2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

Opinion fluctuations and persistent disagreement in social networks

Daron Acemoglu, Giacomo Como, Fabio Fagnani, Asuman Ozdaglar

***Abstract*— We study a tractable opinion dynamics model that generates long-run disagreements and persistent opinion fluctuations. Our model involves an inhomogeneous stochastic gossip process of continuous opinion dynamics in a society consisting of two types of agents: *regular agents*, who update their beliefs according to information that they receive from their social neighbors; and *stubborn agents*, who never update their opinions and might represent leaders, political parties or media sources attempting to influence the beliefs in the rest of the society. When the society contains stubborn agents with different opinions, the belief dynamics never lead to a consensus (among the regular agents). Instead, beliefs in the society almost surely fail to converge, the belief profile keeps on oscillating in an ergodic fashion, and it converges in law to a non-degenerate random vector.**

The structure of the graph describing the social network and the location of the stubborn agents within it shape the opinion dynamics. The expected stationary beliefs vector is proved to be harmonic on the graph, with every regular agent's value being the weighted average of its neighbors' values, and boundary conditions corresponding to the stubborn agents' beliefs.

We prove that, in large-scale societies which are *highly fluid*, meaning that the product of the mixing time of the random walk on the graph describing the social network and the relative size of the linkages to stubborn agents vanishes as the population size grows large, a condition of *homogeneous influence* emerges, whereby the stationary beliefs' marginal distributions of most of the regular agents have approximately equal first and second moment. Homogeneous influence in a highly fluid societies need not imply approximate consensus among the agents, whose beliefs may well oscillate in an essentially uncorrelated way.

I. INTRODUCTION

Disagreement among individuals in a society, even on central questions that have been debated for centuries, is the norm; agreement is the rare exception. How can disagreement of this sort persist for so long? Notably, such disagreement is not a consequence of lack of communication or some other factors leading to fixed opinions. Disagreement remains even as individuals communicate and sometimes change their opinions.

Existing models of communication and learning, based on Bayesian or non-Bayesian updating mechanisms, typically lead to consensus provided that communication takes place over a strongly connected network (e.g., Smith and Sorensen [32], Banerjee and Fudenberg [7], Acemoglu, Dahleh, Lobel and Ozdaglar [2], Bala and Goyal [6], Gale and Kariv [18], DeMarzo, Vayanos and Zwiebel [14], Golub and Jackson [19], Acemoglu, Ozdaglar and ParandehGheibi [3]), and are thus unable to explain persistent disagreements. One notable exception is provided by models that incorporate a form of *homophily* mechanism in communication, whereby individuals are more likely to exchange opinions or communicate with others that have similar beliefs, and fail to interact with agents whose beliefs differ from theirs by more than some given confidence threshold. This mechanism was first proposed by Axelrod [5] in the discrete opinion dynamics setting, and then by Krause [21], and Deffuant and Weisbuch [13], in the continuous opinion dynamics framework. Such beliefs dynamics typically lead to the emergence of different asymptotic opinion clusters (see, e.g., [23], [9], [11]); however, they are unable to explain persistent opinion fluctuations in the society.

In this paper, we investigate a tractable opinion dynamics model that generates both long-run disagreement and opinion fluctuations. We consider an *inhomogeneous stochastic gossip model* of communication wherein there is a fraction of *stubborn agents* in the society who never change their opinions. We show that the presence of stubborn agents with competing opinions leads to persistent opinion fluctuations and disagreement among the rest of the society.

More specifically, we consider a society envisaged as a social network of n interacting agents (or individuals), communicating and exchanging information. Each agent a starts with an opinion (or belief) $X_a(0) \in \mathbb{R}$ and is then activated according to a Poisson process in continuous time. Following this event, she meets one of the individuals in her *social neighborhood* according to a pre-specified stochastic process. This process represents an underlying *social network*. We distinguish between two types of individuals, stubborn and regular. Stubborn agents, which are typically few in number, never change their opinions: they might thus correspond to media sources, opinion leaders, or political parties wishing to influence the rest of the society, and, in a first approximation, not getting any feedback from it. In contrast, regular agents, which make up the great majority of the agents in the social network, update their beliefs to some weighted average of their pre-meeting belief and the belief of the agent they met. The opinions generated through this information exchange process form a Markov process whose long-run behavior is the focus of our analysis.

First, we show that, under general conditions, these opinion dynamics never lead to a consensus (among the regular agents). In fact, regular agents' beliefs almost surely fail to converge, and keep on oscillating in an ergodic fashion.

D. Acemoglu is with the Department of Economics at the Massachusetts Institute of Technology. daron@mit.edu.

G. Como is with the Department of Automatic Control, Lund University, Sweden. giacomo.como@control.lth.se.

F. Fagnani is with the Department of Mathematics at the Politecnico di Torino. fabio.fagnani@polito.it.

A. Ozdaglar is with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. asuman@mit.edu.

Instead, the belief of each regular agent converges in law to a non-degenerate stationary random variable, and, similarly, the vector of beliefs of all agents jointly converge to a non-degenerate stationary random vector. This model therefore provides a new approach to understanding persistent disagreements and opinion fluctuations.

Second, we investigate how the structure of the graph describing the social network and the location of stubborn agents within it shape the behavior of the opinion dynamics. The expected belief vector is proved to evolve according to an ordinary differential equation coinciding with the Kolmogorov backward equation of a continuous-time Markov chain on the graph with absorbing states corresponding to the stubborn agents, and hence to converge to a harmonic vector, with every regular agent's value being the weighted average of its neighbors' values, and boundary conditions corresponding to the stubborn agents' beliefs. Expected cross-products of the agents' beliefs allow for a similar characterization in terms of coupled random walks on the graph describing the social network. The characterization of the expected stationary beliefs as harmonic functions is then used in order to find explicit solutions for some social networks with particular structure or symmetries.

Third, in what we consider the most novel contribution of our analysis, we study the behavior of the stationary beliefs in large-scale *highly fluid* social networks, defined as networks where the product between the fraction of edges incoming in the stubborn agent set times the mixing time of the associated random walk is small. We show that in highly fluid social networks, the expected value and variance of the stationary beliefs of most of the agents concentrate around certain values as the population size grows large. We refer to this result as *homogeneous influence* of stubborn agents on the rest of the society—meaning that their influence on most of the agents in the society are approximately the same. The applicability of this result is then proved by providing several examples of large-scale random networks, including the Erdös–Rényi graph in the connected regime, power law networks, and small-world networks. Finally, we argue that homogeneous influence in a highly fluid societies need not imply approximate consensus among the agents, whose beliefs may well oscillate in an essentially uncorrelated way, leaving a deeper understanding of this topic as a matter for future work.

Our main contribution partly stems from novel applications of several techniques of applied probability in the study of opinion dynamics. In particular, convergence in law and ergodicity of the agents' beliefs is established by first rewriting the dynamics in the form of an iterated affine function system and then using techniques developed in this field [15]. On the other hand, our estimates of the behavior of the expected values and variances of the stationary beliefs in large-scale highly fluid networks are based on techniques from the theory of reversible Markov chains, including approximate exponentiality of the hitting times and fast mixing [4], [22], as well as on results in modern random graph theory [16].

In addition to the aforementioned works on learning and opinion dynamics, this paper is related to some of the literature in the statistical physics of social dynamics: see [10] and references therein for an overview of such research line. More specifically, our model is closely related to some work by Mobilia and co-authors [25], [26], [27], who study a variation of the discrete opinion dynamics model, also called the *voter model*, with inhomogeneities, there referred to as *zealots*: such zealots are agents which tend to favor one opinion in [25], [26], or are in fact equivalent to our stubborn agents in [27]. These works generally present analytical results for some regular graphical structures (such as regular lattices [25], [26], or complete graphs [27]), and are then complemented by numerical simulations. In contrast, we prove convergence in distribution and characterize the properties of the limiting distribution for general finite graphs. Even though our model involves continuous belief dynamics, the voter model with zealots of [27] can be recovered as a special case of our general framework.

Our work is also related to work on consensus and gossip algorithms, which is motivated by different problems, but typically leads to a similar mathematical formulation (Tsitsiklis [33], Tsitsiklis, Bertsekas and Athans [34], Jadbabaie, Lin and Morse [20], Olfati-Saber and Murray [30], Olshevsky and Tsitsiklis [31], Fagnani and Zampieri [17], Nedić and Ozdaglar [28]). In consensus problems, the focus is on whether the beliefs or the values held by different units (which might correspond to individuals, sensors, or distributed processors) converge to a common value. Our analysis here does not focus on limiting consensus of values, but in contrast, characterizes the stationary fluctuations in values.

The rest of this paper is organized as follows: In Section II, we introduce our model of interaction between the agents, describing the resulting evolution of individual beliefs, and we discuss two special cases, in which the arguments simplify particularly, and some fundamental features of the general case are highlighted. Section III presents convergence results on the evolution of agent beliefs over time, for a given social network: the beliefs are shown to converge in distribution, and to be an stationary process, while in general almost surely they do not converge sample-path-wise. In Section IV, we first provide a characterization of the expected stationary beliefs in terms of the hitting probabilities of a random walks on the graph describing the social network. Then, we exploit this characterization in order to provide bounds on the level of dispersion of the expected stationary beliefs' vector: it is shown that, in highly fluid networks, most of the agents have almost the same stationary belief and variance. Section V presents some concluding remarks. All the statements will be presented without proof, which can be found in the journal version of this work [1], along with a more general formulation, and more detailed results and examples.

Before proceeding, we establish some notational conventions and terminology to be followed throughout the paper. We shall typically label the entries of vectors by elements

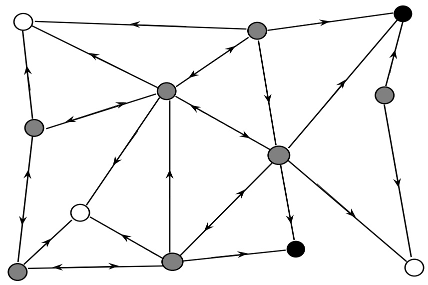

Fig. 1. A social network with seven regular agents (colored in grey), and five stubborn agents (colored in white, and black, respectively). Links are only incoming to the stubborn agents, while links between pairs of regular agents may be uni- or bi-directional.

of finite alphabets, rather than non-negative integers, hence $\mathbb{R}^{\mathcal{I}}$ will stand for the set of vectors with entries labeled by elements of the finite alphabet $\mathcal{I}$. An index denoted by a lower-case letter will implicitly be assumed to run over the finite alphabet denoted by the corresponding calligraphic upper-case letter (e.g. $\sum_i$ will stand for $\sum_{i\in\mathcal{I}}$). For two non-negative sequences $\{a_n\}$, $\{b_n\}$, we will write $a_n = O(b_n)$ if for some positive constant K, $a_n \leq K b_n$ for all sufficiently large n.

II. Belief evolution model

We consider a finite population $\mathcal{V}$ of interacting agents, of possibly very large size $n := |\mathcal{V}|$. The connectivity among the agents is described by a simple undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, whose node set is identified with the agent population, and where $\mathcal{E}$ stands for the set of links among the agents.

At time $t \geq 0$, each agent $v \in \mathcal{V}$ holds a *belief* (or *opinion*) about an underlying state of the world, denoted by $X_v(t) \in \mathbb{R}$. The full vector of beliefs at time t will be denoted by $X(t) := \{X_v(t) : v \in \mathcal{V}\}$. We distinguish between two types of agents: regular and stubborn. Regular agents repeatedly update their own beliefs, based on the observation of the beliefs of their neighbors in $\mathcal{G}$. *Stubborn agents* never change their opinions. Agents which are not stubborn are called *regular*. We shall denote the set of regular agents by $\mathcal{A}$, the set of stubborn agents by $\mathcal{S}$, so that the set of all agents is $\mathcal{V} = \mathcal{A} \cup \mathcal{S}$ (see Figure 1).

More specifically, the agents' beliefs evolve according to the following stochastic update process. At time $t = 0$, each agent $v \in \mathcal{V}$ starts with an initial belief $X_v(0)$. The beliefs of the stubborn agents stay constant in time:

$$X_s(t) = X_s(0) =: x_s\,, \qquad s \in \mathcal{S}\,.$$

In contrast, the beliefs of the regular agents are updated as follows. To every ordered pair of agents of the form (a, v), where necessarily $a \in \mathcal{A}$, $v \in \mathcal{V}$, and $\{a, v\} \in \mathcal{E}$, a clock is associated, ticking at the times of an independent Poisson process of rate $1/d_a$, where d_a is the degree of a in $\mathcal{G}$. If the (a, v)-th clock ticks at time t, agent a meets agent v and updates her belief to a convex combination of her own current belief and the current belief of agent v:

$$X_a(t) = (1-\theta)X_a(t^-) + \theta X_v(t^-)\,, \tag{1}$$

where $X_v(t^-)$ stands for the left limit $\lim_{u\uparrow t} X_v(u)$. Here, the scalar $\theta \in (0, 1]$ is a *trust parameter* that represents the confidence that each regular agent $a \in \mathcal{A}$ puts on her neighbors' beliefs.[1] For every regular agent $a \in \mathcal{A}$, let $\mathcal{S}_a \subseteq \mathcal{S}$ be the subset of stubborn agents which are reachable from a by a path in $\mathcal{G}$ with no intermediate steps in $\mathcal{S}$. We refer to $\mathcal{S}_a$ as the set of stubborn agents *influencing* a.

The pair $\mathcal{N} = (\mathcal{G}, \theta)$ contains the entire information about patterns of interaction among the agents, and will be referred to as the *social network*. Together with an assignment of a probability law for the initial belief vector, the social network designates a *society*. Throughout the paper, we make the following assumptions regarding the underlying social network.

Assumption 1: Every regular agent is influenced by some stubborn agent, i.e., $\mathcal{S}_a$ is non-empty for every a in $\mathcal{A}$.

For a given social network, we associate the stochastic matrix $P \in \mathbb{R}^{\mathcal{V}\times\mathcal{V}}$, and a probability vector π whose entries are defined by

$$P_{vw} = \begin{cases} 1/d_v & \text{if} \quad \{v, w\} \in \mathcal{E} \\ 0 & \text{if} \quad \{v, w\} \notin \mathcal{E} \end{cases} \qquad \pi_v := d_v / \sum_w d_w\,. \tag{2}$$

Observe that P is a reversible matrix, and π is its unique stationary probability vector.

III. Convergence in distribution and ergodicity of the beliefs

This section is devoted to studying the convergence properties of the random belief vector $X(t)$ for the general update model described in Sect. II. Figure 2 reports the typical sample-path behavior of the agents' beliefs for a simple social network with population size $n = 4$, and line graph topology, in which the two stubborn agents are positioned in the extremes and hold beliefs $x_0 < x_3$. As shown in Fig. 2(b), the beliefs of the two regular agents, $X_1(t)$, and $X_2(t)$, oscillate in the interval $[x_0, x_3]$, in an apparently chaotic way. On the other hand, the time averages of the two regular agents' beliefs rapidly approach a limit value, of $2x_0/3 + x_3/3$ for agent 1, and $x_0/3 + 2x_3/3$ for agent 2.

As we shall see below, such behavior is rather general. In our model of social network with at least two stubborn agents having non-coincident constant beliefs, the regular agent beliefs almost surely fail to converge. On the other hand, we shall prove that, regardless of the initial regular agents' beliefs, the belief vector $X(t)$ is convergent in distribution

[1]Note that the one described above can be interpreted as inhomogeneous gossip model of opinion dynamics, where all the regular agents use the trust parameter θ in their beliefs' updates, whereas the stubborn agents use trust parameter equal to 0. The case in which the stubborn agents have positive, though very small trust parameter, so that they get some feedback from their neighbors, is studied in [3], where it is shown that the society does converge to an asymptotic consensus around a random common Z, whose expected value is a convex combination of the initial beliefs of all the agents with higher relative weight to the stubborn (there called 'forceful') agents.

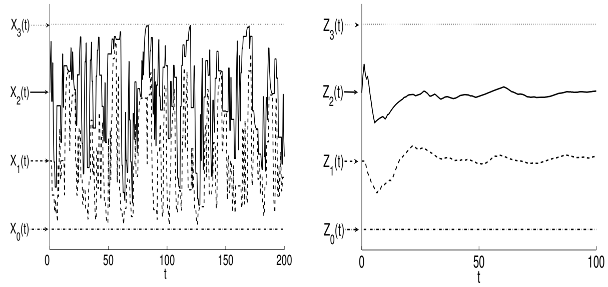

Fig. 2. Typical sample-path behavior of the beliefs, and their ergodic averages for a social network with population size $n = 4$. The topology is a line graph. The stubborn agents corresponds to the two extremes of the line, $\mathcal{S} = \{0, 3\}$, while their constant opinions are $x_0 = 0$, and $x_3 = 1$. The regular agent set is $\mathcal{A} = \{1, 2\}$. The confidence parameters, and the interaction rates are chosen to be $\theta_{av} = 1/2$, and $r_{av} = 1/3$, for all $a = 1, 2$, and $v = a \pm 1$. In picture (b), the trajectories of the actual beliefs $X_v(t)$, for $v = 0, 1, 2, 3$, are reported, whereas picture (c) reports the trajectories of their ergodic averages $\{Z_v(t) := t^{-1} \int_0^t X_v(u)\mathrm{d}u\}$.

to a random asymptotic belief vector X, and in fact it is an ergodic process.

Theorem 1: Let Assumption 1 hold. Then, for every value of the stubborn agents' beliefs $\{x_s : s \in \mathcal{S}\}$, there exists a stationary random belief vector X, whose probability law is invariant for the system and such that, for every initial distribution satisfying $\mathbb{P}(X_s(0) = x_s, \forall s \in \mathcal{S}) = 1$,

$$\lim_{t\to\infty} \mathbb{E}[\varphi(X(t))] = \mathbb{E}[\varphi(X)],$$

and, with probability one,

$$\lim_{t\to\infty} \frac{1}{t} \int_0^t \varphi(X(u))\mathrm{d}u = \mathbb{E}[\varphi(X)],$$

where $\varphi : \mathbb{R}^{\mathcal{V}} \to \mathbb{R}$ is any continuous test function such that $\mathbb{E}[\varphi(X)]$ exists and is finite.

Theorem 1 states that the beliefs of all the agents converge in distribution, and that their empirical averages converge almost surely, to a random asymptotic belief vector X. In contrast, the following theorem shows that the asymptotic belief of a regular agent which is connected to at least two stubborn agents with different beliefs is a non-degenerate random variable. As a consequence, the belief of every such regular agent keeps on oscillating with probability one. Moreover, the theorem shows that, with probability one, the difference between any pair of distinct regular agents which are influenced by more than one stubborn agent does not converge to zero, so that disagreement between them persists in time. For $a \in \mathcal{A}$, let $\mathcal{X}_a = \{x_s : s \in \mathcal{S}\}$ denote the set of stubborn agents' belief values influencing agent a.

Theorem 2: Let Assumption 1 hold, and let $a \in \mathcal{A}$ be such that $|\mathcal{X}_a| \geq 2$. Then, the asymptotic belief X_a is a non-degenerate random variable. Furthermore, if $a, a' \in \mathcal{A}$, with $a' \neq a$ are such that $|\mathcal{X}_a \cap \mathcal{X}_{a'}| \geq 2$, then $\mathbb{P}(X_a \neq X_{a'}) > 0$.

Theorem 1 and Theorem 2 are two of the central results of our paper. Even though beliefs converge in distribution, the presence of stubborn agents with different beliefs ensures that almost surely they fail to converge sample-path-wise. Moreover there will not be a consensus of beliefs in this society. Both of these are a consequence of the fact that each regular agent is continuously being influenced –directly or indirectly– by stubborn agents with different beliefs.

IV. Homogeneous influence in highly fluid social networks

In this section, we study the behavior of the expected stationary belief vector $\{\mathbb{E}[X_v : v \in \mathcal{V}\}$ as a function of the geometry of the underlying social network. Our estimates will prove to be particularly relevant for large-scale social networks satisfying the following condition.

As a first step, we derive a characterization of the expected stationary belief vector in terms of hitting probabilities of a random walk $V(t)$ on $\mathcal{G}$ with transition probability matrix P defined as in (2). We use the notation $T_{\mathcal{S}} := \inf\{t \geq 0 : V(t) \in \mathcal{S}\}$ for the hitting time of $V(t)$ on $\mathcal{S}$, and define the *hitting probability distributions* γ^v over $\mathcal{S}$, by

$$\gamma_s^v := \mathbb{P}(V(T_{\mathcal{S}}) = s | V(0) = v), \qquad s \in \mathcal{S}.$$

Then, one has the following result:

Theorem 3: Let Assumption 1 hold. Then, for every value of the stubborn agents' beliefs $\{x_s\}$, $\{\mathbb{E}[X_v]\}$ coincides with the unique vector $y \in \mathbb{R}^{\mathcal{V}}$ satisfying

$$y_a = \sum_v P_{av} y_v, \qquad y_s = x_s, \qquad a \in \mathcal{A},\ s \in \mathcal{S}. \quad (3)$$

Moreover,

$$\mathbb{E}[X_v] = \sum_s \gamma_s^v x_s, \quad v \in \mathcal{V}. \quad (4)$$

Equation (3) states that the expected stationary belief of each regular agent coincides with the arithmetic average of those all her neighbors'. In other words, the vector of the agents' expected stationary beliefs is harmonic on the graph with boundary conditions corresponding to the stubborn agents' beliefs. Equation (4) shows that the expected stationary belief of an agent can be written as a convex combination of the stubborn agents' beliefs, with weights given by the hitting probabilities on the set $\mathcal{S}$ of the random walk on the graph started at v. Such a characterization turns out to be fundamental in deriving the results presented below.

First, we introduce the notion of fluidity of a social network.

Definition 1: Given a social network, let P and π be as in (2). Let $d_* := \min_v d_v$ be the minimum degree, and let τ denote the (variational distance) mixing time [4], [22] of the standard random walk on $\mathcal{V}$. The *fluidity* of the social network is the ratio

$$\Phi := \frac{n d_*}{\tau \sum_s d_s}. \quad (5)$$

A sequence of social networks (or, more briefly, a social network) of increasing population size n is *highly fluid* if Φ diverges as n grows large.

Our estimates will show that for large-scale highly fluid social networks, the expected stationary beliefs of most of the regular agents in the population are very close to the value

$$\overline{x} := \sum_{s} \overline{\gamma}_s x_s \,, \qquad \overline{\gamma}_s := \sum_{v} \pi_v \gamma_s^v \,, \qquad s \in \mathcal{S} \,. \qquad (6)$$

Observe that $\overline{\gamma}_s$ coincides with probability that the random walk $V(t)$, started from the stationary distribution π, hits the stubborn agent s before any other stubborn agent $s' \in \mathcal{S}$. In fact, one may interpret $\overline{\gamma}_s$ as a relative measure of the influence of the stubborn agent s on the society compared to the rest of the stubborn agents $s' \in \mathcal{S}$.

Theorem 4: Let Assumption 1 hold, and assume that $4\sum_s d_s \leq \sum_v d_v$. Then, for all $\varepsilon > 0$,

$$\frac{1}{n}\left|\left\{v : \left|\mathbb{E}[X_v] - \overline{z}\right| \geq \Delta_* \varepsilon\right\}\right| \leq \frac{\psi(\varepsilon)}{\Phi} \,, \qquad (7)$$

where Φ is the fluidity, $\psi(\varepsilon) := 16\varepsilon^{-1}\log(2e^2/\varepsilon)$, and $\Delta_* := \max_{s,s'}\{x_s - x_{s'}\}$.

This theorem implies that in large-scale highly fluid social networks, as the population size n grows large, the expected stationary beliefs of regular agents concentrate around the value $\overline{z}$. We refer to this as an *homogeneous influence* of the stubborn agents on the rest of the society—meaning that their influence on most of the agents in the society is approximately the same. Indeed, it amounts to homogeneous (at least in their first moment) marginals of the agents' ergodic beliefs. This shows that in highly fluid social networks, most of the regular agents feel the presence of the stubborn agents in approximately the same way.

It is worth stressing how the condition of homogeneous influence may significantly differ from an approximate consensus. In fact, the former only involves the (the first and second moments of) the marginal distributions of the agents' stationary beliefs, and does not have any implication for their joint probability law. A chaotic distribution in which the agents' ergodic beliefs are all mutually independent would be compatible with the condition of approximately equal influence, as well as an approximate consensus condition, which would require the ergodic beliefs of most of the agents to be close to each other with high probability. We will address the investigation of this topic in another work.

Intuitively, if the set $\mathcal{S}$ and the mixing time τ are both small, then the influence of the stubborn agents will be felt by most of the regular agents much later then the time it takes them to influence each other. Hence, their expected beliefs will converge to values very close to each other. The proof of Theorem 4 relies on the characterization (4) of the stationary expected beliefs in terms of the hitting probabilities of the random walk. The definition of highly fluid network implies that the (expected) time it takes $V(t)$ to hit $\mathcal{S}$, when started from most of the nodes of $\mathcal{G}$, is much larger than the mixing time τ. Hence, before hitting $\mathcal{S}$, $V(t)$ looses memory of where it started from, and approaches $\mathcal{S}$ almost as if started from the stationary distribution π.

We now present some examples of family of social networks that are highly fluid in the limit of large population size n. Following a common terminology, we say that some property of such graphs holds with high probability, if the probability that it holds approaches one in the limit of large population size n.

Example 1: (**Connected Erdös-Renyi**) Consider the Erdös-Renyi random graph $\mathcal{G} = \mathcal{ER}(n,p)$, i.e., the random undirected graph with n vertices, in which each pair of distinct vertices is an edge with probability p, independently from the others. We focus on the regime $p = cn^{-1}\log n$, with $c > 1$, where the Erdös-Renyi graph is known to be connected with high probability [16, Thm. 2.8.2]. In this regime, results by Cooper and Frieze [12] ensure that, with high probability, $\tau = O(\log n)$, and that there exists a positive constant δ such that $\delta c \log n \leq d_v \leq 4c\log n$ for each node v [16, Lemma 6.5.2]. In particular, it follows that, with high probability, $4nd^* \geq \delta\sum_v d_v$. Therefore, the resulting social network is highly fluid, provided that $|\mathcal{S}| = o(n/\log n)$, as n grows large.

Example 2: (**Fixed degree distribution**) Consider a random graph $\mathcal{G} = \mathcal{FD}(n,\lambda)$, with n vertices, whose degree d_v are independent and identically distributed random variables with $\mathbb{P}(d_v = k) = \lambda_k$, for $k \in \mathbb{N}$. We assume that $\lambda_1 = \lambda_2 = 0$, that $\lambda_{2k} > 0$ for some $k \geq 2$, and that the first two moments $\sum_k \lambda_k k$, and $\sum_k \lambda_k k^2$ are finite. Then, the probability of the event $E_n := \{\sum_v d_v \text{ is even}\}$ converges to $1/2$ as n grows large, and we may assume that $\mathcal{G} = \mathcal{FD}(n,\lambda)$ is generated by randomly matching the vertices. Results in [16, Ch. 6.3] show that $\tau = O(\log n)$. Therefore, one finds that the resulting social network is highly fluid with high probability provided that $\sum_s d_s = o(n/\log n)$.

Example 3: (**Preferential attachment**) The preferential attachment model was introduced by Barabasi and Albert [8] to model real-world networks which typically exhibit a power law degree distribution. We follow [16, Ch. 4] and consider the random graph $\mathcal{G} = \mathcal{PA}(n,m)$ with n vertices, generated by starting with two vertices connected by m parallel edges, and then subsequently adding a new vertex and connecting it to m of the existing nodes with probability proportional to their current degree. As shown in [16, Th. 4.1.4], the degree distribution converges in probability to the power law $\mathbb{P}(d_v = k) = \lambda_k = 2m(m+1)/k(k+1)(k+2)$, and the graph is connected with high probability [16, Th. 4.6.1]. In particular, it follows that, with high probability, the average degree remains bounded, while the second moment of the degree distribution diverges an n grows large. On the other hand, results by Mihail et al. [24] (see also [16, Th. 6.4.2]) imply that the mixing time $\tau = O(\log n)$. Therefore, the resulting social network is highly fluid with high probability if $\sum_s d_s = o(n/\log n)$.

Example 4: (**Watts & Strogatz's small world**) Watts and Strogatz [35], and then Newman and Watts [29] proposed simple models of random graphs to explain the empirical evidence that most social networks contain a large number of triangles and have a small diameter (the latter has become known as the small-world phenomenon). We consider Newman and Watts' model, which is a random graph $\mathcal{G} = \mathcal{NW}(n,k,p)$, with n vertices, obtained starting from a Cayley graph on the ring $\mathbb{Z}_n$ with generator

$\{-k, -k+1, \ldots, -1, 1, \ldots, k-1, k\}$, and adding to it a Poisson number of shortcuts with mean pkn, and attaching them to randomly chosen vertices. In this case, the average degree remains bounded with high probability as n grows large, while results by Durrett [16, Th. 6.6.1] show that the mixing time $\tau = O(\log^3 n)$. Therefor, the network is highly fluid provided that $\sum_s d_s = o(n/\log^3 n)$.

V. Conclusion

In this paper, we have studied a possible mechanism explaining persistent disagreement and opinion fluctuations in social networks. We have considered a stochastic gossip model of continuous opinion dynamics, combined with the assumption that there are some stubborn agents in the network who never change their opinions. We have shown that the presence of these stubborn agents leads to persistent oscillations and disagreements among the rest of the society: the beliefs of regular agents almost surely do not converge sample-path-wise, and keep on oscillating in an ergodic fashion. First and second moments of the stationary beliefs distribution can be characterized in terms of the hitting probabilities of a random walk on the graph describing the social network. We have shown that in highly fluid social networks, whose associated random walks have mixing times which are sufficiently smaller than the inverse of the stubborn agents' set size, the vectors of the stationary expected beliefs and variances are almost constant, so that the stubborn agents have homogeneous influence on the society.

References

[1] D. Acemoglu, G. Como, F. Fagnani, and A. Ozdaglar, *Opinion fluctuations and disagreement in social networks*, (2010), submitted, available at http://mit.edu/giacomo/www/disagreementsubmittedMOR.pdf.

[2] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, *Bayesian learning in social networks*, Review of Economic Studies (2011), to appear.

[3] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, *Spread of (mis)information in social networks*, Games and Economic Behavior **70** (2010), no. 2, 194–227.

[4] D. Aldous and J. Fill, *Reversible Markov chains and random walks on graphs*, monograph in preparation, 2002.

[5] R. Axelrod, *The dissemination of culture*, Journal of Conflict Resolution **42** (1997), no. 2, 203–226.

[6] V. Bala and S. Goyal, *Learning from neighbours*, Review of Economic Studies **65** (1998), no. 3, 595–621.

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

[8] A.L. Barabasi and R. Albert, *Emergence of scaling in random networks*, Science **268** (1999), 509–512.

[9] V.D. Blondel, J. Hendickx, and J. N. Tsitsiklis, *On krause's consensus formation model with state-dependent connectivity*, IEEE Transactions on Automatic Control **54** (2009), no. 11, 2586–2597.

[10] C. Castellano, S. Fortunato, and V. Loreto, *Statistical physics of social dynamics*, Reviews of Modern Physics **81** (2009), 591–644.

[11] G. Como and F. Fagnani, *Scaling limits for continuous opinion dynamics systems*, The Annals of Applied Probability **21** (2011), no. 4, 1537–1567.

[12] C. Cooper and A. Frieze, *The cover time of sparse random graphs*, Random Structures and Algorithms **30** (2006), no. 1-2, 1–16.

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

[14] P.M. DeMarzo, D. Vayanos, and J. Zwiebel, *Persuasion bias, social influence, and unidimensional opinions*, The Quarterly Journal of Economics **118** (2003), no. 3, 909–968.

[15] P. Diaconis and D. Friedman, *Iterated random functions*, SIAM Review **41** (1999), no. 1, 45–76.

[16] R. Durrett, *Random graph dynamics*, Cambridge University Press, 2006.

[17] F. Fagnani and S. Zampieri, *Randomized consensus algorithms over large scale networks*, IEEE Journal on Selected Areas of Communications **26** (2008), no. 4, 634–649.

[18] D. Gale and S. Kariv, *Bayesian learning in social networks*, Games and Economic Behavior **45** (2003), no. 2, 329–346.

[19] B. Golub and M.O. Jackson, *Naive learning in social networks and the wisdom of crowds*, American Economic Journal: Microeconomics **2** (2010), no. 1, 112–149.

[20] A. Jadbabaie, J. Lin, and S. Morse, *Coordination of groups of mobile autonomous agents using nearest neighbor rules*, IEEE Transactions on Automatic Control **48** (2003), no. 6, 988–1001.

[21] U. Krause, *A discrete non-linear and non-autonomous model of consensus formation*, pp. 227–236, Gordon and Breach, Amsterdam, 2000.

[22] D.A. Levin, Y. Peres, and E.L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, 2010.

[23] J. Lorenz, *A stabilization theorem for continuous opinion dynamics*, Physica A **355** (2005), no. 1, 217–223.

[24] M. Mihail, C. Papadimitriou, and A. Saberi, *On certain connectivity properties of the internet topology*, Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science 2003, 2004.

[25] M. Mobilia, *Does a single zealot affect an infinite group of voters?*, Physical Review Letters **91** (2003), no. 2, 028701.

[26] M. Mobilia and I.T. Georgiev, *Voting and catalytic processes with inhomogeneities*, Physical Review E **71** (2005), 046102.

[27] M. Mobilia, A. Petersen, and S. Redner, *On the role of zealotry in the voter model*, J. Statistical Mechanics: Theory and Experiments **128** (2007), 447–483.

[28] A. Nedić and A. Ozdaglar, *Distributed subgradient methods for multi-agent optimization*, IEEE Transactions on Automatic Control **54** (2009), no. 1, 48–61.

[29] M.E.J. Newman and D.J. Watts, *Renormalization group analysis of the small-world network model*, Physics Letters A **263** (1999), 341–346.

[30] R. Olfati-Saber and R.M. Murray, *Consensus problems in networks of agents with switching topology and time-delays*, IEEE Transactions on Automatic Control **49** (2004), no. 9, 1520–1533.

[31] A. Olshevsky and J.N. Tsitsiklis, *Convergence speed in distributed consensus and averaging*, SIAM Journal on Control and Optimization **48** (2009), no. 1, 33–55.

[32] L. Smith and P. Sorensen, *Pathological outcomes of observational learning*, Econometrica **68** (2000), no. 2, 371–398.

[33] J.N. Tsitsiklis, *Problems in decentralized decision making and computation*, Ph.D. thesis, Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1984.

[34] J.N. Tsitsiklis, D.P. Bertsekas, and M. Athans, *Distributed asynchronous deterministic and stochastic gradient optimization algorithms*, IEEE Transactions on Automatic Control **31** (1986), no. 9, 803–812.

[35] D.J. Watts and S.H. Strogatz, *Collective dynamics of 'small-world' networks*, Nature **393** (1998), 440–442.