



LUND UNIVERSITY

Opinion fluctuations and persistent disagreement in social networks

Acemoglu, Daron; Como, Giacomo; Fagnani, Fabio; Ozdaglar, Asuman

2011

[Link to publication](#)

Citation for published version (APA):

Acemoglu, D., Como, G., Fagnani, F., & Ozdaglar, A. (2011). *Opinion fluctuations and persistent disagreement in social networks*. 2347-2352. Paper presented at 50th IEEE Conference on Decision and Control and European Control Conference, 2011, Orlando, Florida, United States.

Total number of authors:

4

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Opinion fluctuations and persistent disagreement in social networks

Daron Acemoglu, Giacomo Como, Fabio Fagnani, Asuman Ozdaglar

Abstract—Disagreement among individuals in a society, even on central questions that have been debated for centuries, is the rule; agreement is the rare exception. How can disagreement of this sort persist for so long? Most 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. These models are thus unable to explain persistent disagreements, and belief fluctuations.

We propose a tractable model that generates long-run disagreements and persistent opinion fluctuations. Our model involves a stochastic gossip model 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 stubborn agents within it shape the long run behavior of the opinion dynamics. We prove that, when the society is *highly fluid*, meaning that the mixing time of the random walk on the graph describing the social network is small relative to the inverse of the relative size of the linkages to stubborn agents, a condition of *homogeneous influence* emerges, whereby the ergodic beliefs of most of the regular agents have approximately equal marginal distributions. This clearly need not imply approximate consensus and in fact we show, under mild conditions, the ergodic belief distribution becomes *approximately chaotic*, meaning that the variance of the aggregate belief of the society vanishes in the large population limit while individual opinions still fluctuate significantly 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.

D. Acemoglu is with the Department of Economics at the Massachusetts Institute of Technology. {giacomo,asuman}@mit.edu.

G. Como and A. Ozdaglar are with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. {giacomo,asuman}@mit.edu.

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

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 [20], [5], [2], [4], [11], [8], [12], [3]. They are thus unable to explain persistent disagreements. In this paper, we investigate a possible source of persistent disagreement in social networks. We propose a tractable model that generates both long-run disagreement and opinion fluctuations so that a consensus fails to emerge even as individuals communicate and sometimes change their opinions.

We consider a stochastic gossip model of communication combined with the assumption that there are some *stubborn* agents in the network who never change their opinions. We show that the presence of these stubborn agents 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 or political leaders wishing to influence the rest of the society). 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 over the graph induced by the social network. Much of our analysis characterizes the long-run behavior of this Markov process.

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. Instead, the belief of each regular agent converges in law to a non-degenerate random variable and thus has a limiting ergodic distribution (and similarly, the vector of beliefs of all regular agents jointly converge to a non-degenerate random vector). This model therefore provides a new approach to understanding persistent disagreements.

We then study the long-run dynamics of opinions in *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 ergodic opinion of most of the agents concentrate around certain values in the large population limit. 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.

Finally, we show that, if the presence of stubborn agents in the society is *significant*, then the variance of the ergodic aggregate belief of the society vanishes in the large population limit, and the ergodic opinion distribution is *approximately chaotic*. If, moreover, the influence of any stubborn agent does not dominate the influences of the rest, then the mean squared disagreement, i.e., the average of the expected squared differences between the agents’ ergodic beliefs, remains bounded away from zero in the large population limit.

Our analysis uses several new approaches to the study of belief dynamics. First, convergence in law of the regular agents’ beliefs is established by first rewriting the dynamics in the form of an iterated affine function system, and studying the corresponding time-reversed process; the latter is converging almost surely and, at each time instant, has the same marginal distribution as the actual beliefs process. Second, we use a characterization of the expected values and correlations of the ergodic beliefs in terms of the hitting probability distribution of a pair of coupled random walks moving on the directed graph describing the communication structure in the social network. Third, we use the characterization of these hitting distributions as solutions of a Laplace equation with boundary conditions on the stubborn agents set in order to find explicit solutions for the expected ergodic beliefs in some social networks with additional structure. Fourth, we derive bounds on the behavior of the expected values and variances of the ergodic beliefs in large population size limit, by showing that, on highly fluid networks, these expectations and variances are almost equal for most of the agents. This is a consequence of the fact that the hitting probabilities on the stubborn agents set of the associated random walk have a weak dependence on the initial state, which is in turn proved by combining properties of fast-mixing chains, including the approximate exponentiality of the hitting times.

In addition to the aforementioned works on learning and opinion dynamics, our model is closely related to the work by Mobilia [15], which propose a variation of the discrete opinion dynamics model, also called the voter model, with “zealots” (equivalent to our stubborn agents). This work generally relies on heuristic mean-field approximations, valid for certain graphical structures, and numerical simulations, to characterize belief dynamics. 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, we shall also show that Mobilia’s model 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 [21], [22], [13], [18], [19], [10], [16]. 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 ergodic 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 ergodic process, while in general they do not converge almost surely. Section IV presents a characterization of the first and second moments of the ergodic beliefs in terms of the hitting probabilities of two coupled random walks on the network. Section V provides bounds on the level of dispersion of the first two moments of the ergodic beliefs: it is shown that, in highly fluid networks, most of the agents have almost the same ergodic belief and variance. Section VI studies the mean square oscillations and disagreement in highly fluid networks: if there is a significant presence of stubborn agents, the variance of the ergodic aggregate belief of the society vanishes in the large population limit, and the joint distribution of the ergodic beliefs is close to a chaotic law. Section VII contains some concluding remarks. All the statements will be presented without proof, which are presented in the longer version of this work [1].

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 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 a probability distribution μ over a finite set \mathcal{I} , and a subset $\mathcal{J} \subseteq \mathcal{I}$ we will write $\mu(\mathcal{J}) := \sum_j \mu_j$.

Let $V(t)$ and $V'(t)$ be continuous-time random walks on a finite set \mathcal{V} , defined on the same probability space, both with marginal transition probability matrix P . We use the notation $\mathbb{P}_v(\cdot)$, and $\mathbb{P}_{v,v'}(\cdot)$, for the conditional probability measures given the events $V(0) = v$, and, respectively, $(V(0), V'(0)) = (v, v')$. Similarly, for some probability distribution π over \mathcal{V} (possibly the stationary one), $\mathbb{P}_\pi(\cdot) := \sum_{v,v'} \pi_v \pi_{v'} \mathbb{P}_{v,v'}(\cdot)$ will denote the conditional probability measure of the Markov chain with initial distribution π , while $\mathbb{E}_v[\cdot]$, $\mathbb{E}_{v,v'}[\cdot]$, and $\mathbb{E}_\pi[\cdot]$ will denote the corresponding conditional expectations. For two non-negative sequences




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.

$\{a_n\}$, $\{b_n\}$, we will write $a_n = O(b_n)$ if for some positive constant K , $a_n \leq Kb_n$ 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:

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

In contrast, the beliefs of the regular agents are updated as follows. To every 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 :

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

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. For every regular agent $a \in \mathcal{A}$, let $\mathcal{S}_a \subseteq \mathcal{S}$




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 $\{Z_v(t) := t^{-1} \int_0^t X_v(u) du\}$.

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 . For every stubborn agent $s \in \mathcal{S}$, $\mathcal{A}_s := \{a : s \in \mathcal{S}_a\} \subseteq \mathcal{A}$ will stand for the set of regular agents *influenced* by s .

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} .

Assumption 2: Every stubborn agent influences some regular agent, i.e., \mathcal{A}_s is non-empty for every s in \mathcal{S} .

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

$$P_{vv'} = \begin{cases} 1/d_v & \text{if } \{v, v'\} \in \mathcal{E} \\ 0 & \text{if } \{v, v'\} \notin \mathcal{E} \end{cases} \quad \pi_v := d_v / \sum_{v'} d_{v'}.$$

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

Theorem 1: Let Assumptions 1 and 2 hold. Then, for every value of the stubborn agents' beliefs $\{x_s : s \in \mathcal{S}\}$, there exists an $\mathbb{R}^{\mathcal{V}}$ -valued random variable X , such that, for every initial belief distribution satisfying $\mathbb{P}(X_s(0) = x_s) = 1$ for every $s \in \mathcal{S}$, and

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

for all bounded and continuous test functions $\varphi : \mathbb{R}^{\mathcal{V}} \rightarrow \mathbb{R}$. Moreover, the probability law of the asymptotic belief vector X is invariant for the system.

Using standard ergodic theorems for Markov chains, an immediate implication of Theorem 1 is the following corollary, which shows that time averages of any continuous bounded function of agent beliefs are given by their expectation over the limiting distribution. Choosing the relevant function properly, this enables us to express the empirical averages of and correlations across agent beliefs in terms of expectations over the limiting distribution, highlighting the ergodicity of agent beliefs.

Corollary 1: For all initial distributions $X(0) \in \mathbb{R}^{\mathcal{V}}$, with probability one,

$$\lim_{t \rightarrow +\infty} \frac{1}{t} \int_0^t \varphi(X(u)) du = \mathbb{E}[\varphi(X)],$$

for all continuous and bounded test functions $\varphi : \mathbb{R}^{\mathcal{V}} \rightarrow \mathbb{R}$.

Motivated by Corollary 1, for any agent $v \in \mathcal{V}$, we refer to the random variable X_v as the *ergodic belief of agent v*.

Theorem 1, and Corollary 1, respectively, show 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 Assumptions 1 and 2 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 $\mathbb{P}(X_a \neq X_{a'}) > 0$.




Fig. 3. In (a), a network topology consisting of a line with three regular agents and two stubborn agents placed in the extremes. In (b), the corresponding graph product \mathcal{G}_{\square} . The latter has 25 nodes, four of which are absorbing states. The coupled random walk $(V(t), V'(t))$ moves on \mathcal{G}_{\square} , its two components jumping independently to neighbor states, unless they are either on the diagonal, or one of them is in \mathcal{S} : in the former case, there is some chance that the two components jump as a unique one, thus inducing a direct connection along the diagonal; in the latter case, the only component that can keep moving is the one which has not hit \mathcal{S} , while the one who hit \mathcal{S} is bound to remain constant from that point on.

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. 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. EMPIRICAL AVERAGES AND CORRELATIONS OF AGENT BELIEFS

In this section, we provide a characterization of the empirical averages and correlations of agent beliefs $\{X_v(t) : v \in \mathcal{V}\}$, i.e., of the almost surely constant limits

$$\lim_{t \rightarrow +\infty} \frac{1}{t} \int_0^t X_v(u) du, \quad \lim_{t \rightarrow +\infty} \frac{1}{t} \int_0^t X_v(u) X_{v'}(u) du.$$

By Corollary 1, these limits are given by the first two moments of the ergodic beliefs, i.e., $\mathbb{E}[X_v]$ and $\mathbb{E}[X_v X_{v'}]$, respectively, independently of the distribution of initial regular agents' beliefs.

We next provide explicit characterizations of these limits in terms of hitting probabilities of a pair of coupled random walks on $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Specifically, we consider a coupling $(V(t), V'(t))$ of continuous-time random walks on \mathcal{V} , such that both $V(t)$, and $V'(t)$, have marginal state transition rates P_{av} , as defined in (II). In fact, one may interpret $(V(t), V'(t))$ as a random walk on the Cartesian power graph \mathcal{G}_{\square} , whose node set is the product $\mathcal{V} \times \mathcal{V}$, and where there is an edge from (v, v') to (w, w') , if and only if either $\{v, w\} \in \mathcal{E}$ and $v' = w'$, or $v = w$ and $\{v', w'\} \in \mathcal{E}$, or $v = v'$ and $w = w'$ (See Figure 3). The transition rates $K_{(v,v')(w,w')}$ of the coupled random walks $(V(t), V'(t))$ are

given by

$$\begin{cases} P_{vw} & \text{if } v \neq v', v \neq w, v' = w' \\ P_{v'w'} & \text{if } v \neq v', v = w, v' \neq w' \\ 0 & \text{if } v \neq v', w \neq v, w' \neq v' \\ \theta P_{vw} & \text{if } v = v', w = w' \\ (1 - \theta)P_{vw} & \text{if } v = v', v' = w' \\ (1 - \theta)P_{v'w'} & \text{if } v = v', v = w \\ 0 & \text{if } v = v', w \neq w', w \neq v, w' \neq v'. \end{cases} \quad (2)$$

The first three lines of (2) state that, conditioned on $(V(t), V'(t))$ being on a pair of non-coincident nodes (v, v') , each of the two components, $V(t)$ (respectively, $V'(t)$), jumps to a neighbor node w , with transition rate P_{vw} (respectively, to a neighbor node w' with transition rate $P_{v'w'}$), whereas the probability that both components jump at the same time is zero. On the other hand, the last four lines of (2) state that, once the two components have met, i.e., conditioned on $V(t) = V'(t) = v$, they have some chance to stick together and jump as a single particle to a neighbor node w , with rate θP_{vw} , while each of the components $V(t)$ (respectively, $V'(t)$) has still some chance to jump alone to a neighbor node w with rate $(1 - \theta)P_{vw}$ (resp., to w' with rate $(1 - \theta)P_{v'w'}$). In the extreme case when $\theta = 1$ for all v, w , the last three lines of the righthand side of (2) equal 0, and in fact one recovers the expression for the transition rates of two coalescing random walks: once $V(t)$ and $V'(t)$ have met, they stick together and move as a single particle, never separating from each other.

We use the notation $T_{\mathcal{S}} := \inf\{t \geq 0 : V(t) \in \mathcal{S}\}$ and $T'_{\mathcal{S}} := \inf\{t \geq 0 : V'(t) \in \mathcal{S}\}$. Further, for all $v, v' \in \mathcal{V}$, we define the *hitting probability distributions* γ^v over \mathcal{S} , and $\eta^{vv'}$ over \mathcal{S}^2 , whose entries are respectively given by

$$\begin{aligned} \gamma_s^v &:= \mathbb{P}_v(V(T_{\mathcal{S}}) = s), & s \in \mathcal{S}, \\ \eta_{ss'}^{vv'} &:= \mathbb{P}_{vv'}(V(T_{\mathcal{S}}) = s, V'(T'_{\mathcal{S}}) = s'), & s, s' \in \mathcal{S}. \end{aligned}$$

The following lemma characterizes $\{\gamma_s^v : v \in \mathcal{V}\}$ and $\{\eta_{ss'}^{vv'} : v, v' \in \mathcal{V}\}$ as solutions of harmonic equations on \mathcal{A} and \mathcal{A}^2 , with boundary conditions on \mathcal{S} and $\mathcal{V}^2 \setminus \mathcal{A}^2$.

Lemma 1: For all $s, s' \in \mathcal{S}$, one has that

$$\begin{aligned} \sum_v P_{av}(\gamma_s^v - \gamma_s^a) &= 0, & \forall a \in \mathcal{A}, \\ \gamma_s^s &= 1, & \gamma_{s'}^s = 0, & \forall s' \in \mathcal{S} \setminus \{s\}, \end{aligned}$$

$$\begin{aligned} \sum_{v, v'} K_{(a, a')(v, v')} \left(\eta_{ss'}^{vv'} - \eta_{ss'}^{aa'} \right) &= 0, & \forall a, a' \in \mathcal{A}, \\ \eta_{ss'}^{vv'} &= \gamma_s^v \gamma_{s'}^{v'}, & \forall (v, v') \notin \mathcal{A}^2. \end{aligned}$$

The next theorem provides a fundamental characterization of the expected values and correlations of ergodic beliefs in terms of the hitting probabilities of the coupled random walks $V(t)$ and $V'(t)$.

Theorem 3: For all $v, v' \in \mathcal{V}$,

$$\mathbb{E}[X_v] = \sum_s \gamma_s^v x_s, \quad \mathbb{E}[X_v X_{v'}] = \sum_{s, s'} \eta_{ss'}^{vv'} x_s x_{s'}.$$

Remark 1: As a consequence of Theorem 3, one gets that, if $\mathcal{X}_a = \{x_*\}$, then $X_a = x_*$, and, by Corollary 1, $X_a(t)$ converges to x_* with probability one. This can be thought of as a sort of complement to Theorem 2.

V. HOMOGENEOUS INFLUENCE IN HIGHLY FLUID SOCIAL NETWORKS

In this section, we present estimates for the ergodic belief expectations and variances as a function of the underlying social network. Our estimates will prove to be particularly relevant for large-scale social networks satisfying the following condition.

Definition 1: Given a reversible social network, let P denote its transition probability matrix, and π denote its stationary distribution. Define $\pi_* := \min_v \pi_v$, and let

$$\tau := \min_{v, v'} \inf \left\{ t \geq 0 : \sum_w |\mathbb{P}_v(V(t) = w) - \mathbb{P}_{v'}(V(t) = w)| \leq \frac{2}{e} \right\}$$

denote the (variational distance) mixing time of the continuous-time random walk $V(t)$ with transition rate matrix P . We say that a sequence of social networks of increasing population size n is *highly fluid* if it satisfies

$$\tau \pi(\mathcal{S}) = o(1), \quad \liminf n \pi_* > 0, \quad \text{as } n \rightarrow +\infty, \quad (3)$$

where

$$\pi(\mathcal{S}) := \sum_s \pi_s = \sum_s d_s / (n \bar{d}). \quad (4)$$

Our estimates will show that for large-scale highly fluid social networks, the ergodic beliefs of most of the regular agents in the population can be approximated (at least in their first and second moments) by a ‘virtual’ random belief Z , whose distribution is given by

$$\mathbb{P}(Z = x_s) = \bar{\gamma}_s, \quad \bar{\gamma}_s := \sum_v \pi_v \gamma_s^v, \quad s \in \mathcal{S}. \quad (5)$$

We refer to the probability distribution $\{\bar{\gamma}_s : s \in \mathcal{S}\}$ as the *stationary stubborn agent distribution*. Observe that $\bar{\gamma}_s = \mathbb{P}_{\pi}(V(T_{\mathcal{S}}) = s)$ 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, as we shall clarify below, one may interpret $\bar{\gamma}_s$ as a relative measure of the influence of the stubborn agent s on the society compared to the rest of the stubborn agents $s' \in \mathcal{S}$.

More precisely, let us denote the expected value and variance of the virtual belief Z by

$$\mathbb{E}[Z] := \sum_s \bar{\gamma}_s x_s, \quad \sigma_Z^2 := \sum_s \bar{\gamma}_s (x_s - \mathbb{E}[Z])^2. \quad (6)$$

Let σ_v^2 denote the variance of the ergodic belief of agent v ,

$$\sigma_v^2 := \mathbb{E}[X_v^2] - \mathbb{E}[X_v]^2.$$

We also use the notation Δ_* to denote the maximum difference between stubborn agents’ beliefs, i.e.,

$$\Delta_* := \max \{x_s - x_{s'} : s, s' \in \mathcal{S}\}. \quad (7)$$

The next theorem presents the main result of this section.

Theorem 4: Let Assumptions 1 and 2 hold, and assume that $\pi(\mathcal{S}) \leq 1/4$. Then, for all $\varepsilon > 0$,

$$\frac{1}{n} \left| \left\{ v : \left| \mathbb{E}[X_v] - \mathbb{E}[Z] \right| \geq \Delta_* \varepsilon \right\} \right| \leq \psi(\varepsilon) \frac{\tau \pi(\mathcal{S})}{n \pi_*}, \quad (8)$$

with $\psi(\varepsilon) := \frac{16}{\varepsilon} \log(2e^2/\varepsilon)$. Furthermore, if the trust parameters satisfy $\theta_{av} = 1$ for all $(a, v) \in \mathcal{E}$, then

$$\frac{1}{n} \left| \left\{ v : \left| \sigma_v^2 - \sigma_Z^2 \right| \geq \Delta_*^2 \varepsilon \right\} \right| \leq \psi(\varepsilon) \frac{\tau \pi(\mathcal{S})}{n \pi_*}. \quad (9)$$

This theorem implies that in large-scale highly fluid social networks, as the population size n grows large, the expected values and variances of ergodic beliefs of regular agents concentrate around fixed values corresponding to the expected virtual belief $\mathbb{E}[Z]$, and, respectively, its variance σ_Z^2 . 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 two moments) 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.

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 than the time it takes them to influence each other. Hence, their beliefs’ empirical averages and variances will converge to values very close to each other. The proof of Theorem 4 relies on the characterization of the mean ergodic beliefs in terms of the hitting probabilities of the random walk $V(t)$. 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} , $V(t)$ loses memory of where it started from, and approaches \mathcal{S} almost as if started from the stationary distribution π .

Before proving Theorem 4, we present some examples of highly fluid social networks in Sect. V-A.

A. Examples of large-scale highly fluid social networks

We now present some examples of family of social networks that are highly fluid in the limit of large population size n . Before proceeding, let us observe that $\pi_* n \leq 1$, with equality if and only if π is the uniform measure over \mathcal{V} . Hence, one has $\pi_* n = 1$ for regular graphs, while, for general undirected graphs $(\pi_* n)^{-1} \leq \bar{d}$, where \bar{d} is the average degree of the graph.

We shall focus on four examples of random graph sequences which have been the object of extensive research. 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 Erdős-Renyi) Consider the Erdős-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 Erdős-Renyi graph is known to be connected with high probability [9, Thm. 2.8.2]. In this regime, results by Cooper and Frieze [7] 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 [9, Lemma 6.5.2]. In particular, it follows that, with high probability, $(\pi_* n)^{-1} \leq 4/\delta$. Hence, using (4), one finds that 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 $\bar{d} := \sum_k \lambda_k k$, 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 [9, 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 $\sum_s d_s = o\left(\frac{n}{\log n}\right)$.

Example 3: (Preferential attachment) The preferential attachment model was introduced by Barabasi and Albert [6] to model real-world networks which typically exhibit a power law degree distribution. We follow [9, 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 [9, 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 [9, Th. 4.6.1]. In particular, it follows that, with high probability, the average degree \bar{d} remains bounded, while the second moment of the degree distribution diverges as n grows large. On the other hand, results by Mihail et al. [14] (see also [9, Th. 6.4.2]) imply that the mixing time $\tau = O(\log n)$. Therefore, thanks to (4), the resulting social network is highly fluid with high probability if $\sum_{s \in \mathcal{S}} d_s = o\left(\frac{n}{\log n}\right)$.

Example 4: (Watts & Strogatz’s small world) Watts and Strogatz [23], and then Newman and Watts [17] 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 $\{-k, -k+1, \dots, -1, 1, \dots, k-1, k\}$, 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 [9, Th. 6.6.1] show that the mixing time $\tau = O(\log^3 n)$. This, and (4) imply that (3) holds provided that $\sum_{s \in \mathcal{S}} d_s = o\left(\frac{n}{\log^3 n}\right)$.

VI. OPINION OSCILLATIONS AND DISAGREEMENT

We have seen in the previous section that in highly fluid social networks a condition of homogeneous influence is achieved, with the expected values and variances of the

ergodic opinions of almost all the agents close to those of the virtual belief. 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' ergodic 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. In this section, under additional assumptions, we show that the ergodic belief distribution in highly fluid social networks is closer to a chaotic distribution than to an approximate consensus. For the sake of simplicity, throughout this section, we restrict our attention to the voter model.

Assumption 3: For every $e \in \mathcal{E}$, $\theta_e = 1$.

We start by introducing two quantities measuring the amplitude of the aggregate population's oscillations and the average disagreement among the agents. Specifically, let us consider the *ergodic aggregate belief* of the system, $\bar{X} := n^{-1} \sum_v X_v$, and let

$$\sigma_{\bar{X}}^2 := \mathbb{E} \left[(\bar{X} - \mathbb{E}[\bar{X}])^2 \right] \quad (10)$$

be its variance. Also, define the *mean squared disagreement* as

$$\Delta^2 := \frac{1}{2n^2} \sum_{v,v'} \mathbb{E} \left[(X_v - X_{v'})^2 \right], \quad (11)$$

the reason for the factor $1/2$ being mere notational convenience. Observe that, if the ergodic distribution of the agents' beliefs is chaotic (i.e., it is the product of its marginals), then \bar{X} is the arithmetic average of independent random variables with finite variance, and thus $\sigma_{\bar{X}}^2 = o(1)$. On the other hand, an approximate consensus condition, with the ergodic beliefs of most of the agents close to each other with high probability, would imply that $\Delta^2 = o(1)$.

In this section, we focus on highly fluid social networks satisfying the following:

Definition 2: Given a family of reversible social networks of increasing population size, we say that there is a *significant presence of stubborn agents* if

$$\frac{\pi(\mathcal{D})\tau_2}{\pi(\mathcal{S})} = o(1), \quad n \rightarrow +\infty, \quad (12)$$

where τ_2 is the relaxation time, i.e., the inverse of the spectral gap, and

$$\pi(\mathcal{D}) := \sum_a \pi_a^2$$

is the invariant measure of the diagonal set $\mathcal{D} := \{(a, a) : a \in \mathcal{A}\}$.

In order to obtain some intuition on Definition 2, one should think of the ratio $\pi(\mathcal{S})/\pi(\mathcal{D})$ as a measure of the relative intensity of the interactions of the regular agents with the stubborn agents (quantified by $\pi(\mathcal{S})$), as compared

to the intensity of the interactions between typical pairs of regular agents (quantified by $\pi(\mathcal{D})$). If such a ratio grows fast enough (precisely, Definition 2 requires it to grow faster than the relaxation time of the network, but in fact, one may expect that in many cases such ratio going to infinity should suffice), then one may expect that the ergodic beliefs of a typical pair of regular agents in the network should be directly influenced by the stubborn agents' beliefs, without a significant coupling between themselves. Hence, in a social network with a significant presence of stubborn agents, the ergodic beliefs of most of the regular agents' pairs are expected to be weakly coupled, so that the variance of the ergodic aggregate belief should vanish in the large population limit. Indeed, this is formalized in the following:

Theorem 5: For any family of highly fluid social networks, satisfying Assumptions 1, 2, and 3, with a significant presence of stubborn agents, it holds

$$\sigma_{\bar{X}}^2 = o(1), \quad \Delta^2 = \sigma_Z^2 + o(1), \quad \text{as } n \rightarrow +\infty.$$

Theorem 5 shows that in highly fluid social network with a significant presence of stubborn agents, the amplitude of the ergodic oscillations of the aggregate belief vanishes, while the mean square disagreement is asymptotically equivalent to the variance of the virtual belief, in the limit of large population size. Hence, under these conditions the ergodic belief distribution achieved in this setting is close to a chaotic distribution.

An immediate consequence of Theorem 5 is that, if σ_Z^2 is bounded away from zero in the large population limit, then so is the mean squared disagreement. Observe that the condition $\sigma_Z^2 = o(1)$ is equivalent to the fact that the probability measure $\sum_s \bar{\gamma}_s \delta_{x_s}$ (where $\bar{\gamma}_s$ is defined in (5), and δ_x stands for the Dirac's measure centered in some $x \in \mathbb{R}$) concentrates in one single point. We can think of this as the presence of a *dominating stubborn agents' belief*. Hence, we may say that Theorem 5 implies that, on highly fluid social networks with a significant presence of stubborn agents, none of whose beliefs is dominating, a significant disagreement persists in the large population limit.

A. Examples of highly fluid social networks with significant presence of stubborn agents

Observe that one has

$$\pi(\mathcal{D}) = \sum_a \pi_a^2 \leq \sum_v \pi_v^2 = \sum_v \frac{d_v^2}{(n\bar{d})^2} = \frac{\bar{d}^2}{(\bar{d})^2} n^{-1},$$

where \bar{d} is the average degree, and $\bar{d}^2 := n^{-1} \sum_v d_v^2$ is the average squared degree, of \mathcal{G} . Notice that the ratio

$$\frac{\bar{d}^2}{(\bar{d})^2} = 1 + \frac{1}{n} \sum_v \left(\frac{d_v}{\bar{d}} - 1 \right)^2$$

is minimal for regular graphs, where it equals 1, and grows with the normalized variance of the degree distribution. In particular, for a family of social networks with bounded first and second moment of the degree distribution, $\pi(\mathcal{D}) = O(n^{-1})$, so that, in order to have a significant presence of

stubborn agents, its is sufficient that $n\pi(\mathcal{S}) = (\bar{d})^{-1} \sum_s d_s$ grows faster than the relaxation time τ_2 .

Let us return to the examples of Sect. V-A.

Example 5: (Connected Erdős-Renyi) Consider the Erdős-Renyi random graph $\mathcal{G} = \mathcal{ER}(n, p)$, in the regime $p = cn^{-1} \log n$, with $c > 1$, as in Example 1. Then, with high probability, $\bar{d}^2/(\bar{d})^2 = O(1)$, while $\tau_2 \leq \tau = O(\log n)$. It follows that the associated social network is highly fluid and with a significant presence of stubborn agents provided that $|\mathcal{S}|$ grows faster than $\log n$, and slower than $n/\log n$.

Example 6: (Fixed degree distribution) Consider $\mathcal{G} = \mathcal{FD}(n, \lambda)$, as in Example 2. Then, with high probability, since the expected square degree is bounded, one has $\pi(\mathcal{D}) = O(n^{-1})$, while $\tau_2 \leq \tau = O(\log n)$. It follows that the associated social network is highly fluid and with a significant presence of stubborn agents provided that $\sum_s d_s$ grows faster than $\log n$, and slower than $n/\log n$.

Example 7: (Preferential attachment) Consider the preferential attachment model of Example 3. Then, with high probability, $\tau_2 \leq \tau = O(\log n)$, while, according to [9, pag. 180], $\pi(\mathcal{D}) \leq n^{-1} \log n$. It follows that the associated social network is highly fluid and with a significant presence of stubborn agents provided that $\sum_s d_s$ grows faster than $\log^2 n$, and slower than $n/\log n$.

Example 8: (Watts & Strogatz's small world) For the small-world model of Example 4, one has that both the average degree and the average square degree are bounded, so that $\pi(\mathcal{D}) = O(n^{-1})$, while $\tau_2 \leq \tau = O(\log^3 n)$, with high probability. This implies that (3) holds provided that $\sum_{s \in \mathcal{S}} d_s$ grows faster than $\log^3 n$ and slower than $n/\log^3 n$.

VII. 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 do not converge almost surely, and keep on oscillating according to an ergodic distribution. First and second moments of the ergodic beliefs distribution can be characterized in terms of the hitting probabilities of a random walk on the network, while the correlation between the ergodic beliefs of any pair of regular agents can be characterized in terms of the hitting probabilities of a pair of coupled random walks. We have shown that in highly fluid, reversible 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 expected ergodic beliefs and of the ergodic variances are almost constant, so that the stubborn agents have approximately the same influence on the society. Finally, we have also shown that in highly fluid social networks in which there is a significant presence of stubborn agents, the variance of the ergodic aggregate belief of the system vanishes in the

limit of large population size, and the ergodic distribution of the agents beliefs approaches an approximately chaotic condition. This implies that, if the influence of any of the stubborn agents' opinions does not dominate the influence of the rest, then the mean square disagreement does not vanish in the large population size.

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*, Preprint, 2008.
- [3] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, *Spread of (mis)information in social networks*, LIDS report 2812, forthcoming in *Games and Economic Behavior*, 2009.
- [4] V. Bala and S. Goyal, *Learning from neighbours*, *Review of Economic Studies* **65** (1998), no. 3, 595–621.
- [5] A. Banerjee and D. Fudenberg, *Word-of-mouth learning*, *Games and Economic Behavior* **46** (2004), 1–22.
- [6] A.L. Barabasi and R. Albert, *Emergence of scaling in random networks*, *Science* **268** (1999), 509–512.
- [7] C. Cooper and A. Frieze, *The cover time of sparse random graphs*, *Random Structures and Algorithms* **30** (2006), no. 1-2, 1–16.
- [8] 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.
- [9] R. Durrett, *Random graph dynamics*, Cambridge University Press, 2006.
- [10] 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.
- [11] D. Gale and S. Kariv, *Bayesian learning in social networks*, *Games and Economic Behavior* **45** (2003), no. 2, 329–346.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] M. Mobilia, *Does a single zealot affect an infinite group of voters?*, *Physical Review Letters* **91** (2003), no. 2, 028701.
- [16] A. Nedić and A. Ozdaglar, *Distributed subgradient methods for multi-agent optimization*, *IEEE Transactions on Automatic Control* **54** (2009), no. 1, 48–61.
- [17] M.E.J. Newman and D.J. Watts, *Renormalization group analysis of the small-world network model*, *Physics Letters A* **263** (1999), 341–346.
- [18] 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.
- [19] A. Olshevsky and J.N. Tsitsiklis, *Convergence speed in distributed consensus and averaging*, forthcoming in *SIAM Journal on Control and Optimization*, 2008.
- [20] L. Smith and P. Sorensen, *Pathological outcomes of observational learning*, *Econometrica* **68** (2000), no. 2, 371–398.
- [21] 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.
- [22] 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.
- [23] D.J. Watts and S.H. Strogatz, *Collective dynamics of 'small-world' networks*, *Nature* **393** (1998), 440–442.