Opinion fluctuations and disagreement in social networks

Daron Acemoglu
Department of Economics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA, 02139
email: daron@mit.edu `http://econ-www.mit.edu/faculty/acemoglu/`

Giacomo Como
Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA, 02139
email: giacomo@mit.edu `http://mit.edu/giacomo/www/`

Fabio Fagnani
Dipartimento di Matematica, Politecnico di Torino, corso Stati Uniti 24, 10129, Torino, Italy
email: fabio.fagnani@polito.it `http://calvino.polito.it/~fagnani/`

Asuman Ozdaglar
Laboratory for Information and Decision Systems and Department of Electrical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA, 02139
email: asuman@mit.edu `http://web.mit.edu/asuman/www/`

We study a tractable opinion dynamics model that generates long-run disagreements and persistent opinion fluctuations. Our model involves a 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 stubborn agents within it shape 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 continuos 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 stationary beliefs allow for a similar characterization in terms of coupled random walks on the graph describing the social network.

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.

Key words: opinion dynamics, multi-agent systems, social networks, persistent disagreement, opinion fluctuations, social influence.

MSC2000 Subject Classification: Primary: 91D30, 60K35; Secondary: 93A15 (*See the MSC2000 codes at* `http://www.ams.org/msc/`)

OR/MS subject classification: Primary: Games/group decisions: stochastic; Secondary: Markov processes, Random walk (*See the OR/MS classification at* `http://mor.pubs.informs.org/ORSubject.pdf`)

---

1. 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 [43], Banerjee and Fudenberg [7], Acemoglu, Dahleh, Lobel and Ozdaglar [1], Bala and Goyal [6], Gale and Kariv [23], DeMarzo, Vayanos and Zwiebel [17], Golub and Jackson [24], Acemoglu, Ozdaglar and ParandehGheibi [2]), 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 [27], and Deffuant and Weisbuch [16], in the continuous opinion dynamics framework. Such beliefs dynamics typically lead to the emergence of different asymptotic opinion clusters (see, e.g., [31, 9, 12]); 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. 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 continuos 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 stationary 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–*gossip* 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 a 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 [18]. 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 [3, 28], as well as on results in modern random graph theory [19].

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 [33, 34, 35], 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 [33, 34], or are in fact equivalent to our stubborn agents in [35]. These works generally present analytical results for some regular graphical structures (such as regular lattices [33, 34], or complete graphs [35]), 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, we will also show that the voter model with zealots of [35] 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 [45], Tsitsiklis, Bertsekas and Athans [46], Jadbabaie, Lin and Morse [26], Olfati-Saber and Murray [39], Olshevsky and Tsitsiklis [40], Fagnani and Zampieri [22], Nedić and Ozdaglar [36]). 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 2, 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 3 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 they do not converge almost surely. Section 4 presents a characterization of the first and second moments of the stationary beliefs in terms of the hitting probabilities of two coupled random walks on the graph describing the social network. Section 5 narrows down the discussion to reversible social networks, and presents explicit computations of the expected stationary beliefs and variances for some special network topologies. Section 6 provides bounds on the level of dispersion of the first two moments of the stationary beliefs: it is shown that, in highly fluid networks, most of the agents have almost the same stationary belief and variance. Section 7 presents some concluding remarks.

Basic Notation and Terminology We will 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 any finite set $\mathcal{J}$, we use the notation $\mathbb{1}_{\mathcal{J}}$ to denote the indicator function over the set $\mathcal{J}$, i.e., $\mathbb{1}_{\mathcal{J}}(j)$ is equal to 1 if $j \in \mathcal{J}$, and equal to 0 otherwise. For a matrix $M \in \mathbb{R}^{\mathcal{I}\times\mathcal{J}}$, $M^{\mathrm{T}} \in \mathbb{R}^{\mathcal{J}\times\mathcal{I}}$ will stand for its transpose, $||M||$ for its 2-norm. For a probability distribution μ over a finite set $\mathcal{I}$, and a subset $\mathcal{J} \subseteq \mathcal{I}$ we will write $\mu(\mathcal{J}) := \sum_j \mu_j$. If ν is another probability distribution on $\mathcal{I}$, we will use the notation $||\mu - \nu||_{TV} := \frac{1}{2}\sum_i |\mu_i - \nu_i| = \sup\{|\mu(\mathcal{J}) - \nu(\mathcal{J})| : \mathcal{J} \subseteq \mathcal{I}\}$, for the total variation distance between μ and ν. The probability law (or distribution) of a random variable Z will be denoted by $\mathcal{L}(Z)$. We will often refer to continuous-time Markov chains on a finite set $\mathcal{V}$ as random walks on $\mathcal{V}$. Such random walks will be characterized by their generator matrix $M \in \mathbb{R}^{\mathcal{V}\times\mathcal{V}}$, which has zero row sums, and whose non-diagonal elements are nonnegative and correspond to the rates at which the chain jumps from a state to another (see [38, Ch.s 2-3]). If $V(t)$ and $V'(t)$ are random walks on $\mathcal{V}$, defined on the same probability space, we will use the notation $\mathbb{P}_v(\,\cdot\,)$, and $\mathbb{P}_{vv'}(\,\cdot\,)$, for the conditional probability measures given the events $V(0) = v$, and, respectively, $(V(0), V'(0)) = (v, v')$. Similarly, for

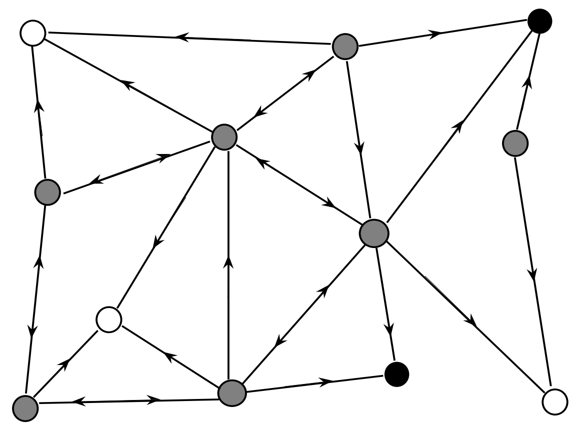

Figure 1: A social network with seven regular agents (colored in grey), and five stubborn agents (colored in white, and black, respectively). Links are only incoming to the stubborn agents, while links between pairs of regular agents may be uni- or bi-directional.

some probability distribution π over $\mathcal{V}$ (possibly the stationary one), $\mathbb{P}_\pi(\,\cdot\,) := \sum_{v,v'} \pi_v \pi_{v'} \mathbb{P}_{vv'}(\,\cdot\,)$ will denote the conditional probability measure of the random walk 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 real-valued sequences $\{a_n : n \in \mathbb{N}\}$, $\{b_n : n \in \mathbb{N}\}$, we will write $a_n = O(b_n)$ if for some positive constant K, $a_n \le K b_n$ for all sufficiently large n, $a_n = \Theta(b_n)$ if $b_n = O(a_n)$, $a_n = o(b_n)$ if $\lim_n a_n/b_n = 0$.

2. 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 directed graph $\overrightarrow{\mathcal{G}} = \big(\mathcal{V}, \overrightarrow{\mathcal{E}}\big)$, whose node set is identified with the agent population, and where $\overrightarrow{\mathcal{E}} \subseteq \mathcal{V} \times \mathcal{V} \setminus \mathcal{D}$, with $\mathcal{D} := \{(v, v) : v \in \mathcal{V}\}$, stands for the set of directed edges (or links) among the agents.[1]

At time $t \ge 0$, 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 out-neighbors in $\overrightarrow{\mathcal{G}}$. *Stubborn agents* never change their opinions, i.e., they do not have any out-neighbors. Agents which are not stubborn are called *regular*. We will 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 directed edge in $\overrightarrow{\mathcal{E}}$ of the form (a, v), where necessarily $a \in \mathcal{A}$, and $v \in \mathcal{V}$, a clock is associated, ticking at the times of an independent Poisson process of rate $r_{av} > 0$. 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:

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

where $X_v(t^-)$ stands for the left limit $\lim_{u \uparrow t} X_v(u)$. Here, the scalar $\theta_{av} \in (0, 1]$ is a *trust parameter* that represents the confidence that the regular agent $a \in \mathcal{A}$ puts on agent v's belief.[2] That r_{av} and θ_{av} are strictly positive for all $(a, v) \in \overrightarrow{\mathcal{E}}$ is simply a convention (since if $r_{av}\theta_{av} = 0$, one can always consider the subgraph of $\overrightarrow{\mathcal{G}}$ obtained by removing the edge (a, v) from $\overrightarrow{\mathcal{E}}$). Similarly, we also adopt the convention

[1]Notice that we don't allows for parallel edges or self-loops.

[2]We have imposed that at each meeting instance, only one agent updates her belief. The model can be easily extended to the case where both agents update their beliefs simultaneously, without significantly affecting any of our general results.

that $r_{vv'} = \theta_{vv'} = 0$ for all $v, v' \in \mathcal{V}$ such that $(v, v') \notin \overrightarrow{\mathcal{E}}$ (hence, including self-loops $v' = v$). 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 directed path in $\overrightarrow{\mathcal{G}}$. 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 tuple $\mathcal{N} = \left(\overrightarrow{\mathcal{G}}, \{\theta_e\}, \{r_e\}\right)$ 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, $\mathcal{L}(X(0))$, the social network designates a *society*. Throughout the paper, we make the following assumptions regarding the underlying social network.

Assumption 2.1 *Every regular agent is influenced by some stubborn agent, i.e., $\mathcal{S}_a$ is non-empty for every a in $\mathcal{A}$.*

Assumption 2.1 may be easily removed. If there are some regular agents which are not influenced by any stubborn agent, then there is no edge in $\mathcal{E}$ connecting the set $\mathcal{R}$ of such regular agents to $\mathcal{V} \setminus \mathcal{R}$. Then, one may decompose the subgraph obtained by restricting $\mathcal{G}$ to $\mathcal{R}$ into its communicating classes, and apply the results in [22] (see Example 3.5 therein), showing that, with probability one, a consensus on a random belief is achieved on every such communicating class.

We denote the *total meeting rate of agent* $v \in \mathcal{V}$ by r_v, i.e., $r_v := \sum_{v' \in \mathcal{V}} r_{vv'}$, and the *total meeting rate of all agents* by r, i.e., $r := \sum_{v \in \mathcal{V}} r_v$. We use $N(t)$ to denote the *total number of agent meetings (or edge activations) up to time* $t \geq 0$, which is simply a Poisson arrival process of rate r. We also use the notation T_k to denote the *time of the k-th belief update*, i.e., $T_k := \inf\{t \geq 0 : N(t) \geq k\}$.

For a given social network, we associate two matrices $Q \in \mathbb{R}^{\mathcal{V} \times \mathcal{V}}$, with entries

$$Q_{vw} := \theta_{vw} r_{vw} \qquad Q_{vv} := -\sum\nolimits_{v' \neq v} Q_{vv'} \,, \qquad v \neq w \in \mathcal{V} \,, \tag{2}$$

and $P \in \mathbb{R}^{\mathcal{A} \times \mathcal{V}}$, whose entries are defined by

$$P_{av} = -Q_{av}/Q_{vv} \,, \qquad a \in \mathcal{A} \,,\; v \in \mathcal{V} \,. \tag{3}$$

The following example describes the canonical construction of a social network from an undirected graph, and will be used often in the rest of the paper.

Example 2.1 *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected connected graph, and $\mathcal{S} \subseteq \mathcal{V}$, $\mathcal{A} = \mathcal{V} \setminus \mathcal{S}$. Define the directed graph $\overrightarrow{\mathcal{G}} = (\mathcal{V}, \overrightarrow{\mathcal{E}})$, where $(a, v) \in \overrightarrow{\mathcal{E}}$ if and only if $a \in \mathcal{A}$, $v \in \mathcal{V}$, and $\{a, v\} \in \mathcal{E}$, i.e., $\overrightarrow{\mathcal{G}}$ is the directed graph obtained by making all edges in $\mathcal{E}$ bidirectional except edges between a regular and a stubborn agent, which are unidirectional (pointing from the regular agent to the stubborn agent). For every node $v \in \mathcal{V}$, let d_v be its degree in $\mathcal{G}$. Let the trust parameter be constant, i.e., $\theta_{av} = \theta \in (0, 1]$ for all $(a, v) \in \overrightarrow{\mathcal{E}}$. Define*

$$r_{av} = 1/d_a \,, \qquad a \in \mathcal{A} \,,\; v \in \mathcal{V} : \; \{a, v\} \in \mathcal{E} \,. \tag{4}$$

This concludes the construction of the social network $\mathcal{N} = \left(\overrightarrow{\mathcal{G}}, \{\theta_e\}, \{r_e\}\right)$. For this social network, one has

$$Q_{av} = \theta/d_a \,, \qquad P_{av} = 1/d_a \,, \qquad \forall (a, v) \in \overrightarrow{\mathcal{E}} \,.$$

We conclude this section by discussing in some detail two special cases whose simple structure sheds light on the main features of the general model. In particular, we consider a social network with a single regular agent and a social network where the trust parameter satisfies $\theta_{av} = 1$ for all $a \in \mathcal{A}$ and $v \in \mathcal{V}$. We show that in both of these cases agent beliefs almost surely fail to converge.

2.1 Single regular agent Consider a society consisting of a single regular agent, i.e., $\mathcal{A} = \{a\}$, and two stubborn agents, $\mathcal{S} = \{s, s'\}$ (see Fig. 2(a)). Assume that $r_{as} = r_{as'} = 1/2$, $\theta_{as} = \theta_{as'} = 1/2$, $x_s = 0$, $x_{s'} = 1$, and $X_a(0) = 0$. Then, one has for all $t \geq 0$,

$$X_a(t) = \sum_{1 \leq k \leq N(t)} 2^{k - N(t) - 1} W(k) \,,$$

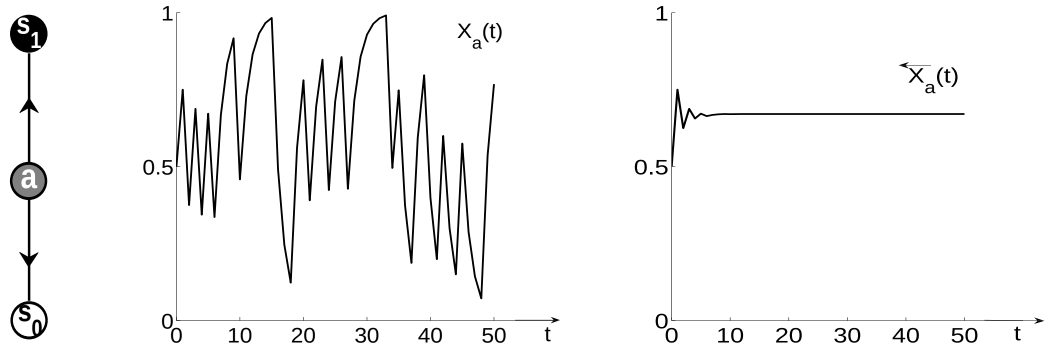

(a) (b) (c)

Figure 2: Typical sample-path behavior of the belief of the regular agent in the simple social network topology depicted in (a). In (b) the actual belief process $X_a(t)$, oscillating ergodically on the interval $[0,1]$; in (c), the time-reversed process, rapidly converging to a stationary random belief X_a.

where $N(t)$ is the total number of agent meetings up to time t (or number of arrivals up to time t of a rate-1 Poisson process), and $\{W(k) : k \in \mathbb{N}\}$ is a sequence of Bernoulli(1/2) random variables, independent mutually and from the process $N(t)$. Observe that, almost surely, arbitrarily long strings of contiguous zeros and ones appear in the i.i.d. sequence $\{W(k)\}$, while the number of meetings $N(t)$ grows unbounded. It follows that, with probability one

$$\liminf_{t\to\infty} X_a(t) = 0\,, \qquad \limsup_{t\to\infty} X_a(t) = 1\,,$$

so that the belief $X_a(t)$ does not converge almost surely.

On the other hand, observe that, since $\sum_{k>n} 2^{-k}|W(k)| \le 2^{-n}$, the series $X_a := \sum_{k\ge1} 2^{-k}W(k)$ is sample-wise converging. It follows that, as t grows large, the time-reversed process

$$\overleftarrow{X_a}(t) := \sum_{1\le k\le N(t)} 2^{-k}W(k)$$

converges to X_a, with probability one, and, a fortiori, in distribution. Notice that, for all positive integer k, the binary k-tuples $\{W(1),\ldots,W(k)\}$ and $\{W(k),\ldots,W(1)\}$ are uniformly distributed over $\{0,1\}^k$, and independent from the Poisson arrival process $N(t)$. It follows that, for all $t \ge 0$, the random variable $\overleftarrow{X}_a(t)$ has the same distribution as $X_a(t)$. Therefore, $X_a(t)$ converges in distribution to X_a as t grows large. Moreover, it is a standard fact (see e.g. [42, pag.92]) that X_a is uniformly distributed over the interval $[0,1]$. Hence, the probability distribution of $X_a(t)$ is asymptotically uniform on $[0,1]$.

The analysis can in fact be extended to any trust parameter $\theta_{is} = \theta_{is'} = \theta \in (0,1)$. In this case, one gets that

$$X_a(t) = \theta \sum_{1\le k\le N(t)} (1-\theta)^{N(t)-k}W(k)$$

converges almost surely to the stationary belief

$$X_a := \theta(1-\theta)^{-1}\sum_{k\ge1}(1-\theta)^kW(k)\,. \tag{5}$$

As explained in [18, Section 2.6], for every value of θ in $(1/2,1)$, the probability law of X_a is singular, and in fact supported on a Cantor set. In contrast, for almost all values of $\theta \in (0,1/2)$, the probability law of X_a is absolutely continuous with respect to Lebesgue's measure.[3] In the extreme case $\theta = 1$, it is not hard to see that $X_a(t) = W(N(t))$ converges in distribution to a random variable X_a with Bernoulli(1/2) distribution. On the other hand, observe that, regardless of the fine structure of the probability law of the stationary belief X_a, i.e., on whether it is absolutely continuous or singular, it is not hard to characterize

[3] See [41]. In fact, explicit counterexamples of values of $\theta \in (0,1/2)$ for which the asymptotic measure is singular are known. For example, Erdös [20, 21] showed that, if $\theta = (3-\sqrt{5})/2$, then the probability law of X_a is singular.

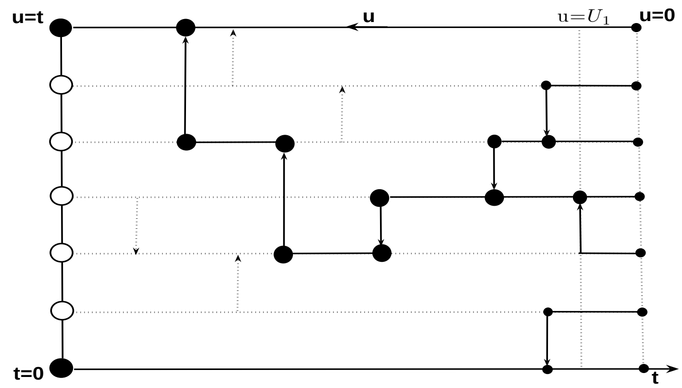

Figure 3: Duality between the voter model with zealots and the coalescing random walks process with absorbing states. The network topology is a line with five regular agents and two stubborn agents placed in the two extremities. The time index for the opinion dynamics, t, runs from left to right, whereas the time index for the coalescing random walks process, u, runs from right to left. Both dotted and solid arrows represent meeting instances. Fixing a time horizon $t > 0$, in order to trace the beliefs $X(t)$, one has to follow coalescing random walks starting at $u = 0$ in the different nodes of the network, and jumping from a state to another one in correspondence to the solid arrows. The particles are represented by bullets at times of their jumps. Clusters of coalesced particles are represented by bullets of increasing size.

its moments for all values of $\theta \in (0, 1]$. In fact, it follows from (5), that the expected value of X_a is given by

$$\mathbb{E}[X_a] = \theta(1-\theta)^{-1}\sum_{k\geq 1}(1-\theta)^k\mathbb{E}[W(k)] = \theta\sum_{k\geq 0}(1-\theta)^k\frac{1}{2} = \frac{1}{2},$$

and, using the mutual independence of the $W(k)$'s, the variance of X_a is given by

$$\mathrm{Var}[X_a] = \theta^2(1-\theta)^{-2}\sum_{k\geq 1}(1-\theta)^{2k}\,\mathrm{Var}[W(k)] = \theta^2\sum_{k\geq 0}(1-\theta)^{2k}\frac{1}{4} = \frac{\theta}{4(2-\theta)}.$$

2.2 Voter model with zealots We now consider the special case when the social network topology $\to G$ is arbitrary, and $\theta_{av} = 1$ for all $(a, v) \in \mathcal{E}$. In this case, whenever an edge $(a, v) \in \overrightarrow{\mathcal{E}}$ is activated, the regular agent a adopts agent v's current opinion as such, completely disregarding her own current opinion.

This opinion dynamics, known as the *voter model*, was introduced independently by Clifford and Sudbury [11], and Holley and Liggett [25]. It has been extensively studied in the framework of interacting particle systems [29, 30]. While most of the research focus has been on the case when the graph is an infinite lattice, the voter model on finite graphs, and without stubborn agents, was considered, e.g., in [13, 15], [3, Ch. 14], and [19, Ch. 6.9]: in this case, consensus is achieved in some finite random time, whose distribution depends on the graph topology only.

In some recent work [35] a variant with one or more stubborn agents (there referred to as *zealots*) has been proposed and analyzed on the complete graph. We wish to emphasize that such voter model with zealots can be recovered as a special case of our model, and hence our general results, to be proven in the next sections, apply to it as well. However, we briefly discuss this special case here, since proofs are much more intuitive, and allow one to anticipate some of the general results.

The main tool in the analysis of the voter model is the *dual process*, which runs backward in time and allows one to identify the source of the opinion of each agent at any time instant. Specifically, let us focus on the belief of a regular agent a at time $t > 0$. Then, in order to trace $X_a(t)$, one has to look at the last meeting instance of agent a that occurred no later than time t. If such a meeting instance occurred at some time $t - U_1 \in [0, t]$ and the agent met was $v \in \mathcal{V}$, then the belief of agent a at time t coincides with the one of agent v at time $t - U_1$, i.e., $X_a(t) = X_v(t - U_1)$. The next step is to look at the last meeting instance of agent v occurred no later than time $t - U_1$; if such an instance occurred at time

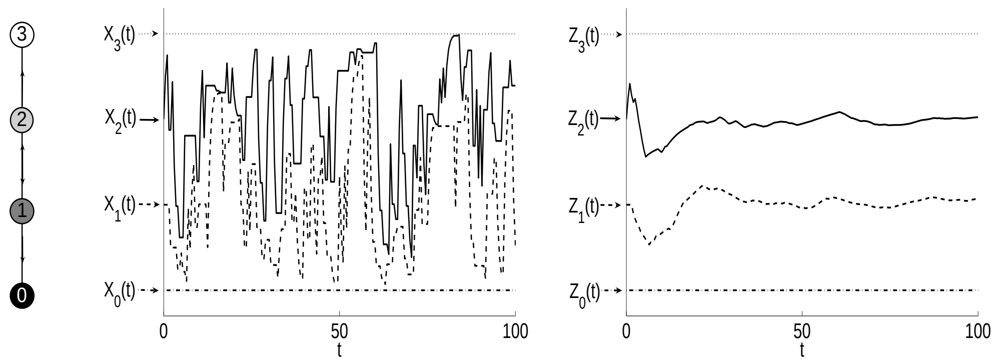

(a) (b) (c)

Figure 4: 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, displayed in (a). 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 $\{Z_v(t) := t^{-1} \int_0^t X_v(u)\mathrm{d}u\}$.

$t - U_2 \in [0, t - U_1]$, and the agent met was w, then $X_a(t) = X_v(t - U_1) = X_w(t - U_2)$. Clearly, one can iterate this argument, going backward in time, until reaching time 0. In this way, one implicitly defines a random walk $V_a(u)$ with state space $\mathcal{V}$, which starts at $V_a(0) = a$ and stays put there until time U_1, when it jumps to node v and stays put there in the time interval $[U_1, U_2)$, then jumps at time U_2 to node w, and so on. It is not hard to see that, thanks to the fact that the meeting instances are independent Poisson processes, the random walk $V_a(u)$ has generator matrix Q. In particular, it halts when it hits some state $s \in \mathcal{S}$. This shows that $\mathcal{L}(X_a(t)) = \mathcal{L}(X_{V_a(t)}(0))$. More generally, if one is interested in the joint probability distribution of the belief vector $X(t)$, then one needs to consider $n - |\mathcal{S}|$ random walks, $\{V_a(t) : a \in \mathcal{A}\}$ each one starting from a different node $a \in \mathcal{A}$, and run simultaneously on $\mathcal{V}$ (see Figure 3). These random walks move independently with generator matrix Q, until the first time that they either meet, or they hit the set $\mathcal{S}$: in the former case, they stick together and continue moving on $\mathcal{V}$ as a single particle, with generator matrix Q; in the second case, they halt. This process is known as the *coalescing random walk process* with absorbing set $\mathcal{S}$. Then, one gets that

$$\mathcal{L}(\{X_a(t) : a \in \mathcal{A}\}) = \mathcal{L}(\{X_{V_a(t)}(0) : a \in \mathcal{A}\}). \tag{6}$$

Equation (6) establishes a *duality* between the voter model with zealots and the coalescing random walks process with absorbing states. In particular, Assumption 2.1 implies that, with probability one, each $V_a(u)$ will hit the set $\mathcal{S}$ in some finite random time $T^a_{\mathcal{S}}$, so that in particular the vector $\{V_a(u) : a \in \mathcal{A}\}$ converges in distribution to an $\mathcal{S}^{\mathcal{A}}$-valued random vector $\{V_a(T^a_{\mathcal{S}}) : a \in \mathcal{A}\}$. It then follows from (6) that $X(t)$ converges in distribution to a stationary belief vector X whose entries are given by $X_s = x_s$ for every stubborn agent $s \in \mathcal{S}$, and $X_a = x_{V_a(T^a_{\mathcal{S}})}$ for every regular agent $a \in \mathcal{A}$.

3. 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 Section 2. Figure 4 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. 4(b), the beliefs of the two regular agents, $X_1(t)$, and $X_2(t)$, oscillate persistently in the interval $[x_0, x_3]$. 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 will 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: we have seen this in the special cases of Section 2.1, while a general result in this sense will be stated as Theorem 3.2. On the other hand, we will prove that, regardless of the initial regular agents' beliefs, the belief vector $X(t)$ is convergent in distribution to a random stationary belief vector X (see Theorem 3.1), and in fact it is an ergodic process (see Corollary 3.1).

In order to prove Theorem 3.1, we will rewrite $X(t)$ in the form of an iterated affine function system and apply standard techniques in this field [18]. In particular, we will consider the so-called time-reversed belief process. This is a stochastic process whose marginal probability distribution, at any time $t \geq 0$, coincides with the one of the actual belief process, $X(t)$. In contrast to $X(t)$, the time-reversed belief process is in general not Markov, whereas it can be shown to converge to a random stationary belief vector with probability one. From this, we recover convergence in distribution of the actual belief vector $X(t)$.

Formally, for any time instant $t \geq 0$, let us introduce the projected belief vector $Y(t) \in \mathbb{R}^{\mathcal{A}}$, where $Y_a(t) = X_a(t)$ for all $a \in \mathcal{A}$. Let $I_{\mathcal{A}} \in \mathbb{R}^{\mathcal{A}\times\mathcal{A}}$ be the identity matrix, and for $a \in \mathcal{A}$, let $e_{(a)} \in \mathbb{R}^{\mathcal{A}}$ be the vector whose entries are all zero, but for the a-th which equals 1. For every positive integer k, consider the random matrix $A(k) \in \mathbb{R}^{\mathcal{A}\times\mathcal{A}}$, and the random vector $B(k) \in \mathbb{R}^{\mathcal{A}}$, defined by

$$A(k) = I_{\mathcal{A}} + \theta_{aa'}\left(e_{(a)}e_{(a')}^{\mathrm{T}} - e_{(a)}e_{(a)}^{\mathrm{T}}\right) \qquad B(k) = 0\,,$$

if the k-th activated edge is $(a, a') \in \overrightarrow{\mathcal{E}}$, with $a, a' \in \mathcal{A}$, and

$$A(k) = I_{\mathcal{A}} - \theta_{as}e_{(a)}e_{(a)}^{\mathrm{T}} \qquad B(k) = e_{(a)}\theta_{as}x_s\,,$$

if the k-th activated edge is $(a, s) \in \overrightarrow{\mathcal{E}}$, with $a \in \mathcal{A}$, and $s \in \mathcal{S}$. Define the matrix product

$$\overrightarrow{A}(k,l) := A(l)A(l-1)\dots A(k+1)A(k)\,, \qquad 1 \leq k \leq l\,, \tag{7}$$

with the convention that $\overrightarrow{A}(k,l) = I_{\mathcal{A}}$ for $k > l$. Then, one has

$$Y(T_n) = A(n)Y(T_n^-) + B(n) = A(n)Y(T_{n-1}) + B(n)\,, \qquad n \geq 1\,,$$

so that, for all $t \geq 0$,

$$Y(t) = \overrightarrow{A}(1, N(t))Y(0) + \sum_{1\leq k\leq N(t)} \overrightarrow{A}(k+1, N(t))B(k)\,, \tag{8}$$

where we recall that $N(t)$ is the total number of agents' meetings up to time t. Now, define the *time-reversed belief process*

$$\overleftarrow{Y}(t) := \overleftarrow{A}(1, N(t))Y(0) + \sum_{1\leq k\leq N(t)} \overleftarrow{A}(1, k-1)B(k)\,, \tag{9}$$

where

$$\overleftarrow{A}(k,l) := A(k)A(k+1)\dots A(l-1)A(l)\,, \qquad k \leq l\,,$$

with the convention that $\overleftarrow{A}(k,l) = I_{\mathcal{A}}$ for $k > l$. The following is a fundamental observation see [18]:

Lemma 3.1 *For all $t \geq 0$, $Y(t)$ and $\overleftarrow{Y}(t)$ have the same probability distribution.*

Proof. Notice that $\{(A(k), B(k)) : k \in \mathbb{N}\}$ is a sequence of independent and identically distributed random variables, independent from the process $N(t)$. This, in particular, implies that, the l-tuple $\{(A(k), B(k)) : 1 \leq k \leq l\}$ has the same distribution as the l-tuple $\{(A(l-k+1), B(l-k+1)) : 1 \leq k \leq l\}$, for all $l \in \mathbb{N}$. From this, and the identities (8) and (9), it follows that the belief vector $Y(t)$ has the same distribution as $\overleftarrow{Y}(t)$, for all $t \geq 0$. ■

The second fundamental result is that, in contrast to the actual regular agents' belief vector $Y(t)$, the time-reversed belief process $\overleftarrow{Y}(t)$ converges almost surely. This is formalized in the next lemma, whose proof relies on showing that the matrices $A(k)$ are contractive in the average, as a consequence of Assumption 2.1.

Lemma 3.2 *Let Assumption 2.1 hold. Then, for every value of the stubborn agents' beliefs $\{x_s\} \in \mathbb{R}^{\mathcal{S}}$, there exists an $\mathbb{R}^{\mathcal{A}}$-valued random variable Y, such that,*

$$\mathbb{P}\left(\lim_{t\to\infty} \overleftarrow{Y}(t) = Y\right) = 1\,,$$

for every initial distribution $\mathcal{L}(Y(0))$ of the regular agents' beliefs.

Proof. Observe that the expected entries of $A(k)$, and $B(k)$, are given by

$$\mathbb{E}[A_{aa'}(k)] = \frac{Q_{aa'}}{r}\,, \qquad \mathbb{E}[A_{aa}(k)] := 1 - \frac{1}{r}\sum_{v\in\mathcal{V}} Q_{av}\,, \qquad \mathbb{E}[B_a(k)] = \frac{1}{r}\sum_{s\in\mathcal{S}} Q_{as}x_s\,,$$

for all $a \neq a' \in \mathcal{A}$. In particular, $\mathbb{E}[A(k)]$ is a substochastic matrix. Moreover, it follows from Assumption 2.1 that there exists j such that $\sum_a \mathbb{E}[A_{ja}](k) < 1$. This implies that the spectral radius ρ of $\mathbb{E}[A(k)]$ is strictly less than one. Standard linear algebra then shows that

$$\left|\left|\mathbb{E}\left[\overleftarrow{A}(1,k)\right]\right|\right| \leq Ck^{n-1}\rho^k\,, \qquad \forall k \geq 0\,,$$

where C is a constant depending on $\mathbb{E}[A(1)]$ only. It follows that

$$\inf_{k\in\mathbb{N}} \frac{1}{k}\mathbb{E}\left[\log||\overleftarrow{A}(1,k)||_1\right] \leq \inf_{k\in\mathbb{N}} \frac{1}{k}\log\mathbb{E}\left[||\overleftarrow{A}(1,k)||_1\right] \leq \lim_{k\to\infty} \frac{\log(Cnk^{n-1}\rho^k)}{k} = \log\rho < 0\,. \tag{10}$$

Hence, it follows from [18, Th. 2.1] that, with probability one, the series

$$Y := \sum_{k\geq 1} \overleftarrow{A}(1, j-1)B(j)$$

is convergent, and $\overleftarrow{A}(1,k)Y(0)$ converges to 0 as k grows large. Since, with probability one, $N(t)$ goes to infinity as t grows large, one has that

$$\lim_{t\to\infty} \overleftarrow{Y}(t) = \lim_{t\to\infty} \overleftarrow{A}(1, N(t))Y(0) + \sum_{1\leq j\leq N(t)} \overleftarrow{A}(1, j-1)B(j) = Y\,,$$

with probability one. This completes the proof. ∎

Lemma 3.1 and Lemma 3.2 allow one to prove convergence in distribution of $X(t)$ to a random belief vector X, as stated in the following result.

Theorem 3.1 *Let Assumption 2.1 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 distribution $\mathcal{L}(X(0))$ satisfying $\mathbb{P}(X_s(0) = x_s) = 1$ for every $s \in \mathcal{S}$,*

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

for all bounded and continuous test functions $\varphi : \mathbb{R}^{\mathcal{V}} \to \mathbb{R}$. Moreover, the probability law of the stationary belief vector X is invariant for the system, i.e., if $\mathcal{L}(X(0)) = \mathcal{L}(X)$, then $\mathcal{L}(X(t)) = \mathcal{L}(X)$ for all $t \geq 0$.

Proof. It follows from Lemma 3.2 $\overleftarrow{Y}(t)$ converges to Y with probability one, and a fortiori in distribution. By Lemma 3.1, $\overleftarrow{Y}(t)$ and $Y(t)$ are identically distributed. Therefore, $Y(t)$ converges to Y in distribution, and the first part of the claim follows by defining $X_a = Y_a$ for all $a \in \mathcal{A}$, and $X_s = x_s$ for all $s \in \mathcal{S}$. For the second part of the claim, it is sufficient to observe that the distribution of $Y = \sum_{k\geq 1} \overleftarrow{A}(1, k-1)B(k)$ is the same as the one of $Y' := A(0)Y + B(0)$, where $A(0)$, and $B(0)$, are independent copies of $A(1)$, and $B(1)$, respectively. ∎

Motivated by Theorem 3.1, for any agent $v \in \mathcal{V}$, we refer to the random variable X_v as the *stationary belief of agent v*. Using standard ergodic theorems for Markov chains, an immediate implication of Theorem 3.1 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 3.1 *Let Assumption 2.1 hold. Then, for every value of the stubborn agents' beliefs $\{x_s : s \in \mathcal{S}\}$, with probability one,*

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

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

Proof. Let $Y(t)$ and Y be the projections of the belief vector at time $t \geq 0$, and of the stationary belief vector A, respectively, to the regular agents set $\mathcal{A}$. Let $\tilde{Y}(0)$ be an $\mathbb{R}^{\mathcal{A}}$-valued random vector, independent from $Y(0)$ and such that $\mathcal{L}(\tilde{Y}(0)) = \mathcal{L}(Y)$. Let $Y(t)$ be as in (8), and

$$\tilde{Y}(t) = \overrightarrow{A}(1, N(t))\tilde{Y}(0) + \sum_{1 \leq k \leq N(t)} \overrightarrow{A}(k+1, N(t))B(k)\,,$$

where $N(t)$ is the total number of agents' meetings up to time t, and $\overrightarrow{A}(k,l)$ is defined as in (7). Then,

$$\tilde{Y}(t) - Y(t) = \overrightarrow{A}(1, N(t))(\tilde{Y}(0) - Y(0))\,.$$

Arguing as in the proof of Lemma 3.2, one shows that $\lim_{t\to\infty} ||\tilde{Y}(t) - Y(t)|| = 0$, with probability one. Now, for $t > 0$, let the vectors $\tilde{X}(t)$ and $X(t)$ be defined by $\tilde{X}_a(t) = \tilde{Y}_a(t)$, $X_a(t) = Y_a(t)$ for $a \in \mathcal{A}$, and $\tilde{X}_s(t) = X_s(t) = x_s$ for $s \in \mathcal{S}$, and observe that, with probability one, $\sup_{t\geq 0} |X(t)| \leq \max_v |X_v(0)| < \infty$, $\sup_{t\geq 0} |\tilde{X}(t)| \leq \max_v |X_v| < \infty$. Then, for every continuous $\varphi : \mathbb{R}^{\mathcal{V}} \to \mathbb{R}$, one has that

$$\lim_{t\to\infty} |\varphi(\tilde{X}(t)) - \varphi(X(t))| = 0\,,$$

with probability one. On the other hand, stationarity of the process $\tilde{X}(t)$ allows one to apply the ergodic theorem (see, e.g., [44, Theorem 6.2.12]), showing that, if $\mathbb{E}[\varphi(X)]$ exists and is finite, then

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

with probability one. Then, for any continuous φ such that $\mathbb{E}[\varphi(X)]$ exists and is finite, one has that

$$\left|\frac{1}{t}\int_0^t \varphi(X(s))\mathrm{d}s - \mathbb{E}[\varphi(X)]\right| \leq \left|\frac{1}{t}\int_0^t \varphi(\tilde{X}(s))\mathrm{d}s - \mathbb{E}[\varphi(X)]\right| + \frac{1}{t}\int_0^t \left|\varphi(X(s)) - \varphi(\tilde{X}(s))\right| \mathrm{d}s \stackrel{t\to\infty}{\longrightarrow} 0\,,$$

with probability one. ■

Theorem 3.1, and Corollary 3.1, respectively, show that the beliefs of all the agents converge in distribution, and that their empirical distributions converge almost surely, to a random stationary belief vector X. In contrast, the following theorem shows that the stationary 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 $\mathcal{X}_a = \{x_s : s \in \mathcal{S}_a\}$ denote the set of stubborn agents' belief values influencing agent a.

Theorem 3.2 *Let Assumption 2.1 hold, and let $a \in \mathcal{A}$ be such that $|\mathcal{X}_a| \geq 2$. Then, the stationary 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 $\mathbb{P}(X_a \neq X_{a'}) > 0$.*

Proof. With no loss of generality, since the distribution of the stationary belief vector X does not depend on the probability law of the initial beliefs of the regular agents, we can assume that such a law is the stationary one, i.e., that $\mathcal{L}(X(0)) = \mathcal{L}(X)$. Then, Theorem 3.1 implies that

$$\mathcal{L}(X(t)) = \mathcal{L}(X)\,, \qquad \forall t \geq 0\,. \tag{11}$$

Let $a \in \mathcal{A}$ be such that X_a is degenerate. Then, almost surely, $X_a(t) = x_a$ for almost all t, for some constant x_a. Then, as we will show below, all the out-neighbors of a will have their beliefs constantly equal to x_a with probability one. Iterating the argument until reaching the set $\mathcal{S}_a$, one eventually finds that $x_s = x_a$ for all $s \in \mathcal{S}$, so that $|\mathcal{X}_a| = 1$. This proves the first part of the claim. For the second part, assume that $X_a = X_{a'}$ almost surely for some $a \neq a'$. Then, one can prove that, with probability one, every out-neighbor of a or a' agrees with a or a' at any time. Iterating the argument until reaching the set $\mathcal{S}_a \cup \mathcal{S}_{a'}$, one eventually finds that $|\mathcal{X}_s \cup \mathcal{X}_{s'}| = 1$.

One can reason as follows in order to see that, if v is an out-neighbor of a, and $X_a = x_a$ is degenerate, then $X_v(t) = x_a$ for all t. Let T_n^{av} be the n-th activation of the edge (a, v). Then, Equation (1) implies that

$$X_v(T_n^{av}) = x_a\,, \qquad \forall n \geq 1\,. \tag{12}$$

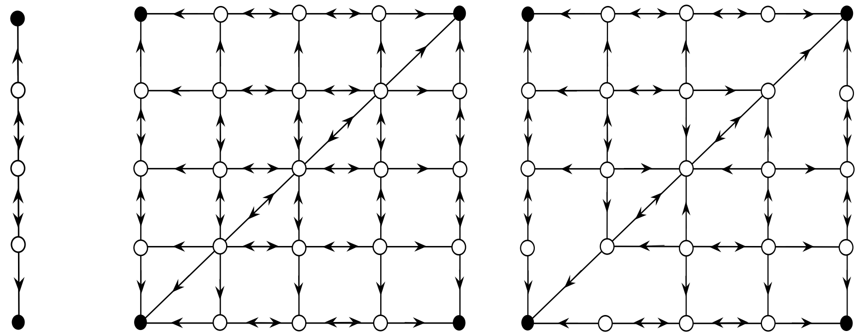

Figure 5: In (a), a network topology consisting of a line with three regular agents and two stubborn agents placed in the extremes. In (b), the corresponding graph product $\mathcal{G}_\square$. The latter has 25 nodes, four of which are absorbing states. The pair $(V(t), V'(t))$ of the coupled random walks 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. In (c), the product graph $\mathcal{G}_\square$ is reported for the extreme case when $\theta_e = 1$ for all $e \in \overrightarrow{\mathcal{E}}$. In this case, the coupled random walks $(V(t), V'(t))$ are coalescing: once they meet, they stick together, moving as a single particle, and never separating from each other. This reflects the fact that there are no outgoing edges from the diagonal set.

Now, define $T^* := \inf\{t \geq 0 : X_v(t) \neq x_a\}$, and assume by contradiction that $\mathbb{P}(T^* < \infty) > 0$. By the strong Markov property, and the property that edge activations are independent Poisson processes, this would imply that

$$\mathbb{P}(\text{first edge activated after } T^* \text{ is } (a, v) | \mathcal{F}_{T^*}) > 0 \qquad \text{on} \qquad \{T^* < \infty\}\,, \tag{13}$$

which would contradict (12). Then, necessarily $T^* = \infty$, and hence $X_v = X_a$, with probability one.

On the other hand, assume that $X_a = X_{a'}$ with probability one, let v be an out-neighbor of a, and T_n^{av} be the n-th activation of edge (a, v). Then,

$$X_v(T_n^{av}) = X_a(T_n^{av})\,, \qquad \forall n \geq 1\,. \tag{14}$$

Let $T^* := \inf\{t \geq 0 : X_v(t) \neq X_a(t)\}$, and assume by contradiction that $\mathbb{P}(T^* < \infty) > 0$. Then, arguing as above, one would find that (13) holds. Upon observing that almost surely two edges are never activated at the same time, this can be easily shown to contradict (14), so that necessarily $T^* = \infty$, and hence $X_v = X_a = X_{a'}$, with probability one. ∎

Even though, by Theorem 3.1, the belief of any agent always converges in distribution, Theorem 3.2 shows that, if a regular agent a is influenced by stubborn agents with different beliefs, then her stationary belief X_a is non-degenerate. By Corollary 3.1, this implies that, with probability one, her belief $X_a(t)$ keeps on oscillating and does not stabilize on a limit. Similarly, Theorem 3.2 and Corollary 3.1 imply that, if two regular agents are influenced by stubborn agents with different beliefs, then, with probability one, their beliefs will not achieve a consensus asymptotically.

4. Empirical averages and correlations of agent beliefs In this section, we provide a characterization of the empirical averages of the agents' beliefs and their cross-products, i.e., of the almost surely constant limits

$$\lim_{t\to\infty} \frac{1}{t}\int_0^t X_v(u)\mathrm{d}u, \qquad \lim_{t\to\infty} \frac{1}{t}\int_0^t X_v(u)X_{v'}(u)\mathrm{d}u\,, \qquad v, v' \in \mathcal{V}\,.$$

Since the stationary beliefs X_v take values in the interval $[\min_s x_s, \max_s x_s]$, one has that both $\mathbb{E}[X_v]$ and $\mathbb{E}[X_v X_{v'}]$ exist and are finite. Hence, Corollary 3.1 implies that the limits above are given by the first two moments of the stationary 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 $\overrightarrow{\mathcal{G}} = (\mathcal{V}, \overrightarrow{\mathcal{E}})$.[4] Specifically, we consider a coupling $(V(t), V'(t))$ of continuous-time Markov chains with state space $\mathcal{V}$, such that both $V(t)$ and $V'(t)$ have generator matrix Q, as defined in (2). In fact, one may interpret the continuous-time Markov chain $(V(t), V'(t))$ 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 $(v, w) \in \overrightarrow{\mathcal{E}}$ and $v' = w'$, or $v = w$ and $(v', w') \in \overrightarrow{\mathcal{E}}$, or $v = v'$ and $w = w'$ (See Figure 5). Such coupled random walks pair $(V(t), V'(t))$ has generator matrix K whose entries are given by

$$K_{(v,v')(w,w')} := \begin{cases} Q_{vw} & \text{if } v \neq v'\,,\ v \neq w\,,\ v' = w' \\ Q_{v'w'} & \text{if } v \neq v'\,,\ v = w\,,\ v' \neq w' \\ 0 & \text{if } v \neq v'\,,\ v \neq w\,,\ v' \neq w' \\ Q_{vv} + Q_{v'v'} & \text{if } v \neq v'\,,\ v = w\,,\ v' = w' \\ \theta_{vw} Q_{vw} & \text{if } v = v'\,,\ w = w'\,,\ v \neq w \\ (1-\theta_{vw}) Q_{vw} & \text{if } v = v'\,,\ v \neq w\,,\ v' = w' \\ (1-\theta_{vw'}) Q_{vw'} & \text{if } v = v'\,,\ v = w\,,\ v' \neq w' \\ 0 & \text{if } v = v'\,,\ w \neq w'\,,\ w \neq v\,,\ w' \neq v' \\ 2Q_{vv} + \sum_{v'' \neq v} \theta_{vv''} Q_{vv''} & \text{if } v = v'\,,\ v = w\,,\ v' = w'\,. \end{cases} \tag{15}$$

The first four lines of (15) 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 Q_{vw} (respectively, to a neighbor node w' with transition rate $Q_{v'w'}$), whereas the probability that both components jump at the same time is zero. On the other hand, the last five lines of (15) 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 $\theta_{vw} Q_{vw}$, while each of the components $V(t)$ (respectively, $V'(t)$) has still some chance to jump alone to a neighbor node w with rate $(1-\theta_{vw})Q_{vw}$ (resp., to w' with rate $(1-\theta_{vw'})Q_{vw'}$). In the extreme case when $\theta_{vw} = 1$ for all v, w, the sixth and seventh line of the righthand side of (15) 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.

For $v, w, v', w' \in \mathcal{V}$, and $t \geq 0$, we will denote by

$$\gamma^v_w(t) := \mathbb{P}_v(V(t) = w) = \mathbb{P}_v(V'(t) = w)\,, \qquad \eta^{vv'}_{ww'}(t) = \mathbb{P}_{vv'}(V(t) = w, V'(t) = w') \tag{16}$$

the marginal and joint transition probabilities of the two random walks at time t. It is a standard fact (see, e.g., [38, Theorem 2.8.3]) that such transition probabilities satisfy the *Kolmogorov backward equations*

$$\frac{\mathrm{d}}{\mathrm{d}t}\gamma^v_w(t) = \sum_{\tilde v} Q_{v\tilde v}\gamma^{\tilde v}_w(t)\,, \qquad \frac{\mathrm{d}}{\mathrm{d}t}\eta^{vv'}_{ww'}(t) = \sum_{\tilde v} K_{(v,v')(\tilde v,\tilde v')}\eta^{\tilde v\tilde v'}_{ww'}(t)\,, \qquad v, v', w, w' \in \mathcal{V}\,, \tag{17}$$

with initial condition

$$\gamma^v_w(0) = \begin{cases} 1 & \text{if} \quad v = w \\ 0 & \text{if} \quad v \neq w\,, \end{cases} \qquad \eta^{vv'}_{ww'}(0) = \begin{cases} 1 & \text{if} \quad (v, v') = (w, w') \\ 0 & \text{if} \quad (v, v') \neq (w, w')\,. \end{cases} \tag{18}$$

The next simple result provides a fundamental link between the belief evolution process introduced in Section 2 and the coupled random walks, by showing that the expected values and expected cross-products of the agents' beliefs satisfy the same linear system (17) of ordinary differential equations as the transition probabilities of $(V(t), V'(t))$.

Lemma 4.1 *For all $v, v' \in \mathcal{V}$, and $t \geq 0$, it holds*

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[X_v(t)] = \sum_w Q_{vw}\mathbb{E}[X_w(t)]\,, \qquad \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[X_v(t)X_{v'}(t)] = \sum_{w,w'} K_{(v,v')(ww')}\mathbb{E}[X_w(t)X_{w'}(t)]\,. \tag{19}$$

[4] Note that the set of states for such a random walk corresponds to the set of agents, therefore we use the terms "state" and "agent" interchangeably in the sequel.

Proof. Recall that, for the belief update model introduced in Section 2, arbitrary agents' interactions occur at the ticking T_k of a Poisson clock of rate r. Moreover, with conditional probability r_{vw}/r, any such interaction involves agent v updating her opinion to a convex combination of her current belief and the one of agent w, with weight θ_{vw} on the latter. It follows that, for all $k \geq 0$, and $v \in \mathcal{V}$,

$$\mathbb{E}[X_v(T_{k+1})|\mathcal{F}_{T_k}] - X_v(T_k) = \frac{1}{r}\sum_w r_{vw}\theta_{vw}(X_w(T_k) - X_v(T_k)) = \sum_w Q_{vw}X_w(T_k)\,.$$

Then, the above and the fact that the Poisson clock has rate r imply the left-most equation in (19).

Similarly, or all $v \neq v' \in \mathcal{V}$, one gets that

$$\begin{array}{rcl}\mathbb{E}[X_v(T_{k+1})X_{v'}(T_{k+1})|\mathcal{F}_{T_k}] - X_v(T_k)X_{v'}(T_k) & = & \frac{1}{r}\sum_w r_{vw}\theta_{vw}\left(X_w(T_k)X_{v'}(T_k) - X_v(T_k)X_{v'}(T_k)\right) \\ & & +\frac{1}{r}\sum_{w'} r_{v'w'}\theta_{v'w'}\left(X_v(T_k)X_{w'}(T_k) - X_v(T_k)X_{v'}(T_k)\right) \\ & = & \frac{1}{r}\sum_{w,w'} K_{(v,v')(w,w')}X_wX_{w'}\end{array}$$

as well as

$$\begin{array}{rcl}\mathbb{E}[X_v^2(T_{k+1})|\mathcal{F}_{T_k}] - X_v^2(T_k) & = & \frac{1}{r}\sum_w r_{vw}\mathbb{E}\left[\left((1-\theta_{vw})X_v + \theta_{vw}X_w\right)^2 - X_v^2\right] \\ & = & \frac{1}{r}\sum_w r_{vw}\theta_{vw}\left(\theta_{vw}X_w^2 + 2(1-\theta_{vw})X_vX_w - (2-\theta_{vw})X_v^2\right) \\ & = & \frac{1}{r}\sum_w Q_{vw}(1-\theta_{av})\left(X_aX_v - X_a^2\right) \\ & & +\frac{1}{r}\sum_{w'} Q_{vw'}(1-\theta_{vw'})\left(X_vX_{w'} - X_v^2\right) \\ & & +\frac{1}{r}\sum_w Q_{vw}\theta_{vw}\left(X_w^2 - X_v^2\right) \\ & = & \frac{1}{r}\sum_{w,w'} K_{(v,v)(w,w')}X_wX_{w'}\,.\end{array}$$

Then, the two identities above and the fact that the Poisson clock has rate r imply the right-most equation in (19). ∎

We are now in a position to prove the main result of this section characterizing the expected values and expected cross-products of the agents' stationary beliefs in terms of the hitting probabilities of the coupled random walks. Let us denote by $T_{\mathcal{S}}$ and $T'_{\mathcal{S}}$ the *hitting times* of the random walks $V(t)$, and respectively $V'(t)$, on the set of stubborn agents $\mathcal{S}$, i.e.,

$$T_{\mathcal{S}} := \inf\{t \geq 0 :\ V(t) \in \mathcal{S}\}\,, \qquad T'_{\mathcal{S}} := \inf\{t \geq 0 :\ V'(t) \in \mathcal{S}\}\,.$$

Observe that Assumption 2.1 implies that both $T_{\mathcal{S}}$ and $T'_{\mathcal{S}}$ are finite with probability one for every initial distribution of the pair $(V(0), V'(0))$. Hence, for all $v, v' \in \mathcal{V}$, we can define the *hitting probability distributions* γ^v over $\mathcal{S}$, and $\eta^{vv'}$ over $\mathcal{S}^2$, whose entries are respectively given by

$$\begin{array}{ll}\gamma_s^v := \mathbb{P}_v(V(T_{\mathcal{S}}) = s)\,, & s \in \mathcal{S}\,, \\ \eta_{ss'}^{vv'} := \mathbb{P}_{vv'}(V(T_{\mathcal{S}}) = s, V'(T'_{\mathcal{S}}) = s')\,, & s, s' \in \mathcal{S}\,.\end{array} \tag{20}$$

Then, we have the following:

Theorem 4.1 *Let Assumption 2.1 hold. Then, for every value of the stubborn agents' beliefs* $\{x_s :\ s \in \mathcal{S}\}$,

$$\mathbb{E}[X_v] = \sum_s \gamma_s^v x_s\,, \qquad \mathbb{E}[X_vX_{v'}] = \sum_{s,s'} \eta_{ss'}^{vv'} x_sx_{s'}\,, \tag{21}$$

for all $v, v' \in \mathcal{V}$. *Moreover,* $\{\mathbb{E}[X_v] :\ v \in \mathcal{V}\}$ *and* $\{\mathbb{E}[X_vX_{v'}] :\ v, v' \in \mathcal{V}\}$ *are the unique vectors in* $\mathbb{R}^{\mathcal{V}}$, *and* $\mathbb{R}^{\mathcal{V}\times\mathcal{V}}$ *respectively, satisfying*

$$\sum_v Q_{av}\mathbb{E}[X_v] = 0\,, \qquad \mathbb{E}[X_s] = x_s\,, \qquad \forall a \in \mathcal{A}\,, \quad \forall s \in \mathcal{S}\,, \tag{22}$$

$$\sum_{w,w'} K_{(a,a'),(w,w')}\mathbb{E}[X_wX_{w'}] = 0\,, \qquad \mathbb{E}[X_vX_{v'}] = \mathbb{E}[X_v]\mathbb{E}[X_{v'}]\,, \qquad \forall a, a' \in \mathcal{A}\,,\ \forall (v, v') \in \mathcal{V}^2 \setminus \mathcal{A}^2\,. \tag{23}$$

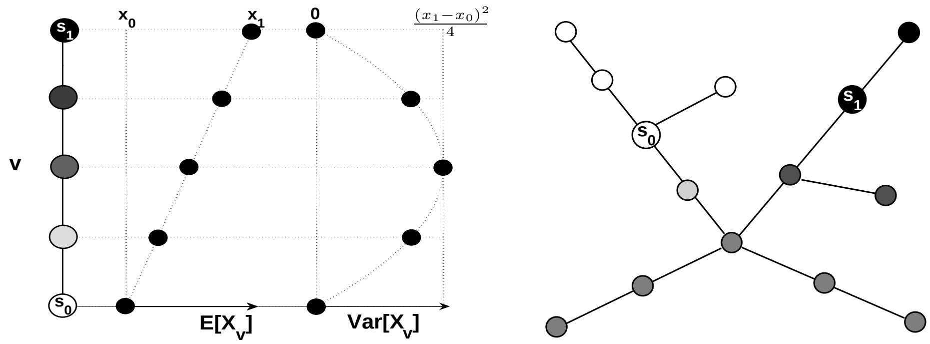

Figure 6: In the left-most figure, expected stationary beliefs and variances in a social network with a line graph topology with $n = 5$, and stubborn agents positioned in the two extremities. The expected stationary beliefs are linear interpolations of the two stubborn agents' beliefs, while their variances follow a parabolic profile with maximum in the central agent, and zero variance for the two stubborn agents s_0, and s_1. In the right-most figure, expected stationary beliefs in a social network with a tree-like topology, represented by different levels of gray. The solution is obtained by linearly interpolating between the two stubborn agents' beliefs, x_0 (white), and x_1 (black), on the vertices lying on the path between s_0 and s_1, and then extended by putting it constant on each of the connected components of the subgraph obtained by removing the edges of such path.

Proof. It follows from (17) and (19) that $\sum_s x_s \gamma_s^v(t)$ and $\mathbb{E}[X_v(t)]$ satisfy the same linear system of differential equations with the same initial condition. On the other hand, Assumption 2.1 implies that $\lim_{t\to\infty} \gamma_s^v(t) = \gamma_s^v$ for every $s \in \mathcal{S}$, and $\lim_{t\to\infty} \gamma_a^v(t) = 0$ for every $a \in \mathcal{A}$. Therefore, one gets that

$$\lim_{t\to\infty} \mathbb{E}[X_v(t)] = \sum_s \gamma_s^v x_s \,, \qquad \forall v \in \mathcal{V} \,. \tag{24}$$

Now, if the initial belief distribution $\mathcal{L}(X(0))$ coincides with the stationary one $\mathcal{L}(X)$, one has that $\mathcal{L}(X(t)) = \mathcal{L}(X)$ for all $t \geq 0$, so that in particular $\mathbb{E}[X_v(t)] = \mathbb{E}[X_v]$, and hence $\lim_{t\to\infty} \mathbb{E}[X_v(t)] = \mathbb{E}[X_v]$ for all v. Substituting in the righthand side of (24), this proves the leftmost identity in (21). The rightmost identity in (21) follows from an analogous argument.

In order to prove the second part of the claim, observe that the expected stationary beliefs and belief cross-products necessarily satisfy (22) and (23), since, by Lemma 4.1, they evolve according to the autonomous differential equations (19), and are convergent by the arguments above. On the other hand, uniqueness of the solutions of (22) and (23) follows from [3, Ch. 2, Lemma 27]. ■

Remark 4.1 *As a consequence of Theorem 4.1, one gets that, if $\mathcal{X}_a = \{x_*\}$, then $X_a = x_*$, and, by Corollary 3.1, $X_a(t)$ converges to x^* with probability one. This can be thought of as a sort of complement to Theorem 3.2.*

5. Explicit computations of expectations and variances of the stationary beliefs We present now a few examples of explicit computations of the stationary expected beliefs and variances for social networks obtained using the construction in Example 2.1, starting from a finite undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Recall that, in this case, $Q_{av} = \theta/d_a$ for all $a \in \mathcal{A}$, and $v \in \mathcal{V}$ such that $\{a, v\} \in \mathcal{E}$. It then follows from Theorem 4.1 that the expected stationary beliefs can be characterized as the unique vectors in $\mathbb{R}^{\mathcal{V}}$ satisfying

$$\mathbb{E}[X_a] = \frac{1}{d_a} \sum_{v:\{v,a\}\in\mathcal{E}} \mathbb{E}[X_v] \,, \qquad \mathbb{E}[X_s] = x_s \,, \qquad \forall a \in \mathcal{A} \,,\ \forall s \in \mathcal{S} \,, \tag{25}$$

and, if $\theta = 1$, the second moments of the stationary beliefs are the unique solutions of

$$\mathbb{E}[Xa^2] = \frac{1}{d_a} \sum_{v:\{v,a\}\in\mathcal{E}} \mathbb{E}[X_v^2] \,, \qquad \mathbb{E}[X_s^2] = x_s^2 \,, \qquad \forall a \in \mathcal{A} \,,\ \forall s \in \mathcal{S} \,. \tag{26}$$

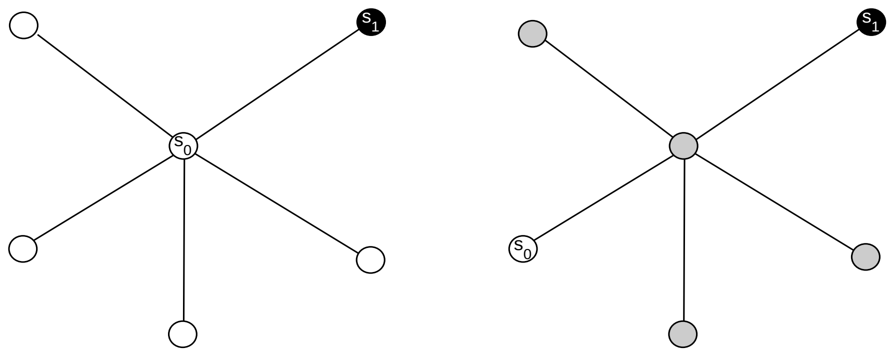

Figure 7: Two social network with a special case of tree-like topology, known as star graph, and two stubborn agents. In social network depicted in left-most figure one of the stubborn agents, s_0, occupies the center, while the other one, s_1, occupies one of the leaves. There, all regular agents' stationary beliefs coincide with the belief x_0 of s_0, represented in white. In social network depicted in right-most figure, none of the stubborn agents, occupy the center. There, all regular agents' stationary beliefs coincide with the arithmetic average (represented in gray) of x_0 (white), and x_1 (black).

Example 5.1 ***(Tree)*** *Let us consider the case when $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a tree. Let the stubborn agent set $\mathcal{S}$ consist of only two elements, s_0 and s_1, with beliefs x_0, and x_1, respectively.*

If $\mathcal{S}_a = \{s_0\}$ (respectively, $\mathcal{S}_a = \{s_1\}$), then Remark 4.1 implies that $\mathbb{E}[X_a] = x_0$ (resp., $\mathbb{E}[X_a] = x_1$), and $\mathrm{Var}[X_a] = 0$. Instead, if $\mathcal{S}_a = \{s_0, s_1\}$, then one has that (25) is satisfied by

$$\mathbb{E}[X_v] = \frac{d(v, s_0)x_1 + d(v, s_1)x_0}{d(v, s_0) + d(v, s_1)},$$

where $d(v, w)$ denotes the distance in $\mathcal{G}$, between (i.e., the length of the shortest path connecting) nodes v and w. Hence, the stationary expected beliefs are linear interpolations of the beliefs of the stubborn agents. Moreover, if the confidence parameters are $\theta_e = 1$ for all e, then (26) is satisfied by

$$\mathbb{E}[X_a^2] = \frac{d(a, s_0)x_1^2 + d(a, s_1)x_0^2}{d(a, s_0) + d(a, s_1)},$$

so that the stationary variance of agent a's belief is given by

$$\mathrm{Var}[X_a] = \mathbb{E}[X_a^2] - \mathbb{E}[X_a]^2 = \frac{d(a, s_0)d(a, s_1)}{(d(a, s_0) + d(a, s_1))^2}\,(x_0 - x_1)^2 .$$

The two equations above show that the belief of each regular agent keeps on oscillating ergodically around a value which depends on the relative distance of the agent from the two stubborn agents. The amplitude of such oscillations is maximal for central nodes, i.e., those which are homogeneously distant from both stubborn agents. This can be given the intuitive explanation that, the closer a regular agent is to a stubborn agent s with respect to the other stubborn agent s', the more frequent her, possibly indirect, interactions are with agent s and the less frequent her interactions are with s', and hence the stronger the influence is from s rather than from s'. Moreover, the more equidistant a regular agent a is from s_0, and s_1, the higher the uncertainty is on whether, in the recent past, agent a has been influenced by either s_0, or s_1.

On its left-hand side, Figure 6 reports the expected stationary beliefs and their variances for a social network with population size $n = 5$, line (a special case of tree-like) topology: the two stubborn agents are positioned in the extremities, and plotted in white, and black, respectively, while regular agents are plotted in different shades of gray corresponding to their relative distance from the extremities, and hence to their expected stationary belief. In the right-hand side of Figure 6, a more complex tree-like topology is reported, again with two stubborn agents colored in white, and black respectively, and with regular agents colored by different shades of gray corresponding to their relative vicinity to the two stubborn agents. Figure 7 reports two social networks with star topology (another special case of tree). In both cases there are two stubborn agents, colored in white, and black, respectively. In the left-most picture, the white stubborn agent occupies the center, so that all the rest of the population will eventually adopt his belief, and is therefore colored in white. In the right-most picture, none of the stubborn agents occupies the center, and hence all the regular agents, hence colored in gray, are equally influenced by the two stubborn agents.

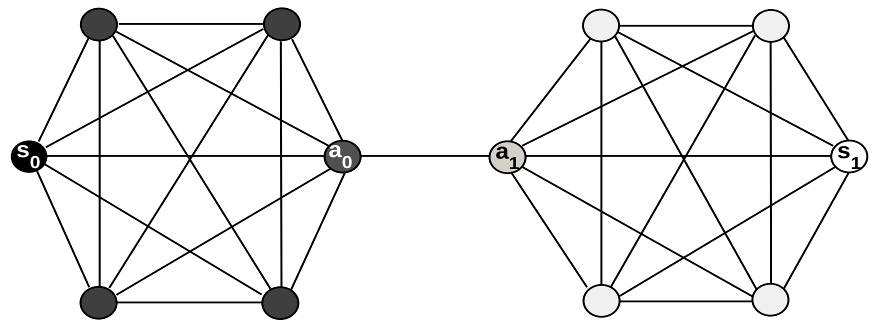

Figure 8: A social network with population size $n = 12$, a barbell-like topology, and two stubborn agents. In each of the two halves of the graph the expected average beliefs concentrate around the beliefs of the stubborn agent in the respective half.

Example 5.2 *(**Barbell**) For even $n \geq 6$, consider a barbell-like topology consisting of two complete graphs with vertex sets $\mathcal{V}_0$, and $\mathcal{V}_1$, both of size $n/2$, and an extra edge $\{a_0, a_1\}$ with $a_0 \in \mathcal{A}_0$, and $a_1 \in \mathcal{A}_1$ (see Figure 8). Let $\mathcal{S} = \{s_0, s_1\}$ with $s_0 \neq a_0 \in \mathcal{V}_0$ and $s_1 \neq a_1 \in \mathcal{V}_1$. Then, (25) is satisfied by*

$$\mathbb{E}[X_a] = \begin{cases} \frac{4}{n+8}x_{s_0} + \frac{n+4}{n+8}x_{s_1} & \text{if} \quad a = a_1 \\ \frac{n+4}{n+8}x_{s_0} + \frac{4}{n+8}x_{s_1} & \text{if} \quad a = a_0 \\ \frac{2}{n+8}x_{s_0} + \frac{n+6}{n+8}x_{s_1} & \text{if} \quad a \in \mathcal{A}_1 \setminus \{a_1\} \\ \frac{n+6}{n+8}x_{s_0} + \frac{2}{n+8}x_{s_1} & \text{if} \quad a \in \mathcal{A}_0 \setminus \{a_0\} \,. \end{cases}$$

In particular, observe that, as n grows large, $\mathbb{E}[X_a]$ converges to x_{s_0} for all $a \in \mathcal{A}_0$, and $\mathbb{E}[X_a]$ converges to x_{s_1} for all $a \in \mathcal{A}_1$. Hence, the network polarizes around the opinions of the two stubborn agents.

Example 5.3 *(**Abelian Cayley graph**) Let us denote by $\mathbb{Z}_m$ the integers modulo m. Put $\mathcal{V} = \mathbb{Z}_m^d$, and let $\Theta \subseteq \mathcal{V} \setminus \{0\}$ be a subset generating $\mathcal{V}$ and such that if $x \in \Theta$, then also $-x \in \Theta$. The Abelian Cayley graph associated with Θ is the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where $\{v, w\} \in \mathcal{E}$ iff $v - w \in \Theta$. Notice that Abelian Cayley graphs are always undirected and regular, with $d_v = |\Theta|$ for any $v \in \mathcal{V}$. Denote by $e_i \in \mathcal{V}$ the vector of all 0's but the i-th component equal to 1. If $\Theta = \{\pm e_1, \ldots, \pm e_d\}$, the corresponding $\mathcal{G}$ is the classical d-dimensional torus of size $n = m^d$. In particular, for $d = 1$, this is a cycle, while, for $d = 2$, this is the torus (see Figure 9).*

Let the stubborn agent set consist of only two elements: $\mathcal{S} := \{s_0, s_1\}$. Then the following formula holds (see [3, Ch. 2, Corollary 10]):

$$\gamma^v_{s_0} = \mathbb{P}_v(T_{s_1} < T_{s_0}) = \frac{E_{vs_0} - E_{vs_1} + E_{s_1 s_0}}{E_{s_0 s_1} + E_{s_1 s_0}} \tag{27}$$

where $E_{vw} := \mathbb{E}_v[T_w]$ denotes the expected time it takes to a random walk started at v to hit for the first time w. On the other hand, average hitting times E_{vw} can be expressed in terms of the Green function of the graph, which is defined as the unique matrix $Z \in \mathbb{R}^{\mathcal{V} \times \mathcal{V}}$ such that

$$Z\mathbb{1} = 0 \,, \quad (I - P)Z = I - n^{-1}\mathbb{1}\mathbb{1}^{\mathrm{T}} \,,$$

where $\mathbb{1}$ stands for the all-1 vector. The relation with the hitting times is given by:

$$E_{vw} = n^{-1}(Z_{ww} - Z_{vw}) \,. \tag{28}$$

Let P be the stochastic matrix corresponding to the simple random walk on $\mathcal{G}$. It is a standard fact that P is irreducible and its unique invariant probability is the uniform one. There is an orthonormal basis of eigenvectors for P good for any Θ: if $l = (l_1, \ldots l_d) \in \mathcal{V}$ define $\phi_l \in \mathbb{R}^{\mathcal{V}}$ by

$$\phi_l(k) = m^{-d/2} \exp\left(\frac{2\pi i}{m} l \cdot k\right) , \quad k = (k_1, \ldots, k_d) \in \mathcal{V} \,,$$

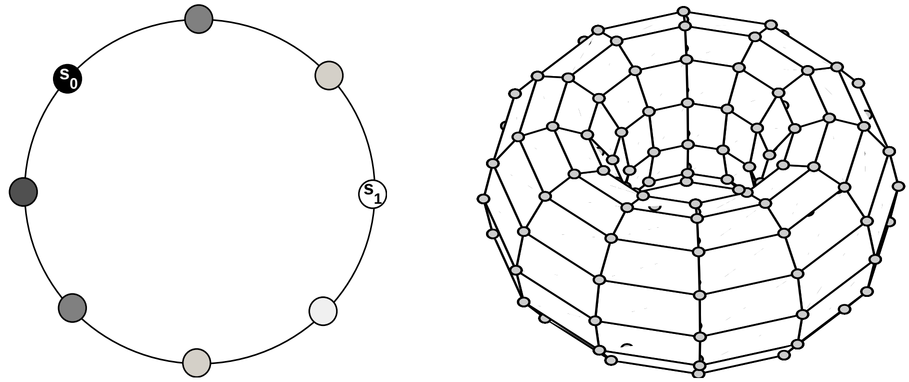

Figure 9: Two social networks with cycle and 2-dimensional toroidal topology, respectively.

(where $l \cdot k = \sum_i l_i k_i$). The corresponding eigenvalues can be expressed as follows

$$\lambda_l = \frac{1}{|\Theta^+|} \sum_{k \in \Theta^+} \cos\left(\frac{2\pi}{m} l \cdot k\right)$$

where Θ^+ is any subset of Θ such that for all $x \in \Theta$, $|\{x, -x\} \cap \Theta^+| = 1$. Hence,

$$Z_{vw} = m^{-d} \sum_{l \neq 0} \frac{\exp\left[\frac{2\pi i}{m} l \cdot (v - w)\right]}{1 - \frac{1}{|\Theta^+|} \sum_{k \in \Theta^+} \cos\left[\frac{2\pi}{m} l \cdot k\right]} \tag{29}$$

From (27), (28), and the fact that $E_{s_0 s_1} = E_{s_1 s_0}$ by symmetry, one obtains

$$\gamma^a_{s_1} = \frac{1}{2} + \frac{m^{-d} \sum_{l \neq 0} \frac{\exp\left[\frac{2\pi i}{m} l \cdot (a - s_1)\right] - \exp\left[\frac{2\pi i}{m} l \cdot (a - s_0)\right]}{1 - \frac{1}{|\Theta^+|} \sum_{k \in \Theta^+} \cos\left[\frac{2\pi}{m} l \cdot k\right]}}{2m^{-d} \sum_{l \neq 0} \frac{1 - \cos\left[\frac{2\pi}{m} l \cdot (s_0 - s_1)\right]}{1 - \frac{1}{|\Theta^+|} \sum_{k \in \Theta^+} \cos\left[\frac{2\pi}{m} l \cdot k\right]}} . \tag{30}$$

The expected stationary beliefs can then be computed using Theorem 4.1 and (30).

6. Homogeneous influence in highly fluid social networks In this section, we present estimates for the stationary expected beliefs and belief variances as a function of the underlying social network. We will confine ourselves to considering social networks satisfying an additional assumption of reversibility, formulated below. Let us recall the definition of the matrix $P \in \mathbb{R}^{\mathcal{A} \times \mathcal{V}}$ associated to a social network.

Assumption 6.1 *The restriction of the matrix P defined by (3) to $\mathcal{A} \times \mathcal{A}$ is irreducible and reversible, i.e., there exist positive values $\tilde{\pi}_a$, for $a \in \mathcal{A}$, such that*

$$\tilde{\pi}_a P_{aa'} = \tilde{\pi}_{a'} P_{a'a} , \qquad \forall a, a' \in \mathcal{A} . \tag{31}$$

Condition (31) entails a certain form of reciprocity between the intensity of influence of a regular agent on another, and vice versa. In particular, it implies that if $(a, a') \in \overrightarrow{\mathcal{E}}$, then also $(a', a) \in \overrightarrow{\mathcal{E}}$, so that the restriction of the di-graph $\overrightarrow{\mathcal{G}}$ to $\mathcal{A}$ is in fact undirected. We will refer to social networks whose associated matrix P satisfies Assumption 6.1 as *reversible (social) networks*.

Observe that, whenever Assumption 6.1 is satisfied, the vector $\{\tilde{\pi}_a : a \in \mathcal{A}\}$ is uniquely defined up to a multiplicative constant: to see this, fix the value $\tilde{\pi}_a$ on some a, then (31) fixes the entries $\tilde{\pi}_{a'}$ for a' in the out-neighborhood of a, and the irreducibility assumption allows one to iterate the argument until covering the whole $\mathcal{A}$. Now, given a reversible social network, it is possible to extend the matrix P on $\mathcal{S} \times \mathcal{V}$ as follows: Put $P_{ss'} := 0$ for all $s, s' \in \mathcal{S}$, and

$$\tilde{\pi}_s := \sum_{a \in \mathcal{A}} \tilde{\pi}_a P_{as} , \qquad P_{sa} := P_{as} \frac{\tilde{\pi}_a}{\tilde{\pi}_s} , \qquad s \in \mathcal{S}, a \in \mathcal{A} . \tag{32}$$

In particular, P_{sa} does not depend on the particular choice of $\tilde{\pi}$, and, extended in this way, P becomes an irreducible and reversible stochastic matrix on $\mathcal{V}$. Furthermore, the probability measure π on $\mathcal{V}$, defined by

$$\pi_v := \left(\sum\nolimits_{v'} \tilde{\pi}_{v'}\right)^{-1} \tilde{\pi}_v \,, \tag{33}$$

is its unique invariant distribution. The measure of the stubborn agents' set under such a distribution is given by

$$\pi(\mathcal{S}) = \sum\nolimits_s \pi_s = \left(\sum\nolimits_v \tilde{\pi}_v\right)^{-1} \sum\nolimits_{a,s} \tilde{\pi}_a P_{as} \tag{34}$$

Observe that (31) and (32) imply that $P_{vw} > 0$ if and only if $P_{vw} > 0$. It is then natural to associate to any social network satisfying Assumption 6.1 an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, in which $\{v, v'\} \in \mathcal{E}$ if and only of $P_{vv'} > 0$. From now on, we refer to this undirected graph $\mathcal{G}$, rather than to the directed graph $\overrightarrow{\mathcal{G}}$ considered so far. Clearly, a given undirected graph $\mathcal{G}$ may be associated to many reversible social networks.

Remark 6.1 *Given a reversible social network, the one proposed above is not its only possible extension of its matrix P that makes it irreducible and reversible, as one may allow, e.g., for non-zero valued $P_{ss'}$. However, our subsequent analysis is valid for all such extensions, while tightness of the estimates obtained may vary with the choice of such an extension.*

Example 6.1 *Let us consider the canonical construction of a social network from a given undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, explained in Example 2.1. Extend P by putting $P_{sv} = 1/d_s$, for all $s \in \mathcal{S}$, and $v \in \mathcal{V}$ such that $\{s, v\} \in \mathcal{E}$, and $P_{sv} = 0$, for all $s \in \mathcal{S}$, and $v \in \mathcal{V}$ such that $\{s, v\} \notin \mathcal{E}$. Then, Assumption 6.1 can be checked to hold, with the invariant measure given by*

$$\pi_v = d_v/(n\overline{d}) \,,$$

where d_v is the degree of node v in $\mathcal{G}$ and $\overline{d} := n^{-1}\sum_v d_v$ is the average degree of $\mathcal{G}$. Observe that, in this construction, $\pi(\mathcal{S}) = (\sum_v d_v)^{-1} \sum_s d_s$ is the fraction of edges incident to the stubborn agents, or, in other words, the relative size of the boundary of $\mathcal{S}$ in $\mathcal{G}$.

We introduce the following definition.

Definition 6.1 *Given a reversible social network, let P be defined as in (3), and extended as in (32), and π denote its stationary distribution defined as in (33). Define $\pi_* := \min_v \pi_v$, and let*

$$\tau := \inf\left\{ t \geq 0 : \sum_w |\mathbb{P}_v(V(t) = w) - \mathbb{P}_{v'}(V(t) = w)| \leq \frac{2}{e} \,, \forall v, v' \in \mathcal{V} \right\}. \tag{35}$$

denote the (variational distance) mixing time of the continuous-time Markov chain $V(t)$ with state space $\mathcal{V}$, and generator matrix $P - I$. The fluidity *of the social network is the ratio*

$$\Phi := \frac{n\pi_*}{\pi(\mathcal{S})\tau} \,. \tag{36}$$

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 marginal distributions of the stationary beliefs of most of the regular agents in the population can be approximated (at least in their first and second moments) by the distribution of a *weighted-mean belief* Z, supported in the finite set $\mathcal{X} := \{x_s : s \in \mathcal{S}\}$, and given by

$$\mathbb{P}(Z = x) = \sum_{s:\, x_s = x} \overline{\gamma}_s \,, \qquad x \in \mathcal{X} \,, \qquad \overline{\gamma}_s := \sum\nolimits_v \pi_v \gamma_s^v \,, \qquad s \in \mathcal{S} \,. \tag{37}$$

We refer to the probability distribution $\{\overline{\gamma}_s : s \in \mathcal{S}\}$ as the *stationary stubborn agent distribution*. Observe that $\overline{\gamma}_s = \mathbb{P}_\pi(V(T_\mathcal{S}) = s)$ 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 will clarify below, 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 $s' \in \mathcal{S}$.

More precisely, let us denote the expected value and variance of the weighted-mean belief Z by

$$\mathbb{E}[Z] := \sum\nolimits_s \overline{\gamma}_s x_s \,, \qquad \sigma_Z^2 := \sum\nolimits_s \overline{\gamma}_s \left(x_s - \mathbb{E}[Z]\right)^2 \,. \tag{38}$$

Let σ_v^2 denote the variance of the stationary 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.,

$$\Delta_* := \max\left\{x_s - x_{s'} \,:\, s, s' \in \mathcal{S}\right\}. \tag{39}$$

The next theorem presents the main result of this section.

Theorem 6.1 *Let Assumptions 2.1 and 6.1 hold, and assume that $\pi(\mathcal{S}) \le 1/4$. Then, for all $\varepsilon > 0$,*

$$\frac{1}{n}\left|\left\{v : \left|\mathbb{E}[X_v] - \mathbb{E}[Z]\right| > \Delta_* \varepsilon\right\}\right| \le \frac{\psi(\varepsilon)}{\Phi} \,, \tag{40}$$

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

$$\frac{1}{n}\left|\left\{v : \left|\sigma_v^2 - \sigma_Z^2\right| > \Delta_*^2 \varepsilon\right\}\right| \le \frac{\psi(\varepsilon)}{\Phi} \,. \tag{41}$$

This theorem implies that in large-scale highly fluid social networks, as the population size n grows large, the stationary expected beliefs and variances of the regular agents concentrate around fixed values corresponding to the expected weighted-mean belief $\mathbb{E}[Z]$, and, respectively, its variance σ_Z^2 (see Fig.s 10 and 11). We refer to this phenomenon as *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' stationary 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 then 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. Theorem 6.1 is proved in Section 6.2. Its proof relies on the characterization of the expected stationary beliefs and variances 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 $\mathcal{S}$, $V(t)$ looses memory of where it started from, and approaches $\mathcal{S}$ almost as if started from the stationary distribution π.

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.

Before proving Theorem 6.1, we present some examples of highly fluid social networks in Section 6.1.

6.1 Examples of large-scale social networks We now present some examples of family of social networks that are highly fluid in the limit of large population size n. All the examples will follow the canonical social network construction of Example 2.1, starting from an undirected graph $\mathcal{G}$. Before proceeding, let us recall that the invariant measure of the stubborn agents set $\pi(\mathcal{S})$ is given by

$$\pi(\mathcal{S}) = \sum\nolimits_s d_s/(n\bar{d}) \,, \tag{42}$$

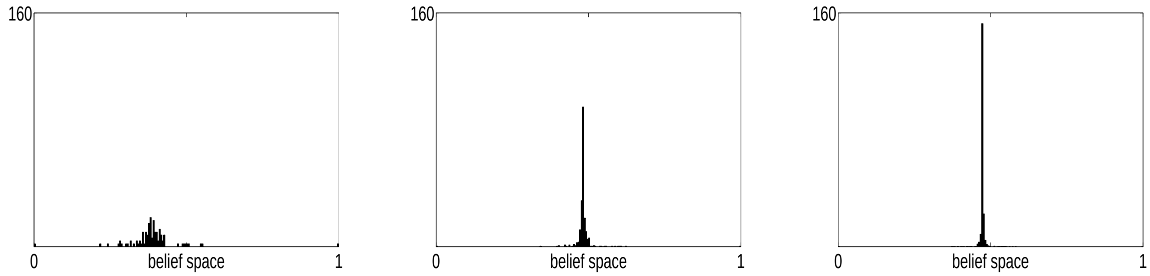

(a) (b) (c)

Figure 10: Homogeneous influence in Erdös-Renyi graphs of increasing sizes. The population sizes are $n = 100$ in (a), $n = 500$ in (b), and $n = 2000$ in (c), while $p = 2n^{-1}\log n$. In each case, there are two stubborn agents randomly selected from the node with uniform probability and holding beliefs equal to 0, and 1, respectively. The figures show the empirical density of the expected stationary beliefs, for typical realizations of such graphs. As predicted by Theorem 6.1, such empirical density tends to concentrate around a single value as the population size grows large.

and 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 $(\pi_* n)^{-1} \leq \overline{d}$, where $\overline{d}$ is the average degree of the graph.

We start with an example of a social network which is not highly fluid.

Example 6.2 ***(Barbell)*** *For even $n \geq 6$, consider the barbell-like topology introduced in Example 5.2. The mixing time of this network can be estimated in terms of the conductance ϕ_* of the graph, which is defined as the minimum over all subsets $\mathcal{V}' \subseteq \mathcal{V}$ with $0 < \sum_{v\in\mathcal{V}'} d_v \leq n\overline{d}/2$, of the ratio between the number of edges connecting $\mathcal{V}'$ with its complement, and the sum of the degrees of the nodes in $\mathcal{V}$. It is not hard to see that such a minimum is achieved by $\mathcal{V}' = \mathcal{V}_0$, so that $\phi_* = \left(\frac{n}{2}(\frac{n}{2}-1)+1\right)^{-1} \leq \frac{4}{(n+1)^2}$. Using [28, Theorem 7.3], it then follows that $\tau \geq (4\phi_*)^{-1} \geq (n+1)^2/16$. Since $d_v \geq n/2 - 1$ for all v, it follows that the barbell-like network is never highly fluid provided that $|\mathcal{S}| \geq 1$. In fact, we have already seen in Example 5.2 that the expected stationary beliefs polarize in this case.*

Let us now consider a standard deterministic family of symmetric graphs.

Example 6.3 *(d**-dimensional tori**) Let us consider the case of a d-dimensional torus of size $n = m^d$, introduced in Example 5.3. Since this is a regular graph, one has $\pi_* n = 1$. Moreover, it was proved by Cox [15] that, as n grows large, $\tau \sim C_d n^{2/d}$, for some constant C_d depending on the dimension d only. Then, $\tau\pi(\mathcal{S}) \sim |\mathcal{S}|n^{2/d-1}$. Hence, if $|\mathcal{S}| = o(n^{2/d-1})$, then the social network with toroidal topology is highly fluid. In fact, using Fourier analysis, one may show that $|\mathcal{S}| = o(n^{1/d-1})$ suffices. In contrast, for the one-dimensional torus (i.e., a ring) of size n, both $\mathbb{E}_\pi[T_\mathcal{S}] \sim n^2$ and $\tau C_2 \sim n^2$; in fact, using the explicit calculations of Example 5.1, one can see that the expected stationary beliefs do not concentrate in this case. Finally, the two-dimensional torus is not highly fluid, hence Theorem 6.1 is not sufficient to prove that the empirical beliefs concentrate around $\mathbb{E}[Z]$. Nevertheless, one could use the explicit expression (30) and Fourier analysis in order to show that the condition $|\mathcal{S}| = o(n^{1/2})$ would suffice for that.*

An intuition for this behavior can be obtained by thinking of a limit continuous model. First recall that the expected stationary beliefs vector solves the Laplace equation on $\mathcal{G}$ with boundary conditions assigned on the stubborn agent set $\mathcal{S}$. Now, consider the Laplace equation on a d-dimensional manifold with boundary conditions on a certain subset. Then, in order for the problem to be well-posed, one needs that the such a subset has dimension $d-1$. Similarly, one needs $|\mathcal{S}| = \Theta(n^{(d-1)/d}) = \Theta(m^{d-1})$ in order to guarantee that the expected stationary beliefs vector is not almost constant in the limit of large n.

We now present 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.

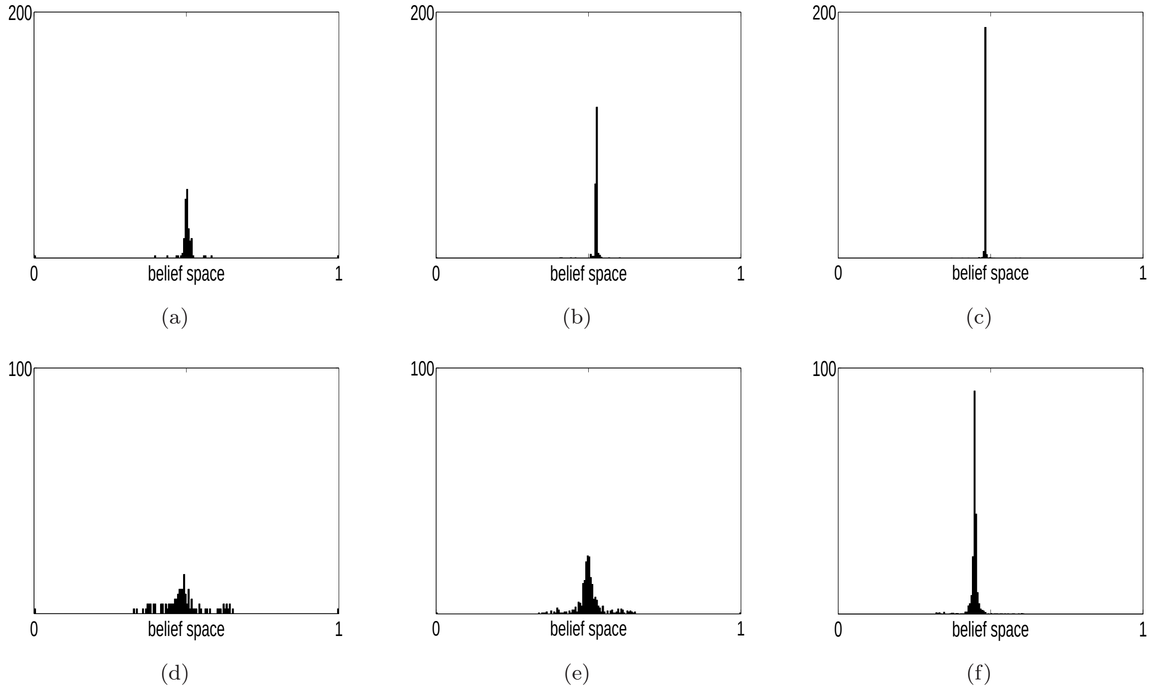

Figure 11: Homogeneous influence in preferential attachment networks of increasing sizes. The figures show the empirical density of the expected stationary beliefs, for typical realizations of such graphs. The population sizes are $n = 100$ in (a) and (d), $n = 500$ in (b) and (e), and $n = 2000$ in (c) and (f), while $m = 4$. In each case, there are two stubborn agents holding beliefs equal to 0, and 1, respectively. In (a), (b), and (c), the stubborn agents are chosen to coincide with the two latest attached nodes, and therefore tend to have the lowest degree. In contrast, in (d), (e), and (f), the stubborn agents are chosen to be the two initial nodes, and therefore tend to have the highest degrees. As predicted by Theorem 6.1, the empirical density of the expected stationary beliefs tends to concentrate around a single value as the population size grows large. The rate at which this concentration occurs is faster in the top three figures, where $\sum_s d_s$ is smaller, and slower in the bottom three figures, where $\sum_s d_s$ is larger.

Example 6.4 ***(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 [19, Thm. 2.8.2]. In this regime, results by Cooper and Frieze [14] ensure that, with high probability, $\tau = O(\log n)$, and that there exists a positive constant δ such that $\delta c \log n \le d_v \le 4c \log n$ for each node v [19, Lemma 6.5.2]. In particular, it follows that, with high probability, $(\pi_* n)^{-1} \le 4/\delta$. Hence, using (42), one finds that the resulting social network is highly fluid, provided that $|\mathcal{S}| = o(n/\log n)$, as n grows large. Figure 10 shows the empirical density of the expected stationary beliefs for typical realizations of Erdös-Renyi graphs of increasing size $n = 100, 500, 2000$, and constant stubborn agents number $|\mathcal{S}| = 2$.*

Example 6.5 ***(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 $k \ge 2$, and that the first two moments $\overline{d} := \sum_k \lambda_k k$, 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 [19, Ch. 6.3] show that $\tau = O(\log n)$. Therefore, using (42), one finds that the resulting social network is highly fluid with high probability provided that $\sum_s d_s = o\big(n/\log n\big)$.*

Example 6.6 ***(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 dis-*

tribution. We follow [19, 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 [19, 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 [19, Th. 4.6.1]. In particular, it follows that, with high probability, the average degree $\overline{d}$ remains bounded, while the second moment of the degree distribution diverges an n grows large. On the other hand, results by Mihail et al. [32] (see also [19, Th. 6.4.2]) imply that the mixing time $\tau = O(\log n)$. Therefore, thanks to (42), the resulting social network is highly fluid with high probability if $\sum_{s\in\mathcal{S}} d_s = o\big(n/\log n\big)$.

Example 6.7 ***(Watts & Strogatz's small world)*** *Watts and Strogatz [47], and then Newman and Watts [37] 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 [19, Th. 6.6.1] show that the mixing time $\tau = O(\log^3 n)$. This, and (42) imply that the network is highly fluid with high probability provided that $\sum_{s\in\mathcal{S}} d_s = o\big(n/\log^3 n\big)$.*

6.2 Proof of Theorem 6.1 Recall that Theorem t4.1 allows one to express the stationary expected beliefs and belief variances can in terms of the hitting probability distributions γ^v on the stubborn agent set $\mathcal{S}$ of a continuous-time Markov chain with state space $\mathcal{V}$, and generator matrix Q. Observe that such hitting probabilities only depend on the restriction of the matrix P to $\mathcal{A}\times\mathcal{V}$, while they are independent of the holding rates of the chain $V(t)$, as well as of the specific extension of P to $\mathcal{S}\times\mathcal{V}$. As a consequence, without any loss of generality, one can assume —as we will— that the extension of P to $\mathcal{S}\times\mathcal{V}$ is given as in (32), and that the holding times of $V(t)$ in every state have exponential distribution of rate 1 (or, equivalently, that the generator matrix of $V(t)$ is $P - I$).

In order to prove Theorem 6.1, we will obtain estimates on the hitting probabilities of the random walk $V(t)$ on $\mathcal{V}$ with generator matrix P. We start by stating a standard result on the distance of transition probability distribution of a random walk from its stationary distribution.

Proposition 6.1 *[3, Ch. 4,Lemma 5] Let $V(t)$ be a continuous-time Markov chain on $\mathcal{V}$ with irreducible jump matrix P, and holding times with rate-1 exponential distribution. For $v \in \mathcal{V}$, and $t \geq 0$, let $\gamma^v(t)$ be the probability distribution of $V(t)$, conditioned on $V(0) = v$. Then,*

$$\max_{v,v'} \left|\left|\gamma^v(t) - \gamma^{v'}(t)\right|\right|_{TV} \leq \exp(1 - t/\tau)\,, \qquad \forall t \geq 0\,,$$

where τ is the mixing time [cf. Eq. (35)].

The following result, whose proof is an application of Proposition 6.1, provides a useful estimate on the total variation distance between the hitting probability distribution γ^v over $\mathcal{S}$ and the stationary stubborn agent distribution $\overline{\gamma}$.

Lemma 6.1 *Let $V(t)$ be a continuous-time Markov chain on $\mathcal{V}$ with irreducible jump matrix P, and holding times with rate-1 exponential distribution. Let γ^v be the hitting distribution on $\mathcal{S} \subseteq \mathcal{V}$ for the chain $V(t)$ started in $v \in \mathcal{V}$, as defined in (20), and $\overline{\gamma} := \sum_v \pi_v \gamma^v$, where π is the unique invariant distribution of P. Then, for all $t \geq 0$, and $v \in \mathcal{V}$,*

$$||\gamma^v - \overline{\gamma}||_{TV} \leq p_v(t) + \exp(-t/\tau + 1)\,, \tag{43}$$

where $p_v(t) := \mathbb{P}_v\left(T_{\mathcal{S}} \leq t\right)$.

Proof. Notice that (43) is trivial when $v \in \mathcal{S}$, for in that case $p_v(t) = 1$. For $a \in \mathcal{A}$, one reason as follows. Let $\tilde{T}_{\mathcal{S}} := \inf\{t' \geq t : V(t) \in \mathcal{S}\}$. Using the identity

$$||\mu - \nu||_{TV} = \frac{1}{2} \sup_{f\in[-1,1]^{\mathcal{V}}} \left\{\sum\nolimits_v (\mu_v - \nu_v) f_v\right\}$$

(see, e.g., [28, Prop. 4.5]), and observing that the event $\{T_{\mathcal{S}} > t\}$ implies $V_{T_{\mathcal{S}}} = V_{\tilde{T}_{\mathcal{S}}}$, one gets that

$$\begin{aligned}\left|\left|\gamma^a - \sum\nolimits_v q_v^a(t)\gamma^v\right|\right|_{TV} &= \frac{1}{2}\sup_{f\in[-1,1]^{\mathcal{V}}}\left\{\mathbb{E}_a\left[f_{V_{T_{\mathcal{S}}}} - f_{V_{\tilde{T}_{\mathcal{S}}}}\right]\right\} \\ &= \frac{1}{2}\sup_{f\in[-1,1]^{\mathcal{V}}}\left\{\mathbb{E}_a\left[1_{\{T_{\mathcal{S}}\le t\}}\left(f_{V_{T_{\mathcal{S}}}} - f_{V_{\tilde{T}_{\mathcal{S}}}}\right)\right]\right\} \quad \le \quad p_a(t)\,.\end{aligned}$$

On the other hand, since every Markov kernel is contractive in total variation distance, one has that

$$||\sum_v \gamma_v^a(t)\gamma^v - \overline{\gamma}||_{TV} = ||\sum_v(\gamma_v^a(t) - \pi_v)\gamma^v - ||_{TV} \le ||\gamma^a(t) - \pi||_{TV}\,.$$

By applying the triangle inequality, the two estimates above, and Proposition 6.1, one shows that

$$\begin{aligned}||\gamma^a - \overline{\gamma}||_{TV} &\le ||\gamma^a - \sum\nolimits_v \gamma_v^a(t)\gamma^v||_{TV} + ||\sum\nolimits_v \gamma_v^a(t)\gamma^v - \overline{\gamma}||_{TV} \\ &\le p_a(t) + ||\gamma^a(t) - \pi||_{TV} \\ &\le p_a(t) + \exp(-t/\tau + 1)\,,\end{aligned}$$

thus proving the claim. ■

Lemma 6.2, stated below, is the main technical result of this section. Its proof relies on the "approximate exponentiality" of the hitting time $T_{\mathcal{S}}$. This is the property that the probability law of $T_{\mathcal{S}}$ is close to the exponential distribution with expectation $\mathbb{E}_\pi[T_{\mathcal{S}}]$ when the initial distribution is the stationary one, and the mixing time τ is small with respect to the expected hitting time $\mathbb{E}_\pi[T_{\mathcal{S}}]$. In particular, we make use of the following result, due to Aldous and Brown:

Proposition 6.2 *([3, Ch. 3,Prop. 23]) Let $V(t)$ be a continuous-time Markov chain on $\mathcal{V}$ with holding times exponentially distributed with rate 1, irreducible jump matrix P, and stationary distribution π. Let τ_2 be its relaxation time, i.e., the inverse of the spectral gap of P. Then, for all $\mathcal{S}\subset\mathcal{V}$,*

$$\sup_{t\ge 0}|\mathbb{P}_\pi(T_{\mathcal{S}} > t) - \exp(-t/\mathbb{E}_\pi[T_{\mathcal{S}}])| \le \tau_2/\mathbb{E}_\pi[T_{\mathcal{S}}]\,.$$

Lemma 6.2 *Let Assumptions 2.1 and 6.1 hold. Then, for all $\varepsilon > 0$,*

$$\frac{1}{n}\left|\left\{v\in\mathcal{V} : ||\gamma^v - \overline{\gamma}||_{TV} \ge \varepsilon\right\}\right| \le \frac{2\log(2e^2/\varepsilon)}{\varepsilon}\frac{\tau}{n\pi_*\mathbb{E}_\pi[T_{\mathcal{S}}]}\,. \tag{44}$$

Proof. Let $p_v(t) := \mathbb{P}_v(T_{\mathcal{S}} \le t)$. From Lemma 6.1, and Proposition 6.2, it follows that, for all $t\ge 0$,

$$\sum_v \pi_v p_v(t) = \mathbb{P}_\pi\left(T_{\mathcal{S}}\le t\right) \le 1 - \exp\left(-\frac{t}{\mathbb{E}_\pi[T_{\mathcal{S}}]}\right) + \frac{\tau_2}{\mathbb{E}_\pi[T_{\mathcal{S}}]} \le \frac{t+\tau_2}{\mathbb{E}_\pi[T_{\mathcal{S}}]}\,,$$

where the last step follows from the inequality $1 - e^{-x} \le x$. Hence, Markov's inequality implies that

$$\frac{1}{n}|\{v : p_v(t)\ge\varepsilon/2\}| \le \frac{2}{\varepsilon}\sum_v\frac{1}{n}p_v(t) \le \frac{2}{\varepsilon n\pi_*}\sum_v \pi_v p_v(t) \le \frac{2}{\varepsilon n\pi_*}\frac{t+\tau_2}{\mathbb{E}_\pi[T_{\mathcal{S}}]}\,. \tag{45}$$

Now, choose $t = \tau\log(2e/\varepsilon)$, so that $\exp(-t/\tau+1) = \varepsilon/2$, and $t+\tau_2 \le \tau\log(2e/\varepsilon) + \tau = \tau(\log(2e^2/\varepsilon))$, thanks to the inequality $\tau_2\le\tau$ (see, e.g., [3, Ch. 4, p. 2]). Then, applying (43) and (45) yields

$$\frac{1}{n}|\{v : ||\gamma^v - \overline{\gamma}||_{TV}\ge\varepsilon\}| \le \frac{1}{n}\left|\left\{v : ||\gamma^v(t) - \pi||_{TV} \ge \frac{\varepsilon}{2}\right\}\right| \le \frac{2\log(2e^2/\varepsilon)\tau}{\varepsilon n\pi_*\mathbb{E}_\pi[T_{\mathcal{S}}]}\,,$$

which proves the claim. ■

Lemma 6.2 is particularly relevant when τ is much smaller than $\mathbb{E}_\pi[T_{\mathcal{S}}]$. Indeed, in this case, it shows that, for all but a negligible fraction of initial states $v\in\mathcal{V}$, the hitting probability distribution γ^v will be close to the stationary stubborn agent distribution $\overline{\gamma}$. The intuition behind this result is rather simple: if the chain $V(t)$ mixes much before hitting the stubborn agents set $\mathcal{S}$, then it will hit some s before any other $s'\in\mathcal{S}$ with probability close to $\overline{\gamma}_s$, independently of the initial state. While the expected hitting time $\mathbb{E}_\pi[T_{\mathcal{S}}]$ may be computable in certain cases, it is often easier to estimate it in terms of the invariant measure of the stubborn agent set, $\pi(\mathcal{S})$, e.g., using the following result:

Proposition 6.3 *[4, Proposition 7.13] For all $\mathcal{S} \subseteq \mathcal{V}$, $\mathbb{E}_\pi[T_\mathcal{S}] \geq 1/(2\pi(\mathcal{S})) - 3/2$.*

Lemma 6.2 and Proposition 6.3 immediately imply the following result:

Lemma 6.3 *Let Assumptions 2.1 and 6.1 hold, and assume that $\pi(\mathcal{S}) \leq 1/4$. Then, for all $\varepsilon > 0$,*

$$\frac{1}{n}\left|\{v : ||\gamma^v - \overline{\gamma}||_{TV} \geq \varepsilon\}\right| \leq \frac{\psi(\varepsilon)}{\Phi},$$

where $\psi(\varepsilon) := 16\varepsilon^{-1}\log(2e^2/\varepsilon)$, and Φ is the fluidity of the network, as defined by (36).

Proof. The assumption $\pi(\mathcal{S}) \leq 1/4$ implies that $3/(8\pi(\mathcal{S})) \geq 3/2$, so that, by Proposition 6.3, $\mathbb{E}_\pi[T_\mathcal{S}] \geq 1/(2\pi(\mathcal{S})) - 3/2 \geq 1/(8\pi(\mathcal{S}))$. The claim now follows by applying this bound to the righthand side of (44). ■

Proof of Theorem 6.1:

Let $y_s := x_s + \Delta_*/2 - \max\{x_{s'} : s' \in \mathcal{S}\}$ for all $s \in \mathcal{S}$. Clearly $|y_s| \leq \Delta_*/2$. Then, it follows from Theorem 4.1 that

$$\Big|\mathbb{E}[X_v] - \mathbb{E}[Z]\Big| = \Big|\sum_s \gamma^v_s x_s - \sum_s \overline{\gamma}_s x_s\Big| = \Big|\sum_s \gamma^v_s y_s - \sum_s \overline{\gamma}_s y_s\Big| \leq \Delta_* ||\gamma^v - \overline{\gamma}||_{TV},$$

so that (40) immediately follows from Lemma 6.3.

On the other hand, in order to show (41), first recall that, if $\theta_e = 1$ for all $e \in \overrightarrow{\mathcal{E}}$, then Eq. (15) provides the transition rates of coalescing random walks. In particular, if $V(0) = V'(0)$, then $V(T_\mathcal{S}) = V'(T'_\mathcal{S})$, so that $\eta^{vv}_{ss'} = \gamma^v_s$ if $s = s'$, and $\eta^{vv}_{ss'} = 0$ otherwise. Then, it follows from Theorem 4.1 that

$$\begin{aligned}
\sigma^2_v &= \mathbb{E}[X^2_v] - \mathbb{E}[X_v]^2 \\
&= \textstyle\sum_{s,s'} \eta^{vv}_{ss'} x_s x_{s'} - \left(\sum_s \gamma^v_s x_s\right)^2 \\
&= \textstyle\sum_s \gamma^v_s x^2_s - \left(\sum_s \gamma^v_s x_s\right)^2 \\
&= \textstyle\frac{1}{2}\sum_s \sum_{s'} \gamma^v_s \gamma^v_{s'} (x_s - x_{s'})^2 .
\end{aligned}$$

Similarly, $\sigma^2_Z = \frac{1}{2}\sum_{s,s'} \overline{\gamma}_s \overline{\gamma}_{s'} (x_s - x_{s'})^2$, so that

$$\begin{aligned}
|\sigma^2_v - \sigma^2_Z| &\leq \textstyle\frac{1}{2}\sum_{s,s'} |\gamma^v_s\gamma^v_{s'} - \overline{\gamma}_s\overline{\gamma}_{s'}|\,(x_s - x_{s'})^2 \\
&\leq \textstyle\frac{1}{2}\sum_{s,s'} |\gamma^v_s\gamma^v_{s'} - \overline{\gamma}_s\overline{\gamma}_{s'}|\,\Delta^2_* \\
&\leq \textstyle\frac{1}{2}\sum_{s,s'} \left(\gamma^v_s\,|\gamma^v_{s'} - \overline{\gamma}_{s'}| + \overline{\gamma}^v_{s'}\,|\gamma^v_s - \overline{\gamma}_s|\right)\Delta^2_* \\
&= ||\gamma^v - \overline{\gamma}||_{TV}\,\Delta^2_* ,
\end{aligned}$$

and (41) follows again from a direct application of Lemma 6.3.

7. 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 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, while the correlation between the stationary 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 stationary expected beliefs and variances are almost constant, so that the stubborn agents have homogeneous influence on the society.

Acknowledgments. This research was partially supported by the NSF grant SES-0729361, the AFOSR grant FA9550-09-1-0420, the ARO grant 911NF-09-1-0556, the Draper UR&D program, and the AFOSR MURI R6756-G2.

References

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

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

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

[4] D.J. Aldous, *Random walks on finite groups and rapidly mixing Markov chains*, Séminaire de probabilités (Strasbourg) **17** (1983), 243–297.

[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] P. Clifford and A. Sudbury, *A model for spatial conflict*, Biometrika **60** (1973), 581–588.

[12] G. Como and F. Fagnani, *Scaling limits for continuous opinion dynamics systems*, The Annals of Applied Probability (2010), to appear.

[13] P. Connely and D. Welsh, *Finite particle systems and infection models*, Mathematical Proceedings of the Cambridge Philosophical Society **94** (1982), 167–182.

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

[15] J.T. Cox, *Coalescing random walks and voter model consensus times on the torus in* $\mathbb{Z}^d$, The Annals of Probability **17** (1989), no. 4, 1333–1366.

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

[17] 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.

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

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

[20] P. Erdös, *On a family of symmetric bernoulli convolutions*, American Journal of Mathematics **61** (1939), 974–975.

[21] ______, *On the smoothness properties of bernoulli convolutions*, American Journal of Mathematics **62** (1940), 180–186.

[22] 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.

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

[24] 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.

[25] R.A. Holley and T.M. Liggett, *Ergodic theorems for weakly interacting infinite systems and the voter model*, The Annals of Probability **3** (1975), no. 4, 643–663.

[26] 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.

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

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

[29] T.M. Liggett, *Interacting particle systems*, Springer-Verlag, 1985.

[30] ______, *Stochastic interacting systems: Contact, voter, and exclusion processes*, Springer Berlin, 1999.

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

[32] 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.

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

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

[35] 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.

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

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

[38] J.R. Norris, *Markov chains*, Cambridge University Press, 1997.

[39] 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.

[40] 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.

[41] Y. Peres and B. Solomyak, *Absolute continuity of bernoulli convolutions, a simple proof*, Mathematics Research Letters **3** (1996), 231–239.

[42] M.M. Rao and R.J. Swift, *Probability theory with applications*, Springer, 2006.

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

[44] D.W. Strook, *Probability theory: An analytic view*, 2nd ed., Cambridge University Press, 2011.

[45] 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.

[46] 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.

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