run a unique lab experiment in the field across 19 villages in rural Karnataka, India, to discriminate between the two leading classes of social learning models – Bayesian learning models versus DeGroot models.[3] They find evidence that the DeGroot model better explains the data than the Bayesian learning model at the network level.[4] At the individual level, they find that the DeGroot model performs much better than Bayesian learning in explaining the actions of an individual given a history of play.[5]

Manipulation is modeled as a communicative or interactional practice, where the manipulating agent exercises some control over the manipulated agent against her will. In this sense, manipulation is illegitimate, see Van Dijk (2006). Notice that manipulating the trust of other agents (instead of the opinions directly) can be seen as an attempt to influence their opinions in the medium- or even long-run since they will continue to use these manipulated weights in the future.[6] Agents only engage in

---

[2]Choi et al. (2012) report an experimental investigation of learning in three-person networks and find that already in simple three-person networks people fail to account for repeated information. They argue that the Quantal Response Equilibrium (QRE) model can account for the behavior observed in the laboratory in a variety of networks and informational settings.

[3]Notice that in order to compare the two concepts, they study DeGroot action models, i.e., agents take an action after aggregating the actions of their neighbors using the DeGroot updating rule.

[4]At the network level (i.e., when the observational unit is the sequence of actions), the Bayesian learning model explains 62% of the actions taken by individuals while the degree weighting DeGroot model explains 76% of the actions taken by individuals.

[5]At the individual level (i.e., when the observational unit is the action of an individual given a history), both the degree weighting and the uniform DeGroot model largely outperform Bayesian learning models.

[6]In our approach, the opinion of the manipulated agent is only affected indirectly through the manipulated trust weights. Therefore, her opinion continues to be affected

manipulation if it is worth the effort. They face a trade-off between their increase in satisfaction with the opinions (and possibly the trust itself) of the other agents and the cost of manipulation. In examples, we will frequently use a utility model where agents prefer each other agent's opinion one step ahead to be as close as possible to their current opinion. This reflects the idea that the support of others is necessary to enforce interests. Agents will only engage in manipulation when it brings the opinion (possibly several steps ahead) of the manipulated agent sufficiently closer to their current opinion compared to the cost of doing so. In our view, this constitutes a natural way to model lobbying incentives.

We first show that manipulation can modify the trust structure. If the society is split up into several disconnected clusters of agents and there are also some agents outside these clusters, then the latter agents might connect different clusters by manipulating the agents therein. Such an agent, previously outside any of these clusters, would not only get influential on the agents therein, but also serve as a bridge and connect them. As we show by means of an example, this can lead to a connected society, and thus, make the society reaching a consensus.

Second, we analyze the long-run beliefs and show that manipulation fosters opinion leadership in the sense that the manipulating agent always increases her influence on the long-run beliefs. For the other agents, this is ambiguous and depends on the social network. Surprisingly, the manipulated agent may thus even gain influence on the long-run opinions. As a consequence, the expected change of influence on the long-run beliefs is ambiguous and depends on the agents' preferences and the social network. We also show that a definitive trust structure evolves in the society and, if the satisfaction of agents only depends on the current and future opinions and not directly on the trust, manipulation will come to an end and they reach a consensus (under some weak regularity condition). At some point, opinions become too similar to be manipulated. Furthermore, we discuss the speed of convergence and note that manipulation can accelerate or slow down convergence. In particular, in sufficiently homophilic societies, i.e., societies where agents tend to trust those agents who are similar to them,

---

by the manipulation in the following periods, but the extent might be diminished by further manipulations.

and where costs of manipulation are rather high compared to its benefits, manipulation accelerates convergence if it decreases homophily and otherwise it slows down convergence.

Finally, we investigate the tension between information aggregation and spread of misinformation. We find that if manipulation is rather costly and the agents underselling their information gain and those overselling their information lose overall influence (i.e., influence in terms of their initial information), then manipulation reduces misinformation and agents converge jointly to more accurate opinions about some underlying true state. In particular, this means that an agent for whom manipulation is cheap can severely harm information aggregation.

There is a large and growing literature on learning in social networks. Models of social learning either use a Bayesian perspective or exploit some plausible rule of thumb behavior.[7] We consider a model of non-Bayesian learning over a social network closely related to DeGroot (1974), DeMarzo et al. (2003), Golub and Jackson (2010) and Acemoglu et al. (2010). De-Marzo et al. (2003) consider a DeGroot rule of thumb model of opinion formation and they show that persuasion bias affects the long-run process of social opinion formation because agents fail to account for the repetition of information propagating through the network. Golub and Jackson (2010) study learning in an environment where agents receive independent noisy signals about the true state and then repeatedly communicate with each other. They find that all opinions in a large society converge to the truth if and only if the influence of the most influential agent vanishes as the society grows.[8]

Acemoglu et al. (2010) investigate the tension between information aggregation and spread of misinformation. They characterize how the presence of forceful agents affects information aggregation. Forceful agents influence the beliefs of the other agents they meet, but do not change their own opinions. Under the assumption that even forceful agents obtain some

---

[7]Acemoglu et al. (2011) develop a model of Bayesian learning over general social networks, and Acemoglu and Ozdaglar (2011) provide an overview of recent research on opinion dynamics and learning in social networks.

[8]Golub and Jackson (2012) examine how the speed of learning and best-response processes depend on homophily. They find that convergence to a consensus is slowed down by the presence of homophily but is not influenced by network density.

information from others, they show that all beliefs converge to a stochastic consensus. They quantify the extent of misinformation by providing bounds on the gap between the consensus value and the benchmark without forceful agents where there is efficient information aggregation.[9] Friedkin (1991) studies measures to identify opinion leaders in a model related to DeGroot. Recently, Büchel et al. (2012) develop a model of opinion formation where agents may state an opinion that differs from their true opinion because agents have preferences for conformity. They find that lower conformity fosters opinion leadership. In addition, the society becomes wiser if agents who are well informed are less conform, while uninformed agents conform more with their neighbors.

Furthermore, Watts (2014) studies the influence of social networks on correct voting. Agents have beliefs about each candidate's favorite policy and update their beliefs based on the favorite policies of their neighbors and on whom the latter support. She finds that political agreement in an agent's neighborhood facilitates correct voting, i.e., voting for the candidate who's favorite policy is closest to his own favorite policy. Our paper is also related to the literature on lobbying as costly signaling, e.g., Austen-Smith and Wright (1994); Esteban and Ray (2006). These papers do not consider networks and model lobbying as providing one-shot costly signals to decision-makers in order to influence a policy decision.[10]

To the best of our knowledge we are the first allowing agents to manipulate the trust of others in social networks and we find that the implications of manipulation are non-negligible for opinion leadership, reaching a consensus, and aggregating dispersed information.

The paper is organized as follows. In Section 2 we introduce the model of opinion formation. In Section 3 we show how manipulation can change the trust structure of society. Section 4 looks at the long-run effects of manipulation. In Section 5 we investigate how manipulation affects the extent of misinformation in society. Section 6 concludes. The proofs are

---

[9]In contrast to the averaging model, Acemoglu et al. (2010) have a model of pairwise interactions. Without forceful agents, if a pair meets two periods in a row, then in the second meeting there is no information to exchange and no change in beliefs takes place.

[10]Notice that we study how (repeated) manipulation and lobbying affect public opinion (and potientially single decision-makers) in the long-run, but do not model explicitly any decision-making process.

presented in Appendix A.

2   Model and Notation

Let $\mathcal{N} = \{1, 2, \ldots, n\}$ be the set of agents who have to take a decision on some issue and repeatedly communicate with their neighbors in the social network. Each agent $i \in \mathcal{N}$ has an initial *opinion* or *belief* $x_i(0) \in \mathbb{R}$ about the issue and an initial vector of *social trust* $m_i(0) = (m_{i1}(0), m_{i2}(0), \ldots, m_{in}(0))$, with $0 \leq m_{ij}(0) \leq 1$ for all $j \in \mathcal{N}$ and $\sum_{j \in \mathcal{N}} m_{ij}(0) = 1$, that captures how much attention agent i pays (initially) to each of the other agents. More precisely, $m_{ij}(0)$ is the initial weight or trust that agent i places on the current belief of agent j in forming her updated belief. For $i = j$, $m_{ii}(0)$ can be interpreted as how much agent i is confident in her own initial opinion.

At period $t \in \mathbb{N}$, the agents' beliefs are represented by the vector $x(t) = (x_1(t), x_2(t), \ldots, x_n(t))' \in \mathbb{R}^n$ and their social trust by the matrix $M(t) = (m_{ij}(t))_{i,j \in \mathcal{N}}$.[11] First, one agent is chosen (probability $1/n$ for each agent) to meet and to have the opportunity to manipulate an agent of her choice. If agent $i \in \mathcal{N}$ is chosen at t, she can decide which agent j to meet and furthermore how much effort $\alpha \geq 0$ she would like to exert on j. We write $E(t) = (i; j, \alpha)$ when agent i is chosen to manipulate at t and decides to exert effort α on j. The decision of agent i leads to the following updated trust weights of agent j:

$$m_{jk}(t+1) = \begin{cases} m_{jk}(t)/(1+\alpha) & \text{if } k \neq i \\ (m_{jk}(t) + \alpha)/(1+\alpha) & \text{if } k = i \end{cases}.$$

The trust of j in i increases with the effort i invests and all trust weights of j are normalized. Notice that we assume for simplicity that the trust of j in an agent other than i decreases by the factor $1/(1+\alpha)$, i.e., the absolute decrease in trust is proportional to its level. If i decides not to invest any effort, the trust matrix does not change. We denote the resulting updated trust matrix by $M(t+1) = [M(t)](i; j, \alpha)$.

Agent i decides on which agent to meet and on how much effort to exert

---

[11] We denote the transpose of a vector (matrix) x by x'.

according to her utility function

$$u_i\big(M(t), x(t); j, \alpha\big) = v_i\big([M(t)](i; j, \alpha), x(t)\big) - c_i(j, \alpha),$$

where $v_i\big([M(t)](i; j, \alpha), x(t)\big)$ represents her *satisfaction* with the other agents' opinions and trust resulting from her decision (j, α) and $c_i(j, \alpha)$ represents its *cost*. We assume that v_i is continuous in all arguments and that for all $j \neq i$, $c_i(j, \alpha)$ is strictly increasing in $\alpha \geq 0$, continuous and strictly convex in $\alpha > 0$, and that $c_i(j, 0) = 0$. Note that these conditions ensure that there is always an optimal level of effort $\alpha^*(j)$ given agent i decided to manipulate j.[12] Agent i's optimal choice is then $(j^*, \alpha^*(j^*))$ such that $j^* \in \text{argmax}_{j \neq i}\, u_i\big(M(t), x(t); j, \alpha^*(j)\big)$.

Secondly, all agents communicate with their neighbors and update their beliefs using the updated trust weights:

$$x(t + 1) = [x(t)](i; j, \alpha) = M(t + 1)x(t) = [M(t)](i; j, \alpha)x(t).$$

In the sequel, we will often simply write $x(t + 1)$ and omit the dependence on the agent selected to manipulate and her choice (j, α). We can rewrite this equation as $x(t + 1) = \overline{M}(t + 1)x(0)$, where $\overline{M}(t + 1) = M(t + 1)M(t)\cdots M(1)$ (and $\overline{M}(t) = I_n$ for $t < 1$, where I_n is the $n \times n$ identity matrix) denotes the *overall trust matrix*.

Now, let us give some examples of satisfaction functions that fulfill our assumptions.

Example 1 (Satisfaction functions)**.**

(i) Let $\gamma \in \mathbb{N}$ and

$$v_i\big([M(t)](i; j, \alpha), x(t)\big) = -\frac{1}{n-1} \sum_{k \neq i} \Big(x_i(t) - \big(M(t+1)^\gamma\, x(t)\big)_k\Big)^2,$$

where $M(t + 1) = [M(t)](i; j, \alpha)$. That is, agent i's objective is that each other agent's opinion γ periods ahead is as close as possible

---

[12]Note that for all j, $v_i(M(i; j, \alpha), x)$ is continuous in α and bounded from above since $v_i(\cdot, x)$ is bounded from above on the compact set $[0, 1]^{n \times n}$ for all $x \in \mathbb{R}^n$. In total, the utility is continuous in $\alpha > 0$ and since the costs are strictly increasing and strictly convex in $\alpha > 0$, there always exists an optimal level of effort, which might not be unique, though.

to her current opinion, disregarding possible manipulations in future periods.

(ii)

$$v_i\big([M(t)](i;j,\alpha), x(t)\big) = -\left(x_i(t) - \frac{1}{n-1}\sum_{k\neq i} x_k(t+1)\right)^2,$$

where $x_k(t+1) = \big([M(t)](i;j,\alpha)x(t)\big)_k$. That is, agent i wants to be close to the average opinion in society one period ahead, but disregards differences on the individual level.

We will frequently choose in examples the first satisfaction function with parameter $\gamma = 1$, together with a cost function that combines fixed costs and quadratic costs of effort.

Remark 1. If we choose satisfaction functions $v_i \equiv v$ for some constant v and all $i \in \mathcal{N}$, then agents do not have any incentive to exert effort and our model reverts to the classical model of DeGroot (1974).

We now introduce the notion of *consensus*. Whether or not a consensus is reached in the limit depends generally on the initial opinions.

Definition 1 (Consensus). We say that a group of agents $G \subseteq \mathcal{N}$ reaches a *consensus* given initial opinions $(x_i(0))_{i\in\mathcal{N}}$, if there exists $x(\infty) \in \mathbb{R}$ such that

$$\lim_{t\to\infty} x_i(t) = x(\infty) \text{ for all } i \in G.$$

3 The Trust Structure

We investigate how manipulation can modify the structure of interaction or trust in society. We first shortly recall some graph-theoretic terminology.[13] We call a group of agents $C \subseteq \mathcal{N}$ *minimal closed* at period t if these agents only trust agents inside the group, i.e., $\sum_{j\in C} m_{ij}(t) = 1$ for all $i \in C$, and if this property does not hold for a proper subset $C' \subsetneq C$. The set of minimal closed groups at period t is denoted $\mathcal{C}(t)$ and is called the *trust structure*.

---

[13]See Golub and Jackson (2010).

A walk at period t of length K is a sequence of agents $i_1, i_2, \ldots, i_{K+1}$ such that $m_{i_k, i_{k+1}}(t) > 0$ for all $k = 1, 2, \ldots, K$. A walk is a *path* if all agents are distinct. A *cycle* is a walk that starts and ends in the same agent. A cycle is *simple* if only the starting agent appears twice in the cycle. We say that a minimal closed group of agents $C \in \mathcal{C}(t)$ is *aperiodic* if the greatest common divisor[14] of the lengths of simple cycles involving agents from C is 1.[15] Note that this is fulfilled if $m_{ii}(t) > 0$ for some $i \in C$.

At each period t, we can decompose the set of agents $\mathcal{N}$ into minimal closed groups and agents outside these groups, the *rest of the world*, $R(t)$:

$$\mathcal{N} = \bigcup_{C \in \mathcal{C}(t)} C \cup R(t).$$

Within minimal closed groups, all agents interact indirectly with each other, i.e., there is a path between any two agents. We say that the agents are *strongly connected*. For this reason, minimal closed groups are also called strongly connected and closed groups, see Golub and Jackson (2010). Moreover, agent $i \in \mathcal{N}$ is part of the rest of the world $R(t)$ if and only if there is a path at period t from her to some agent in a minimal closed group $C \not\ni i$.

We say that a manipulation at period t does not change the trust structure if $\mathcal{C}(t+1) = \mathcal{C}(t)$. It also implies that $R(t+1) = R(t)$. We find that manipulation changes the trust structure when the manipulated agent belongs to a minimal closed group and additionally the manipulating agent does not belong to this group, but may well belong to another minimal closed group. In the latter case, the group of the manipulated agent is disbanded since it is not anymore closed and its agents join the rest of the world. However, if the manipulating agent does not belong to a minimal closed group, the effect on the group of the manipulated agent depends on the trust structure. Apart from being disbanded, it can also be the case that the manipulating agent and possibly others from the rest of the world join the group of the manipulated agent.

Proposition 1. *Suppose that $E(t) = (i; j, \alpha)$, $\alpha > 0$, at period t.*

---

[14]For a set of integers $S \subseteq \mathbb{N}$, $\gcd(S) = \max \{k \in \mathbb{N} \mid m/k \in \mathbb{N} \text{ for all } m \in S\}$ denotes the greatest common divisor.

[15]Note that if one agent in a simple cycle is from a minimal closed group, then so are all.

(i) *Let $i \in \mathcal{N}, j \in R(t)$ or $i, j \in C \in \mathcal{C}(t)$. Then, the trust structure does not change.*

(ii) *Let $i \in C \in \mathcal{C}(t)$ and $j \in C' \in \mathcal{C}(t) \backslash \{C\}$. Then, C' is disbanded, i.e.,*
$$\mathcal{C}(t+1) = \mathcal{C}(t) \backslash \{C'\}.$$

(iii) *Let $i \in R(t)$ and $j \in C \in \mathcal{C}(t)$.*

    (a) *Suppose that there exists no path from i to k for any $k \in \cup_{C' \in \mathcal{C}(t) \backslash \{C\}} C'$. Then, $R' \cup \{i\}$ joins C, i.e.,*
$$\mathcal{C}(t+1) = \mathcal{C}(t) \backslash \{C\} \cup \{C \cup R' \cup \{i\}\},$$

    where $R' = \{l \in R(t) \backslash \{i\} \mid \text{there is a path from } i \text{ to } l\}$.

    (b) *Suppose that there exists $C' \in \mathcal{C}(t) \backslash \{C\}$ such that there exists a path from i to some $k \in C'$. Then, C is disbanded.*

All proofs can be found in Appendix A. The following example shows that manipulation can enable a society to reach a consensus due to changes in the trust structure.

Example 2 (Consensus due to manipulation). Take $\mathcal{N} = \{1, 2, 3\}$ and assume that

$$u_i\big(M(t), x(t); j, \alpha\big) = -\frac{1}{2} \sum_{k \neq i} \big(x_i(t) - x_k(t+1)\big)^2 - \big(\alpha^2 + 1/10 \cdot \mathbf{1}_{\{\alpha > 0\}}(\alpha)\big)$$

for all $i \in \mathcal{N}$. Notice that the first part of the utility is the satisfaction function in Example 1 part (i) with parameter $\gamma = 1$, while the second part, the costs of effort, combines fixed costs, here $1/10$, and quadratic costs of effort. Let $x(0) = (10, 5, -5)'$ be the vector of initial opinions and

$$M(0) = \begin{pmatrix} .8 & .2 & 0 \\ .4 & .6 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

be the initial trust matrix. Hence, $\mathcal{C}(0) = \{\{1, 2\}, \{3\}\}$. Suppose that first agent 1 and then agent 3 are drawn to meet another agent. Then, at period

11

0, agent 1's optimal decision is to exert $\alpha = 2.54$[16] effort on agent 3. The trust of the latter is updated to

$$m_3(1) = (.72, 0, .28),$$

while the others' trust does not change, i.e., $m_i(1) = m_i(0)$ for $i = 1, 2$, and the updated opinions become

$$x(1) = M(1)x(0) = (9, 7, 5.76)'.$$

Notice that the group of agent 3 is disbanded (see part (ii) of Proposition 1). In the next period, agent 3's optimal decision is to exert $\alpha = .75$ effort on agent 1. This results in the following updated trust matrix:

$$M(2) = \begin{pmatrix} .46 & .11 & .43 \\ .4 & .6 & 0 \\ .72 & 0 & .28 \end{pmatrix}.$$

Notice that agent 3 joins group $\{1, 2\}$ (see part (iii,a) of Proposition 1) and therefore, $\mathcal{N}$ is minimal closed, which implies that the group will reach a consensus, as we will see later on.

However, notice that if instead of agent 3 another agent is drawn in period 1, then agent 3 never manipulates since when finally she would have the opportunity, her opinion is already close to the others' opinions and therefore, she stays disconnected from them. Nevertheless, the agents would still reach a consensus in this case due to the manipulation at period 0. Since agent 3 trusts agent 1, she follows the consensus that is reached by the first two agents.

4 The Long-Run Dynamics

We now look at the long-run effects of manipulation. First, we study the consequences of a single manipulation on the long-run opinions of minimal closed groups. In this context, we are interested in the role of manipulation in opinion leadership. Secondly, we investigate the outcome of the

---

[16]Stated values are rounded to two decimals for clarity reasons.

influence process. Finally, we discuss how manipulation affects the speed of convergence of minimal closed groups and illustrate our results by means of an example.

4.1   Opinion Leadership

Typically, an agent is called *opinion leader* if she has substantial influence on the long-run beliefs of a group. That is, if she is among the most influential agents in the group. Intuitively, manipulating others should increase her influence on the long-run beliefs and thus foster opinion leadership.

To investigate this issue, we need a measure for how remotely agents are located from each other in the network, i.e., how directly agents trust other agents. For this purpose, we can make use of results from Markov chain theory. Let $(X_s^{(t)})_{s=0}^{\infty}$ denote the homogeneous Markov chain induced by the transition matrix $M(t)$. The agents are then interpreted as states of the Markov chain and the trust of i in j, $m_{ij}(t)$, is interpreted as the transition probability from state i to state j. Then, the *mean first passage time* from state i to state j is defined as $\mathbb{E}[\inf\{s \geq 0 \mid X_s^{(t)} = j\} \mid X_0^{(t)} = i]$. Given the current state of the Markov chain is i, the mean first passage time to j is the expected time it takes for the chain to reach state j.

In other words, the mean first passage time from i to j corresponds to the average (expected) length of a random walk on the weighted network $M(t)$ from i to j that takes each link with probability equal to the assigned weight.[17] This average length is small if the weights along short paths from i to j are high, i.e., if agent i trusts agent j rather directly. We therefore call this measure *weighted remoteness* of j from i.

Definition 2 (Weighted remoteness)**.** Take $i, j \in \mathcal{N}$, $i \neq j$. The *weighted remoteness* at period t of agent j from agent i is given by

$$r_{ij}(t) = \mathbb{E}[\inf\{s \geq 0 \mid X_s^{(t)} = j\} \mid X_0^{(t)} = i],$$

where $(X_s^{(t)})_{s=0}^{\infty}$ is the homogeneous Markov chain induced by $M(t)$.

---

[17]More precisely, it is a random walk on the state space $\mathcal{N}$ that, if currently in state k, travels to state l with probability $m_{kl}(t)$. The length of this random walk to j is the time it takes for it to reach state j.

The following remark shows that the weighted remoteness attains its minimum when i trusts solely j.

Remark 2. Take $i, j \in \mathcal{N}$, $i \neq j$.

(i) $r_{ij}(t) \geq 1$,

(ii) $r_{ij}(t) < +\infty$ if and only if there is a path from i to j, and, in particular, if $i, j \in C \in \mathcal{C}(t)$,

(iii) $r_{ij}(t) = 1$ if and only if $m_{ij}(t) = 1$.

To provide some more intuition, let us look at an alternative (implicit) formula for the weighted remoteness. Suppose that $i, j \in C \in \mathcal{C}(t)$ are two distinct agents in a minimal closed group. By part (ii) of Remark 2, the weighted remoteness is finite for all pairs of agents in that group. The unique walk from i to j with (average) length 1 is assigned weight (or has probability, when interpreted as a random walk) $m_{ij}(t)$. And the average length of walks to j that first pass through $k \in C\backslash\{j\}$ is $r_{kj}(t) + 1$, i.e., walks from i to j with average length $r_{kj}(t) + 1$ are assigned weight (have probability) $m_{ik}(t)$. Thus,

$$r_{ij}(t) = m_{ij}(t) + \sum_{k \in C\backslash\{j\}} m_{ik}(t)(r_{kj}(t) + 1) \ .$$

Finally, applying $\sum_{k \in C} m_{ik}(t) = 1$ leads to the following remark.

Remark 3. Take $i, j \in C \in \mathcal{C}(t)$, $i \neq j$. Then,

$$r_{ij}(t) = 1 + \sum_{k \in C\backslash\{j\}} m_{ik}(t)r_{kj}(t).$$

Note that computing the weighted remoteness using this formula amounts to solving a linear system of $|C|(|C| - 1)$ equations, which has a unique solution.

We denote by $\pi(C; t)$ the probability vector of the agents' influence on the final consensus of their group $C \in \mathcal{C}(t)$ at period t, given that the group is aperiodic and the trust matrix does not change any more.[18] In this case,

---

[18]In the language of Markov chains, $\pi(C; t)$ is known as the unique stationary distribution of the aperiodic communication class C. Without aperiodicity, the class might fail to converge to consensus.

the group converges to

$$x(\infty) = \pi(C;t)'\, x(t)|_C = \sum_{i \in C} \pi_i(C;t) x_i(t),$$

where $x(t)|_C = (x_i(t))_{i \in C}$ is the restriction of $x(t)$ to agents in C. In other words, $\pi_i(C;t)$, $i \in C$, is the influence weight of agent i's opinion at period t, $x_i(t)$, on the consensus of C. Notice that the influence vector $\pi(C;t)$ depends on the trust matrix $M(t)$ and therefore it changes with manipulation. A higher value of $\pi_i(C;t)$ corresponds to more influence of agent i on the consensus. Each agent in a minimal closed group has at least some influence on the consensus: $\pi_i(C;t) > 0$ for all $i \in C$.[19]

We now turn back to the long-run consequences of manipulation and thus, opinion leaders. We restrict our analysis to the case where both the manipulating and the manipulated agent are in the same minimal closed group. Since in this case the trust structure is preserved we can compare the influence on the long-run consensus of the group before and after manipulation.

Proposition 2. *Suppose that at period t, group $C \in \mathcal{C}(t)$ is aperiodic and $E(t) = (i;j,\alpha)$, $i,j \in C$. Then, aperiodicity is preserved and the influence of agent $k \in C$ on the final consensus of her group changes as follows,*

$$\pi_k(C;t+1) - \pi_k(C;t) =$$
$$\begin{cases} \alpha/(1+\alpha)\pi_i(C;t)\pi_j(C;t+1)\sum_{l \in C\setminus\{i\}} m_{jl}(t)r_{li}(t) & \text{if } k = i \\ \alpha/(1+\alpha)\pi_k(C;t)\pi_j(C;t+1)\left(\sum_{l \in C\setminus\{k\}} m_{jl}(t)r_{lk}(t) - r_{ik}(t)\right) & \text{if } k \neq i \end{cases}.$$

Corollary 1. *Suppose that at period t, group $C \in \mathcal{C}(t)$ is aperiodic and $E(t) = (i;j,\alpha)$, $i,j \in C$. If $\alpha > 0$, then*

(i) *agent i strictly increases her long-run influence, $\pi_i(C;t+1) > \pi_i(C;t)$,*

(ii) *any other agent $k \neq i$ of the group can either gain or lose influence, depending on the trust matrix. She gains if and only if*

$$\sum_{l \in C\setminus\{k,i\}} m_{jl}(t)\big(r_{lk}(t) - r_{ik}(t)\big) > m_{jk}(t)r_{ik}(t),$$

---

[19]See Golub and Jackson (2010).

(iii) agent $k \neq i, j$ loses influence for sure if j trusts solely her, i.e., $m_{jk}(t) = 1$.

Proposition 2 tells us that the change in long-run influence for any agent k depends on the effort agent i exerts to manipulate agent j, agent k's current long-run influence and the future long-run influence of the manipulated agent j. In particular, the magnitude of the change increases with i's effort, and it is zero if agent i does not exert any effort. Furthermore, notice that dividing both sides by agent k's current long-run influence, $\pi_k(C; t)$, yields the relative change in her long-run influence.

When agent $k = i$, we find that this change is strictly positive whenever she exerts some effort. In this sense, manipulation fosters opinion leadership. It is large if the weighted remoteness of i from agents (other than i) that are significantly trusted by j is large. To understand this better, notice that the long-run influence of an agent depends on how much she is trusted by agents that are trusted. Or, in other words, an agent is influential if she is influential on other influential agents. Thus, there is a direct gain of influence due to an increase of trust from j and an indirect loss of influence (that is always dominated by the direct gain) due to a decrease of trust from j faced by agents that (indirectly) trust i. This explains why it is better for i if agents facing a large decrease of trust from j (those trusted much by j) do not (indirectly) trust i much, i.e., i has a large weighted remoteness from them.

For any other agent $k \neq i$, it turns out that the change can be positive or negative. It is positive if, broadly speaking, j does not trust k a lot, the weighted remoteness of k from i is small and furthermore the weighted remoteness of k from agents (other than i) that are significantly trusted by j is larger than that from i. In other words, it is positive if the manipulating agent, who gains influence for sure, (indirectly) trusts agent k significantly (small weighted remoteness of k from i), k does not face a large decrease of trust from j and those agents facing a large decrease from j (those trusted much by j) (indirectly) trust k less than i does.

Notice that for any agent $k \neq i, j$, this is a trade-off between an indirect gain of trust due to the increase of trust that i obtains from j, on the one hand, and an indirect loss of influence due to a decrease of trust from j

faced by agents that (indirectly) trust k as well as the direct loss of influence due to a decrease of trust from j, on the other hand. In the extreme case where j only trusts k, the direct loss of influence dominates the indirect gain of influence for sure.

In particular, it means that even the manipulated agent j can gain influence. In a sense, such an agent would like to be manipulated because she trusts the "wrong" agents. For agent j, being manipulated is positive if her weighted remoteness from agents she trusts significantly is large and furthermore, her weighted remoteness from i is small. Hence, it is positive if the manipulating agent (indirectly) trusts her significantly (small weighted remoteness from i) and agents facing a large decrease of trust from her (those she trusts) do not (indirectly) trust her much. Here, the trade-off is between the indirect gain of trust due to the increase of trust that i obtains from her and the indirect loss of influence due to a decrease of trust from her faced by agents that (indirectly) trust her. Note that the gain of influence is particularly large if the manipulating agent trusts j significantly.

The next example shows that indeed in some situations an agent can gain from being manipulated in the sense that her influence on the long-run beliefs increases.

Example 3 (Being manipulated can increase influence)**.** Take $\mathcal{N} = \{1, 2, 3\}$ and assume that

$$M(0) = \begin{pmatrix} .25 & .25 & .5 \\ .5 & .5 & 0 \\ .4 & .5 & .1 \end{pmatrix}$$

is the initial trust matrix. Notice that $\mathcal{N}$ is minimal closed. Suppose that agent 1 has the opportunity to meet another agent and decides to exert effort $\alpha > 0$ on agent 3. Then, from Proposition 2, we get

$$\pi_3(\mathcal{N}; 1) - \pi_3(\mathcal{N}; 0) = \frac{\alpha}{1 + \alpha} \pi_3(\mathcal{N}; 0) \pi_3(\mathcal{N}; 1) \sum_{l=1,2} m_{3l}(0) r_{l3}(0) - r_{13}(0)$$

$$= \frac{\alpha}{1 + \alpha} \pi_3(\mathcal{N}; 0) \pi_3(\mathcal{N}; 1) \frac{7}{10} > 0,$$

since $\pi_3(\mathcal{N}; 0), \pi_3(\mathcal{N}; 1) > 0$. Hence, being manipulated by agent 1 increases agent 3's influence on the long-run beliefs. The reason is that, initially, she trusts too much agent 2 – an agent that does not trust her at

17

all. She gains influence from agent 1's increase of influence on the long-run beliefs since this agent trusts her. In other words, after being manipulated she is trusted by an agent that is trusted more.

Furthermore, we can use Proposition 2 to compare the expected influence on the long-run consensus of society before and after manipulation when all agents are in the same minimal closed group.[20] For this result we need to slightly change our notation. We denote the decision of agent $i \in \mathcal{N}$ when she is selected to meet another agent by $\big(j(i), \alpha(i; j(i))\big)$, i.e., agent i decides to exert effort $\alpha(i; j(i))$ on agent $j(i)$.

Corollary 2. *Suppose that at period t, $\mathcal{C}(t) = \{\mathcal{N}\}$ and that $\mathcal{N}$ is aperiodic. Then, aperiodicity is preserved and, in expectation, the influence of agent $k \in \mathcal{N}$ on the final consensus of the society changes as follows from period t to $t+1$,*

$$
\mathbb{E}[\pi_k(\mathcal{N}; t+1) - \pi_k(\mathcal{N}; t) \mid M(t), x(t)] =
$$
$$
\frac{\pi_k(\mathcal{N}; t)}{n} \left[\sum_{i \in \mathcal{N}} \left(\frac{\alpha(i; j(i))}{1 + \alpha(i; j(i))} \pi_{j(i)}(\mathcal{N}; t+1) \sum_{l \neq k} m_{j(i)l}(t) r_{lk}(t) \right) - \right.
$$
$$
\left. \sum_{i \neq k} \frac{\alpha(i; j(i))}{1 + \alpha(i; j(i))} \pi_{j(i)}(\mathcal{N}; t+1) r_{ik}(t) \right].
$$

Notice that an agent gains long-run influence in expectation if and only if the term in the square brackets is positive. For this to hold, it is necessary that $\alpha(i; j(i)) > 0$ for some $i \in \mathcal{N}$ at period t. Moreover, it follows from Corollary 1 part (i) that $\alpha(k; j(k)) > 0$ and $\alpha(i; j(i)) = 0$ for all $i \neq k$ at period t (i.e., only agent k would manipulate if she was selected at t) is a sufficient condition for that she gains influence in expectation. The reason is that agent k gains influence for sure when she manipulates herself, and since no other agent manipulates when selected, she gains in expectation. Notice that by dividing both sides by agent k's current long-run influence, $\pi_k(C; t)$, we get the expected relative change in her long-run influence.

---

[20]Notice that if not all agents are in the same minimal closed group, then the group in question could be disbanded with some probability and hence would not anymore reach a consensus.

4.2 Convergence

We now determine where the process finally converges to. First, we look at the case where all agents are in the same minimal closed group. Given the group is aperiodic, we show that if the satisfaction level only depends on the opinions (before and after manipulation), i.e., a change in trust that does not affect opinions does not change the satisfaction of an agent, and if there is a fixed cost for exerting effort, then manipulation comes to an end, eventually. At some point, opinions in the society become too similar to be manipulated. Second, we determine the final consensus the society converges to.

Lemma 1. *Suppose that $\mathcal{C}(0) = \{\mathcal{N}\}$ and that $\mathcal{N}$ is aperiodic. If for all $i, j \in \mathcal{N}$ and $\alpha > 0$,*

(i) $v_i\big(M(i; j, \alpha), x\big) - v_i\big(M(i; j, 0), x\big) \to 0$ *if* $\|x(i; j, \alpha) - x(i; j, 0)\| \to 0$, *and*

(ii) $c_i(j, \alpha) \geq \underline{c} > 0$,

then, there exists an almost surely finite stopping time τ such that from period $t = \tau$ on there is no more manipulation, where $\|\cdot\|$ is any norm on $\mathbb{R}^n$.[21] The society converges to the random variable

$$x(\infty) = \pi(\mathcal{N}; \tau)' \, \overline{M}(\tau - 1) \, x(0).$$

Now, we turn to the general case of any trust structure. We show that after a finite number of periods, the trust structure settles down. Then, it follows from the above result that, under the beforementioned conditions, manipulation within the minimal closed groups that have finally been formed comes to an end. We also determine the final consensus opinion of each aperiodic minimal closed group.

Proposition 3.

(i) *There exists an almost surely finite stopping time τ such that for all $t \geq \tau$, $\mathcal{C}(t) = \mathcal{C}(\tau)$.*

---

[21]In our context, this means that τ is a random variable such that the event $\tau = t$ only depends on which agents were selected to meet another agent at periods $1, 2, \ldots, t$, and furthermore τ is almost surely finite, i.e., the event $\tau < +\infty$ has probability 1.

(*ii*) *If $C \in \mathcal{C}(\tau)$ is aperiodic and for all $i, j \in C$, $\alpha > 0$,*

(1) $v_i\big(M(i; j, \alpha), x\big) - v_i\big(M(i; j, 0), x\big) \to 0$ *if* $\|x(i; j, \alpha) - x(i; j, 0)\| \to 0$, *and*

(2) $c_i(j, \alpha) \geq \underline{c} > 0$,

then, there exists an almost surely finite stopping time $\widehat{\tau} \geq \tau$ such that at all periods $t \geq \widehat{\tau}$, agents in C are not manipulated. Moreover, they converge to the random variable

$$x(\infty) = \pi(C; \widehat{\tau})' \, M(\widehat{\tau} - 1)\big|_C \, M(\widehat{\tau} - 2)\big|_C \cdots M(1)\big|_C \, x(0)\big|_C \, .$$

In what follows we use τ and $\widehat{\tau}$ in the above sense. We denote by $\overline{\pi}_i(C; t)$ the *overall influence* of agent i's initial opinion on the consensus of group C at period t given no more manipulation affecting C takes place. The overall influence is implicitly given by Proposition 3.

Corollary 3. *The overall influence of the initial opinion of agent $i \in \mathcal{N}$ on the consensus of an aperiodic group $C \in \mathcal{C}(\tau)$ is given by*

$$\overline{\pi}_i(C; \widehat{\tau}) = \begin{cases} \big(\pi(C; \widehat{\tau})' \, M(\widehat{\tau} - 1)\big|_C \, M(\widehat{\tau} - 2)\big|_C \cdots M(1)\big|_C\big)_i & \text{if } i \in C \\ 0 & \text{if } i \notin C \end{cases} \, .$$

It turns out that an agent outside a minimal closed group that has finally formed can never have any influence on its consensus opinion.

4.3   Speed of Convergence

We have seen that within an aperiodic minimal closed group $C \in \mathcal{C}(t)$ agents reach a consensus given that the trust structure does not change anymore. This means that their opinions converge to a common opinion. By *speed of convergence* we mean the time that this convergence takes. That is, it is the time it takes for the expression

$$|x_i(t) - x_i(\infty)|$$

to become small. It is well known that this depends crucially on the second largest eigenvalue $\lambda_2(C; t)$ of the trust matrix $M(t)\big|_C$, where $M(t)\big|_C =$

$(m_{ij}(t))_{i,j \in C}$ denotes the restriction of $M(t)$ to agents in C. Notice that $M(t)|_C$ is a stochastic matrix since C is minimal closed. The smaller the eigenvalue in absolute value, the faster the convergence to consensus (see Jackson, 2008).

Thus, the change in the second largest eigenvalue due to manipulation tells us whether the speed of convergence has increased or decreased. In this context, the concept of *homophily* is important, that is, the tendency of people to interact relatively more with those people who are similar to them.[22]

Definition 3 (Homophily). The *homophily* of a group of agents $G \subseteq \mathcal{N}$ at period t is defined as

$$\text{Hom}(G; t) = \frac{1}{|G|} \left(\sum_{i,j \in G} m_{ij}(t) - \sum_{i \in G, j \notin G} m_{ij}(t) \right).$$

The homophily of a group of agents is the normalized difference of their trust in agents inside and outside the group. Notice that a minimal closed group $C \in \mathcal{C}(t)$ attains the maximum homophily, $\text{Hom}(C; t) = 1$. Consider a *cut of society* $(S, \mathcal{N} \backslash S)$, $S \subseteq \mathcal{N}, S \neq \emptyset$, into two groups of agents S and $\mathcal{N} \backslash S$.[23] The next lemma establishes that manipulation across the cut decreases homophily, while manipulation within a group increases it.

Lemma 2. *Take a cut of society* $(S, \mathcal{N} \backslash S)$. *If $i \in \mathcal{N}$ manipulates $j \in S$ at period t, then*

(i) *the homophily of S (strictly) increases if $i \in S$ (and $\sum_{k \in S} m_{jk}(t) < 1$), and*

(ii) *the homophily of S (strictly) decreases if $i \notin S$ (and $\sum_{k \in S} m_{jk}(t) > 0$).*

Now, we come back to the speed of convergence. Given the complexity of the problem for $n \geq 3$, we consider an example of a two-agent soci-

---

[22]Notice that we do not model explicitly the characteristics that lead to homophily.

[23]There exist many different notions of homophily in the literature. Our measure is similar to the one used in Golub and Jackson (2012). We can consider the average homophily $(\text{Hom}(S; t) + \text{Hom}(\mathcal{N} \backslash S; t))/2$ with respect to the cut $(S, \mathcal{N} \backslash S)$ as a generalization of degree-weighted homophily to general weighted averages.

ety that suggests that homophily helps to explain the change in speed of convergence.

Example 4 (Speed of convergence with two agents). Take $\mathcal{N} = \{1, 2\}$ and suppose that at period t, $\mathcal{N}$ is minimal closed and aperiodic. Then, we have that $\lambda_2(\mathcal{N}; t) = m_{11}(t) - m_{21}(t) = m_{22}(t) - m_{12}(t)$. Therefore, we can characterize the change in the second largest eigenvalue as follows:

$$
\begin{aligned}
|\lambda_2(\mathcal{N}; t+1)| \leq |\lambda_2(\mathcal{N}; t)| &\Leftrightarrow |m_{11}(t+1) - m_{21}(t+1)| \leq |m_{11}(t) - m_{21}(t)| \\
&\Leftrightarrow |m_{22}(t+1) - m_{12}(t+1)| \leq |m_{22}(t) - m_{12}(t)|.
\end{aligned}
$$

It means that convergence is faster after manipulation if afterwards agents behave more similar, i.e., the trust both agents put on agent 1's opinion is more similar (which implies that also the trust they put on agent 2's opinion is more similar). Thus, if for instance

$$
m_{22}(t) > (1 + \alpha)m_{12}(t), \tag{1}
$$

then agent 1 manipulating agent 2 accelerates convergence. However, if $m_{22}(t) < m_{12}(t)$, it slows down convergence since manipulation increases the already existing tendency of opinions to oscillate. The more interesting case is the first one, though. We can write (1) as

$$
(1 + \alpha)\text{Hom}(\{1\}, t) + \text{Hom}(\{2\}, t) > \alpha,
$$

that is, manipulation accelerates convergence if there is sufficient aggregated homophily in the society and agent 1 does not exert too much effort.

The example shows that manipulation can speed up or slow down the convergence process. More important, it suggests that in a sufficiently homophilic society where exerting effort is rather costly, manipulation reducing homophily (i.e., across the cut, see Lemma 2) increases the speed of convergence. Notice that manipulation increasing homophily (i.e., within one of the groups separated by the cut) is not possible in this simple setting since both groups are singletons. However, it seems plausible that it would slow down convergence in homophilic societies.[24]

---

[24]In the above example, increasing homophily is attained by increasing the weight

4.4 Three-agents Example

Finally, let us consider an example with three agents to illustrate the results of this section. We use a utility model that is composed of the satisfaction function in Example 1 (i) and a cost function that combines fixed costs and quadratic costs of effort.

Example 5 (Three-agents society). Take $\mathcal{N} = \{1, 2, 3\}$ and assume that

$$u_i\big(M(t), x(t); j, \alpha\big) = -\frac{1}{2} \sum_{k \neq i} \big(x_i(t) - x_k(t+1)\big)^2 - \big(\alpha^2 + 1/10 \cdot \mathbf{1}_{\{\alpha > 0\}}(\alpha)\big)$$

for all $i \in \mathcal{N}$. Let $x(0) = (10, 5, 1)'$ be the vector of initial opinions and

$$M(0) = \begin{pmatrix} .6 & .2 & .2 \\ .1 & .4 & .5 \\ 0 & .6 & .4 \end{pmatrix}$$

be the initial trust matrix. Notice that this society is connected. The vector of initial long-run influence – and of long-run influence in the classical model without manipulation – is $\pi(\mathcal{N}; 0) = \pi_{\mathrm{cl}} = (.12, .46, .42)'$ and the initial speed of convergence is measured by $\lambda_2(\mathcal{N}; 0) = \lambda_{2,\mathrm{cl}} = .55$. At period 0, any agent selected to exert effort would do so. It is either $E(0) = (1; 3, 1.46), (2; 1, .6)$ or $(3; 1, 1.4)$. In expectation, we get $\mathbb{E}[\pi(\mathcal{N}; 1)] = (.2, .41, .39)'$ and $\mathbb{E}[\lambda_2(\mathcal{N}; 1)] = .21$. So, on average agent 1 profits from manipulation. Since initially the other agents almost did not listen to her and also her opinion was far apart from the others' opinions, she exerts significant effort when selected. In particular, the society is homophilic: taking the cut $(\{1\}, \{2, 3\})$, we get

$$\mathrm{Hom}(\{1\}, 0) = .2 \text{ and } \mathrm{Hom}(\{2, 3\}, 0) = .9.$$

So, since with probability one the manipulation is across the cut, the strong decrease in the (expected) second largest eigenvalue supports our suggestion from Section 4.3 that manipulation reducing homophily (i.e., across the cut) increases the speed of convergence.

---

of an agent on herself, which leads to an increase of the second largest eigenvalue in sufficiently homophilic societies.

At the next period, there is only manipulation if at the last period an agent other than agent 3 was selected to manipulate. In expectation, we get $\mathbb{E}[\pi(\mathcal{N}; 2)] = (.22, .41, .38)'$ and $\mathbb{E}[\lambda_2(\mathcal{N}; 2)] = .17$. Again, agent 1 profits on average from manipulation, but only slightly since opinions are already closer and since she is not as isolated as in the beginning. The convergence gets, on average, slightly faster as well.

Manipulation ends here, that is, with probability one no agent exerts effort from period 2 on, i.e $M(t) = M(2)$ for all $t \geq 2$. Hence, the expected influence of the agents' initial opinions on the consensus is

$$\mathbb{E}[\overline{\pi}(\mathcal{N}; 2)'] = \mathbb{E}[\pi(\mathcal{N}; 2)' \, \overline{M}(1)] = \mathbb{E}[\pi(\mathcal{N}; 2)' \, M(1)] = (.21, .41, .38).$$

Thus, the expected consensus that society reaches is

$$\mathbb{E}[x(\infty)] = \mathbb{E}[\overline{\pi}(\mathcal{N}; 2)']x(0) = 4.53.$$

Compared to this, the classical model gives $x_{\mathrm{cl}}(\infty) = \pi_{\mathrm{cl}}'x(0) = 3.88$ and hence, our model leads to an average long-run belief of society that is closer to the initial opinion of agent 1 since she is the one who (on average) gains influence due to manipulation.

5   The Wisdom of Crowds

We now investigate how manipulation affects the extent of misinformation in society. In this section, we assume that the society forms one minimal closed and aperiodic group. Clearly, societies that are not connected fail to aggregate information.[25] We use an approach similar to Acemoglu et al. (2010) and assume that there is a *true state* $\mu = (1/n) \sum_{i \in \mathcal{N}} x_i(0)$ that corresponds to the average of the initial opinions of the n agents in the society. Information about the true state is dispersed, but can easily be aggregated by the agents: uniform overall influence on the long-run beliefs leads to perfect aggregation of information.[26] Notice that, in general, agents

---

[25]However, as in Example 2, we can observe that manipulation leads to a connected society and thus such an event can also be viewed as reducing the extent of misinformation in the society.

[26]We can think of the initial opinions as being drawn independently from some distribution with mean μ. Then, uniform overall influence leads as well to optimal aggre-

cannot infer the true state from the initial information since they only get to know the information of their neighbors.

At a given period t, the *wisdom* of the society is measured by the difference between the true state and the consensus they would reach in case no more manipulation takes place:

$$\overline{\pi}(\mathcal{N};t)'\, x(0) - \mu = \sum_{i \in \mathcal{N}} \left(\overline{\pi}_i(\mathcal{N};t) - \frac{1}{n} \right)\, x_i(0).$$

Hence, $\|\overline{\pi}(\mathcal{N};t) - (1/n)\mathbb{I}\|_2$ measures the extent of misinformation in the society, where $\mathbb{I} = (1,1,\ldots,1)' \in \mathbb{R}^n$ is a vector of 1s and $\|x\|_2 = \sqrt{\sum_{k \in \mathcal{N}} |x_k|^2}$ is the standard Euclidean norm of $x \in \mathbb{R}^n$. We say that an agent i *undersells* (*oversells*) her information at period t if $\overline{\pi}_i(\mathcal{N};t) < 1/n$ ($\overline{\pi}_i(\mathcal{N};t) > 1/n$). In a sense, an agent underselling her information is, compared to her overall influence, (relatively) well informed.

Definition 4 (Extent of misinformation). A manipulation at period t reduces the *extent of misinformation* in society if

$$\|\overline{\pi}(\mathcal{N};t+1) - (1/n)\mathbb{I}\|_2 < \|\overline{\pi}(\mathcal{N};t) - (1/n)\mathbb{I}\|_2,$$

otherwise, it (weakly) increases the *extent of misinformation*.

The next lemma describes, given some agent manipulates another agent, the change in the overall influence of an agent from period t to period $t+1$.

Lemma 3. *Suppose that $\mathcal{C}(0) = \{\mathcal{N}\}$ and that $\mathcal{N}$ is aperiodic. For $k \in \mathcal{N}$, at period t,*

$$\overline{\pi}_k(\mathcal{N};t+1) - \overline{\pi}_k(\mathcal{N};t) = \sum_{l=1}^{n} \overline{m}_{lk}(t)\big(\pi_l(\mathcal{N};t+1) - \pi_l(\mathcal{N};t)\big).$$

In case there is manipulation at period t, the overall influence of the initial opinion of an agent increases if the agents that overall trust her gain (on average) influence from the manipulation. Next, we provide conditions ensuring that a manipulation reduces the extent of misinformation in the

gation, the difference being that it is not perfect in this case due to the finite number of samples.

society. First, manipulation should not be too cheap for the agent who is manipulating. Second, only agents underselling their information should gain overall influence. We say that $\overline{\pi}(\mathcal{N};t)$ is *generic* if for all $k \in \mathcal{N}$ it holds that $\overline{\pi}_k(\mathcal{N};t) \neq 1/n$.

Proposition 4. *Suppose that $\mathcal{C}(0) = \{\mathcal{N}\}$, $\mathcal{N}$ is aperiodic and that $\overline{\pi}(\mathcal{N};t)$ is generic. Then, there exists $\overline{\alpha} > 0$ such that $E(t) = (i;j,\alpha)$, $\alpha > 0$, reduces the extent of misinformation if*

(i) $\alpha \leq \overline{\alpha}$, *and*

(ii) $\sum_{l=1}^{n} \overline{m}_{lk}(t)\big(\pi_l(\mathcal{N};t+1) - \pi_l(\mathcal{N};t)\big) \geq 0$ *if and only if k undersells her information at period t.*

Intuitively, condition (ii) says that (relatively) well informed agents (those that undersell their information) should gain overall influence, while (relatively) badly informed agents (those that oversell their information) should lose overall influence. Then, this leads to a distribution of overall influence in the society that is more equal and hence reduces the extent of misinformation in the society – but only if i does not exert too much effort on j (condition (i)). Otherwise, manipulation makes some agents too influential, in particular the manipulating agent, and leads to a distribution of overall influence that is even more unequal than before. In other words, information aggregation can be severely harmed when for some agents manipulation is rather cheap.

We now introduce a true state of the world into Example 5. On average, manipulation reduces the extent of misinformation in each period and the society converges to a more precise consensus.

Example 6 (Three-agents society, cont'd)**.** Recall that $\mathcal{N} = \{1,2,3\}$ and that

$$u_i\big(M(t), x(t); j, \alpha\big) = -\frac{1}{2} \sum_{k \neq i} \big(x_i(t) - x_k(t+1)\big)^2 - \big(\alpha^2 + 1/10 \cdot \mathbf{1}_{\{\alpha>0\}}(\alpha)\big)$$

for all $i \in \mathcal{N}$. Furthermore, $x(0) = (10, 5, 1)'$ and

$$M(0) = \begin{pmatrix} .6 & .2 & .2 \\ .1 & .4 & .5 \\ 0 & .6 & .4 \end{pmatrix}.$$

Hence, $\mu = (1/3) \sum_{i \in \mathcal{N}} x_i(0) = 5.33$ is the true state. The vector of initial overall influence is $\overline{\pi}(\mathcal{N}; 0) = \pi(\mathcal{N}; 0) = (.12, .46, .42)'$. Recall that in expectation, we obtain $\mathbb{E}[\overline{\pi}(\mathcal{N}; 1)] = \mathbb{E}[\pi(\mathcal{N}; 1)] = (.2, .41, .39)'$, $\mathbb{E}[\overline{\pi}(\mathcal{N}; 2)] = (.21, .41, .38)'$ and that there is no more manipulation from period 2 on. Thus,

$$\|\overline{\pi}(\mathcal{N}; 0) - (1/3)\mathbb{I}\|_2 = .268 > \|\mathbb{E}[\overline{\pi}(\mathcal{N}; 1)] - (1/3)\mathbb{I}\|_2 = .161$$
$$> \|\mathbb{E}[\overline{\pi}(\mathcal{N}; 2)] - (1/3)\mathbb{I}\|_2 = .158.$$

So, in terms of the expected long-run influence, manipulation reduces the extent of misinformation in society. And indeed, the agents reach the expected consensus $\mathbb{E}[x(\infty)] = 4.53$, which is closer to the true state $\mu = 5.33$ than the consensus they would have reached in the classical model of DeGroot, $x_{\mathrm{cl}}(\infty) = 3.88$.

This confirms the intuition that manipulation has the most bite in the beginning, before potentially misleading opinions have spread. Furthermore, this example suggests that manipulation can have positive effects on information aggregation if agents have homogeneous preferences for manipulation.

6   Conclusion

We investigated the role of manipulation in a model of opinion formation where agents have beliefs about some question of interest and update them taking weighted averages of neighbors' opinions. Our analysis focused on the consequences of manipulation for the trust structure and long-run beliefs in the society, including learning.

We showed that manipulation can modify the trust structure and lead to a connected society, and thus, to consensus. Furthermore, we found

that manipulation fosters opinion leadership in the sense that the manipulating agent always increases her influence on the long-run beliefs. And more surprisingly, this may even be the case for the manipulated agent. The expected change of influence on the long-run beliefs is ambiguous and depends on the agents' preferences and the social network.

We also showed that the trust structure of the society settles down and, if the satisfaction of agents does not directly depend on the trust, manipulation will come to an end and they reach a consensus (under some weak regularity condition). To obtain insights on the relation of manipulation and the speed of convergence, we provided examples and argued that in sufficiently homophilic societies where manipulation is rather costly, manipulation accelerates convergence if it decreases homophily and otherwise it slows down convergence.

Regarding learning, we were interested in the question whether manipulation is beneficial or harmful for information aggregation. We used an approach similar to Acemoglu et al. (2010) and showed that manipulation reduces the extent of misinformation in the society if manipulation is rather costly and the agents underselling their information gain and those overselling their information lose overall influence. Not surprisingly, agents for whom manipulation is cheap can severely harm information aggregation. Furthermore, our main example suggests that homogeneous preferences for manipulation favor a reduction of the extent of misinformation in society.

We should notice that manipulation has no bite if we use the approach of Golub and Jackson (2010). They studied large societies and showed that opinions converge to the true state if the influence of the most influential agent in the society is vanishing as the society grows. Under this condition, manipulation does not change convergence to the true state since its consequences are negligible compared to the size of the society. In large societies, information is aggregated before manipulation (and possibly a series of manipulations) can spread misinformation. The only way manipulation could have consequences for information aggregation in large societies would be to enable agents to manipulate a substantial proportion of the society instead of only one agent. Relaxing the restriction to manipulation of a single agent at a time is left for future work.

We view our paper as first attempt in studying manipulation and mis-

information in social networks. Our approach incorporated strategic considerations in a model of opinion formation à la DeGroot. We made several simplifying assumptions and derived results that apply to general societies. We plan to address some of the open issues in future work, e.g., extending manipulation to groups and allowing for more sophisticated agents.

A   Appendix

Proof of Proposition 1

(i) Follows immediately since all minimal closed groups remain unchanged.

(ii) If agent i manipulates agent j, then $m_{ji}(t+1) > 0$ and thus, since $C' \ni j$ is minimal closed at period t, there exists a path at $t+1$ from l to i for all $l \in C'$. Since C is still minimal closed, it follows that $R(t+1) = R(t) \cup C'$, i.e., $\mathcal{C}(t+1) = \mathcal{C}(t) \backslash \{C'\}$.

(iii) (a) If agent i manipulates agent j, then it follows that $\sum_{l \in C \cup \{i\}} m_{kl}(t+1) = 1$ for all $k \in C$ since C is closed at t. Furthermore, since by assumption there is no path from i to k for any $k \in \cup_{C' \in \mathcal{C}(t) \backslash \{C\}} C'$ and by definition of R', $\sum_{l \in C \cup R' \cup \{i\}} m_{kl}(t+1) = 1$ for all $k \in R' \cup \{i\}$. Hence, it follows that $\sum_{l \in C \cup R' \cup \{i\}} m_{kl}(t+1) = 1$ for all $k \in C \cup R' \cup \{i\}$, i.e., $C \cup R' \cup \{i\}$ is closed.

Note that moreover, since by assumption there is no path from i to k for any $k \in \cup_{C' \in \mathcal{C}(t) \backslash \{C\}} C'$, there is a path from i to j (otherwise $R' \cup \{i\}$ was closed at t). Thus, since C is minimal closed and i manipulates j, there is a path from k to l for all $k, l \in C \cup \{i\}$ at $t+1$. Then, by definition of R', there is also a path from k to l for all $k \in C \cup \{i\}$ and $l \in R'$. Moreover, again by assumption and definition of R', there exists a path from k to l for all $k \in R'$ and all $l \in C$ (otherwise a subset of R' was closed at t).

Combined, this implies that the same holds for all $k, l \in C \cup R' \cup \{i\}$. Hence, $C \cup R' \cup \{i\}$ is minimal closed, i.e., $\mathcal{C}(t+1) = \mathcal{C}(t) \backslash \{C\} \cup \{C \cup R' \cup \{i\}\}$.

(b) If agent i manipulates agent j, then $m_{ji}(t+1) > 0$ and thus, since $C \ni j$ is minimal closed at period t, there exists a path at $t+1$ from l to i for all $l \in C$. Hence, by assumption there exists a path from agent j to k, but not vice versa since $C' \ni k$ is minimal closed. Thus, $R(t+1) = R(t) \cup C$, which finishes the proof.

Proof of Proposition 2

Suppose w.l.o.g. that $\mathcal{C}(t) = \{\mathcal{N}\}$. First, note that aperiodicity is preserved since manipulation can only increase the number of simple cycles. We can write

$$M(t+1) = M(t) + e_j z(t)',$$

where e_j is the j-th unit vector, and

$$
z_k(t) = \begin{cases} (m_{ji}(t) + \alpha)/(1+\alpha) - m_{ji}(t) & \text{if } k = i \\ (m_{jk}(t))/(1+\alpha) - m_{jk}(t) & \text{if } k \neq i \end{cases}
$$

$$
= \begin{cases} \alpha(1 - m_{ji}(t))/(1+\alpha) & \text{if } k = i \\ -\alpha m_{jk}(t)/(1+\alpha) & \text{if } k \neq i \end{cases}.
$$

From Hunter (2005), we get

$$\pi_k(\mathcal{N}; t+1) - \pi_k(\mathcal{N}; t) = -\pi_k(\mathcal{N}; t)\pi_j(\mathcal{N}; t+1) \sum_{l \neq k} z_l(t) r_{lk}(t)$$

$$
= \begin{cases} \alpha/(1+\alpha)\, \pi_i(\mathcal{N}; t)\pi_j(\mathcal{N}; t+1) \sum_{l \neq i} m_{jl}(t) r_{li}(t) & \text{if } k = i \\ \alpha/(1+\alpha)\, \pi_k(\mathcal{N}; t)\pi_j(\mathcal{N}; t+1) \left(\sum_{l \neq k} m_{jl}(t) r_{lk}(t) - r_{ik}(t) \right) & \text{if } k \neq i \end{cases},
$$

which finishes the proof.

Proof of Corollary 1

We know that $\pi_k(C; t), \pi_k(C; t+1) > 0$ for all $k \in C$. Note that if i manipulates j, i.e., $\alpha > 0$, then it must be that $m_{ji}(t) < 1$ since otherwise $[M(t)](i; j, \alpha) = [M(t)](i; j, 0)$ and thus the agent would not have exerted effort. Thus, by Remark 2, $\sum_{l \in C \setminus \{i\}} m_{jl}(t) r_{li}(t) > 0$ and hence $\pi_i(\mathcal{N}; t+$

$1) > \pi_i(\mathcal{N}; t)$, which proves part (i). Part (ii) is obvious. Part (iii) follows since $m_{jk}(t) = 1$ implies $\sum_{l \in C \setminus \{k\}} m_{jl}(t) r_{lk}(t) = 0$, which finishes the proof.

Proof of Lemma 1

By Proposition 1, we know that $\mathcal{C}(t) = \{\mathcal{N}\}$ for all $t \geq 0$, and furthermore, also aperiodicity is preserved. First, we show that the opinions converge to a consensus $x(\infty)$. Therefore, suppose to the contrary that the opinions (with positive probability) do not converge. This implies that there exists a periodic trust matrix $M^* \in \mathbb{R}^{n \times n}$ such that for some sequence of agents $\{i^*(t)\}_{t \geq 0}$ chosen to manipulate, $M(t) \to M^*$ for $t \to \infty$. Denote the decision of $i^*(t)$ at period t by $(j^*(t), \alpha^*(t))$. Notice that since $M(t)$ is aperiodic for all $t \geq 0$, i.e., $M(t) \neq M^*$ for all $t \geq 0$, this is only possible if there are infinitely many manipulations. $\qquad(2)$

Denoting by $x^*(t)$ the opinions and by $M^*(t)$ the trust matrix at period t in the above case, we get

$$\left\| [x^*(t)] \big(i^*(t); j^*(t), \alpha^*(t) \big) - [x^*(t)] \big(i^*(t); j^*(t), 0 \big) \right\|$$
$$= \left\| [M^*(t)] \big(i^*(t); j^*(t), \alpha^*(t) \big) x^*(t) - M^*(t) x^*(t) \right\|$$
$$\to 0 \text{ for } t \to \infty,$$

and thus, by assumption,

$$v_{i^*} \big([M^*(t)] \big(i^*(t); j^*(t), \alpha^*(t) \big), x^*(t) \big) - v_{i^*} \big([M^*(t)] \big(i^*(t); j^*(t), 0 \big), x^*(t) \big)$$
$$\to 0 < \underline{c} \leq c_{i^*}(j^*(t), \alpha^*(t)) \text{ for } t \to \infty,$$

which is a contradiction to (2). Having established the convergence of opinions, it follows directly that $\| [x(t)](i; j, \alpha) - [x(t)](i; j, 0) \| \to 0$ for $t \to \infty$, any i selected at t and her choice (j, α). Hence, by assumption, $v_i \big([M(t)](i; j, \alpha), x(t) \big) - v_i \big([M(t)](i; j, 0), x(t) \big) \to 0 < \underline{c} \leq c_i(j, \alpha)$ for $t \to \infty$, any i selected at t and her choice (j, α), which shows that there exits an almost surely finite stopping time τ such that for all $t \geq \tau$, $E(t) = (i; \cdot, 0)$ for any i chosen to manipulate at t.

Furthermore, since $M(\tau)$ is aperiodic and no more manipulation takes

place, agents reach a (random) consensus that can be written as

$$x(\infty) = \pi(\mathcal{N}; \tau)' x(\tau) = \pi(\mathcal{N}; \tau)' M(\tau) x(\tau - 1)$$
$$= \pi(\mathcal{N}; \tau)' M(\tau - 1) M(\tau - 2) \cdots M(1) x(0)$$
$$= \pi(\mathcal{N}; \tau)' \overline{M}(\tau - 1) x(0),$$

where the second equality follows from the fact that $\pi(\mathcal{N}; \tau)$ is a left eigenvector of $M(\tau)$ corresponding to eigenvalue 1, which finishes the proof.

Proof of Proposition 3

Suppose that the sequence $(\tau_k)_{k=1}^{\infty}$ of stopping times denotes the periods where the trust structure changes, i.e., at $t = \tau_k$ the trust structure changes the k-th time. Notice that $\tau_k = +\infty$ if the k-th change never happens. By Proposition 1, it follows that when $\tau_k < +\infty$, either

(a) $1 \leq |\mathcal{C}(\tau_k + 1)| < |\mathcal{C}(\tau_k)|$ and $|R(\tau_k + 1)| > |R(\tau_k)|$, or

(b) $|\mathcal{C}(\tau_k + 1)| = |\mathcal{C}(\tau_k)|$ and $0 \leq |R(\tau_k + 1)| < |R(\tau_k)|$

holds. This implies that the maximal number of changes in the trust structure is finite, i.e., there exists $K < +\infty$ such that there are at most K changes in the structure and thus, almost surely $\tau_{K+1} = +\infty$. Hence, $\tau = \max\{\tau_k + 1 \mid \tau_k < +\infty\} < +\infty$, where $\tau_0 \equiv 0$, is the desired almost surely finite stopping time, which finishes part (i). Part (ii) follows from Lemma 1. The restriction to C of the matrices $M(t)$ in the computation of the consensus belief is due to the fact that $M(t)|_C$ is a stochastic matrix for all $t \geq 0$ since C is minimal closed at $t = \widehat{\tau}$, which finishes the proof.

Proof of Lemma 2

Suppose that $i \in S$. Since $\sum_{k \in S} m_{jk}(t) - \sum_{k \notin S} m_{jk}(t) \leq (<)1$, it follows that

$$\sum_{k \in S} m_{jk}(t) - \sum_{k \notin S} m_{jk}(t)$$

$$\leq (<) \left(\sum_{k \in S} m_{jk}(t) - \sum_{k \notin S} m_{jk}(t) \right) /(1 + \alpha) + \alpha/(1 + \alpha)$$

$$= \left(\sum_{k \in S \setminus \{i\}} m_{jk}(t) - \sum_{k \notin S} m_{jk}(t) \right) /(1 + \alpha) + (m_{ji}(t) + \alpha)/(1 + \alpha)$$

$$= \sum_{k \in S} m_{jk}(t+1) - \sum_{k \notin S} m_{jk}(t+1)$$

and hence $\mathrm{Hom}(S; t+1) \geq (>)\mathrm{Hom}(S; t)$, which finishes part (i). Part (ii) is analogous, which finishes the proof.

Proof of Lemma 3

We can write

$$\overline{\pi}_k(\mathcal{N}; t+1) = \sum_{l=1}^{n} \overline{m}_{lk}(t) \pi_l(\mathcal{N}; t+1)$$

$$= \sum_{l=1}^{n} \overline{m}_{lk}(t) \big(\pi_l(\mathcal{N}; t+1) - \pi_l(\mathcal{N}; t) \big) + \sum_{l=1}^{n} \overline{m}_{lk}(t) \pi_l(\mathcal{N}; t)$$

$$= \sum_{l=1}^{n} \overline{m}_{lk}(t) \big(\pi_l(\mathcal{N}; t+1) - \pi_l(\mathcal{N}; t) \big) + \underbrace{\sum_{l=1}^{n} \overline{m}_{lk}(t-1) \pi_l(\mathcal{N}; t)}_{=\overline{\pi}_k(\mathcal{N}; t)},$$

where the last equality follows since $\pi(\mathcal{N}; t)$ is a left eigenvector of $M(t)$, which finishes the proof.

Proof of Proposition 4

Let $N_* \subseteq \mathcal{N}$ denote the set of agents that undersell their information at period t. Then, the agents in $N^* = \mathcal{N} \backslash N_*$ oversell their information and additionally, $N_*, N^* \neq \emptyset$. By Proposition 2, we have $\pi_k(\mathcal{N}; t+1) -$

$\pi_k(\mathcal{N}; t) \to 0$ for $\alpha \to 0$ and all $k \in \mathcal{N}$ and thus by Lemma 3 we have

$$\overline{\pi}_k(\mathcal{N}; t+1) - \overline{\pi}_k(\mathcal{N}; t) \to 0 \text{ for } \alpha \to 0 \text{ and all } k \in \mathcal{N}. \qquad (3)$$

Let $k \in N_*$, then by (ii) and Lemma 3, $\overline{\pi}_k(\mathcal{N}; t+1) \geq \overline{\pi}_k(\mathcal{N}; t)$. Hence, by (3), there exists $\overline{\alpha}(k) > 0$ such that

$$1/n \geq \overline{\pi}_k(\mathcal{N}; t+1) \geq \overline{\pi}_k(\mathcal{N}; t) \text{ for all } 0 < \alpha \leq \overline{\alpha}(k).$$

Analogously, for $k \in N^*$, there exists $\overline{\alpha}(k) > 0$ such that

$$1/n \leq \overline{\pi}_k(\mathcal{N}; t+1) < \overline{\pi}_k(\mathcal{N}; t) \text{ for all } 0 < \alpha \leq \overline{\alpha}(k).$$

Therefore, setting $\overline{\alpha} = \min_{k \in \mathcal{N}} \overline{\alpha}(k)$, we get for $0 < \alpha \leq \overline{\alpha}$

$$
\begin{aligned}
\|\overline{\pi}(\mathcal{N}; t) - (1/n) \cdot \mathbb{I}\|_2^2 &= \sum_{k \in \mathcal{N}} |\overline{\pi}_k(\mathcal{N}; t) - 1/n|^2 \\
&= \sum_{k \in N_*} \underbrace{|\overline{\pi}_k(\mathcal{N}; t) - 1/n|^2}_{\geq |\overline{\pi}_k(\mathcal{N}; t+1) - 1/n|^2} + \sum_{k \in N^*} \underbrace{|\overline{\pi}_k(\mathcal{N}; t) - 1/n|^2}_{> |\overline{\pi}_k(\mathcal{N}; t+1) - 1/n|^2} \\
&> \sum_{k \in \mathcal{N}} |\overline{\pi}_k(\mathcal{N}; t+1) - 1/n|^2 \\
&= \|\overline{\pi}(\mathcal{N}; t+1) - (1/n) \cdot \mathbb{I}\|_2^2,
\end{aligned}
$$

which finishes the proof.

Acknowledgements

References

Acemoglu, D., M. Dahleh, I. Lobel, and A. Ozdaglar (2011). Bayesian learning in social networks. *The Review of Economic Studies 78*(4), 1201–36.

Acemoglu, D. and A. Ozdaglar (2011). Opinion dynamics and learning in social networks. *Dynamic Games and Applications 1*(1), 3–49.

Acemoglu, D., A. Ozdaglar, and A. ParandehGheibi (2010). Spread of (mis)information in social networks. *Games and Economic Behavior 70*(2), 194–227.

Austen-Smith, D. and J. R. Wright (1994). Counteractive lobbying. *American Journal of Political Science 38*(1), 25–44.

Büchel, B., T. Hellmann, and S. Klößner (2012). Opinion dynamics under conformity. Center for Mathematical Economics Working Papers 469, Bielefeld University.

Chandrasekhar, A., H. Larreguy, and J. Xandri (2012). Testing models of social learning on networks: evidence from a framed field experiment. Mimeo, Massachusetts Institute of Technology.

Choi, S., D. Gale, and S. Kariv (2012). Social learning in networks: a quantal response equilibrium analysis of experimental data. *Review of Economic Design 16*(2-3), 135–57.

DeGroot, M. (1974). Reaching a consensus. *Journal of the American Statistical Association 69*(345), 118–21.

DeMarzo, P., D. Vayanos, and J. Zwiebel (2003). Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics 118*(3), 909–68.

Esteban, J. and D. Ray (2006). Inequality, lobbying, and resource allocation. *The American Economic Review*, 257–79.

Friedkin, N. E. (1991). Theoretical foundations for centrality measures. *American journal of Sociology 96*(6), 1478–504.

Golub, B. and M. O. Jackson (2010). Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics 2*(1), 112–49.

Golub, B. and M. O. Jackson (2012). How homophily affects the speed of learning and best-response dynamics. *The Quarterly Journal of Economics 127*(3), 1287–338.

Gullberg, A. T. (2008). Lobbying friends and foes in climate policy: the case of business and environmental interest groups in the European Union. *Energy Policy 36*(8), 2964–72.

Hunter, J. J. (2005). Stationary distributions and mean first passage times of perturbed Markov chains. *Linear Algebra and its Applications 410*, 217–43.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.

Van Dijk, T. A. (2006). Discourse and manipulation. *Discourse Society 17*(3), 359–83.

Watts, A. (2014). The influence of social networks and homophily on correct voting. *Network Science 2*(01), 90–106.