Strategic Communication in Social Networks[†]

Manuel Förster[*]

CEREC, Saint-Louis University – Brussels, Belgium

January 26, 2015

Abstract

We study the role of conflicting interests in boundedly rational belief dynamics. Agents meet pairwise with their neighbors in the social network and exchange information strategically. They hold beliefs about an issue of common interest and differ in their preferences about the action to take. The sender of information would like to spread his belief about the action to take, while the receiver would like to learn the sender's belief about the issue. In equilibrium the sender only communicates a noisy message containing information about his belief; the receiver interprets the sent message and updates her belief by taking the average of the interpretation and her pre-meeting belief. With conflicting interests, the belief dynamics generically fails to converge almost surely: each agent's belief converges to some interval and keeps fluctuating on it forever. In particular, our results suggest that the classical consensus result is not stable with respect to conflicts of interest.

JEL classification: C72, D74, D83, D85, Z13.

Keywords: Social networks, strategic communication, conflict of interest, persistent disagreement, belief fluctuations.

1 Introduction

Individuals form their beliefs and opinions on various economic, political and social issues based on information they receive from their social environment. This may

[†]We thank Tom Truyts, Dunia López-Pintado, Jean-Jacques Herings, Ana Mauleon, Vincent Vannetelbosch, Michel Grabisch, Joel Sobel and participants of seminars at CORE, University of Saint-Louis – Brussels, Université libre de Bruxelles, Bielefeld University and the UECE Lisbon Meetings 2014 for helpful comments. This research project is supported by the doctoral program EDE-EM ("European Doctorate in Economics – Erasmus Mundus") of the European Commission, which is gratefully acknowledged.

[*]E-Mail: manuel.forster.uni@gmail.com. Correspondence to: Saint-Louis University – Brussels, Boulevard du Jardin botanique 43, 1000 Brussels, Belgium.

include friends, neighbors and coworkers as well as political actors and news sources, among others. Typically, all these individuals have widely diverging interests, views and tastes, as can be seen in daily political discussions or in all kinds of bargaining situations. In election campaigns, politicians have incentives to argue solutions or proposals that differ from their beliefs. In budget allocation problems, the recipients of capital, e.g., ministries, local governments or departments of companies or universities, have incentives to overstate their capital requirement, while the other side is concerned with efficiency. Another example are court trials, where the accused has clearly incentives to misreport the events in question. And in marketing, firms might overstate the product quality to attract costumers.

When interests are conflicting, individuals will find it more advantageous not to reveal their true belief for strategic reasons. However, in the literature on communication in social networks, it is usually assumed that agents report their beliefs truthfully, see, e.g., DeGroot (1974); Golub and Jackson (2010); DeMarzo et al. (2003); Acemoglu et al. (2010); Förster et al. (2013). DeMarzo et al. (2003) state that this assumption is for simplicity, but that "[n]onetheless, in many persuasive settings, (e.g., political campaigns and court trials) agents clearly do have incentives to strategically misreport their beliefs."

In this paper, individuals differ in their preferences about the action to take with respect to some issue of common interest. Each individual holds a belief about the issue, henceforth simply *belief*, and a belief about the action to take, henceforth *action-belief*.[1] We assume that when two individuals communicate, the receiver of information would like to get to know the belief of the sender about the issue as precisely as possible in order to refine her own belief, while the sender wants to spread his action-belief, i.e., he would like the receiver to adopt his action-belief.

To illustrate this approach, consider an international meeting of politicians, e.g., the United Nations climate change conferences. The common issue of the decision-makers at these meetings is to find and to agree on the measures or actions to take in order to limit global warming. Each decision-maker holds a belief about which measures are to be taken by the global community to achieve this goal. However, the measures they intend to support in front of the other decision-makers (their action-belief) might differ from this belief due to strategic reasons that depend on the local environment within their country. These reasons include local costs of adaption of the measures, the risk profile of the country, and the local public opinion.[2] During

[1] We can interpret her belief as to be about the fundamentals of the issue, while her action-belief is a personal judgement about the issue for strategic reasons or taste considerations.

[2] The 2009 United Nations climate change conference that took place in Copenhagen, Denmark,

these meetings, politicians interact repeatedly with each other. When receiving information, they would like to do so as precisely as possible since the ideal action for each country depends on the fundamentals of global warming, while they intend to support their action-belief when sending information in order to reach an outcome close to the ideal measure for their country.

An important question for society is how the presence of these conflicts influences information aggregation, long-run beliefs and opinions in society. We develop a framework of belief dynamics where individuals with conflicting interests communicate strategically in a social network and update their beliefs naïvely with the obtained information.

We show that these conflicts lead to persistent disagreement among the agents and fluctuating beliefs. In particular, our results suggest that the classical consensus result is not stable with respect to conflicts of interest. Thus, we provide a rationale for why disagreement among individuals in our societies is the norm on many central issues that have been debated for decades.

More precisely, we consider a society represented by a social network of n agents. At time $t \geq 0$, each agent holds a belief $x_i(t) \in [0, 1]$ about some issue of common interest.[3] Furthermore, each agent holds an action-belief $x_i(t) + b_i$ about the action to take, where $b_i \in \mathbb{R}$ is a *bias* relative to her belief.[4] Each agent starts with an initial belief $x_i(0) \in [0, 1]$ and meets (communicates with) agents in her social neighborhood according to a Poisson process in continuous time that is independent of the other agents.[5] When an agent is selected by her associated Poisson process, she receives information from one of her neighbors (called the sender of information) according to a stochastic process that forms her social network.[6] We assume that the sender wants to spread his action-belief, while the receiver wants to infer his belief in order to update her own belief.

In equilibrium, this conflict of interest leads to noisy communication à la Craw-

---

led to a political agreement on the goal of limiting global warming to no more than two degrees Celsius over the pre-industrial average. However, views on the measures to take remained widely diverging depending on local environments and therefore prevented a full-fledged legal agreement, see Bodansky (2010).

[3]We refer to DeMarzo et al. (2003) for a discussion about the representation of beliefs by a unidimensional structure.

[4]Thus, if θ was the true answer to the issue, her ideal action would be $\theta + b_i$.

[5]See Acemoglu et al. (2010, 2013), who use this timing in related models.

[6]Note that we model communication as directed. We want to allow for asymmetric communication since, e.g., an agent might obtain a lot of information from another agent, but this might not be the case vice versa. We can think of journalists whose information reach a large audience, who themselves only receive information from few people, though.

ford and Sobel (1982, henceforth CS): the sender sends one of finite messages that contains information about his belief, which is then interpreted by the receiver.[7] In optimal equilibrium, communication is as informative as possible given the conflict of interest, i.e., the sender uses as many messages as possible and discriminates as finely as possible between different beliefs.[8] The receiver updates her belief by taking the average of the interpretation of the sent message and her pre-meeting belief. Although simple, this updating rule reflects the idea that agents fail to adjust for repetitions and dependencies in information they hear several times due to the complexity of social networks, as argued by DeMarzo et al. (2003).[9] In other words, agents are assumed to be only boundedly rational: they are rational when communicating (with respect to the conflicts of interest), but naïve or boundedly rational when updating their beliefs as they fail to take into account the history of beliefs.

Our framework induces a belief dynamics process as well as an action-belief dynamics process. As a first observation, we note that we can concentrate our analysis on the belief dynamics process since both processes have the same convergence properties. We say that an agent's belief *fluctuates* on an interval if her belief will (almost surely) never leave the interval and if this does not hold for any subinterval. In other words, the belief "travels" the whole interval, but not beyond.

In our main result, we show that for any initial beliefs, the belief dynamics process converges with probability 1 to a set of intervals that is *minimal mutually confirming*. Given each agent's belief lies in her corresponding interval, these intervals are the convex combinations of the interpretations the agents use when communicating. Furthermore, we show that the belief of an agent eventually fluctuates on her corresponding interval whenever the interval is proper, i.e., whenever it contains infinitely many elements (beliefs). As a consequence, the belief dynamics has a steady state if and only if there exists a minimal mutually confirming set such that all its intervals are degenerate, i.e., contain only a single point.

[7]The receiver faces (unmeasurable) uncertainty (also called Knightian uncertainty, see Knight (1921)) about the sender's belief after receiving the message. We assume that agents are maxmin utility maximizers in the sense that upon receiving a message, they choose an interpretation that maximizes their worst-case utility.

[8]Note that CS argue that the optimal equilibrium is particularly plausible in a situation like ours, where communication is repeated.

[9]Note that this updating rule has another appealing interpretation: if the initial beliefs were drawn independently from a normal distribution with equal mean and equal variance and if there was no conflict of interest, then this updating rule would be optimal. In view of this, we should think about the conflicts of interest as being rather small.

However, we remark that such a situation must be constructed explicitly by choosing specific biases and network configurations and thus, we conclude that the belief dynamics generically fails to converge almost surely. On the contrary, in absence of conflicting interests, we recover the classical result that a connected social network makes society reaching a consensus. Thus, it is conflicts of interest that drive our main result and moreover, our results suggest that the classical consensus result is not stable with respect to conflicts of interest.

We then investigate more closely the pattern of the fluctuations and show that the belief dynamics – although failing to converge almost surely – converges in distribution to a random vector. And moreover, we find that the beliefs fluctuate in an ergodic way, i.e., the empirical averages of the agents' beliefs converge to their long-run expectations. We illustrate our results by several examples.

The introduction of conflict of interest leads not only to persistent disagreement among the agents, but also to fluctuating beliefs. These phenomena are common in our societies, where disagreement among individuals and constantly changing beliefs and opinions are the norm on many central issues, and agreement is the rare exception even though these issues have been debated for decades. Persistent disagreement and in particular belief fluctuations have been frequently observed in social sciences, see, e.g., Kramer (1971) who documents large swings in US voting behavior within short periods, and works in social and political psychology that study how political parties and other organizations influence political beliefs, e.g., Cohen (2003); Zaller (1992). Furthermore, DiMaggio et al. (1996) show that the variance of several opinion dimensions has not changed significantly in the US between the 70s and the 90s, and Evans (2003) finds that the variance has actually increased on moral issues. At the same time, our result is surprising in view of the literature on communication in social networks: in most models, a strongly connected network leads to mutual consensus among the agents in the long-run. To this respect, Acemoglu et al. (2013) is the closest to our work, where the authors introduce stubborn agents that never change their belief, which leads to fluctuating beliefs when the other agents update regularly from different stubborn agents.

There exists a large literature on communication in social networks, using both Bayesian and non-Bayesian updating rules.[10] Apart from the various works that

[10] In Bayesian and observational learning models communication is typically assumed to be truthful and agents have the tendency to converge to a mutual consensus, e.g., Banerjee and Fudenberg (2004); Gale and Kariv (2003); Acemoglu et al. (2011). Another stream of literature studies how observable behaviors spread in a population, e.g., López-Pintado (2008, 2012); Jackson and Yariv (2007); Morris (2000).

assume truthful communication, Büchel et al. (2012) study a model where agents act strategically in the sense that their stated belief differs from their true belief depending on their preferences for conformity. Acemoglu et al. (2014) study a model of Bayesian learning where the agents' objective is to form beliefs (acquire information) about an irreversible decision that each agent has to make, eventually. In this setting, agents might want to misreport their information in order to delay the decisions of other agents. The authors show that it is an equilibrium to report truthfully whenever truthful communication leads to asymptotic learning, i.e., the fraction of agents taking the right decision converges to 1 (in probability) as the society grows. They also show that in some situations, misreporting can lead to asymptotic learning while truthful communication would not. However, also these models lead to mutual consensus under the condition that the underlying social network is strongly connected (and some regularity condition).

Several authors have proposed models to explain non-convergence of beliefs, usually incorporating some kind of homophily that leads to segregated societies and polarized beliefs.[11] Axelrod (1997) proposed such a model in a discrete belief setting, and later on Hegselmann and Krause (2002) and Deffuant et al. (2000) studied the continuous case, see also Lorenz (2005); Blondel et al. (2009); Como and Fagnani (2011). Golub and Jackson (2012) argue that the presence of homophily can substantially slow down convergence and thus lead to a high persistence of disagreement. While being able to explain persistent disagreement, these models fail to explain belief fluctuations in society.

Furthermore, our work is related to contributions on cheap-talk games. Hagenbach and Koessler (2010), Galeotti et al. (2013) and Ambrus and Takahashi (2008) extend the framework of CS to a multi-player (-sender) environment, but maintain the one-shot nature of the model.

The paper is organized as follows. In Section 2 we introduce the model and notation. Section 3 concerns the equilibrium in the communication stage. In Section 4 we study the long-run belief dynamics. In Section 5 we conclude. The proofs are presented in the Appendix.

[11]An exception being Friedkin and Johnsen (1990), who study a variation of the model by DeGroot (1974) where agents can adhere to their initial beliefs to some degree. This leads as well to persistent disagreement among the agents.

2 Model and Notation

We consider a set $\mathcal{N} = \{1, 2, \ldots, n\}$, with $n \geq 2$, of agents who repeatedly communicate with their neighbors in a social network. At time $t \geq 0$, each agent $i \in \mathcal{N}$ holds a *belief* $x_i(t) \in [0,1]$ about some issue of common interest. Furthermore, agent i holds an *action-belief* $x_i(t) + b_i$ about the action to take, where $b_i \in \mathbb{R}$ is a commonly known *bias* relative to her belief $x_i(t)$.

The *social network* is given by a stochastic matrix $P = (p_{ij})_{i,j\in\mathcal{N}}$, i.e., $p_{ij} \geq 0$ for all $i, j \in \mathcal{N}$ and $\sum_{j\in\mathcal{N}} p_{ij} = 1$ for all $i \in \mathcal{N}$. For agent i, p_{ij} is the probability to meet agent j, and $\mathcal{N}_i = \{j \in \mathcal{N} \mid p_{ij} > 0\}$ denotes i's *neighborhood*. Let $(\mathcal{N}, g)$ denote the directed graph where $g = \{(i,j) \mid p_{ij} > 0\}$ is the set of directed links induced by meeting probabilities $p_{ij} > 0$. Throughout the paper we will make the following assumption.

Assumption 1. (i) (Self-communication) Agents do not communicate with themselves, i.e., $p_{ii} = 0$ for all $i \in \mathcal{N}$.

(ii) (Connectivity) The graph $(\mathcal{N}, g)$ is strongly connected, i.e., for all $i, j \in \mathcal{N}$ there exists a directed path connecting i to j with links in g.

The first part states that "self-communication" is not possible. We make this assumption for simplicity, but it could be included as a possibility to abstain from communication. The second part guarantees that every agent communicates indirectly with every other agent, possibly through several links. We make this assumption for several reasons. First, it seems to be natural as evidence suggests that our societies are indeed connected, see, e.g., Watts (2003). And second, it is known to be a necessary condition for convergence of beliefs to a consensus. We want to exclude that beliefs fail to converge because agents are not connected.

Each agent $i \in \mathcal{N}$ starts with an initial belief $x_i(0) \in [0,1]$. Agents meet (communicate) and update their beliefs according to an asynchronous continuous-time model. Each agent is chosen to meet another agent at instances defined by a rate one Poisson process independent of the other agents. Therefore, over all agents, the meetings occur at time instances t_s, $s \geq 1$, according to a rate n Poisson process. Note that by convention, at most one meeting occurs at a given time $t \geq 0$. Hence, we can discretize time according to the agent meetings and refer to the interval $[t_s, t_{s+1})$ as the s^{th} time slot. There are on average n meetings per unit of absolute time, see Boyd et al. (2006) for a detailed relation between the meetings and absolute time. At time slot s, we represent the beliefs of the agents by the vector $x(s) = (x_1(s), x_2(s), \ldots, x_n(s))$.

If agent $i \in \mathcal{N}$ is chosen at time slot s, $s \geq 1$ (probability $1/n$), she meets agent $j \in \mathcal{N}$ with probability p_{ij} and communicates with him. We assume that agent i updates her belief with information she receives from agent j.[12] Agent j sends a *message* (or *signal*) $m \in \mathcal{M} := \{m_1, m_2, \ldots, m_L\}$ containing information about his belief $x_j(s-1)$, where $L \in \mathbb{N}$ is very large but finite, and which is interpreted by i as an estimate $y^{ij}(m)$ of $x_j(s-1)$.[13] Agent i then updates her belief by taking the average of this interpretation and her pre-meeting belief:

$$x_i(s) = \frac{x_i(s-1) + y^{ij}(m)}{2}.$$

If not stated otherwise, agent i will denote the agent that updates her belief (the *receiver* of information), and agent j will denote the agent with whom she communicates (the *sender* of information). We write $g(s) = ij$ if link (i, j) is chosen at time slot s.

Next, we specify how communication between agents takes place. We adapt the framework of Jäger et al. (2011) to conflicting interests and repeated communication. Suppose that $g(s) = ij$; we make the following assumption about the objectives of the agents.

Assumption 2 (Objectives)**.** Agent i's objective is to infer agent j's belief $x_j(s-1)$, while agent j's objective is to spread his action-belief $x_j(s-1) + b_j$.

We represent Assumption 2 by the (communication) preferences

$$u_i(x_j(s-1), y^{ij}(m)) = h(|x_j(s-1) - y^{ij}(m)|)$$

and

$$u_j(x_j(s-1), y^{ij}(m)) = h(|x_j(s-1) + (b_j - b_i) - y^{ij}(m)|),$$

where $h : \mathbb{R}_+ \to \mathbb{R}$ is a continuous, concave and strictly decreasing function. Agent j wants to send a message m such that $i's$ interpretation is as close as possible to $x_j(s-1)+(b_j-b_i)$, the ideal interpretation from his point of view, while agent i wants to choose an interpretation that is as close as possible to j's belief $x_j(s-1)$.[14] Notice that, first, in absence of conflict of interest ($b_i = b_j$) the ideal interpretation from

[12] Note that agent j does not update his belief. Together with the directed social network, this assumption allows for asymmetric communication.

[13] We know from CS that assuming a (sufficiently) large but finite number of messages represents only a restriction in absence of conflict of interest. Since we focus on conflicting interests, we take this assumption for analytical convenience.

[14] We can also interpret $-u_k(x_j(s-1), y^{ij}(m))$ as the loss from communication, see Jäger et al. (2011).

agent j's point of view is equal to his belief, i.e., preferences are aligned; second, if action-beliefs coincide ex-ante (at time slot $s-1$), then agent j's ideal interpretation $x_j(s-1)+(b_j-b_i)$ is the unique interpretation that leaves i's action-belief unchanged and that is hence consistent with Assumption 2.[15] Furthermore, the belief dynamics is well-defined since agent i optimally chooses an interpretation in $[0,1]$ whatever message she receives. A simple example are quadratic preferences.

Example 1 (Quadratic preferences)**.**

$$u_i(x,y) = -(x-y)^2 \text{ and } u_j(x,y) = -(x+(b_j-b_i)-y)^2.$$

In this signaling game, a strategy for the sender j is a measurable function

$$m^{ij} : [0,1] \to \mathcal{M}$$

that assigns a message to each possible belief and for the receiver i, it is a function

$$y^{ij} : \mathcal{M} \to [0,1]$$

that assigns an interpretation to each possible message. We refer to the interpretation of message m_l as $y_l = y^{ij}(m_l)$ and to the set of beliefs that induces m_l as $C_l = (m^{ij})^{-1}(m_l) = \{x \in [0,1] : m^{ij}(x) = m_l\}$ when there is no confusion. An agent receiving message m_l faces (unmeasurable) uncertainty (or Knightian uncertainty) about the location of the sender's belief within C_l. We assume that she is a maxmin utility maximizer in the following sense: upon receiving m_l, she maximizes her worst-case utility $\min_{x\in C_l} u_i(x,y)$.[16]

Then, an equilibrium of the game consists of strategies (m^{ij}, y^{ij}) such that

(i) for each message $m_l \in \mathcal{M}$,

$$y_l \in \operatorname*{argmax}_{y\in\mathbb{R}} \min_{x\in C_l} u_i(x,y), \text{ and}$$

[15]Suppose that action-beliefs coincide ex-ante, i.e., $x_i(s-1)+b_i = x_j(s-1)+b_j$. Then, the ideal interpretation y^* for agent j should not change i's action-belief, i.e., by assumption,

$$x_j(s-1)+b_j = x_i(s-1)+b_i \overset{!}{=} x_i(s)+b_i = \frac{x_i(s-1)+y^*}{2}+b_i = \frac{x_j(s-1)+b_j+y^*+b_i}{2}$$
$$\Leftrightarrow y^* = x_j(s-1)+b_j-b_i.$$

[16]Notice that due to this assumption, the meanings of messages stay the same over time, but nevertheless differ between pairs of agents depending on their particular conflict of interest.

(ii) for each belief $x \in [0,1]$,

$$m^{ij}(x) \in \operatorname*{argmax}_{m \in \mathcal{M}} u_j\big(x, y^{ij}(m)\big).$$

Condition (i) says that agent i uses an interpretation that maximizes her worst-case utility for each message she receives. Condition (ii) says that for each belief agent j chooses a message that maximizes his utility. We exclude the possibility of any prior commitment of the agents and assume without loss of generality that whenever two messages lead to the same interpretation, then agent j only sends the message with the lower index. We say that a message m_l is *induced* (*used*) in equilibrium if $C_l = (m^{ij})^{-1}(m_l) \neq \emptyset$. Thus, we can restrict our attention to the messages that are induced in equilibrium and their interpretations, which are distinct.[17] Throughout the paper we assume that Assumption 1 and 2 hold.

3 Communication Stage

In this section we characterize, given $g(s) = ij$, how agent j communicates with agent i. Suppose j uses messages $m \in \mathcal{M}|_{L(ij)} := \{m_1, m_2, \ldots, m_{L(ij)}\}$ in equilibrium that lead to distinct interpretations $(y_l)_{l=1}^{L(ij)}$. Then, given j holds belief $x_j(s-1)$, he sends a message that maximizes his utility, i.e.,

$$\begin{aligned} m^{ij}(x_j(s-1)) &\in \operatorname*{argmax}_{m \in \mathcal{M}|_{L(ij)}} u_j\big(x_j(s-1), y^{ij}(m)\big) \\ &= \operatorname*{argmax}_{m \in \mathcal{M}|_{L(ij)}} h(|x_j(s-1) + (b_j - b_i) - y^{ij}(m)|) \\ &= \operatorname*{argmin}_{m \in \mathcal{M}|_{L(ij)}} |x_j(s-1) + (b_j - b_i) - y^{ij}(m)|, \end{aligned}$$

where the last equality follows since h is strictly decreasing. Note that this choice is not uniquely defined if $x_j(s-1) + (b_j - b_i)$ has equal distance to two interpretations; we assume without loss of generality that j sends the message with the lowest index in this case. Hence, we can identify j's strategy in equilibrium with a partition $(C_l)_{l=1}^{L(ij)}$ of $[0,1]$, where

$$C_l = (m^{ij})^{-1}(m_l) = \{x \in [0,1] : m^{ij}(x) = m_l\} = [c_{l-1}, c_l)$$

is such that $0 = c_0 < c_1 < \cdots < c_{L(ij)} = 1$. Note that c_l refers to the belief where j is indifferent between sending message m_l and m_{l+1}. So, in equilibrium he partitions

[17] Notice that we restrict our attention to pure strategies.

the unit interval and only communicates the element of the partition his belief is from. Upon receiving message m_l, i will choose an interpretation that maximizes her worst-case utility, i.e.,

$$y_l = \operatorname*{argmax}_{y\in\mathbb{R}} \min_{x\in C_l} u_i(x,y) = \operatorname*{argmin}_{y\in\mathbb{R}} \max_{x\in C_l}|x-y| = \frac{c_l - c_{l-1}}{2}.$$

The number of messages induced in equilibrium is bounded under conflict of interest: we show that the distance between any two interpretations induced in equilibrium is larger than the distance $|b_j - b_i|$ between the ideal interpretations of the agents. Only the equilibrium with one message always exists: in this equilibrium, j's strategy is given by $C_1 = [0,1]$ and i uses the interpretation $y_1 = \operatorname{argmax}_{y\in\mathbb{R}} \min_{x\in C_1} u_i(x,y) = 1/2$. We refer to the finite upper bound on the number of messages (or the "size" of the partition) induced in equilibrium by $L(ij)$. We call the equilibrium using $L(ij)$ messages *optimal equilibrium* since it is *most informative* in the sense that it uses the finest partition. Furthermore, this equilibrium is *essentially unique* in the sense that all equilibria using $L(ij)$ messages induce the same partition and as the receiver's interpretation of a given partition element is unique. And, following the argumentation of CS, we assume that agents coordinate on this equilibrium.[18]

In absence of conflicting interests, the same result holds since we only allow for a finite number of messages. Agents use the maximum number of messages $L(ij) = L$ in optimal equilibrium. Since we do not want to restrict the game under conflict of interest, we assume $L \geq \max\{L(ij) \mid b_i \neq b_j\}$. The following proposition summarizes our findings and explicitly determines the optimal equilibrium.

Proposition 1. *Suppose that $g(s) = ij$.*

(i) If $b_i \neq b_j$, then there exists a finite upper bound

$$L(ij) = \max\{l \in \mathbb{N} \mid 1/(2l) \geq |(l-1)(b_j - b_i)|\}$$

on the number of messages in equilibrium.

(ii) The game has an essentially unique optimal equilibrium (m^{ij}, y^{ij}) in which agent j uses $L(ij)$ (L if $b_i = b_j$) messages and his strategy is given by a partition $(C_l)_{l=1}^{L(ij)}$, where $C_l = (m^{ij})^{-1}(m_l) = [c_{l-1}, c_l)$ is such that

$$c_l = l/L(ij) - 2l(L(ij) - l)(b_j - b_i).$$

[18] They argue that this equilibrium seems to be particularly plausible in situations where communication is repeated, that is, in our case.

Furthermore, agent i's strategy is given by interpretations

$$y_l = y^{ij}(m_l) = (2l-1)/(2L(ij)) - \big((2l-1)L(ij) - 2(l^2-l) - 1\big)(b_j - b_i)$$

for $l = 1, 2, \ldots, L(ij)$.

All proofs can be found in the Appendix. We denote the optimal equilibrium when $g(s) = ij$ by the triple $\mathcal{E}^{ij} = (L(ij), C^{ij}, Y^{ij})$, where $C^{ij} = (c_1, c_2, \ldots, c_{L(ij)-1})$ denotes j's strategy and $Y^{ij} = (y_1, y_2, \ldots, y_{L(ij)})$ denotes i's strategy. The next example illustrates how such equilibria can look like.

Example 2. Consider $\mathcal{N} = \{1, 2\}$ and the vector of biases $b = (0, 1/20)$. The first agent is not biased, while the second is biased to the right.

When $g(s) = 12$, then $L(12) = 3$ messages are induced in optimal equilibrium and strategies are $C^{12} = (4/30, 14/30)$ and $Y^{12} = (2/30, 9/30, 22/30)$. This means that if, for instance, agent 2's belief is below $c_1 = 4/30$, then he sends message m_1 and agent 1 interprets this as $y_1 = 2/30$. When $g(s) = 21$, then as well $L(21) = 3$ messages are induced in optimal equilibrium and strategies are $C^{21} = (16/30, 26/30)$ and $Y^{21} = (8/30, 21/30, 28/30)$. Both equilibria are depicted in Figure 1.

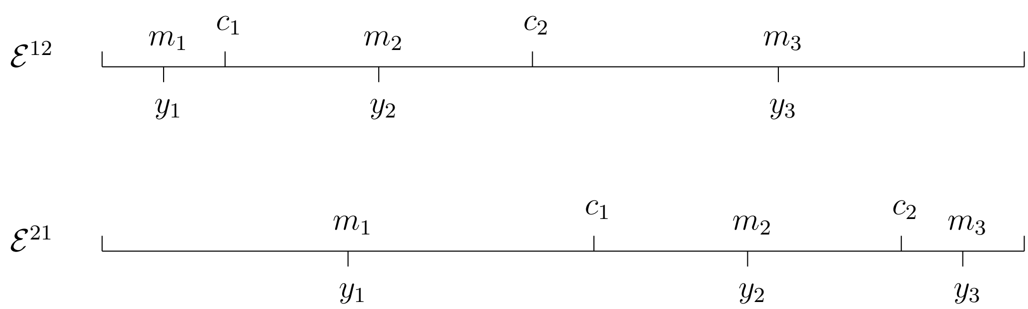

Figure 1: Optimal equilibria in Example 2.

4 Belief Dynamics

In this section we study the long-run behavior of the belief dynamics. At each time slot s, a pair of agents $g(s) = ij$ is selected according to the social network and communicates by employing the optimal equilibrium $\mathcal{E}^{ij}$. Agent i adopts the average of her pre-meeting belief and the equilibrium outcome of communication (her interpretation) as her updated belief. Hence, the belief dynamics $\{x(t)\}_{t\geq 0}$

defines a Markovian stochastic process. Note that we can define as well the action-belief dynamics process $\{x(t)+b\}_{t\geq 0}$, where $b=(b_1,b_2,\ldots,b_n)$ denotes the vector of biases.

Remark 1. The action-belief dynamics $\{x(t)+b\}_{t\geq 0}$ is obtained by a translation of the state space of the belief dynamics $\{x(t)\}_{t\geq 0}$. Hence, both processes have the same properties in terms of convergence.

In the following, we will focus on the belief dynamics. The next example suggests that conflicting interests might prevent society from finding a consensus and instead lead to fluctuating beliefs.

Example 3 (Belief fluctuation)**.** Consider $\mathcal{N}=\{1,2,3\}$ and the vector of biases $b=(0,1/25,-1/15)$. Furthermore, all agents hold the same initial belief $x_i(0)=1/2$ and the social network is given by

$$P=\begin{pmatrix} & 1/2 & 1/2\\ 1/2 & & 1/2\\ 1/2 & 1/2 & \end{pmatrix},$$

i.e., each possible pair of agents is chosen with probability 1/6 at a given time slot. This leads to the following equilibria in the communication stage:

- $\mathcal{E}^{12}=(4,(6/600,108/600,306/600),(3/600,57/600,207/600,453/600))$,
- $\mathcal{E}^{13}=(3,(360/600,560/600),(180/600,460/600,580/600))$,
- $\mathcal{E}^{21}=(4,(294/600,492/600,594/600),(147/600,393/600,543/600,597/600))$,
- $\mathcal{E}^{23}=(2,(428/600),(214/600,514/600))$,
- $\mathcal{E}^{31}=(3,(40/600,240/600),(20/600,140/600,420/600))$,
- $\mathcal{E}^{32}=(2,(172/600),(86/600,386/600))$.

The number of messages induced in equilibrium varies depending on the pair of agents selected to communicate. Agents 1 and 2 use four messages when communicating. The agents with the largest conflict of interest, 2 and 3, only use two messages in equilibrium, though. When looking at the long-run belief dynamics, we find that beliefs do not converge although agents started with identical beliefs. Instead, the beliefs keep fluctuating forever. In particular, each belief fluctuates on some subinterval of $[0,1]$. Agent 1's belief fluctuates on $[180/600,460/600]$, agent 2's belief on $[147/600,393/600]$, and agent 3's belief on $[86/600,420/600]$. Figure 2 depicts one outcome of the long-run belief dynamics.

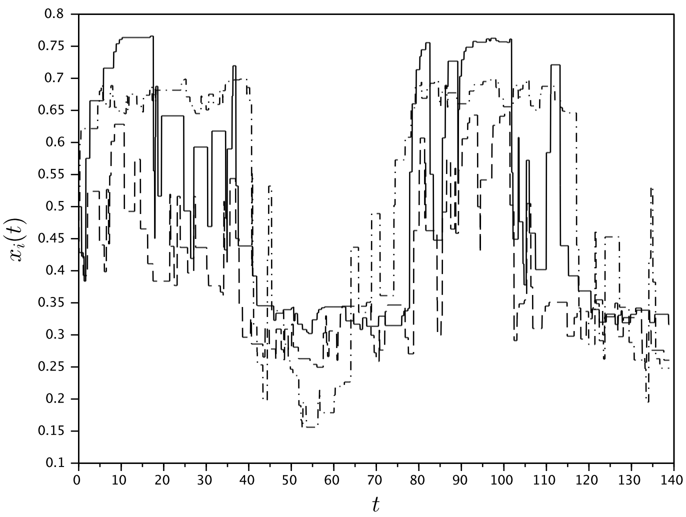

Figure 2: Long-run belief dynamics in Example 3. The solid line represents agent 1, the dashed line agent 2 and the dashed-dotted line agent 3.

Note that the boundaries of the subintervals on which the beliefs fluctuate in the above example are related to the interpretations used by the agents when receiving information. In the following, we want to characterize the asymptotic behavior of the belief dynamics. First, we formalize what we mean by *fluctuation*. We say that an interval is *proper* if it contains infinitely many elements (beliefs).

Definition 1 (Fluctuation)**.** We say that the belief of an agent $i \in \mathcal{N}$ *fluctuates* on the closed and proper interval $\mathcal{I} \subseteq [0,1]$ at time slot s if almost surely $x_i(s') \in \mathcal{I}$ for all $s' \geq s$, but for any closed subinterval $\mathcal{I}' \subsetneq \mathcal{I}$ this does not hold.

In other words, fluctuation on some interval means that the agent's belief never leaves the interval again, but still it "travels" the whole interval. Next, we define the concept of *mutually confirming intervals*. For $j \in \mathcal{N}_i$, let

$$Y^{ij}\big|_{\mathcal{I}_j} = \{y \in Y^{ij} \mid y = y^{ij}(m^{ij}(x)) \text{ for some } x \in \mathcal{I}_j\}$$

denote the restriction of Y^{ij} to the interpretations that correspond to messages sent when j's belief is in $\mathcal{I}_j$.

Definition 2 (Mutually confirming intervals)**.** We say that a set of intervals $\{\mathcal{I}_i\}_{i\in\mathcal{N}}$ is *mutually confirming* if, for all $i \in \mathcal{N}$,

$$\mathcal{I}_i = \text{conv}\left(\cup_{j\in\mathcal{N}_i} Y^{ij}\big|_{\mathcal{I}_j}\right).$$

We say that a set of intervals $\{\mathcal{I}_i\}_{i\in\mathcal{N}}$ is *minimal mutually confirming* if it is mutually confirming and there does not exist a mutually confirming set $\{\mathcal{I}'_i\}_{i\in\mathcal{N}}$ such that $\mathcal{I}'_i \subseteq \mathcal{I}_i$ for all $i \in \mathcal{N}$ and $\mathcal{I}'_i \subsetneq \mathcal{I}_i$ for at least one $i \in \mathcal{N}$.

Mutually confirming intervals are the convex combinations of the interpretations of the messages sent when communicating, given each agent's belief lies in her corresponding interval. The next theorem shows that the belief dynamics converges to a minimal mutually confirming set of intervals. Furthermore, we show that the belief of an agent eventually fluctuates on her corresponding interval whenever the interval is proper.

Theorem 1. *(i) For any vector of initial beliefs $x(0) \in [0,1]^n$, the belief dynamics $\{x(t)\}_{t\geq 0}$ converges almost surely to a minimal mutually confirming set of intervals $\{\mathcal{I}_i\}_{i\in\mathcal{N}}$, and*

(ii) there exists an almost surely finite stopping time τ on the probability space induced by the belief dynamics process such that the belief of agent $i \in \mathcal{N}$ fluctuates on $\mathcal{I}_i$ at time slot s under the event $\{\tau = s\}$ if $\mathcal{I}_i$ is proper.

Theorem 1 implies that if all intervals of a minimal mutually confirming set are *degenerate*, i.e., contain only a single point, then the belief dynamics process has a steady state.

Corollary 1. *The belief dynamics $\{x(t)\}_{t\geq 0}$ has a steady state x^* if and only if there exists a minimal mutually confirming set of intervals $\{\mathcal{I}_i\}_{i\in\mathcal{N}}$ such that $\mathcal{I}_i$ is degenerate for all $i \in \mathcal{N}$. In this case, $x^* = \{\mathcal{I}_i\}_{i\in\mathcal{N}}$.*

When each agent communicates only with one other agent, there is a steady state for sure. The next example shows that this is also possible if some agent communicates with several agents.

Example 4 (Steady state)**.** Consider $\mathcal{N} = \{1,2,3\}$ and the vector of biases $b = (0, 37/600, -26/600)$. Furthermore, all agents hold the same initial belief $x_i(0) = 1/2$ and the social network is given by

$$P = \begin{pmatrix} & 1/2 & 1/2 \\ 1 & & \\ 1 & & \end{pmatrix},$$

i.e., agent 1 is connected to all other agents, while these agents only listen to agent 1. This leads to the following equilibria in the communication stage:

- $\mathcal{E}^{12} = (3, (26/300, 126/300), (13/300, 76/300, 213/300))$,
- $\mathcal{E}^{13} = (3, (152/300, 252/300), (76/300, 202/300, 276/300))$,
- $\mathcal{E}^{21} = (3, (174/300, 274/300), (87/300, 224/300, 287/300))$,
- $\mathcal{E}^{31} = (3, (48/300, 148/300), (24/300, 98/300, 224/300))$.

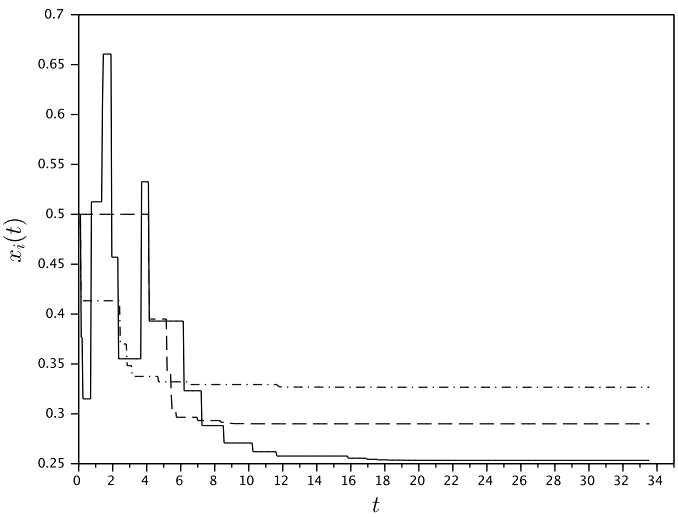

Figure 3: Long-run belief dynamics in Example 4. The solid line represents agent 1, the dashed line agent 2 and the dashed-dotted line agent 3.

All equilibria induce three messages in equilibrium. The vector of beliefs $x^* = (76/300, 87/300, 98/300)$ is a steady state of the process. Note that since agent 1 communicates with two different agents, it is key that the interpretation $y = 76/300$ is part of both equilibria when she is selected to update her belief. Figure 3 depicts an outcome where beliefs converge to this steady state.

The above example shows that the belief dynamics might converge in certain cases. However, such an outcome must be constructed explicitly by choosing specific biases and network configurations. The network needs to be sparse since each time an agent communicates with several agents, we need to find biases such that some interpretation is part of all equilibria. And additionally, we must ensure that these common interpretations are mutually confirming. In particular, a steady state is not stable with respect to the biases.

Remark 2. If conflicts of interest are small enough such that in optimal equilibrium (some) agents send more than one message (no "babbling"), then outcomes with a steady state are non-generic, i.e., the belief dynamics process generically fails to converge almost surely.

Furthermore, in absence of conflicts of interest, i.e., $b_i = \bar{b}$ for all $i \in \mathcal{N}$ so that all equilibria are the same, our model leads to mutual consensus among the agents. Thus, we recover the classical result that a connected social network (and typically some weak regularity condition) makes society reaching a consensus. In particular, it follows then from Remark 2 that this result is not stable with respect to conflicts of interest. When introducing small conflicts, the behavior of the dynamics changes drastically from approaching a consensus to fluctuations. The following remark summarizes these observations.

Remark 3. (i) The belief dynamics approaches a consensus value in absence of conflicts of interest ($b_i = \bar{b}$ for all $i \in \mathcal{N}$), and

(ii) this behavior is not stable with respect to conflicts of interest as introducing small conflicts generically leads to fluctuations.

Thus, we have identified conflicts of interest as the channel that drives our main result. The next example shows how the belief dynamics changes when we introduce small conflicts of interest.

Example 5 (Non-Stability of consensus). Recall that in Example 3, $\mathcal{N} = \{1, 2, 3\}$ and

$$P = \begin{pmatrix} & 1/2 & 1/2 \\ 1/2 & & 1/2 \\ 1/2 & 1/2 & \end{pmatrix}.$$

Consider initial beliefs $x(0) = (3/20, 1/2, 3/4)$ and that the agents are not biased, i.e., $b_i = 0$ for all $i \in \mathcal{N}$.[19] As there are no conflicts of interest, the agents' beliefs quickly approach a consensus value.

Second, consider the vector of biases $b = (1/200, 1/2500, -1/200)$, then agents 1 and 3 use 10 messages in optimal equilibrium when communicating with agent 2, and 7 messages otherwise (as their conflict is larger compared to that each of them has with agent 2). Already these small conflicts lead to substantial fluctuations. Notice however that the amplitude of the fluctuations is smaller compared to Example 3 as the conflicts are smaller.[20] Figure 4 depicts how the belief dynamics changes when

[19] We use a maximum number of messages of $L = 20$ to compute the optimal equilibria in this case. Using more messages would not change the outcome qualitatively.

[20] Nevertheless, the mutually confirming intervals might be larger with small conflicts.

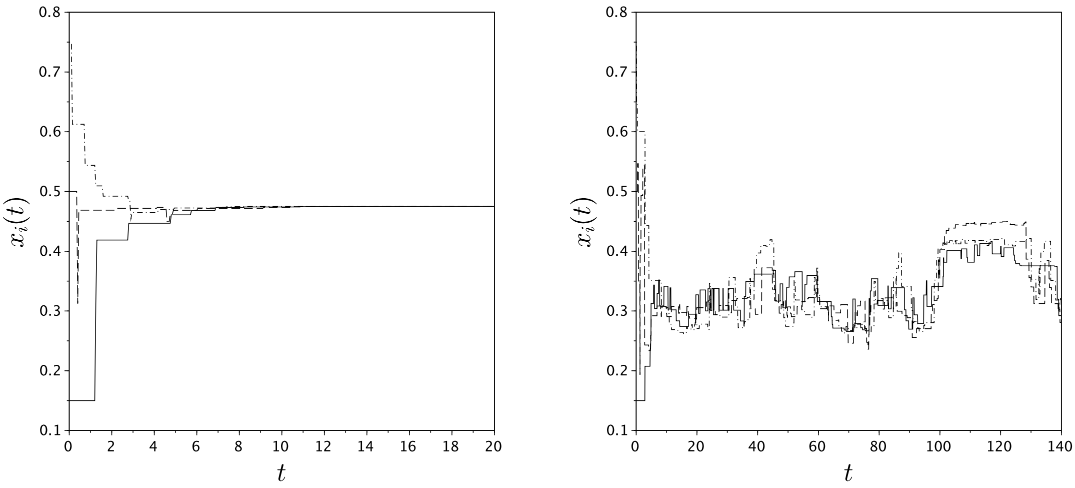

Figure 4: Long-run belief dynamics before (left) and after (right) the introduction of small biases in Example 5. The solid line represents agent 1, the dashed line agent 2 and the dashed-dotted line agent 3.

introducing these small biases.

Next, we investigate more closely the pattern of the fluctuations. So far, we have shown that the beliefs converge almost surely to mutually confirming intervals and each belief fluctuates on its corresponding interval forever. We find that – although failing to converge almost surely – the belief dynamics converges in distribution to a random vector.

Proposition 2. *Suppose the belief dynamics* $\{x(t)\}_{t\geq 0}$ *has minimal mutually confirming sets of intervals* $\{\mathcal{I}_i^1\}_{i\in\mathcal{N}}, \{\mathcal{I}_i^2\}_{i\in\mathcal{N}}, \dots, \{\mathcal{I}_i^r\}_{i\in\mathcal{N}}$. *For any vector of initial beliefs* $x(0) \in [0,1]^n$,

(i) there exists a probability vector $(q_1, q_2, \dots, q_r)$ *such that* $\{x(t)\}_{t\geq 0}$ *converges to the minimal mutually confirming set of intervals* $\{\mathcal{I}_i^k\}_{i\in\mathcal{N}}$ *(henceforth* $x(t) \Rightarrow \mathcal{I}^k$*) with probability* $q_k = q_k(x(0))$, $k = 1, 2, \dots, r$, *and furthermore*

(ii) $\{x(t)\}_{t\geq 0}$ *converges in distribution, i.e., there exist* $[0,1]^n$*-valued stationary random vectors* $\widehat{x}^1, \widehat{x}^2, \dots, \widehat{x}^r$ *such that*

$$\lim_{t\to\infty} \mathcal{L}(x(t) \mid x(t) \Rightarrow \mathcal{I}^k) = \mathcal{L}(\widehat{x}^k), \text{ and in particular}$$

$$\lim_{t\to\infty} \mathbb{E}[\varphi(x(t)) \mid x(t) \Rightarrow \mathcal{I}^k] = \mathbb{E}[\varphi(\widehat{x}^k)]$$

for all bounded and continuous test functions $\varphi : [0,1]^n \to \mathbb{R}$.

Notice that we can obtain the limit of the conditional expected belief of agent $i \in \mathcal{N}$ by choosing the projection $\varphi_i : x \mapsto x_i$.

In the above result, the limiting distribution depends on the minimal mutually confirming set of intervals the belief dynamics converges to. However, when the number of agents is large and the network dense, it will usually be the case that there exists a unique set of those intervals. The following corollary states the result for this case and furthermore, we find that the belief dynamics obeys an ergodic property. We show that the empirical average of the belief dynamics process approaches its long-run expected value with probability 1.[21]

Corollary 2. *Suppose the belief dynamics* $\{x(t)\}_{t\geq 0}$ *has a unique minimal mutually confirming set of intervals* $\{\mathcal{I}_i\}_{i\in\mathcal{N}}$. *For any vector of initial beliefs* $x(0) \in [0,1]^n$,

(i) $\{x(t)\}_{t\geq 0}$ *converges in distribution to a* $[0,1]^n$*-valued stationary random vector* $\widehat{x}$, *i.e.,*

$$\lim_{t\to\infty} \mathcal{L}(x(t)) = \mathcal{L}(\widehat{x}) \text{ and } \lim_{t\to\infty} \mathbb{E}[\varphi(x(t))] = \mathbb{E}[\varphi(\widehat{x})],$$

and furthermore, with probability 1,

(ii)

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t \varphi(x(u))du = \mathbb{E}[\varphi(\hat{x})],$$

for all bounded and continuous test functions $\varphi : [0,1]^n \to \mathbb{R}$.

To illustrate this result, we revisit Example 3 and look at the empirical averages of the agents.

Example 6 (Empirical average)**.** Recall the setting of Example 3: $\mathcal{N} = \{1,2,3\}$ and $b = (0, 1/25, -1/15)$. Furthermore, $x_i(0) = 1/2$ for all $i \in \mathcal{N}$ and

$$P = \begin{pmatrix} & 1/2 & 1/2 \\ 1/2 & & 1/2 \\ 1/2 & 1/2 & \end{pmatrix}.$$

Although the agents' beliefs keep fluctuating forever on their corresponding intervals, their empirical averages do converge. Figure 5 depicts one outcome of the long-run belief dynamics.

[21] Notice that this result can be extended to the case with several minimal mutually confirming sets of intervals. Since in this case the long-run expected value is conditional on the minimal mutually confirming set of intervals the process converges to, the same holds for the empirical averages, i.e., with probability $q_k(x(0))$, we have $\lim_{t\to\infty} 1/t \int_0^t \varphi(x(u))du = \mathbb{E}[\varphi(\hat{x}^k)]$ for all bounded and continuous test functions $\varphi : [0,1]^n \to \mathbb{R}$.

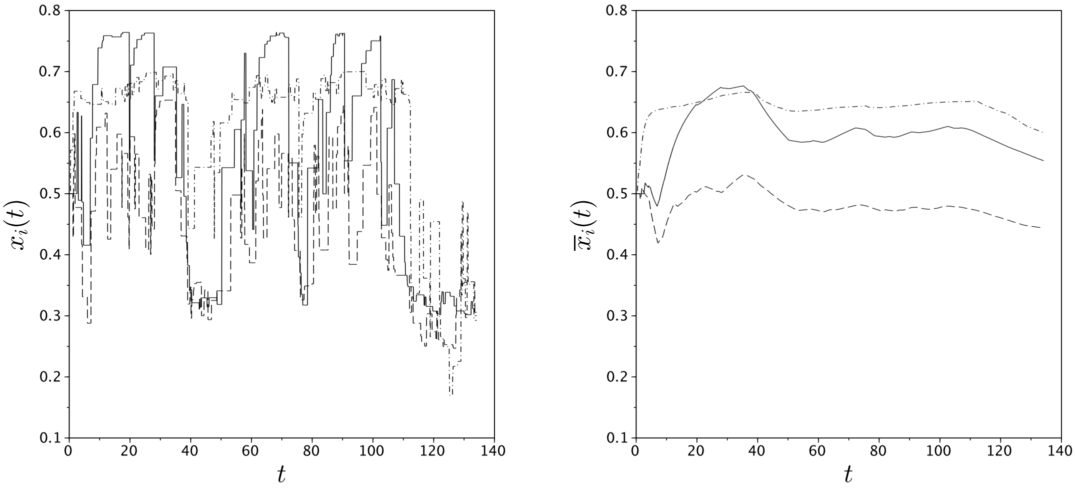

Figure 5: Long-run belief dynamics (left) and its empirical averages (right) in Example 6. The solid line represents agent 1, the dashed line agent 2 and the dashed-dotted line agent 3.

5 Discussion and Conclusion

We study the role of conflicting interests in boundedly rational belief dynamics. Our analysis is motivated by numerous examples such as political campaigns or court trials, where conflicts between different individuals are clearly present. We consider a society represented by a strongly connected network, agents meet (communicate) pairwise with their neighbors and exchange information strategically.

Agents have different preferences about the action to take with respect to some issue of common interest. When two individuals communicate, the receiver of information would like to get to know the belief of the sender about the issue as precisely as possible in order to refine her own belief, while the sender wants to spread his action-belief (about the action to take), i.e., he would like the receiver to adopt his action-belief.

This conflict of interest prevents the agents from revealing their true belief in equilibrium, and instead it leads to noisy communication à la CS: the sender sends one of finite messages that contains information about his belief, which is then interpreted by the receiver. In optimal equilibrium, communication is as informative as possible given the conflict of interest, i.e., the sender uses as many messages as possible. The receiver updates her belief by taking the average of the interpretation of the sent message and her pre-meeting belief.

In our main result, we show that the belief dynamics process converges with probability 1 to a set of intervals that is *minimal mutually confirming*. Given each agent's belief lies in her corresponding interval, these intervals are the convex combinations of the interpretations the agents use when communicating. Furthermore, we show that the belief of an agent eventually *fluctuates* on her corresponding interval whenever the interval is proper. As a consequence, the belief dynamics has a steady state if and only if there exists a minimal mutually confirming set such that all its intervals are degenerate.

We remark that outcomes with a steady state are non-generic as long as conflicts of interest are not too large. Hence, we can conclude that the introduction of conflicts of interest generically leads not only to persistent disagreement among the agents, but also to fluctuating beliefs. And as the belief dynamics approaches a consensus value in absence of conflicting interests, we identify conflicts of interest as the channel that drives our results, which hence suggests that the classical consensus result is not stable with respect to conflicts of interest.

Furthermore, we show that the belief dynamics – although failing to converge almost surely – converges in distribution to a random vector. And moreover, we find that the beliefs fluctuate in an ergodic way, i.e., the empirical averages of the agents' beliefs converge to their long-run expectations.

Though frequently observed in social science, the phenomenon of fluctuation is barely studied in the literature on communication in social networks, the only exception being Acemoglu et al. (2013). While their results are very close to ours, they achieve them with a different approach. Instead of conflicting interests, they introduce stubborn agents that never change their belief into a model of belief dynamics. This also leads to fluctuating beliefs when the other agents update regularly from different stubborn agents. In our model, a stubborn agent would be an agent that only communicates with herself.

Finally, we would like to comment briefly on some of our model choices. Our paper presents a first attempt to enrich a model of belief dynamics with a framework of communication that incorporates conflicting interests. We assume that agents are rational when communicating, but fail to account for the history of beliefs. This allows us to identify the role of conflicting interests in boundedly rational belief dynamics, in particular with respect to the classical result that a consensus emerges when the social network is connected. However, it would be interesting to see whether (partially) Bayesian models would generate similar results under conflicting interests. We leave this issue for future work.

A Appendix

Proof of Proposition 1

Suppose that agent j uses L' messages in equilibrium. We know that agent j's strategy is given by a partition $(C_l)_{l=1}^{L'}$, where $C_l = (m^{ij})^{-1}(m_l) = [c_{l-1}, c_l)$ is such that $0 = c_0 < c_1 < \cdots < c_{L'} = 1$ and

$$|c_l + (b_j - b_i) - y_l| = |c_l + (b_j - b_i) - y_{l+1}| \text{ for } l = 1, 2, \dots, L' - 1.$$

And furthermore, agent i's strategy is given by interpretations

$$y_l = y^{ij}(m_l) = \underset{y \in \mathbb{R}}{\operatorname{argmax}} \min_{x \in C_l} u_i(x, y) = \frac{c_l - c_{l-1}}{2} \text{ for } l = 1, 2, \dots, L'.$$

Thus, j's strategy satisfies

$$\begin{aligned}
&|c_l + (b_j - b_i) - y_l| = |c_l + (b_j - b_i) - y_{l+1}| \\
\Leftrightarrow &|c_l + (b_j - b_i) - (c_l + c_{l-1})/2| = |c_l + (b_j - b_i) - (c_{l+1} + c_l)/2| \\
\Leftrightarrow &|(c_l - c_{l-1})/2 + (b_j - b_i)| = |(c_l - c_{l+1})/2 + (b_j - b_i)|.
\end{aligned}$$

By monotonicity of the c_l, this yields

$$c_{l+1} = 2c_l - c_{l-1} + 4(b_j - b_i) \text{ for } l = 1, 2, \dots, L(ij) - 1.$$

And by the boundary condition $c_0 = 0$, it follows that

$$c_l = lc_1 + 2l(l-1)(b_j - b_i) \text{ for } l = 1, 2, \dots, L(ij).$$

The other boundary condition, $c_{L(ij)} = 1$, implies that $c_1 = 1/L(ij) - 2(L(ij) - 1)(b_j - b_i)$ and hence,

$$\begin{aligned}
c_l &= lc_1 + 2l(l-1)(b_j - b_i) \\
&= l\big(1/L(ij) - 2(L(ij) - 1)(b_j - b_i)\big) + 2l(l-1)(b_j - b_i) \\
&= l/L(ij) - 2l(L(ij) - l)(b_j - b_i). \qquad (1)
\end{aligned}$$

Hence,

$$\begin{aligned}
y_l &= (c_l + c_{l-1})/2 \\
&= (2l-1)/(2L(ij)) - l(L(ij) - l)(b_j - b_i) - (l-1)(L(ij) - l + 1)(b_j - b_i) \\
&= (2l-1)/(2L(ij)) - \big(l(2L(ij) - 2l + 1) - (L(ij) - l + 1)\big)(b_j - b_i) \\
&= (2l-1)/(2L(ij)) - \big((2l-1)L(ij) - 2(l^2 - l) - 1\big)(b_j - b_i).
\end{aligned}$$

Next, we show that there is an upper bound on the number of messages induced in equilibrium under conflict of interest, i.e., $b_i \neq b_j$. The equilibrium with L' messages exists if the strategy determined by (1) is feasible, which, by monotonicity, is the case if and only if

$$\begin{cases} c_1 = 1/L' - 2(L'-1)(b_j - b_i) \geq 0 \\ c_{L'-1} = (L'-1)/L' - 2(L'-1)(b_j - b_i) \leq 1 \end{cases}$$
$$\Leftrightarrow 1/(2L') \geq |(L'-1)(b_j - b_i)|. \tag{2}$$

Thus,

$$\bar{L} = \max\{l \in \mathbb{N} \mid 1/(2l) \geq |(l-1)(b_j - b_i)|\} \tag{3}$$

is the upper bound on the number of messages induced in equilibrium. Note that (2) has only finitely many positive integer solutions, among them $L' = 1$, and thus, (3) is well-defined.

If $b_i = b_j$, then there is no bound due to the biases and thus, the number of messages in equilibrium is bounded by $\bar{L} = L$. Altogether, agent j uses $1 \leq L(ij) = \bar{L}$ messages in optimal (i.e., most informative) equilibrium and moreover, this equilibrium is essentially unique since i's interpretations and j's partition are unique, which finishes the proof.

Proof of Theorem 1

To prove the theorem, we first construct a homogeneous Markov chain $\{\widetilde{x}(s)\}_{s\in\mathbb{N}} = \{(\widetilde{x}_i(s))_{i\in\mathcal{N}}\}_{s\in\mathbb{N}}$ in discrete time with finite n-dimensional state space $\mathcal{A} = \times_{i\in\mathcal{N}} A_i$, where A_i denotes the set of states for agent i. We know that we can replace the time-continuous belief dynamics process $\{x(t)\}_{t\geq 0}$ by the time-discrete process $\{x(s)\}_{s\in\mathbb{N}}$, where $x(s)$ is the vector of beliefs at time slot s. In the following, we also simplify the state space of the process. We find a partition of the unit interval such that it is enough to know in which element of the partition each agent's belief is.

Let $i \in \mathcal{N}$ and $C^i = \cup_{j\in\mathcal{N}_i} C^{ij}$ denote the set of points for which some agent $j \in \mathcal{N}_i$ is indifferent between two messages when communicating with i. Furthermore, $Y^i = \cup_{j\in\mathcal{N}_i} Y^{ij}$ denotes the set of agent i's interpretations. We assume without loss of generality that the set $C^i \cup Y^i$ consists of rational numbers for all $i \in \mathcal{N}$.[22] Then, there exists a lowest common denominator d of the set $\cup_{i\in\mathcal{N}} C^i \cup Y^i$.

[22]If some number is irrational, then we can approximate it arbitrarily well by a rational number, e.g., using the method of continued fractions.

This allows us to define the partition $C_d = \{k/d \mid 0 \leq k \leq d\}$ of $[0,1]$, where each partition element (without loss of generality) is an interval $[(k-1)/d, k/d)$, $k = 1, 2, \dots, d$. This partition distinguishes the beliefs of the agents finely enough to keep track of how the belief dynamics process evolves as we will show. Take $i \in \mathcal{N}$, $j \in \mathcal{N}_i$ and suppose that

$$x_i(s-1) \in [(k_i - 1)/d, k_i/d) \text{ and } x_j(s-1) \in [(k_j - 1)/d, k_j/d),$$

$1 \leq k_i, k_j \leq d$. By construction of the partition, there exists $1 \leq l \leq L(ij)$ such that $x_j(s-1) \in [c_{l-1}, c_l)$, i.e., C_d is fine enough to determine the message $m^{ij}(x_j(s-1))$ sent in equilibrium by agent j. Moreover, also by construction, there exists $1 \leq \bar{k} \leq d-1$ such that the interpretation of this message is $y^{ij}\big(m^{ij}(x_j(s-1))\big) = \bar{k}/d$. And since $x_i(s-1) \in [(k_i - 1)/d, k_i/d)$, it follows that

$$\begin{aligned} x_i(s) = 1/2(x_i(s-1) + \bar{k}/d) &\in [(k_i - 1 + \bar{k})/(2d), (k_i + \bar{k})/(2d)) \\ &\subseteq [(\lceil (k_i + \bar{k})/2 \rceil - 1)/d, \lceil (k_i + \bar{k})/2 \rceil /d), \end{aligned}$$

i.e., C_d is also fine enough to determine i's updated belief and altogether, it is fine enough to keep track of the belief dynamics process.

Therefore, we can identify the continuous state space $[0,1]^n$ of $\{x(s)\}_{s\in\mathbb{N}}$ with the finite state space $\mathcal{A} = A^n = \{a_1, a_2, \dots, a_d\}^n$ of $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$, where $a_k \equiv [(k-1)/d, k/d)$, $k = 1, 2, \dots, d$. In other words, a state $a \in \mathcal{A}$ specifies for each agent the partition element of C_d her belief is in at time slot s.

Let $\bar{x}(a_k) = (2k-1)/(2d)$ denote the average value of $[(k-1)/d, k/d)$ and furthermore, let $\tilde{y}^{ij}(a_k) = y^{ij}\big(m^{ij}(\bar{x}(a_k))\big)$ denote i's interpretation of j's message when j's belief is in $[(k-1)/d, k/d)$. We define the transition probabilities of $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$ as follows:

$$\mathbb{P}[\widetilde{x}(s) = (a^{-i}, a_l) \mid \widetilde{x}(s-1) = a] = 1/n \sum_{\substack{j\in\mathcal{N}_i: \\ (a^i, a^j)\in B^{ij}(l)}} p_{ij} \tag{4}$$

for all $a \in \mathcal{A}$ and $l \in \{1, 2, \dots, d\}$, where

$$B^{ij}(l) = \{(a_k, a_{k'}) \in A^2 \mid 1/2[\bar{x}(a_k) + \tilde{y}^{ij}(a_{k'})] \in [(l-1)/d, l/d)\}$$

is the set of all pairs of individual states (a^i, a^j) such that agent i changes from state a^i to state $a^{i'} = a_l$ given that she updates from agent j who is in state a^j. All other transition probabilities (i.e., those where more than one component changes) are assumed to be equal to zero. By construction, the following result holds.

Lemma 1. *$\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$ is a homogeneous Markov chain with finite state space $\mathcal{A}$ and transition probabilities given by* (4)*, and, in particular, at any time slot s,*

$$\widetilde{x}(s) = (a_{k_1}, a_{k_2}, \ldots, a_{k_n}) \text{ if and only if } x(s) \in \times_{i\in\mathcal{N}}[(k_i-1)/d, k_i/d).$$

Furthermore, for a set of states $Z \subseteq \mathcal{A}$, let $Z|_k = \{a \in A \mid \exists z \in Z : z_k = a\}$ denote the set of all possible values the k^{th} component of states in Z can take. Then, the following holds.

Lemma 2. *If $Z \subseteq \mathcal{A}$ is a recurrent communication class of $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$, then $\{\mathcal{I}_i\}_{i\in\mathcal{N}}$ is a minimal mutually confirming set of $\{x(t)\}_{t\geq 0}$, where*

(i) $\mathcal{I}_i = \bigcup_{k:\, a_k\in Z|_i}[(k-1)/d, k/d]$ if $|Z|_i| \geq 2$, and

(ii) $\mathcal{I}_i = (k-1)/d$ or $\mathcal{I}_i = k/d$ if $Z|_i = \{a_k\}$.

Proof. Suppose that Z is a recurrent communication class of $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$, i.e., the Markov chain will never leave this class and each state $z \in Z$ is visited infinitely often by $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$. We show that $\{\mathcal{I}_i\}_{i\in\mathcal{N}}$ is a minimal mutually confirming set of $\{x(t)\}_{t\geq 0}$.

Note that for $i \in \mathcal{N}$ and each individual state $z^i \in Z|_i$, it is $\widetilde{x}_i(s) = z^i$ for infinitely many time slots s. Let

$$Y_Z^i = \bigcup_{j\in\mathcal{N}_i} \bigcup_{z^j\in Z|_j} \tilde{y}^{ij}(z^j)$$

denote the set of all interpretations of agent i when $\widetilde{x}(s) \in Z$. Note that if $a_k, a_{k'} \in Z|_i$ for $k < k'$, then also $a_{k''} \in Z|_i$ for all $k < k'' < k'$. Thus, if $|Z|_i| \geq 2$,

$$\mathcal{I}_i = \bigcup_{k:\, a_k\in Z|_i} [(k-1)/d, k/d] = \operatorname{conv}(Y_Z^i)$$

since all intervals $[(k-1)/d, k/d]$ in the union are visited by i. On the other hand, if $Z|_i = \{a_k\}$, then i always uses the same interpretation when in Z, either $(k-1)/d$ or k/d. Hence, $\mathcal{I}_i = (k-1)/d = \operatorname{conv}(Y_Z^i)$ or $\mathcal{I}_i = k/d = \operatorname{conv}(Y_Z^i)$. Altogether, we have $\mathcal{I}_i = \operatorname{conv}(Y_Z^i)$ for all $i \in \mathcal{N}$. And furthermore, note that

$$\begin{aligned} Y_Z^i = \bigcup_{j\in\mathcal{N}_i} \bigcup_{z^j\in Z|_j} \tilde{y}^{ij}(z^j) &= \bigcup_{j\in\mathcal{N}_i} \bigcup_{a_k\in Z|_j} \tilde{y}^{ij}(a_k) \\ &= \bigcup_{j\in\mathcal{N}_i} \bigcup_{a_k\in Z|_j} y^{ij}\big(m^{ij}(\bar{x}(a_k))\big) \\ &= \bigcup_{j\in\mathcal{N}_i} \bigcup_{k:\, a_k\in Z|_j} y^{ij}\big(m^{ij}((2k-1)/(2d))\big) \\ &= \bigcup_{j\in\mathcal{N}_i} Y^{ij}\big|_{\mathcal{I}_j}, \end{aligned}$$

where the last equality follows from the definition of $\mathcal{I}_j$. Hence, we get $\mathcal{I}_i = \text{conv}\left(\bigcup_{j\in\mathcal{N}_i} Y^{ij}|_{\mathcal{I}_j}\right)$ for all $i \in \mathcal{N}$, i.e., we have shown that $\{\mathcal{I}_i\}_{i\in\mathcal{N}}$ is a mutually confirming set of $\{x(t)\}_{t\geq 0}$ and furthermore, it is also minimal since by assumption Z is a recurrent communication class of $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$, which finishes the proof. □

Since the state space of $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$ is finite, there exists an almost surely finite stopping time τ such that for any initial state $\widetilde{x}(0) \in \mathcal{A}$,

$$\widetilde{x}(s) \in \{a \in \mathcal{A} \mid \exists Z \ni a \text{ recurrent communication class of } \{\widetilde{x}(s)\}_{s\in\mathbb{N}}\}^{23}$$

under the event $\{\tau = s\}$. So, suppose that $\widetilde{x}(s) \in Z$. We show that this implies that the original chain converges almost surely to the minimal mutually confirming set $\{\mathcal{I}_i\}_{i\in\mathcal{N}}$ defined in Lemma 2.

If $|Z|_i| \geq 2$, then part (i) of Lemma 2 implies that $x_i(s) \in \mathcal{I}_i$ almost surely. Furthermore, since Z is a recurrent communication class of $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$, the boundaries of $\mathcal{I}_i$ are used infinitely often as interpretations by i and thus, $x_i(\tau)$ fluctuates on $\mathcal{I}_i$. On the other hand, if $Z|_i = \{a_k\}$, agent i uses only a single interpretation when updating since $(k-1)/d$ and k/d cannot be both interpretations by choice of the partition C_d. This implies that, without loss of generality, almost surely $x_i(t) \to k/d = \mathcal{I}_i$ for $t \to \infty$, and hence, almost surely $x(t) \to \{\mathcal{I}_i\}_{i\in\mathcal{N}}$ for $t \to \infty$, which finishes the proof.

Proof of Proposition 2

Recall from the proof of Theorem 1 that $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$ is a homogeneous Markov chain in discrete time and with finite state space $\mathcal{A}$ that keeps track of how the belief dynamics process $\{x(t)\}_{t\geq 0}$ evolves (Lemma 1). Notice that for each minimal mutually confirming set of intervals $\{\mathcal{I}_i^k\}_{i\in\mathcal{N}}$ of $\{x(t)\}_{t\geq 0}$, there exists a unique recurrent communication class Z^k of $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$ such that $\{x(t)\}_{t\geq 0}$ converges to $\{\mathcal{I}_i^k\}_{i\in\mathcal{N}}$ if and only if $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$ converges to Z^k. Therefore, the following result proves part (i).

Lemma 3. *(Brémaud, 1999, p. 157, Theorem 6.2)*

Suppose $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$ has recurrent communication classes $Z^1, Z^2, \ldots, Z^r$. Then, for any initial state $a \in \mathcal{A}$, there exists a vector of probabilities $(q_1, q_2, \ldots, q_r)$ such that $\{\widetilde{x}(s)\}_{s\in\mathbb{N}}$ converges to Z^k with probability $q_k = q_k(a)$, k=1,2,...,r.

In order to prove part (ii), let us first introduce some notation. Recall from the proof of Theorem 1 that d denotes the lowest common denominator of the set

[23] We refer, e.g., to Brémaud (1999) for this result.

$\cup_{i\in\mathcal{N}} C^i \cup Y^i$. Let $\{d_l\}_{l=0}^{\infty}$ denote a sequence of common denominators of $\cup_{i\in\mathcal{N}} C^i \cup Y^i$ such that $d_0 = d$ and $d_l \to \infty$ for $l \to \infty$ and denote $\{\widetilde{x}(d_l; s)\}_{s\in\mathbb{N}}$ the corresponding finite state Markov chain and $\mathcal{A}(d_l)$ its state space (see Theorem 1).[24] Notice that if $x_i(s) \in [(k-1)/d_l, k/d_l)$ (and thus $\widetilde{x}_i(d_l; s) = a_k \in \mathcal{A}(d_l)|_i$), $i \in \mathcal{N}$, $k = 1, 2, \ldots, d_l$, then

$$|x_i(s) - \bar{x}[d_l](a_k)| \le \frac{1}{2d_l} \to 0 \text{ for } l \to \infty,$$

where $\bar{x}[d_l](a_k) = (2k-1)/(2d_l)$ denotes the average value of $[(k-1)/d_l, k/d_l)$. Thus, we can approximate the belief dynamics process arbitrarily well by the finite state Markov chains.

We know from part (i) that, for each d_l, $\{\widetilde{x}(d_l; s)\}_{s\in\mathbb{N}}$ converges to the recurrent communication class $Z^k(d_l)$ with probability $q_k = q_k(a)$ if $\widetilde{x}(d_l; 0) = a$. Notice that the probabilities are independent of d_l and that these classes are aperiodic since the agent updating her belief is chosen randomly. It follows from the latter that $\{\widetilde{x}(d_l; s)\}_{s\in\mathbb{N}}$ has a unique stationary distribution $\Pi^k[d_l]$ on $Z^k(d_l)$, i.e., $\Pi^k[d_l](a) > 0$ if and only if $a \in Z^k(d_l)$.

Suppose that $\{\widetilde{x}(d_l; s)\}_{s\in\mathbb{N}}$ converges to $Z^k(d_l)$ (henceforth $\widetilde{x}(d_l; s) \Rightarrow Z^k(d_l)$). This implies that its distribution converges to $\Pi^k[d_l]$, i.e.,

$$\lim_{s\to\infty} \mathcal{L}(\widetilde{x}(d_l; s) \mid \widetilde{x}(d_l; s) \Rightarrow Z^k(d_l)) = \Pi^k[d_l].$$ [25]

To simplify the notation, we omit in the following the condition $\widetilde{x}(d_l; s) \Rightarrow Z^k(d_l)$. For $x \in [0, 1]^n$, define

$$\underline{\mathcal{A}}(d_l; x) := \{(a_{k_1}, a_{k_2}, \ldots, a_{k_n}) \in \mathcal{A}(d_l) \mid k_i/d_l \le x_i \ \forall i \in \mathcal{N}\} \text{ and}$$
$$\overline{\mathcal{A}}(d_l; x) := \{(a_{k_1}, a_{k_2}, \ldots, a_{k_n}) \in \mathcal{A}(d_l) \mid (k_i - 1)/d_l < x_i \ \forall i \in \mathcal{N}\}.$$

Notice that, for all $x \in [0, 1]^n$,

$$\Pi^k[d_l](\overline{\mathcal{A}}(d_l; x)) - \Pi^k[d_l](\underline{\mathcal{A}}(d_l; x)) \to 0 \text{ as } l \to \infty.$$

Therefore, let $\hat{x}^k$ be a random belief vector with cumulative distribution function

$$\mathcal{F}_{\hat{x}^k}(x) := \lim_{l\to\infty} \Pi^k[d_l](\overline{\mathcal{A}}(d_l; x)) = \lim_{l\to\infty} \Pi^k[d_l](\underline{\mathcal{A}}(d_l; x)), x \in [0, 1]^n. \tag{5}$$

[24] For simplicity, we write $\mathbb{P} = \mathbb{P}_{d_l}$ for the corresponding probability measure. This is without loss of generality since the state space $\mathcal{A}(d_l)$ can be identified with the continuous state space of the belief dynamics process, see Theorem 1.

[25] We say that the distribution of $\{\widetilde{x}(d_l; s)\}_{s\in\mathbb{N}}$ converges to the distribution $\Pi^k[d_l]$ if $\lim_{s\to\infty} \mathbb{P}(\widetilde{x}(d_l; s) = a) = \Pi^k[d_l](a)$ for all $a \in \mathcal{A}(d_l)$.

Next, we show that the belief dynamics process converges in distribution to $\hat{x}^k$. By construction, we have

$$\mathbb{P}(\widetilde{x}(d_l;s) \in \underline{\mathcal{A}}(d_l;x)) \leq \mathcal{F}_{x(s)}(x) \leq \mathbb{P}(\widetilde{x}(d_l;s) \in \overline{\mathcal{A}}(d_l;x)), \text{ for all } s, l \in \mathbb{N},$$

where $\mathcal{F}_{x(s)}$ denotes the cumulative distribution function of the belief dynamics process at time slot s. And thus, for all $l \in \mathbb{N}$,

$$\begin{aligned}\Pi^k[d_l](\underline{\mathcal{A}}(d_l;x)) = \lim_{s\to\infty} \mathbb{P}(\widetilde{x}(d_l;s) \in \underline{\mathcal{A}}(d_l;x)) &\leq \lim_{s\to\infty} \mathcal{F}_{x(s)}(x) \\ &\leq \lim_{s\to\infty} \mathbb{P}(\widetilde{x}(d_l;s) \in \overline{\mathcal{A}}(d_l;x)) \\ &= \Pi^k[d_l](\overline{\mathcal{A}}(d_l;x)).\end{aligned}$$

Hence, together with (5), it follows that

$$\lim_{s\to\infty} \mathcal{F}_{x(s)}(x) = \lim_{l\to\infty} \Pi^k[d_l](\overline{\mathcal{A}}(d_l;x)) = \lim_{l\to\infty} \Pi^k[d_l](\underline{\mathcal{A}}(d_l;x)) = \mathcal{F}_{\hat{x}^k}(x),$$

i.e., $\lim_{s\to\infty} \mathcal{L}(x(s) \mid x(t) \Rightarrow \mathcal{I}^k) = \mathcal{L}(\widehat{x}^k)$ for all $k = 1, 2, \ldots, r$. In particular, this is equivalent to

$$\lim_{s\to\infty} \mathbb{E}[\varphi(x(s)) \mid x(s) \Rightarrow \mathcal{I}^k] = \mathbb{E}[\varphi(\widehat{x}^k)]$$

for all $k = 1, 2, \ldots, r$ and for all bounded and continuous test functions $\varphi : [0,1]^n \to \mathbb{R}$ (see Klenke (2007, p. 257, Corollary 13.24)), which finishes the proof.

Proof of Corollary 2

The first part follows immediately from Proposition 2. For part (ii), notice first that the sequence of the lengths of intervals between meetings, $\{t_{s+1} - t_s\}_{s=0}^{\infty}$, is independent and identically distributed with mean $1/n$. Therefore, by the strong law of large numbers, with probability 1,

$$\lim_{u\to\infty} \frac{1}{u} t_u = \lim_{u\to\infty} \frac{1}{u} \sum_{s=0}^{u-1} (t_{s+1} - t_s) = \frac{1}{n}.$$

This implies that

$$\lim_{u\to\infty} \mathbb{E}\left[\frac{t_{s+1} - t_s}{1/u \cdot t_u} - 1\right] = 0.$$

It follows that, for any bounded and continuous test function $\varphi : [0,1]^n \to \mathbb{R}$ and any $\varepsilon > 0$, using the continuity of $\mathbb{P}$ with respect to monotonous sets,

$$\begin{aligned} &\lim_{v\to\infty} \mathbb{P}\left(\sup_{u\geq v} \left| \frac{1}{t_u}\sum_{s=0}^{u-1}(t_{s+1}-t_s)\varphi(x(t_s)) - \frac{1}{u}\sum_{s=0}^{u-1}\varphi(x(t_s))\right| > \varepsilon\right) \\ =&\lim_{v\to\infty} \mathbb{P}\left(\sup_{u\geq v} \left| \frac{1}{u}\sum_{s=0}^{u-1}\left(\frac{t_{s+1}-t_s}{1/u\cdot t_u} - 1\right)\varphi(x(t_s))\right| > \varepsilon\right) \\ =&\mathbb{P}\left(\limsup_{u\to\infty} \left| \frac{1}{u}\sum_{s=0}^{u-1}\left(\frac{t_{s+1}-t_s}{1/u\cdot t_u} - 1\right)\varphi(x(t_s))\right| > \varepsilon\right) = 0 \end{aligned} \tag{6}$$

since $\{x(t_s)\}_{s\in\mathbb{N}}$ is independent of $\{t_{s+1}-t_s\}_{s=0}^{\infty}$ and converges in distribution and almost surely $\sup_{s\geq 0}\|x(t_s)\|_1 \leq 1$. Notice that (6) is equivalent to

$$\lim_{u\to\infty} \left| \frac{1}{t_u}\sum_{s=0}^{u-1}(t_{s+1}-t_s)\varphi(x(t_s)) - \frac{1}{u}\sum_{s=0}^{u-1}\varphi(x(t_s))\right| = 0 \text{ almost surely.} \tag{7}$$

Furthermore, recall from the proof of Proposition 2 that for all $i \in \mathcal{N}$,

$$|x_i(t_s) - \bar{x}[d_l](\widetilde{x}_i(d_l;s))| \leq \frac{1}{2d_l} \to 0 \text{ for } l \to \infty.$$

Hence, also

$$\|x(t_s) - (\bar{x}[d_l](\widetilde{x}_i(d_l;s)))_{i\in\mathcal{N}}\|_\infty \leq \frac{1}{2d_l} \to 0 \text{ for } l \to \infty.$$

And thus, by continuity of φ, there exists a sequence $\{\varepsilon_l\}_{l=0}^{\infty}$, $\varepsilon_l \to 0$ for $l \to \infty$, such that

$$|\varphi(x(t_s)) - \varphi\left((\bar{x}[d_l](\widetilde{x}_i(d_l;s)))_{i\in\mathcal{N}}\right)| \leq \varepsilon_l.$$

This implies that

$$\begin{aligned} &\lim_{u\to\infty} \left| \frac{1}{u}\sum_{s=0}^{u-1}\varphi(x(t_s)) - \frac{1}{u}\sum_{s=0}^{u-1}\varphi\left((\bar{x}[d_l](\widetilde{x}_i(d_l;s)))_{i\in\mathcal{N}}\right)\right| \\ \leq &\lim_{u\to\infty}\frac{1}{u}\sum_{s=0}^{u-1}\left|\varphi(x(t_s)) - \varphi\left((\bar{x}[d_l](\widetilde{x}_i(d_l;s)))_{i\in\mathcal{N}}\right)\right| \\ \leq &\varepsilon_l. \end{aligned} \tag{8}$$

Next, recall from the proof of Proposition 2 that the distribution of $\{\widetilde{x}_i(d_l;s)\}_{s\in\mathbb{N}}$ converges to its unique stationary distribution $\Pi[d_l]$. Therefore, the ergodic theorem for finite state Markov chains (see Brémaud (1999, p. 111, Theorem 4.1)) yields, with probability 1,

$$\lim_{u\to\infty}\frac{1}{u}\sum_{s=0}^{u-1}\varphi\left((\bar{x}[d_l](\widetilde{x}_i(d_l;s)))_{i\in\mathcal{N}}\right) = \sum_{a\in\mathcal{A}(d_l)}\varphi((\bar{x}[d_l](a^i))_{i\in\mathcal{N}})\Pi[d_l](a). \tag{9}$$

Applying (7)–(9), we get, with probability 1,

$$\begin{aligned}
&\lim_{t\to\infty}\left|\frac{1}{t}\int_0^t \varphi(x(t)) - \mathbb{E}[\varphi(\hat{x})]\right| \\
=&\lim_{u\to\infty}\left|\frac{1}{t_u}\sum_{s=0}^{u-1}(t_{s+1}-t_s)\varphi(x(t_s)) - \mathbb{E}[\varphi(\hat{x})]\right| \\
\leq&\lim_{u\to\infty}\left|\frac{1}{t_u}\sum_{s=0}^{u-1}(t_{s+1}-t_s)\varphi(x(t_s)) - \frac{1}{u}\sum_{s=0}^{u-1}\varphi(x(t_s))\right| \\
&+\lim_{u\to\infty}\left|\frac{1}{u}\sum_{s=0}^{u-1}\varphi(x(t_s)) - \frac{1}{u}\sum_{s=0}^{u-1}\varphi\left((\bar{x}[d_l](\widetilde{x}_i(d_l;s)))_{i\in\mathcal{N}}\right)\right| \\
&+\lim_{u\to\infty}\left|\frac{1}{u}\sum_{s=0}^{u-1}\varphi\left((\bar{x}[d_l](\widetilde{x}_i(d_l;s)))_{i\in\mathcal{N}}\right) - \mathbb{E}[\varphi(\hat{x})]\right| \\
\leq&\varepsilon_l + \left|\sum_{a\in\mathcal{A}(d_l)}\varphi((\bar{x}[d_l](a^i))_{i\in\mathcal{N}})\Pi[d_l](a) - \mathbb{E}[\varphi(\hat{x})]\right|.
\end{aligned}$$

Finally, since $\lim_{l\to\infty}\left|\sum_{a\in\mathcal{A}(d_l)}\varphi((\bar{x}[d_l](a^i))_{i\in\mathcal{N}})\Pi[d_l](a) - \mathbb{E}[\varphi(\hat{x})]\right| = 0$ by definition of $\hat{x}$, letting $l\to\infty$ finishes the proof.

References

Acemoglu, D., K. Bimpikis, and A. Ozdaglar (2014). Dynamics of information exchange in endogenous social networks. *Theoretical Economics 9*(1), 41–97.

Acemoglu, D., G. Como, F. Fagnani, and A. Ozdaglar (2013). Opinion fluctuations and disagreement in social networks. *Mathematics of Operations Research 38*(1), 1–27.

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

Ambrus, A. and S. Takahashi (2008). Multi-sender cheap talk with restricted state spaces. *Theoretical Economics 3*(1), 1–27.

Axelrod, R. (1997). The dissemination of culture a model with local convergence and global polarization. *Journal of conflict resolution 41*(2), 203–26.

Banerjee, A. and D. Fudenberg (2004). Word-of-mouth learning. *Games and Economic Behavior 46*(1), 1–22.

Blondel, V. D., J. M. Hendrickx, and J. N. Tsitsiklis (2009). On Krause's multi-agent consensus model with state-dependent connectivity. *Automatic Control, IEEE Transactions on 54*(11), 2586–97.

Bodansky, D. (2010). The Copenhagen climate change conference: a postmortem. *American Journal of International Law 104*(2), 230–40.

Boyd, S., A. Ghosh, B. Prabhakar, and D. Shah (2006). Randomized gossip algorithms. *Information Theory, IEEE Transactions on 52*(6), 2508–30.

Brémaud, P. (1999). *Markov Chains*. Springer.

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

Cohen, G. (2003). Party over policy: the dominating impact of group influence on political beliefs. *Journal of personality and social psychology 85*(5), 808–22.

Como, G. and F. Fagnani (2011). Scaling limits for continuous opinion dynamics systems. *The Annals of Applied Probability 21*(4), 1537–67.

Crawford, V. and J. Sobel (1982). Strategic information transmission. *Econometrica 50*(6), 1431–51.

Deffuant, G., D. Neau, F. Amblard, and G. Weisbuch (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems 3*, 87–98.

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.

DiMaggio, P., J. Evans, and B. Bryson (1996). Have american's social attitudes become more polarized? *American journal of Sociology*, 690–755.

Evans, J. (2003). Have americans' attitudes become more polarized?—an update. *Social Science Quarterly 84*(1), 71–90.

Förster, M., M. Grabisch, and A. Rusinowska (2013). Anonymous social influence. *Games and Economic Behavior 82*, 621–35.

Friedkin, N. E. and E. C. Johnsen (1990). Social influence and opinions. *Journal of Mathematical Sociology 15*(3-4), 193–206.

Gale, D. and S. Kariv (2003). Bayesian learning in social networks. *Games and Economic Behavior 45*(2), 329–46.

Galeotti, A., C. Ghiglino, and F. Squintani (2013). Strategic information transmission networks. *Journal of Economic Theory 148*(5), 1751–69.

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.

Hagenbach, J. and F. Koessler (2010). Strategic communication networks. *The Review of Economic Studies 77*(3), 1072–99.

Hegselmann, R. and U. Krause (2002). Opinion dynamics and bounded confidence – models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation 5*(3).

Jackson, M. O. and L. Yariv (2007). Diffusion of behavior and equilibrium properties in network games. *The American Economic Review 97*(2), 92–8.

Jäger, G., L. Metzger, and F. Riedel (2011). Voronoi languages: equilibria in cheap-talk games with high-dimensional types and few signals. *Games and Economic Behavior 73*(2), 517–37.

Klenke, A. (2007). *Probability theory: a comprehensive course*. Springer.

Knight, F. H. (1921). Risk, uncertainty and profit. *New York: Hart, Schaffner and Marx*.

Kramer, G. (1971). Short-term fluctuations in US voting behavior, 1896–1964. *American Political Science Review 65*(1), 131–43.

López-Pintado, D. (2008). Diffusion in complex social networks. *Games and Economic Behavior 62*(2), 573–90.

López-Pintado, D. (2012). Influence networks. *Games and Economic Behavior 75*(2), 776–87.

Lorenz, J. (2005). A stabilization theorem for dynamics of continuous opinions. *Physica A: Statistical Mechanics and its Applications 355*(1), 217–23.

Morris, S. (2000). Contagion. *The Review of Economic Studies 67*(1), 57–78.

Watts, D. (2003). Six degrees: the science of a connected age. *New York City, NY: WW Norton and Company*.

Zaller, J. (1992). *The nature and origins of mass opinion.* Cambridge University Press.