Submitted to the Annals of Applied Probability

OPINION FLUCTUATIONS AND DISAGREEMENT IN SOCIAL NETWORKS

By Daron Acemoglu, Giacomo Como, Fabio Fagnani and Asuman Ozdaglar,

Massachusetts Institute of Technology and Politecnico di Torino

We study a stochastic gossip model of continuous opinion dynamics in a society consisting of two types of agents: regular agents, who update their beliefs according to information that they receive from their social neighbors; and stubborn agents, who never update their opinions and might represent leaders, political parties or media sources attempting to influence the beliefs in the rest of the society. When the society contains stubborn agents with different opinions, opinion dynamics never lead to a consensus (among the regular agents). Instead, beliefs in the society almost surely fail to converge, and the belief of each regular agent converges in law to a non-degenerate random variable. The model thus generates long-run disagreement and continuous opinion fluctuations. The structure of the social network and the location of stubborn agents within it shape opinion dynamics. When the society is "highly fluid", meaning that the mixing time of the random walk on the graph describing the social network is small relative to (the inverse of) the relative size of the linkages to stubborn agents, the ergodic beliefs of most of the agents concentrate around a certain common value. We also show that under additional conditions, the ergodic beliefs distribution becomes "approximately chaotic", meaning that the variance of the aggregate belief of the society vanishes in the large population limit while individual opinions still fluctuate significantly.

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

*This research was partially supported by the NSF grant SES-0729361, the AFOSR grant FA9550-09-1-0420, the ARO grant W911NF-09-1-0556, the Draper UR&D program, and the AFOSR MURI R6756-G2.

AMS 2000 subject classifications: 60K35, 91D30, 93A15

Keywords and phrases: multi-agent systems, social networks, opinion dynamics, persistent disagreement, opinion fluctuations, social influence

non-Bayesian updating mechanisms, typically lead to consensus provided that communication takes place over a strongly connected network (e.g., Smith and Sorensen [42], 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.[1]

In this paper, we investigate a possible source of persistent disagreement in social networks. We propose a tractable model that generates both long-run disagreement and opinion fluctuations so that a consensus fails to emerge even as individuals communicate and sometimes change their opinions.

We consider a stochastic gossip model of communication combined with the assumption that there are some "stubborn" agents in the network who never change their opinions. We show that the presence of these stubborn agents leads to persistent opinion fluctuations and disagreement among the rest of the society.

More specifically, we consider a society envisaged as a social network of n interacting agents (or individuals), communicating and exchanging information. Each agent a starts with an opinion (or belief) $X_a(0) \in \mathbb{R}$ and is then "activated" according to a Poisson process in continuous time. Following this event, she meets one of the individuals in her *social neighborhood* according to a pre-specified stochastic process. This process represents an underlying *social network*. We distinguish between two types of individuals, *stubborn* and *regular*. Stubborn agents, which are typically few in number, never change their opinions (they might thus correspond to media sources or political leaders wishing to influence the rest of the society). In contrast, regular agents, which make up the great majority of the agents in the social network, update their beliefs to some weighted average of their pre-meeting belief and the belief of the agent they met. The opinions generated through this information exchange process form a Markov process over the graph induced by the social network. Much of our analysis characterizes the long-run behavior of this Markov process.

[1] 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.

We show that, under general conditions, these opinion dynamics never lead to a consensus (among the regular agents). In fact, regular agents' beliefs almost surely fail to converge, and keep on oscillating. Instead, the belief of each regular agent converges in law to a non-degenerate random variable and thus has a limiting ergodic distribution (and similarly, the vector of beliefs of all regular agents jointly converge to a non-degenerate random vector). This model therefore provides a new approach to understanding persistent disagreements.

We then study the long-run dynamics of opinions in "highly fluid" social networks, defined as networks where the product between the fraction of edges incoming in the stubborn agent set times the mixing time of the associated random walk is small. We show that in highly fluid social networks, the expected value and variance of the ergodic opinion of most of the agents concentrate around certain values in the large population limit. We refer to this result as "approximately equal influence" of stubborn agents on the rest of the society—meaning that their influence on most of the agents in the society are approximately the same.

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

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

In addition to the aforementioned works on learning and opinion dynamics, our model is closely related to the work by Mobilia and co-authors [33, 34, 35], which propose a variation of the discrete opinion dynamics model, also called the voter model, with "zealots" (equivalent to our stubborn agents). This work generally relies on heuristic mean-field approximations, valid for certain graphical structures, and numerical simulations, to characterize belief dynamics. In contrast, we prove convergence in distribution and characterize the properties of the limiting distribution for general finite graphs. Even though our model involves continuous belief dynamics, we shall also show that Mobilia's model can be recovered as a special case of our general framework.

Our work is also related to work on consensus and gossip algorithms, which is motivated by different problems, but typically leads to a similar mathematical formulation (Tsitsiklis [43], Tsitsiklis, Bertsekas and Athans [44], Jadbabaie, Lin and Morse [26], Olfati-Saber and Murray [38], Olshevsky and Tsitsiklis [39], 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 ergodic 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 ergodic process, while in general they do not converge almost surely. Section 4 presents a characterization of the first and second moments of the ergodic beliefs in terms of the hitting probabilities of two coupled random walks on the network. Section 5 narrows down the discussion to reversible social networks, and presents explicit computations of the expected ergodic beliefs and variances for some special network topologies. Section 6 provides bounds on the level of dispersion of the first two moments of the ergodic beliefs: it is shown that, in highly fluid networks, most of the agents have almost the same ergodic belief and variance. Section 7 studies the mean square oscillations and disagreement in highly fluid networks: if there is a significant presence of stubborn agents, the variance of the ergodic aggregate belief of the society vanishes in the large population limit, and the joint distribution of the ergodic beliefs is close to a chaotic law. Section 8 contains some concluding remarks.

Basic Notation and Terminology

Before proceeding, we establish some notational conventions and terminology to be followed throughout the paper. We shall typically label the entries of vectors by elements of finite alphabets, rather than non-negative integers, hence $\mathbb{R}^{\mathcal{I}}$ will stand for the set of vectors with entries labeled by elements of the finite alphabet $\mathcal{I}$. An index denoted by a lower-case letter will implicitly be assumed to run over the finite alphabet denoted by the corresponding calligraphic upper-case letter (e.g. $\sum_i$ will stand for $\sum_{i\in\mathcal{I}}$). For 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 $M \in \mathbb{R}^{\mathcal{I}\times\mathcal{J}}$, $||M||_1 := \max_j \sum_i M_{ij}$ and $||M||_\infty := \max_i \sum_j M_{ij}$ will denote its 1-norm, and ∞-norm, respectively, as an operator from $\mathbb{R}^{\mathcal{I}}$ to $\mathbb{R}^{\mathcal{J}}$. 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 shall use the notation

$$||\mu - \nu||_{TV} := \frac{1}{2}\sum\nolimits_i |\mu_i - \nu_i| = \sup\left\{\mu(\mathcal{J}) - \nu(\mathcal{J}) :\, \mathcal{J} \subseteq \mathcal{I}\right\},$$

for the total variation distance between μ and ν. The probability law (or distribution) of a random variable Z will be denoted by $\mathcal{L}(Z)$.

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

For two non-negative 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, $a_n \le Kb_n$ for all sufficiently large n, $a_n = \Theta(b_n)$ if $b_n = O(a_n)$, $a_n = o(b_n)$ if $\lim_n a_n/b_n = 0$, $a_n \approx b_n$ if $\lim_n a_n/b_n = 1$, and $a_n = \omega(b_n)$ if $b_n = o(a_n)$.

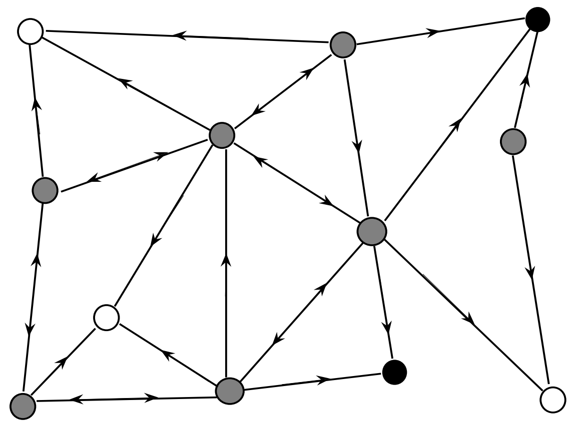

Fig 1. *A social network with seven regular agents (colored in grey), and five stubborn agents (colored in white, and black, respectively). Links are only incoming to the stubborn agents, while links between pairs of regular agents may be uni- or bi-directional.*

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 $\overrightarrow{\mathcal{G}} = \left(\mathcal{V}, \overrightarrow{\mathcal{E}}\right)$, whose node set is identified with the agent population, and where $\overrightarrow{\mathcal{E}} \subseteq \mathcal{V} \times \mathcal{V}$ stands for the set of directed edges (or links) among the agents.

At time $t \geq 0$, each agent $v \in \mathcal{V}$ holds a *belief* (or *opinion*) about an underlying state of the world, denoted by $X_v(t) \in \mathbb{R}$. The full vector of beliefs at time t will be denoted by $X(t) = \{X_v(t) : v \in \mathcal{V}\}$. We distinguish between two types of agents: regular and stubborn. Regular agents repeatedly update their own beliefs, based on the observation of the beliefs of their 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 shall denote the set of regular agents by $\mathcal{A}$, the set of stubborn agents by $\mathcal{S}$, so that the set of all agents is $\mathcal{V} = \mathcal{A} \cup \mathcal{S}$ (see Figure 1).

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

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

In contrast, the beliefs of the regular agents are updated as follows. To every 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 that $r_{vv'} = \theta_{vv'} = 0$ for all $v, v' \in \mathcal{V}$ such that $(v,v') \notin \overrightarrow{\mathcal{E}}$. 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 1. *Every regular agent is influenced by some stubborn agent, i.e., $\mathcal{S}_a$ is non-empty for every a in $\mathcal{A}$.*

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

Notice that both assumptions 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. On the other hand, if a stubborn agent does not influence any agent, it can clearly be neglected.

[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.

In the subsequent analysis, it is convenient to consider a rate matrix $R \in \mathbb{R}^{\mathcal{A}\times\mathcal{V}}$ whose entries coincide with the edge activation rates r_{av}. 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 a transition rate matrix $H \in \mathbb{R}^{\mathcal{A}\times\mathcal{V}}$, whose entries are defined by

$$H_{av} := \theta_{av} r_{av}\,, \qquad a \in \mathcal{A}\,,\ v \in \mathcal{V}\,, \tag{2}$$

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

$$P_{av} = H_{av} / \sum\nolimits_{v'} H_{av'}\,, \qquad a \in \mathcal{A}\,,\ v \in \mathcal{V}\,. \tag{3}$$

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 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

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

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 $v \in \mathcal{V}$, which corresponds to the classical voter model with zealots. We show that in both of these cases agent beliefs almost surely fail to converge.

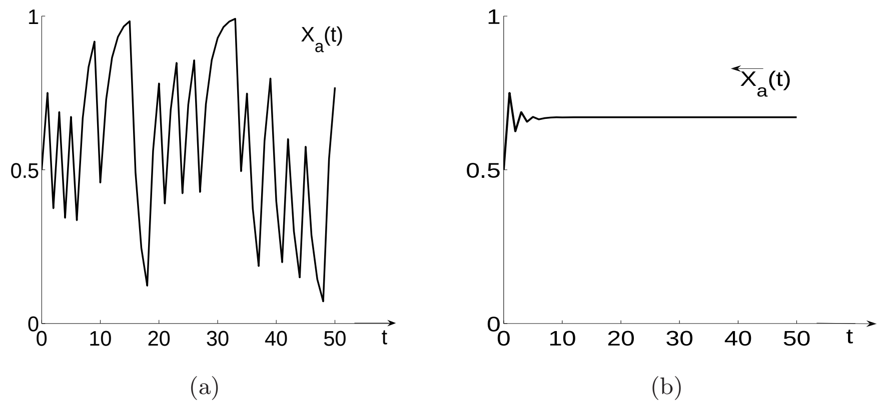

FIG 2. *Typical sample-path behavior: in (a) the actual belief process $X_a(t)$, oscillating ergodically on the interval $[0,1]$; in (b), the time-reversed process, rapidly converging to an asymptotic random belief X_a.*

2.1. *Single regular agent.* Consider a society consisting of a single regular agent, i.e., $\mathcal{A} = \{a\}$, and two stubborn agents, $\mathcal{S} = \{s, s'\}$. 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 $t \geq 0$,

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

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 an independent sequence of Bernoulli(1/2) random variables. Consider the events $E_k := \{W(3k-2) = 0, W(3k-1) = 0, W(3k) = 1\}$, for $k \geq 1$, which are independent and all have probability 1/8. Then, an application of the Borel-Cantelli lemma implies that E_k occurs for infinitely many values of $k \geq 1$. Notice that, when E_k occurs, one has

$$\begin{aligned} X_a(T_{3k}) &= \tfrac{1}{2} + \sum_{1 \leq j \leq 3k-3} 2^{j-3k-1} W(j) \geq \tfrac{1}{2}\,, \\ X_a(T_{3k-1}) &= \sum_{1 \leq j \leq 3k-3} 2^{j-3k} W(j) \leq \tfrac{1}{4}\,, \end{aligned} \tag{5}$$

where T_k is the time of the k-th belief update. It follows from (5), and the fact that $N(t)$ grows unbounded almost surely, 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

$$X_a := \sum_{k\geq 1} 2^{-k} W(k)$$

is sample-wise converging. It follows that, as t grows large, the time-reversed process

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

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 turns out that, for all $t\geq 0$, 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. [41, 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

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

converges in distribution to the asymptotic belief

$$X_a := \theta(1-\theta)^{-1} \sum_{k\geq 1} (1-\theta)^k W(k)\,. \tag{6}$$

As explained in [18, Sect. 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] The extreme case $\theta=1$ falls within the framework of Sect. 2.2. On the other hand, observe that, regardless of the fine structure of the probability law of the asymptotic belief X_a, i.e., on whether it is absolutely continuous or singular, it is not hard to characterize its moments for all values of $\theta\in(0,1]$. In fact, it follows from (6), that the expected value of X_a is given by

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

[3] See [40]. 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.

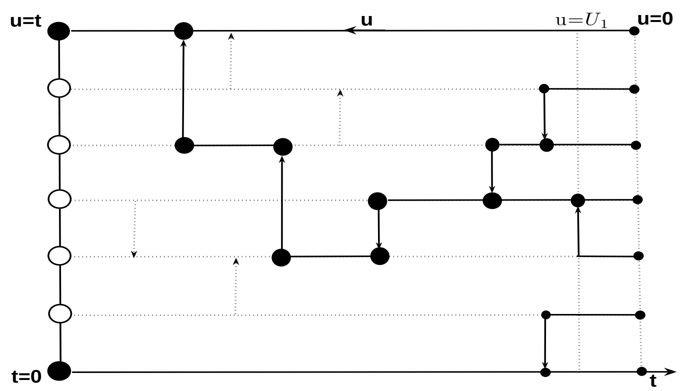

FIG 3. *Duality between the voter model with zealots and the coalescing random walks process with absorbing states. The network topology is a line with five regular agents and two stubborn agents placed in the two extremities. The time index for the opinion dynamics, t, runs from left to right, whereas 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.*

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.* Let us consider the case when $\theta_{av} = 1$ 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 works of Mobilia and others [33, 34, 35] a variant with one or more stubborn agents (there referred to as *zealots*) has been proposed and analyzed mainly through simulations. We wish to emphasize that the 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 prefer to discuss this case here in some detail, 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 $t - U_2 \in [0, t - U_1]$, 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 $V_a(u)$ on $\mathcal{V}$, which starting 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 transition rate matrix R (recall that R is the matrix of edge activation rates $R = [r_{av}]_{a \in \mathcal{A},\ v \in \mathcal{V}}$), and it halts when it hits some state $s \in \mathcal{S}$. In particular, 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, $\{V_a(t) : a \in \mathcal{A}\}$ starting one from each node $a \in \mathcal{A}$, and run simultaneously on $\mathcal{V}$ (see Figure 3). These random walks move independently with transition rate matrix R, 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 transition rate matrix R; 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

$$\mathcal{L}(\{X_a(t) : a \in \mathcal{A}\}) = \mathcal{L}(\{X_{V_a(t)}(0) : a \in \mathcal{A}\}) . \tag{7}$$

Equation (7) establishes a *duality* between the voter model with zealots and

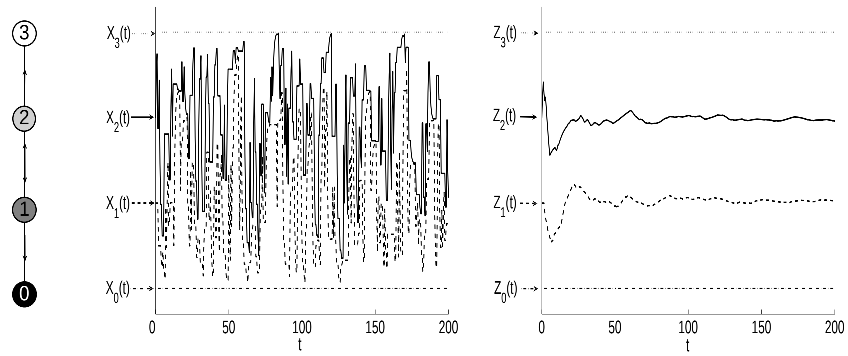

(a) (b) (c)

FIG 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 $\{Z_v(t) := t^{-1}\int_0^t X_v(u)\mathrm{d}u\}$.*

the coalescing random walks process with absorbing states. In particular, Assumptions 1 and 2 imply 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 asymptotic $\mathcal{S}^{\mathcal{A}}$-valued random vector $\{V_a(T^a_{\mathcal{S}}) : a \in \mathcal{A}\}$. It then follows from (7) that $X(t)$ converges in distribution to an asymptotic 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 Sect. 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 in the interval $[x_0, x_3]$, in an apparently chaotic way. On the other hand, the time averages of the two regular agents' beliefs rapidly approach a limit value, of $2x_0/3 + x_3/3$ for agent 1, and $x_0/3 + 2x_3/3$ for agent 2.

As we shall see below, such behavior is rather general. In our model of social network with at least two stubborn agents having non-coincident constant beliefs, the regular agent beliefs almost surely fail to converge: we have seen this in the special cases of Sect. 2.1, while a general result in this sense will be stated as Theorem 2. On the other hand, we shall prove that, regardless of the initial regular agents' beliefs, the belief vector $X(t)$ is convergent in distribution to a random asymptotic belief vector X (see Theorem 1), and in fact it is an ergodic process (see Corollary 1).

In order to prove Theorem 1, we shall rewrite $X(t)$ in the form of an iterated affine function system [18], as explained below. We shall then 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 asymptotic 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. Similarly, for $a, a' \in \mathcal{A}$, let $e_{aa'} \in \mathbb{R}^{\mathcal{A}\times\mathcal{A}}$ be the matrix whose entries are all 0, but for the (a, 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

$$A(k) = I_{\mathcal{A}} + \theta_{aa'}(e_{aa'} - e_{aa}) \qquad B(k) = 0\,,$$

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

$$A(k) = I_{\mathcal{A}} - \theta_{as}e_{aa} \qquad B(k) = e_a\theta_{as}x_s\,,$$

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

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

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

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

with the convention that $\overleftarrow{A}(k,l) = I_{\mathcal{A}}$ for $k > l$. The following is our first, but fundamental, observation:

LEMMA 1. *For all $t \ge 0$, $Y(t)$ 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 $\{(A(k), B(k)) : 1 \le k \le l\}$ has the same distribution as the l-tuple $\{(A(l-k+1), B(l-k+1)) : 1 \le k \le l\}$, 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 $t \ge 0$. ■

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, as formalized in the next lemma.

LEMMA 2. *Let Assumptions 1 and 2 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

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

for all $a \ne a' \in \mathcal{A}$. In particular, it follows from Assumption 1 that $\mathbb{E}[A(k)]$ is a strictly substochastic matrix, with no invariant subset, i.e., such that if y is a non-negative vector supported in some $\mathcal{J} \subseteq \mathcal{A}$, then $\sum_j \sum_a \mathbb{E}[A_{ja}(k)]y_a < \sum_j y_j$. Hence, its spectrum is contained in the disk centered in 0 of radius ρ, where $\rho \in (0,1)$ is the largest eigenvalue of $\mathbb{E}[A(k)]$. Then, using the Jordan canonical decomposition, one can show that

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

where C is a constant depending on $\mathbb{E}[A(1)]$ only. Now, upon observing that the $\overleftarrow{A}(1,k)$ has non-negative entries, and using the inequality $\mathbb{E}[\max\{Z,W\}] \leq \mathbb{E}[Z]+\mathbb{E}[W]$ valid for all nonnegative-valued random variables W and Z, one gets that

$$\begin{array}{rcl} \mathbb{E}\left[\left|\left|\overleftarrow{A}(1,k)\right|\right|_1\right] & = & \mathbb{E}\left[\max_{a'}\sum_a \overleftarrow{A}_{aa'}(1,k)\right] \\ & \leq & \sum_{a,a'}\mathbb{E}\left[\overleftarrow{A}_{a,a'}(1,k)\right] \\ & \leq & n\left|\left|\mathbb{E}\left[\overleftarrow{A}(1,k)\right]\right|\right|_\infty \\ & \leq & Cnk^{n-1}\rho^k\,. \end{array} \tag{10}$$

It follows that

$$\begin{array}{rcl} \inf_{k\in\mathbb{N}}\frac{1}{k}\mathbb{E}\left[\log||\overleftarrow{A}(1,k)||_1\right] & \leq & \inf_{k\in\mathbb{N}}\frac{1}{k}\log\mathbb{E}\left[||\overleftarrow{A}(1,k)||_1\right] \\ & \leq & \lim_{k\to+\infty}\frac{\log(Cnk^{n-1}\rho^k)}{k} \\ & = & \log\rho \\ & < & 0\,, \end{array} \tag{11}$$

Then, it follows from [18, Th. 2.1] that the series

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

is convergent with probability one. Since, with probability one, $\lim_{t\to+\infty} N(t) = +\infty$, one has that

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

with probability one. This completes the proof. ∎

Lemma 1 and Lemma 2 allow one to prove convergence in distribution of $X(t)$ to a random belief vector X, as stated in the following result.

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

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

for all bounded and continuous test functions $\varphi : \mathbb{R}^{\mathcal{V}} \to \mathbb{R}$. Moreover, the probability law of the asymptotic 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 2 $\overleftarrow{Y}(t)$ converges to Y with probability one, and a fortiori in distribution. By Lemma 2, $\overleftarrow{Y}(t)$ 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. ■

REMARK 1. In fact, a more detailed proof of Lemma 2 (based on the estimate (10), and using directly Markov's inequality, without appealing to [18, Th. 2.1]) would have shown that

$$\lim_{t \to +\infty} W_1\left(\mathcal{L}(X(t)), \mathcal{L}(X)\right)^{1/t} \leq \rho\,,$$

where W_1 denotes the so-called order-1 Wasserstein distance. The latter is a metric between probability measures on $\mathbb{R}^{\mathcal{V}}$ which metrizes weak convergence, and has been made popular by optimal transportation theory: we refer to [45] for definition, and an extensive survey of its properties.

Now, we observe that, if α is the unique probability distribution on $\mathcal{A}$ such that $\mathbb{E}[A(k)]\alpha = \rho\alpha$, and $\mathbb{E}_\alpha[T_{\mathcal{S}}]$ is the expected time to hit the stubborn agents set $\mathcal{S}$ for the continuous-time random walk $V(t)$, with transition rates H_{av} defined in (2), and initial state distribution α, then

$$(1-\rho)^{-1} = \mathbb{E}_\alpha[T_{\mathcal{S}}]$$

(see [3, Ch. 3, Sect. 6.5]). This shows that, as already observed in the voter model with zealots of Sect. 2.2, the speed of convergence of the beliefs' distribution to the asymptotic belief vector is related to the hitting time $T_{\mathcal{S}}$ of the random walk $V(t)$ on the stubborn agents' set $\mathcal{S}$. In the following section, we shall see that also the probability law of the asymptotic beliefs, $\mathcal{L}(X)$, can be related to properties of the random walk $V(t)$, and specifically to the probability distribution of $V(T_{\mathcal{S}})$.

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

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

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

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

Proof. Consider the joint discrete-time process $\{Z_k := (X(T_k), U_k) : k \in \mathbb{N}\}$, where $U_k := T_k - T_{k-1}$, and T_k is the time of the k-th belief update (with the convention that $T_0 = 0$). Since $\{U_k : k \in \mathbb{N}\}$ is an independent identically distributed sequence of exponential random variables with rate r, independent from $X(T_k)$, Theorem 1 implies that Z_k converges in distribution to $Z := (X, U)$, where X is as in Theorem 1, U is an independent random variable with exponential distribution of rate r. It follows from [10] that the process $\{Z_k\}$ s ergodic. Hence, for every test function φ continuous and bounded over $\mathbb{R}^{\mathcal{V}}$,

$$\begin{array}{rcl}
\frac{1}{k}\int_0^{T_k}\varphi(X(u))\mathrm{d}u & = & \frac{1}{k}\sum\limits_{1\le j\le k}\varphi(X(T_j^-))U_j \\
& \stackrel{k\to+\infty}{\longrightarrow} & \mathbb{E}[\varphi(X)U] \\
& = & \mathbb{E}[\varphi(X)]\mathbb{E}[U] \\
& = & r\mathbb{E}[\varphi(X)]\,,
\end{array}$$

with probability one. Moreover, one has

$$\frac{1}{k}T_k = \frac{1}{k}\sum_{1\le j\le k}(T_k - T_{k-1}) \stackrel{k\to+\infty}{\longrightarrow} \mathbb{E}[T] = r\,.$$

Hence,

$$\frac{1}{T_{N(t)}}\int_0^{T_{N(t)}}\varphi(X(u))\mathrm{d}u = \frac{N(t)^{-1}\int_0^{T_{N(t)}}\varphi(X(u))\mathrm{d}u}{N(t)^{-1}T_{N(t)}} \overset{t\to+\infty}{\longrightarrow} \mathbb{E}[\varphi(X)]$$

with probability one, and the claim follows from the observation that, as t grows large, $t^{-1}(t-T_{N(t)})$ converges to zero almost surely. ■

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

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

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

PROOF. With no loss of generality, since the distribution of the asymptotic 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 asymptotic one, i.e., that $\mathcal{L}(X(0)) = \mathcal{L}(X)$. Then, Theorem 1 implies that

$$\mathcal{L}(X(t)) = \mathcal{L}(X)\,, \qquad \forall t \geq 0\,. \tag{12}$$

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 shall 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 and a' agrees with a and 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 is degenerate, then $X_v(t) = X_a$ for all t. Fix some scalars x_a, $\zeta > 0$, $\varepsilon > 0$, and, for every time instant $t \geq 0$, consider the events

$$C_t := \{X_a(t) = x_a\}\,, \qquad D_t := \{|X_v(t) - x_a| \geq \zeta\}\,, \qquad E_t := C_t \cap D_t\,,$$
$$E'_t := \{\text{ agent } v \text{ is never active in } (t, t+\varepsilon)\}\,,$$
$$E''_t := \{\text{edge } (a, v) \text{ is activated at some time } u \in (t, t+\varepsilon)\}\,,$$
$$F_t := \{|X_a(u) - x_a| \geq \theta_{av}\zeta\,, \text{ for some time } u \in (t, t+\varepsilon)\}\,.$$

Observe that $E_t \cap E'_t \cap E''_t$ implies F_t. Moreover, E_t, E'_t, and E''_t are mutually independent, and if $r_{av} > 0$ is the activation rate of edge (a, v), and $r_v = \sum_{v'} r_{vv'} \geq 0$ is the meeting rate of agent v, then $\mathbb{P}(E'_t) = e^{-r_v\varepsilon}$, and $\mathbb{P}(E''_t) = 1 - e^{-r_a\varepsilon}$, respectively. It follows that

$$\mathbb{P}(F_t) \geq \mathbb{P}(E'_t \cap E''_t \cap E_t) = \alpha\mathbb{P}(E_t)\,, \qquad \alpha := e^{-r_v\varepsilon}(1 - e^{-r_a\varepsilon})\,. \tag{13}$$

Let $\mathbb{Q}^+$ be the set on non-negative rationals, define $C := \bigcap_{t\in\mathbb{Q}^+} C_t$, $D := \bigcup_{t\in\mathbb{Q}^+} D_t$, and observe that, since F_t implies C^c, the complementary event of C, one has $\mathbb{P}(F_t) \leq 1 - \mathbb{P}(C)$ for every $t \geq 0$. The foregoing, combined with (13), gives

$$\mathbb{P}(E_t) \leq \alpha^{-1}(1 - \mathbb{P}(C))\,, \qquad \forall t \geq 0\,. \tag{14}$$

On the other hand, since the event $C \cap D$ implies that E_t occurs for some nonnegative rational t, the union bound gives

$$\mathbb{P}(C \cap D) \leq \mathbb{P}\left(\bigcup_{t\in\mathbb{Q}^+} E_t\right) \leq \sum_{t\in\mathbb{Q}^+} \mathbb{P}(E_t)\,. \tag{15}$$

Then, (12) implies $\mathbb{P}(X_a(t) = x_a) = 1$ for all $t \geq 0$, and then C has probability one since it is a countable intersection of events of probability one. Then, it follows from (14) that $\mathbb{P}(E_t) = 0$ for all $t \geq 0$, and thus (15) implies that $\mathbb{P}(D) \leq 1 - \mathbb{P}(C) + \mathbb{P}(C \cap D) = 0$.

In order to prove that, if $X_a = X_{a'}$ almost surely, then $X_v(t) = X_a(t)$ for all $t \geq 0$, and every out-neighbor v of either a or a', one can argue in a very similar fashion. Assume without loss of generality that v is an out-neighbor of a. Fix some $\varepsilon > 0$, $\zeta > 0$, and, for $t \geq 0$, consider the events

$$G_t := \{X_a(t) = X_{a'}(t)\}\,, \qquad H_t := \{|X_v(t) - X_a(t)| \geq \zeta\}\,, \qquad L_t := G_t \cap H_t\,,$$
$$E'''_t := \{\text{agent } a' \text{ is never active in } (t, t+\varepsilon)\}\,,$$
$$M_t := \{|X_a(u) - X_{a'}(u)| \geq \theta_{av}\zeta\,, \text{ for some time } u \in (t, t+\varepsilon)\}$$
$$G := \bigcap_{t\in\mathbb{Q}^+} G_t\,, \qquad H := \bigcup_{t\in\mathbb{Q}^+} H_t\,.$$

Arguing as before, one finds that $E'_t \cap E''_t \cap E'''_t \cap L_t$ implies M_t, and that E'_t, E''_t, and E'''_t are conditionally independent given L_t, with conditioned probabilities $\mathbb{P}(E'_t|L_t) = e^{-r_v\varepsilon}$, $\mathbb{P}(E''_t|L_t) = 1 - e^{-r_a\varepsilon}$, and $\mathbb{P}(E'''_t|L_t) = 1 - e^{-r_{a'}\varepsilon}$, respectively, so that

$$\mathbb{P}(M_t \cap L_t) \geq \mathbb{P}(E'_t \cap E''_t \cap E'''_t \cap L_t) = \beta \mathbb{P}(L_t)\,, \tag{16}$$

where $\beta := e^{-r_v\varepsilon}(1 - e^{-r_a\varepsilon})(1 - e^{-r_{a'}\varepsilon})$. Moreover, since M_t implies G^c, for all $t \geq 0$, one has that $\mathbb{P}(L_t \cap M_t) \leq \mathbb{P}(M_t) \leq 1 - \mathbb{P}(G)$, which, together with (16), implies that $\mathbb{P}(L_t) \leq \beta^{-1}(1 - \mathbb{P}(G))$ for all $t \geq 0$. Arguing as before, one finds that $\mathbb{P}(G) = 1$, so that $\mathbb{P}(L_t) = 0$ for all $t \geq 0$, and thus $\mathbb{P}(H) = \mathbb{P}(G \cap H) = 0$. Then, from the arbitrariness of $\zeta > 0$, it follows that $X_v(t) = X_a(t)$ for all $t \geq 0$. ■

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

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

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

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

We next provide explicit characterizations of these limits in terms of hitting probabilities of a pair of coupled random walks on $\overrightarrow{\mathcal{G}} = (\mathcal{V}, \overrightarrow{\mathcal{E}})$.[4] Specifically, we consider a coupling $(V(t), V'(t))$ of continuous-time random walks on $\mathcal{V}$, such that both $V(t)$, and $V'(t)$, have marginal state transition rates H_{av}, as defined in (2). In fact, one may interpret $(V(t), V'(t))$ as a random walk on the Cartesian power graph $\mathcal{G}_\square$, whose node set is the product $\mathcal{V} \times \mathcal{V}$, and where there is an edge from (v, v') to (w, w'), if and only if either $(v, w) \in \overrightarrow{\mathcal{E}}$ and $v' = w'$, or $v = w$ and $(v', w') \in \overrightarrow{\mathcal{E}}$, or $v = v'$ and

[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.

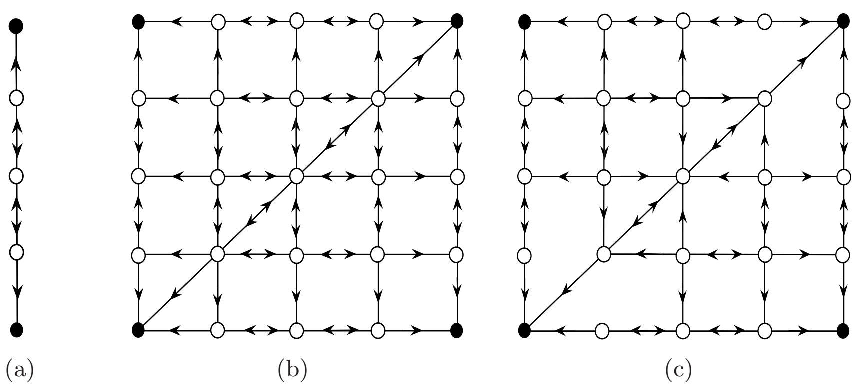

FIG 5. *In (a), a network topology consisting of a line with three regular agents and two stubborn agents placed in the extremes. In (b), the corresponding graph product $\mathcal{G}_\square$. The latter has 25 nodes, four of which are absorbing states. The coupled random walk $(V(t), V'(t))$ moves on $\mathcal{G}_\square$, its two components jumping independently to neighbor states, unless they are either on the diagonal, or one of them is in $\mathcal{S}$: in the former case, there is some chance that the two components jump as a unique one, thus inducing a direct connection along the diagonal; in the latter case, the only component that can keep moving is the one which has not hit $\mathcal{S}$, while the one who hit $\mathcal{S}$ is bound to remain constant from that point on. 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.*

$w = w'$ (See Figure 5). The transition rates of the coupled random walks $(V(t), V'(t))$ are given by

(17)
$$K_{(v,v')(w,w')} := \begin{cases} H_{vw} & \text{if } v \neq v'\,,\ v \neq w\,,\ v' = w' \\ H_{v'w'} & \text{if } v \neq v'\,,\ v = w\,,\ v' \neq w' \\ 0 & \text{if } v \neq v'\,,\ w \neq v\,,\ w' \neq v \\ \theta_{vw} H_{vw} & \text{if } v = v'\,,\ w = w' \\ (1-\theta_{vw}) H_{vw} & \text{if } v = v'\,,\ v' = w' \\ (1-\theta_{vw'}) H_{vw'} & \text{if } v = v'\,,\ v = w \\ 0 & \text{if } v = v'\,,\ w \neq w'\,,\ w \neq v\,,\ w' \neq v'\,. \end{cases}$$

The first three lines of (17) 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 H_{vw} (respectively, to a neighbor node w' with transition rate $H_{v'w'}$), whereas the probability that both components jump at the same time is zero. On the other hand, the last four lines of (17) 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 $\theta_{vw} H_{vw}$, while each of the components $V(t)$ (respectively, $V'(t)$) has still some chance to jump alone to a neighbor node w with rate $(1-\theta_{vw})H_{vw}$ (resp., to w' with rate $(1-\theta_{vw'})H_{vw'}$). In the extreme case when $\theta_{vw} = 1$ for all v, w, the last three lines of the righthand side of (17) equal 0, and in fact one recovers the expression for the transition rates of two coalescing random walks: once $V(t)$ and $V'(t)$ have met, they stick together and move as a single particle, never separating from each other.

We use the notation $T_{\mathcal{S}}$ and $T'_{\mathcal{S}}$ to denote the *hitting times* of the random walks $V(t)$, and respectively $V'(t)$, to the set of stubborn agents $\mathcal{S}$, i.e.,

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

Further, for all $v, v' \in \mathcal{V}$, we define the *hitting probability distributions* γ^v over $\mathcal{S}$, and $\eta^{vv'}$ over $\mathcal{S}^2$, whose entries are respectively given by

$$\begin{aligned} \gamma^v_s &:= \mathbb{P}_v(V(T_{\mathcal{S}}) = s)\,, \qquad s \in \mathcal{S}\,, \\ \eta^{vv'}_{ss'} &:= \mathbb{P}_{vv'}(V(T_{\mathcal{S}}) = s, V'(T'_{\mathcal{S}}) = s')\,, \qquad s, s' \in \mathcal{S}\,. \end{aligned} \tag{18}$$

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

LEMMA 3. *For all $s, s' \in \mathcal{S}$, one has that*

$$\begin{aligned}
&\textstyle\sum_v H_{av}(\gamma_s^v - \gamma_s^a) = 0\,, && \forall a \in \mathcal{A}\,,\\
&\gamma_s^s = 1\,, \qquad \gamma_{s'}^s = 0\,, && \forall s' \in \mathcal{S}\setminus\{s\}\,,
\end{aligned} \tag{19}$$

$$\begin{aligned}
&\textstyle\sum_{v,v'} K_{(a,a')(v,v')}\left(\eta_{ss'}^{vv'} - \eta_{ss'}^{aa'}\right) = 0\,, && \forall a, a' \in \mathcal{A}\,,\\
&\eta_{ss'}^{vv'} = \gamma_s^v \gamma_{s'}^{v'}\,, && \forall (v,v') \in \mathcal{V}^2 \setminus \mathcal{A}^2\,.
\end{aligned} \tag{20}$$

PROOF. Observe that the second line of (19) is trivial since, if $V(0) = s$, then $T_{\mathcal{S}} = 0$, and thus $\gamma_{s'}^s = \mathbb{P}_s(V(T_{\mathcal{S}}) = s') = \mathbb{P}_s(V(0) = s')$ is 1 if $s' = s$ and 0 otherwise. On the other hand, the first line of (19) follows by conditioning on the first state $v \in \mathcal{V}$ hit by a random walk $V(t)$ started from $V(0) = a \in \mathcal{A}$: the probability γ_s^a that such a walk hits some s before any other $s' \in \mathcal{S}$ equals the sum over all neighbors v of a of the probability γ_s^v that a walk started from v hits some s before any other $s' \in \mathcal{S}$, times the probability that the first neighbor hit by the random walk started in a is actually v, which is proportional to H_{av}.

As far as the first line of (20) is concerned, it follows from conditioning on the first pair of vertices (v, v') hit by the joint random walk $(V(t), V'(t))$, arguing as before. To see why the second line of (20) holds, first assume $v' \in \mathcal{S}$ (the alternative case when $v \in \mathcal{S}$ follows from a symmetric argument): in this case, necessarily $V'(T'_{\mathcal{S}}) = V'(0) = s'$, so that $\eta_{ss'}^{vv'} = 0$ if $v' \neq s'$, while, if $v = s'$, one has

$$\eta_{ss'}^{vv'} = \mathbb{P}(V(T_{\mathcal{S}}) = s, V'(T'_{\mathcal{S}}) = s') = \mathbb{P}(V(T_{\mathcal{S}})) = \gamma_s^v\,,$$

thus concluding the proof. ∎

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

THEOREM 3. For all $v, v' \in \mathcal{V}$,

$$\mathbb{E}[X_v] = \sum\nolimits_s \gamma_s^v x_s\,, \qquad \mathbb{E}[X_v X_{v'}] = \sum\nolimits_{s,s'} \eta_{ss'}^{vv'} x_s x_{s'}\,.$$

PROOF. With no loss of generality, since the distribution of the asymptotic 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 asymptotic one, i.e., that $\mathcal{L}(X(0)) = \mathcal{L}(X)$. Then, Theorem 1 implies that

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}\left[\varphi(X(t)\right] = 0\,, \qquad \forall t \geq 0\,, \tag{21}$$

for all continuous bounded test function $\varphi : \mathbb{R}^{\mathcal{V}} \to \mathbb{R}$.

In order to show the first part of the claim, we consider the Laplace equation with boundary conditions

$$\begin{aligned} &\textstyle\sum_v H_{av}(h_v - h_a) = 0\,, && \forall a \in \mathcal{A}\,, \\ &h_s = x_s\,, && \forall s \in \mathcal{S}\,. \end{aligned} \tag{22}$$

It is a standard fact (see e.g. [3, Ch. 2, Lemma 27]) that the system of equations (22) admits a unique solution $\{h_v : v \in \mathcal{V}\}$, and, thanks to (19), it is easy to verify that (22) is satisfied by $h_v := \sum_s \gamma_s^v x_s$. Now, (21) with the specific choice of the test function $\varphi(x) = x_a$ yields

$$0 = \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[X_a(t)] = \sum_v r_{av}\theta_{av}\left(\mathbb{E}[X_v(t)] - \mathbb{E}[X_a(t)]\right) = \sum_v H_{av}\left(\mathbb{E}[X_v] - \mathbb{E}[X_a]\right)\,,$$

for every $a \in \mathcal{A}$. On the other hand, $\mathbb{E}[X_s(t)] = x_s$ for all $s \in \mathcal{S}$. Hence, the expected asymptotic beliefs vector $\{\mathbb{E}[X_v] : v \in \mathcal{V}\}$ solves (22), and the first part of the claim follows by the aforementioned uniqueness of solutions.

In order to prove the second part of the claim, we proceed in a similar fashion. First, observe that thanks to (20), the vector $\{h_{vv'} := \sum_{s,s'} \eta_{ss'}^{vv'} x_s x_{s'}\}$ is the unique solution of the Laplace equation with boundary conditions

$$\begin{aligned} &\textstyle\sum_{v,v'} K_{(a,a')(v,v')}(h_{(v,v')} - h_{(a,a')}) = 0\,, && \forall a, a' \in \mathcal{A}\,, \\ &h_{(v,v')} = \mathbb{E}[X_v]\mathbb{E}[X_{v'}]\,, && \forall (v, v') \in \mathcal{V}^2 \setminus \mathcal{A}^2\,. \end{aligned} \tag{23}$$

Then, for all $a \neq a' \in \mathcal{A}$, (21) applied to the test function $\varphi(x) = x_a x_{a'}$ implies that

$$\begin{aligned} 0 &= \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[X_a(t)X_{a'}(t)] \\ &= \textstyle\sum_v r_{av}\theta_{av}\left(\mathbb{E}[X_v(t)X_{a'}(t)] - \mathbb{E}[X_a(t)X_{a'}(t)]\right) \\ &\quad + \textstyle\sum_{v'} r_{a'v'}\theta_{a'v'}\left(\mathbb{E}[X_a(t)X_{v'}(t)] - \mathbb{E}[X_a(t)X_{a'}(t)]\right) \\ &= \textstyle\sum_{v,v'} K_{(a,a')(v,v')}\left(\mathbb{E}[X_v X_{v'}] - \mathbb{E}[X_a X_{a'}]\right)\,, \end{aligned}$$

while (21) applied to $\varphi(x) = x_a^2$ gives

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[X_a^2(t)] &= \textstyle\sum_v r_{av}\mathbb{E}\left[\left((1-\theta_{av})X_a + \theta_{av}X_v\right)^2 - X_a^2\right] \\ &= \textstyle\sum_v r_{av}\theta_{av}\left(\theta_{av}\mathbb{E}[X_v^2] + 2(1-\theta_{av})\mathbb{E}[X_a X_v] - (2-\theta_{av})\mathbb{E}[X_a^2]\right) \\ &= \textstyle\sum_v H_{av}(1-\theta_{av})\left(\mathbb{E}[X_a X_v] - \mathbb{E}[X_a^2]\right) \\ &\quad + \textstyle\sum_{v'} H_{av'}(1-\theta_{av'})\left(\mathbb{E}[X_a X_{v'}] - \mathbb{E}[X_a^2]\right) \\ &\quad + \textstyle\sum_v H_{av}\theta_{av}\left(\mathbb{E}[X_v^2] - \mathbb{E}[X_a^2]\right) \\ &= \textstyle\sum_{v,v'} K_{(a,a')(v,v')}\left(\mathbb{E}[X_v X_{v'}] - \mathbb{E}[X_a X_{a'}]\right)\,. \end{aligned}$$

On the other hand, it is easy to see that $\mathbb{E}[X_v X_{v'}] = \mathbb{E}[X_v]\mathbb{E}[X_{v'}]$ for all $(v, v') \in \mathcal{V}^2 \setminus \mathcal{A}^2$, thus proving that $\mathbb{E}[X_v X_{v'}]$ solves (23). Hence, the second part of the claim follows by the uniqueness of solutions of (23). ■

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

5. Reversible social networks. From now on, we restrict to considering social networks whose transition probability matrix P satisfies an additional assumption of reversibility. Recall from Sect. 2 that the stochastic matrix P was defined only on $\mathcal{A} \times \mathcal{V}$, and not on $\mathcal{S} \times \mathcal{V}$.

ASSUMPTION 3. *The restriction of P 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*

$$\tilde{\pi}_a P_{aa'} = \tilde{\pi}_{a'} P_{a'a}\,, \qquad a, a' \in \mathcal{A}\,. \tag{24}$$

The irreducibility assumption will be satisfied when the graph $\overrightarrow{\mathcal{G}}$ restricted to the regular agent set $\mathcal{A}$ is strongly connected. In fact, the irreducibility assumption causes no significant loss of generality, as one can always separately study the different communicating classes in the network. On the other hand, the reversibility assumption is more stringent, as it entails a sort of reciprocity between the intensity of influence of a regular agent on another, and vice versa. In particular, it implies that if $(a, a') \in \overline{\mathcal{E}}$, then also $(a', a) \in \overline{\mathcal{E}}$, so that the restriction of the di-graph $\overline{\mathcal{G}}$ to $\mathcal{A}$ is in fact undirected.

We shall refer to social networks whose transition probability matrix P satisfies Assumption 3 as *reversible (social) networks*. Observe that, whenever Assumption 3 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 (24) 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, it is possible to extend P on $\mathcal{S} \times \mathcal{V}$ as follows: Put $P_{ss'} := 0$ for all $s, s' \in \mathcal{S}$, and

$$\tilde{\pi}_s := \sum_{a \in \mathcal{A}} \tilde{\pi}_a P_{as}\,, \qquad P_{sa} := P_{as}\tilde{\pi}_a / \tilde{\pi}_s\,, \qquad s \in \mathcal{S},\, a \in \mathcal{A}\,. \tag{25}$$

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 of dimension $\mathcal{V} \times \mathcal{V}$.[5] Furthermore, the probability measure π on $\mathcal{V}$, defined by

$$\pi_v := \left(\sum\nolimits_{v'} \tilde{\pi}_{v'}\right)^{-1} \tilde{\pi}_v \,, \tag{26}$$

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

$$\pi(\mathcal{S}) = \sum\nolimits_s \pi_s = \left(\sum\nolimits_v \tilde{\pi}_v\right)^{-1} \sum\nolimits_{a,s} \tilde{\pi}_a P_{as} \tag{27}$$

Observe that (24) and (25) imply that $P_{vv'} > 0$ if and only if $P_{vv'} > 0$. It is then natural to associate to any social network satisfying Assumption 3 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.

Example 2. Let us consider the canonical construction of a social network from a given undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, explained in Example 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 3 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,

$$\pi(\mathcal{S}) = \left(\sum\nolimits_v d_v\right)^{-1} \sum\nolimits_s d_s$$

is the fraction of edges incident to the stubborn agents, or, in other words, the relative size of the boundary of $\mathcal{S}$ in $\mathcal{G}$.

[5]This is not the only possible extension that makes P 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 may vary with the choice of such an extension.

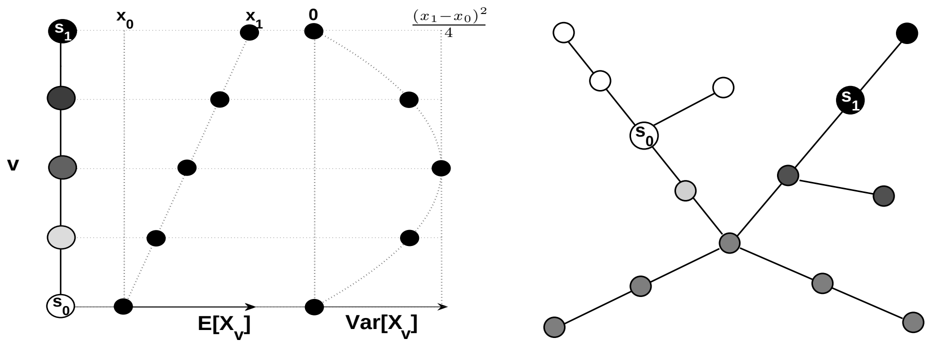

FIG 6. *In the left-most figure, expected ergodic beliefs and variances in a social network with a line graph topology with $n = 5$, and stubborn agents positioned in the two extremities. The expected ergodic beliefs are linear interpolations of the two stubborn agents' beliefs, while their variances 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 ergodic 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.*

5.1. *Explicit computations for some reversible networks.* We present now a few examples of explicit computations of the ergodic average beliefs and variances for reversible social networks, obtained using the construction in Example 1.

EXAMPLE 3. (**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 2 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

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

where $d(a, v)$ denotes the distance in $\mathcal{G}$, between (i.e., the length of the shortest path connecting) nodes a and v. Hence, the ergodic average beliefs are linear interpolations of the beliefs of the stubborn agents. Moreover, if the confidence parameters are $\theta_e = 1$ for all e, then the ergodic variance of agent a' belief is given by

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

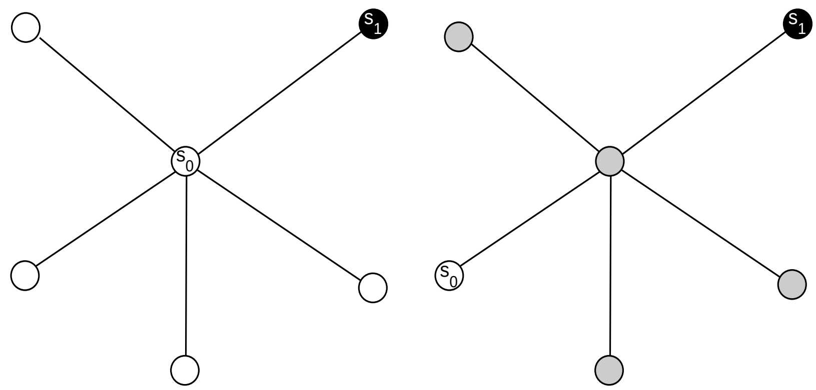

Fig 7. *Two social network with a special case of tree-like topology, known as star graph, and two stubborn agents. In social network depicted in left-most figure one of the stubborn agents, s_0, occupies the center, while the other one, s_1, occupies one of the leaves. There, all regular agents' ergodic 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' ergodic beliefs coincide with the arithmetic average (represented in gray) of x_0 (white), and x_1 (black).*

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 approximately equally 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 ergodic 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 ergodic 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 cor-

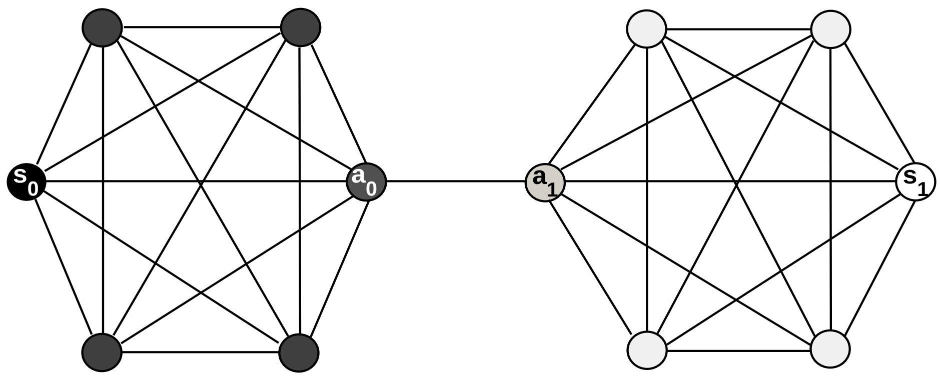

FIG 8. *A social network with population size $n = 12$, a barbell-like topology, and two stubborn agents. In each of the two halves of the graph the expected average beliefs concentrate around the beliefs of the stubborn agent in the respective half.*

responding 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.

EXAMPLE 4. (**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, the expected ergodic beliefs satisfy

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

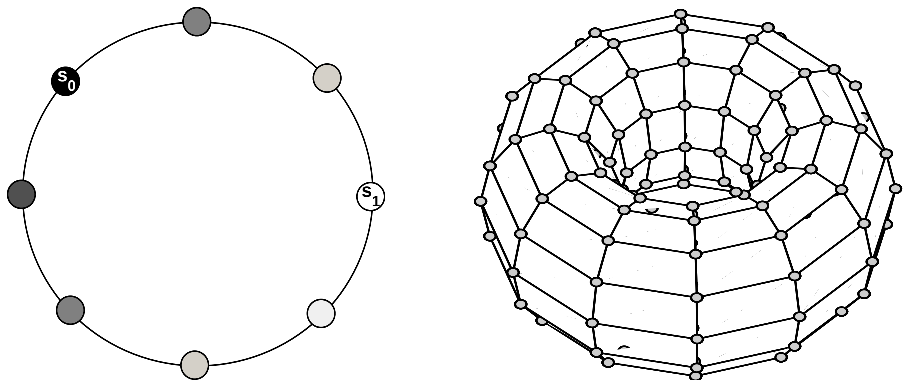

FIG 9. *Two social networks with cycle and 2-dimensional toroidal topology, respectively.*

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. (**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 $\Theta = \{\pm e_1, \dots, \pm e_d\}$, 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]):

$$\gamma_{s_0}^v = \mathbb{P}_v(T_{s_1} < T_{s_0}) = \frac{E_{vs_0} - E_{vs_1} + E_{s_1 s_0}}{E_{s_0 s_1} + E_{s_1 s_0}} \tag{28}$$

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, mean 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

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

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

$$E_{vw} = n^{-1}(Z_{ww} - Z_{vw})\,. \tag{29}$$

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 $l = (l_1, \dots l_d) \in \mathcal{V}$ define $\phi_l \in \mathbb{R}^{\mathcal{V}}$ by

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

(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,

$$Z_{vw} = m^{-d} \sum_{l \neq 0} \frac{\exp\left[\frac{2\pi i}{m} l \cdot (v - w)\right]}{1 - \frac{1}{|\Theta^+|} \sum\limits_{k \in \Theta^+} \cos\left[\frac{2\pi}{m} l \cdot k\right]} \tag{30}$$

From (28), (29), and the fact that $E_{s_0 s_1} = E_{s_1 s_0}$ by symmetry, one obtains

$$\gamma^a_{s_1} = \frac{1}{2} + \frac{m^{-d} \sum\limits_{l \neq 0} \dfrac{\exp\left[\frac{2\pi i}{m} l \cdot (a - s_1)\right] - \exp\left[\frac{2\pi i}{m} l \cdot (a - s_0)\right]}{1 - \frac{1}{|\Theta^+|} \sum\limits_{k \in \Theta^+} \cos\left[\frac{2\pi}{m} l \cdot k\right]}}{2m^{-d} \sum\limits_{l \neq 0} \dfrac{1 - \cos\left[\frac{2\pi}{m} l \cdot (s_0 - s_1)\right]}{1 - \frac{1}{|\Theta^+|} \sum\limits_{k \in \Theta^+} \cos\left[\frac{2\pi}{m} l \cdot k\right]}} \tag{31}$$

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

DEFINITION 1. Given a reversible social network, let P denote its transition probability matrix, extended as in (25), and π denote its stationary distribution defined as in (26). Define $\pi_* := \min_v \pi_v$, and let

$$\tau := \inf\left\{ t \geq 0 : \sum_w |\mathbb{P}_v(V(t) = w) - \mathbb{P}_{v'}(V(t) = w)| \leq \frac{2}{e}\,, \forall v, v' \in \mathcal{V} \right\}. \tag{32}$$

denote the (variational distance) mixing time of the continuous-time random walk $V(t)$ with transition rate matrix P. We say that a sequence[6] of social networks of increasing population size n is *highly fluid* if it satisfies

$$\tau\pi(\mathcal{S}) = o(1)\,, \qquad n\pi_* = \Theta(1)\,, \qquad \text{as } n \to +\infty\,, \tag{33}$$

where $\pi(\mathcal{S})$ is the size of the stubborn agents' set, defined in (27).

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

$$\mathbb{P}(Z = x_s) = \overline{\gamma}_s\,, \qquad \overline{\gamma}_s := \sum\nolimits_v \pi_v \gamma_s^v\,, \qquad s \in \mathcal{S}\,. \tag{34}$$

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 shall 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 virtual belief Z by

$$\mathbb{E}[Z] := \sum\nolimits_s \overline{\gamma}_s x_s\,, \qquad \sigma_Z^2 := \sum\nolimits_s \overline{\gamma}_s \left(x_s - \mathbb{E}[Z]\right)^2\,. \tag{35}$$

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

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

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

$$\Delta_* := \max\left\{x_s - x_{s'} : s, s' \in \mathcal{S}\right\}. \tag{36}$$

The next theorem presents the main result of this section.

THEOREM 4. Let Assumptions 1, 2, and 3 hold, and assume that $\pi(\mathcal{S}) \le 1/4$. Then, for all $\varepsilon > 0$,

$$\frac{1}{n}\left|\left\{v : \left|\mathbb{E}[X_v] - \mathbb{E}[Z]\right| \ge \Delta_*\varepsilon\right\}\right| \le \psi(\varepsilon)\frac{\tau\pi(\mathcal{S})}{n\pi_*}\,, \tag{37}$$

[6]With a slight abuse of notation, in the following we shall sometimes refer to a sequence of social networks of increasing population size n simply as a social network.

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

$$\frac{1}{n}\left|\left\{v : \left|\sigma_v^2 - \sigma_Z^2\right| \geq \Delta_*^2\varepsilon\right\}\right| \leq \psi(\varepsilon)\frac{\tau\pi(\mathcal{S})}{n\pi_*}\,. \tag{38}$$

This theorem implies that in large-scale highly fluid social networks, as the population size n grows large, the expected values and variances of ergodic beliefs of regular agents concentrate around fixed values corresponding to the expected virtual belief $\mathbb{E}[Z]$, and, respectively, its variance σ_Z^2. We refer to this as an *approximately equal 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 approximately equal (at least in their first two moments) marginals of the agents' ergodic beliefs. This shows that in highly fluid social networks, most of the regular agents feel the presence of the stubborn agents in approximately the same way.

Intuitively, if the set $\mathcal{S}$ and the mixing time τ are both small, then the influence of the stubborn agents will be felt by most of the regular agents much later 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 4 is proved in Sect. 6.2. Its proof relies on the characterization of the mean ergodic beliefs in terms of the hitting probabilities of the random walk $V(t)$. The definition of highly fluid network implies that the (expected) time it takes $V(t)$ to hit $\mathcal{S}$, when started from most of the nodes of $\mathcal{G}$, is much larger than the mixing time τ. Hence, before hitting $\mathcal{S}$, $V(t)$ looses memory of where it started from, and approaches $\mathcal{S}$ almost as if started from the stationary distribution π.

Before proving Theorem 4, we present some examples of highly fluid social networks in Sect. 6.1.

6.1. *Examples of large-scale highly fluid social networks.* We now present some examples of family of social networks that are highly fluid in the limit of large population size n. All the examples will follow the canonical social network construction of Example 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

$$\pi(\mathcal{S}) = \sum\nolimits_s d_s/(n\bar{d})\,, \tag{39}$$

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. (**Barbell**) For even $n \geq 6$, consider the barbell-like topology introduced in Example 4. 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

$$\Phi_* = \frac{1}{\frac{n}{2}(\frac{n}{2}-1)+1} \leq \frac{4}{(n+1)^2}\,.$$

Using [28, Theorem 7.3], it then follows that

$$\tau \geq \frac{1}{4\Phi_*} \geq \frac{(n+1)^2}{16}\,.$$

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 4 that the expected ergodic beliefs polarize in this case.

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

EXAMPLE 7. (**d-dimensional tori**) Let us consider the case of a d-dimensional torus of size $n = m^d$, introduced in Example 5. 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 $|\mathcal{S}| = o(n)$, then the social network with toroidal topology is highly fluid.[7] 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 3, that the expected asymptotic opinions do not concentrate in this case. Finally, the two-dimensional torus is not highly fluid, hence Theorem 4 is not sufficient to prove that the empirical beliefs concentrate around $\mathbb{E}[Z]$. Nevertheless, one could use the explicit expression (31) 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 ergodic beliefs vector solves

[7] In fact, using Fourier analysis, one may show that $|\mathcal{S}| = o(n^{1/d-1})$ suffices.

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 ergodic 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.

Example 8. (**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 $\delta c \log n \leq d_v \leq 4c \log n$ for each node v [19, Lemma 6.5.2]. In particular, it follows that, with high probability, $(\pi_* n)^{-1} \leq 4/\delta$. Hence, using (39), one finds that the resulting social network is highly fluid, provided that $|\mathcal{S}| = o(n/\log n)$, as n grows large.

Example 9. (**Fixed degree distribution**) Consider a random graph $\mathcal{G} = \mathcal{FD}(n,\lambda)$, with n vertices, whose degree d_v are independent and identically distributed random variables with $\mathbb{P}(d_v = k) = \lambda_k$, for $k \in \mathbb{N}$. We assume that $\lambda_1 = \lambda_2 = 0$, that $\lambda_{2k} > 0$ for some $k \geq 2$, and that the first two moments $\overline{d} := \sum_k \lambda_k k$, 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 (39), one finds that the resulting social network is highly fluid with high probability provided that $\sum_s d_s = o\big(\frac{n}{\log n}\big)$.

Example 10. (**Preferential attachment**) The preferential attachment model was introduced by Barabasi and Albert [8] to model real-world networks which typically exhibit a power law degree distribution. We follow [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 (39), the resulting social network is highly fluid with high probability if $\sum_{s\in\mathcal{S}} d_s = o\big(\frac{n}{\log n}\big)$.

Example 11. (**Watts & Strogatz's small world**) Watts and Strogatz [46], 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 $\{-k, -k+1, \ldots, -1, 1, \ldots, k-1, k\}$, and adding to it a Poisson number of shortcuts with mean pkn, and attaching them to randomly chosen vertices. In this case, the average degree remains bounded with high probability as n grows large, while results by Durrett [19, Th. 6.6.1] show that the mixing time $\tau = O(\log^3 n)$. This, and (39) imply that (33) holds provided that $\sum_{s\in\mathcal{S}} d_s = o\big(\frac{n}{\log^3 n}\big)$.

6.2. *Proof of Theorem 4.* In order to prove Theorem 4, we shall obtain estimates on the hitting probabilities of the random walk. We start by stating a standard result on the distance of transition probability distribution of a random walk from its stationary distribution. For a random walk $V(t)$ on set $\mathcal{V}$, let $q^v(t)$ be its probability distribution at time $t \geq 0$ when started from some $v \in \mathcal{V}$[8], i.e.,

$$q^v_{v'}(t) := \mathbb{P}_v(V(t) = v')\,, \qquad v, v' \in \mathcal{V}\,.$$

Proposition 1. *[3, Ch. 4,Lemma 5] Let $V(t)$ be a random walk on set $\mathcal{V}$. For all $t \geq 0$, we have*

$$\max_{v,v'} \left|\left| q^v(t) - q^{v'}(t) \right|\right|_{TV} \leq \exp(1 - t/\tau)\,,$$

[8]We use here the same notation for the generic random walk as the random walk induced by a social network for convenience.

where τ is the mixing time of the random walk $V(t)$ [cf. Eq. (32)].

The following result, whose proof is an application of Proposition 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 4. *Let Assumptions 1, and 2 hold. Then, for all $t \geq 0$, and $v \in \mathcal{V}$,*

$$||\gamma^v - \overline{\gamma}||_{TV} \leq p_v(t) + \exp(-t/\tau + 1)\,, \tag{40}$$

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

PROOF. Observe that, with no loss of generality, we may assume $\sum_{v'} H_{vv'} = 1$ for all v, since having different rates $\sum_{v'} H_{vv'}$ does not alter the values of the hitting probability distributions γ^v. Then, notice that (40) is trivial when $v \in \mathcal{S}$, for in that case $p_v(t) = 1$.

On the other hand, for every $a \in \mathcal{A}$, $v \in \mathcal{V}$, and $s \in \mathcal{S}$, let us define

$$\tilde{\gamma}^a_s := \mathbb{P}_a\left(V(T_{\mathcal{S}}) = s, T_{\mathcal{S}} > t\right)\,, \qquad \rho^a_s := \mathbb{P}_a(V(T_{\mathcal{S}}) = s|\, T_{\mathcal{S}} \leq t)\,.$$

$$\tilde{q}^a_v := \mathbb{P}_a(V(t) = v, T_{\mathcal{S}} > t)\,, \qquad \chi^a_s := \sum_v (q^a_v - \tilde{q}^a_v)\gamma^s_v / p_a\,.$$

Clearly, ρ^a is a probability measure over $\mathcal{S}$, and the same is true for χ^a, since $\chi^a_v \geq 0$, and

$$\sum\nolimits_s \chi^a_s = \sum\nolimits_v \mathbb{P}_a(V(t) = v|\, T_{\mathcal{S}} \leq t) \sum\nolimits_s \gamma^s_v = 1\,.$$

On the other hand, neither $\tilde{\gamma}^a$ nor $\tilde{q}^a$ are generally probability measures over $\mathcal{V}$. Let $\tilde{T}_{\mathcal{S}} := \inf\{t' \geq t : \, V(t) \in \mathcal{S}\}$. For $s \in \mathcal{S}$, one has

$$\begin{aligned}
\gamma^a_s - p_a(t)\rho^a_s &= \mathbb{P}_a\left(V(T_{\mathcal{S}}) = s, T_{\mathcal{S}} > t\right) \\
&= \textstyle\sum_v \mathbb{P}_a\left(V(T_{\mathcal{S}}) = s, V(t) = v, T_{\mathcal{S}} > t\right) \\
&= \textstyle\sum_v \mathbb{P}_a\left(V(\tilde{T}_{\mathcal{S}}) = s, V(t) = v, T_{\mathcal{S}} > t\right) \\
&= \textstyle\sum_v \tilde{q}^a_v \mathbb{P}_a\left(V(\tilde{T}_{\mathcal{S}}) = s|\, V(t) = v, T_{\mathcal{S}} > t\right) \\
&= \textstyle\sum_v \tilde{q}^a_v \gamma^v_s \\
&= \textstyle\sum_v q^a_v(t)\gamma^v_s - p_a(t)\chi^a_s\,,
\end{aligned}$$

the last equality following from the strong Markov property of $V(t)$. Hence,

$$||\gamma^a - \textstyle\sum_v q^a_v(t)\gamma^v||_{TV} = p_a(t)||\rho^a - \chi^a||_{TV} \leq p_a(t)\,.$$

Using (34), one has that

$$\begin{aligned}||\textstyle\sum_v q^a_v(t)\gamma^v-\overline{\gamma}||_{TV} &= \tfrac{1}{2}\textstyle\sum_s|\sum_v q^a_v(t)\gamma^v_s-\overline{\gamma}_s|\\ &= \tfrac{1}{2}\textstyle\sum_s|\sum_v \gamma^v_s(q^a_v(t)-\pi_v)|\\ &\le \tfrac{1}{2}\textstyle\sum_v\sum_s \gamma^v_s\,|q^a_v(t)-\pi_v|\\ &= \tfrac{1}{2}\textstyle\sum_v |q^a_v(t)-\pi_v|\\ &= ||q^a(t)-\pi||_{TV}\,.\end{aligned}$$

By applying the triangle inequality, the two estimates above, and Proposition 1, one shows that

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

thus proving the claim. ■

Lemma 5, 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 2. *([3, Ch. 3,Prop. 23]) Let $V(t)$ be a continuous-time reversible random walk on $\mathcal{V}$ with irreducible transition probability 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 $\mathcal{S}\subset\mathcal{V}$,*

$$\sup_{t\ge0}|\mathbb{P}_\pi(T_{\mathcal{S}}>t)-\exp(-t/\mathbb{E}_\pi[T_{\mathcal{S}}])|\le\tau_2/\mathbb{E}_\pi[T_{\mathcal{S}}]\,.$$

LEMMA 5. *Let Assumptions 1, 2, and 3 hold. Then, for all $\varepsilon>0$,*

$$\frac{1}{n}\left|\left\{v\in\mathcal{V}:\,||\gamma^v-\overline{\gamma}||_{TV}\ge\varepsilon\right\}\right|\le\frac{2\log(2e^2/\varepsilon)}{\varepsilon}\frac{\tau}{n\pi_*\mathbb{E}_\pi[T_{\mathcal{S}}]}\,. \tag{41}$$

PROOF. From Lemma 4, and Proposition 2, it follows that, for all $t\ge0$,

$$\sum_v\pi_vp_v(t)=\sum_v\pi_v\mathbb{P}_v\left(T_{\mathcal{S}}\le t\right)\le1-\exp\left(-\frac{t}{\mathbb{E}_\pi[T_{\mathcal{S}}]}\right)+\frac{\tau_2}{\mathbb{E}_\pi[T_{\mathcal{S}}]}\le\frac{t+\tau_2}{\mathbb{E}_\pi[T_{\mathcal{S}}]}\,,$$

where the last step follows from the inequality $e^x \leq 1+x$. Hence, Markov's inequality implies that

$$\frac{1}{n}\left|\{v : p_v(t) \geq \varepsilon/2\}\right| \leq \frac{2}{\varepsilon}\sum_v \frac{1}{n}p_v(t) \leq \frac{2}{\varepsilon n \pi_*}\sum_v \pi_v p_v(t) \leq \frac{2}{\varepsilon n\pi_*}\frac{t+\tau_2}{\mathbb{E}_\pi[T_{\mathcal{S}}]}\,. \tag{42}$$

Now, by applying (40) and (42) with $t = \tau\log(2e/\varepsilon)$, and using the inequality $\tau_2 \leq \tau$, one gets

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

which proves the claim. ■

Lemma 5 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 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 3. *[4, Proposition 7.13] For all $\mathcal{S} \subseteq \mathcal{V}$, $\mathbb{E}_\pi[T_{\mathcal{S}}] \geq \frac{1}{2\pi(\mathcal{S})} - \frac{3}{2}$.*

Lemma 5 and Proposition 3 immediately imply the following result:

LEMMA 6. *Let Assumptions 1, 2, and 3 hold, and assume that $\pi(\mathcal{S}) \leq 1/4$. Then, for all $\varepsilon > 0$,*

$$\frac{1}{n}\left|\{v : ||\gamma^v - \overline{\gamma}||_{TV} \geq \varepsilon\}\right| \leq \psi(\varepsilon)\frac{\tau\pi(\mathcal{S})}{n\pi_*}\,,$$

with $\psi(\varepsilon) := \frac{16}{\varepsilon}\log(2e^2/\varepsilon)$.

Proof of Theorem 4:

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 3 that

$$\left|\mathbb{E}[X_v] - \mathbb{E}[Z]\right| = \left|\sum\nolimits_s \gamma_s^v x_s - \sum\nolimits_s \overline{\gamma}_s x_s\right| = \left|\sum\nolimits_s \gamma_s^v y_s - \sum\nolimits_s \overline{\gamma}_s y_s\right| \leq \Delta_*||\gamma^v - \overline{\gamma}||_{TV}\,,$$

so that (37) immediately follows from Lemma 6.

On the other hand, in order to show (38), first recall that, if $\theta_e = 1$ for all $e \in \overrightarrow{\mathcal{E}}$, then Eq. (17) 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 3 that

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

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

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

and (38) follows again from a direct application of Lemma 6.

7. Opinion oscillations and disagreement. We have seen in the previous section that in highly fluid social networks a condition of approximately equal influence is achieved, with the expected values and variances of the ergodic opinions of almost all the agents close to those of the virtual belief. It is worth stressing how the condition of approximately equal influence may significantly differ from an approximate consensus. In fact, the former only involves the (the first and second moments of) the marginal distributions of the agents' ergodic beliefs, and does not have any implication for their joint probability law. A chaotic distribution in which the agents' ergodic beliefs are all mutually independent would be compatible with the condition of approximately equal influence, as well as an approximate consensus condition, which would require the ergodic beliefs of most of the agents to be close to each other with high probability. In this section, under additional assumptions, we show that the ergodic belief distribution in highly fluid social networks is closer to a chaotic distribution than to an approximate consensus. For the sake of simplicity, throughout this section, we restrict our attention to the voter model.

ASSUMPTION 4. *For every $e \in \overrightarrow{\mathcal{E}}$, $\theta_e = 1$.*

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

$$\sigma^2_{\overline{X}} := \mathbb{E}\left[\left(\overline{X} - \mathbb{E}[\overline{X}]\right)^2\right] \tag{43}$$

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

$$\Delta^2 := \frac{1}{2n^2}\sum_{v,v'}\mathbb{E}\left[(X_v - X_{v'})^2\right], \tag{44}$$

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

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

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

$$\frac{\pi(\mathcal{S})}{\pi(\mathcal{D})} = \omega(\tau_2), \qquad n \to +\infty, \tag{45}$$

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

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

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

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

THEOREM 5. For any family of highly fluid social networks, satisfying Assumptions 1-4, with a significant presence of stubborn agents, it holds

$$\sigma^2_{\overline{X}} = o(1)\,, \qquad \Delta^2 = \sigma^2_Z + o(1)\,, \qquad \text{as } n \to +\infty\,.$$

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

An immediate consequence of Theorem 5 is that, if σ^2_Z is bounded away from zero in the large population limit, then so is the mean squared disagreement. Observe that the condition $\sigma^2_Z = o(1)$ is equivalent to the fact that the probability measure $\sum_s \overline{\gamma}_s \delta_{x_s}$ (where $\overline{\gamma}_s$ is defined in (34), and δ_x stands for the Dirac's measure centered in some $x \in \mathbb{R}$) concentrates in one single point. We can think of this as the as the presence of a *dominating stubborn agents' belief*. Hence, we may say that Theorem 5 implies that, on highly fluid social networks with a significant presence of stubborn agents, none of whose beliefs is dominating, a significant disagreement persists in the large population limit.

7.1. *Examples of highly fluid social networks with significant presence of stubborn agents.* Observe that, for the canonical construction of a social network from an indirected graph $\mathcal{G}$, as described in Example 1, one has

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

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

$$\frac{\overline{d^2}}{(\overline{d})^2} = 1 + \frac{1}{n} \sum\nolimits_v \left(\frac{d_v}{\overline{d}} - 1 \right)^2$$

is minimal for regular graphs, where it equals 1, and grows with the normalized variance of the degree distribution. In particular, for a family of social networks with bounded first and second moment of the degree distribution, $\pi(\mathcal{D}) = O(n^{-1})$, so that, in order to have a significant presence of stubborn agents, its is sufficient that $n\pi(\mathcal{S}) = (\overline{d})^{-1} \sum_s d_s$ grows faster than the relaxation time τ_2.

Let us return to the five examples of Sect. 6.1.

Example 12. (**d-dimensional tori**) Let us consider the case of a d-dimensional torus of size n, introduced in Example 5, and discussed in Example 12. Then, $\overline{d^2} = 2^d$, and $\tau_2 \leq \tau = O(n^{2/d})$ for $d \geq 2$. Thus, if $d > 4$, and $|\mathcal{S}|$ grows faster than $n^{2/d}$, and slower than $n^{1-2/d}$, then the associated social network is highly fluid and with a significant presence of stubborn agents.

Example 13. (**Connected Erdös-Renyi**) Consider the Erdös-Renyi random graph $\mathcal{G} = \mathcal{ER}(n, p)$, in the regime $p = cn^{-1} \log n$, with $c > 1$, as in Example 8. Then, with high probability, $\overline{d^2}/(\overline{d})^2 = O(1)$, while $\tau_2 \leq \tau = O(\log n)$. It follows that the associated social network is highly fluid and with a significant presence of stubborn agents provided that $|\mathcal{S}|$ grows faster than $\log n$, and slower than $n/\log n$.

Example 14. (**Fixed degree distribution**) Consider $\mathcal{G} = \mathcal{FD}(n, \lambda)$, as in Example 9. Then, with high probability, since the expected square degree is bounded, one has $\pi(\mathcal{D}) = O(n^{-1})$, while $\tau_2 \leq \tau = O(\log n)$. It follows that the associated social network is highly fluid and with a significant presence of stubborn agents provided that $\sum_s d_s$ grows faster than $\log n$, and slower than $n/\log n$.

Example 15. (**Preferential attachment**) Consider the preferential attachment model of Example 10. Then, with high probability, $\tau_2 \leq \tau = O(\log n)$, while, according to [19, pag. 180], $\pi(\mathcal{D}) \leq n^{-1} \log n$. It follows that the associated social network is highly fluid and with a significant presence of stubborn agents provided that $\sum_s d_s$ grows faster than $\log^2 n$, and slower than $n/\log n$.

Example 16. (**Watts & Strogatz's small world**) For the small-world model of Example 11, one has that both the average degree and the average square degree are bounded, so that $\pi(\mathcal{D}) = O(n^1)$, while $\tau_2 \leq \tau = O(\log^3 n)$, with high probability. This implies that (33) holds provided that $\sum_{s \in \mathcal{S}} d_s$ grows faster than $\log^3 n$ and slower than $n/\log^3 n$.

7.2. *Proof of Theorem 5.* The following result is a consequence of Theorem 4:

LEMMA 7. *For any family of highly fluid social networks, satisfying Assumptions 1-4,*

$$\Delta^2 + \sigma^2_{\overline{X}} = \sigma^2_Z + o(1)\,, \qquad as\ n \to +\infty\,. \tag{46}$$

PROOF. Observe that

$$\begin{aligned}
\Delta^2 &= \tfrac{1}{2n^2}\textstyle\sum_{v,v'} \mathbb{E}\left[(X_v - X_{v'})^2\right] \\
&= \tfrac{1}{2n^2}\textstyle\sum_{v,v'} \left(\mathbb{E}\left[X_v^2\right] + \mathbb{E}\left[X_{v'}^2\right] - 2\mathbb{E}\left[X_v X_{v'}\right]\right) \\
&= \tfrac{1}{2n^2}\textstyle\sum_{v,v'} \big(\mathbb{E}\left[X_v^2\right] + \mathbb{E}\left[X_{v'}^2\right] - \mathbb{E}\left[X_v\right]^2 - \mathbb{E}\left[X_{v'}\right]^2 \\
&\qquad\qquad + \mathbb{E}\left[X_v\right]^2 + \mathbb{E}\left[X_{v'}\right]^2 - 2\mathbb{E}[X_v]\mathbb{E}[X_{v'}] \\
&\qquad\qquad + 2\mathbb{E}[X_v]\mathbb{E}[X_{v'}] - 2\mathbb{E}[X_v X_{v'}]\big) \\
&= \tfrac{1}{2n^2}\textstyle\sum_{v,v'} \left(\sigma_v^2 + \sigma_{v'}^2 + (\mathbb{E}[X_v] - \mathbb{E}[X_{v'}])^2 - 2\,\mathrm{Cov}(X_v, X_{v'})\right) \\
&= n^{-1}\textstyle\sum_v \sigma_v^2 + \tfrac{1}{2}n^{-2}\sum_{v,v'}(\mathbb{E}[X_v] - \mathbb{E}[X_{v'}])^2 - n^{-2}\sum_{v,v'}\mathrm{Cov}(X_v, X_{v'})\,.
\end{aligned} \tag{47}$$

Similarly, one gets that

$$\begin{aligned}
\sigma^2_{\overline{X}} &= \mathbb{E}\left[(n^{-1}\textstyle\sum_v X_v)^2\right] - \left(\mathbb{E}\left[n^{-1}\sum_v X_v\right]\right)^2 \\
&= n^{-2}\textstyle\sum_{v,v'}\left(\mathbb{E}[X_v X_{v'}] - \mathbb{E}[X_v]\mathbb{E}[X_{v'}]\right) \\
&= n^{-2}\textstyle\sum_{v,v'}\mathrm{Cov}(X_v, X_{v'})\,.
\end{aligned} \tag{48}$$

It follows from Theorem 4 that

$$n^{-2}\sum_{v,v'}(\mathbb{E}[X_v] - \mathbb{E}[X_{v'}])^2 = o(1)\,, \qquad n^{-1}\sum_v \sigma_v^2 = \sigma_Z^2 + o(1)\,. \tag{49}$$

Then the claim follows from (47), (48), and (49). ■

Corollary 7 implies that, the sum of the mean squared disagreement Δ^2 and the variance of the ergodic aggregate belief , $\sigma^2_{\overline{X}}$, remains bounded away from 0 in the limit of large population size, provided that there is no dominating stubborn agents' belief. In fact, we now show that $\sigma^2_{\overline{X}}$ vanishes in the large population limit of highly social networks with significant presence of stubborn agents.

To argue that, some considerations are in order on the coupled random walk $(V(t), V'(t))$ of transition rates $K_{(v,v'),(w,w')}$ defined in (17). Recall that, under Assumption 4, such rates reduce to the ones of a pair of coalescing random walks which stick together once they meet, and never separate from each other. It is then of particular interest to consider the diagonal set

$$\mathcal{D} := \{(a,a) :\, a \in \mathcal{A}\} \subseteq \mathcal{V} \times \mathcal{V}\,,$$

and the boundary set

$$\mathcal{B} := (\mathcal{S} \times \mathcal{V}) \cup (\mathcal{V} \times \mathcal{S})\,.$$

Let $\mathcal{C} := \mathcal{B} \cup \mathcal{D}$. Let $T_{\mathcal{D}}, T_{\mathcal{B}}, T_{\mathcal{C}}$ denote the hitting times of the random walk on the sets $\mathcal{D}$, $\mathcal{B}$, and $\mathcal{C}$, respectively. We shall study the probability

$$p_{\mathcal{D}} := \mathbb{P}_\pi(T_{\mathcal{D}} < T_{\mathcal{B}})\,,$$

that, when started from the stationary distribution, $V(t)$ and $V'(t)$ meet before any of them hits the stubborn agents' set. In particular, one has the following result:

LEMMA 8. *For every social network satisfying Assumptions 1–4,*

$$\sigma^2_{X} \leq \frac{\Delta_*^2}{2(n\pi_*)^2} p_{\mathcal{D}}\,.$$

PROOF. Let us consider another pair of random walks on $\mathcal{V}$, $\tilde{V}(t)$ and $\tilde{V}'(t)$, such that $\tilde{V}(t) = V(t)$, and $\tilde{V}'(t) = V'(t)$, for all $t \leq T_{\mathcal{C}}$. If $T_{\mathcal{C}} = T_{\mathcal{D}}$, then, for $t \geq T_{\mathcal{C}}$, $\tilde{V}(t)$ and $\tilde{V}'(t)$ continue to move on $\mathcal{V}$ with transition rates $P_{vv'}$, independent of each other, and independent from $V(t)$ and $V'(t)$. Otherwise, if $T_{\mathcal{C}} = T_{\mathcal{S}}$, then, for all $t \geq T_{\mathcal{C}}$, $\tilde{V}'(t) = V'(t)$, while $\tilde{V}(t)$ continues to move on $\mathcal{V}$ with transition rates $P_{vv'}$, independently from $V(t)$, and $V'(t)$. In the symmetric case when $T_{\mathcal{C}} = T'_{\mathcal{S}}$, for all $t \geq T_{\mathcal{C}}$, $\tilde{V}(t) = V(t)$, while $\tilde{V}'(t)$ continues to move on $\mathcal{V}$ with transition rates $P_{vv'}$, independently from $V(t)$, and $V'(t)$. In particular, $(\tilde{V}(t), \tilde{V}'(t))$ is a pair of independent random walks on $\mathcal{V}$ both with transition rates $P_{vv'}$.

Observe that, if $T_{\mathcal{C}} = T_{\mathcal{B}}$, then either $T_{\mathcal{C}} = T_{\mathcal{S}}$, or $T_{\mathcal{C}} = T'_{\mathcal{S}}$. In both cases, it is not hard to verify that $\tilde{V}(\tilde{T}_{\mathcal{S}}) = V(T_{\mathcal{S}})$, and $\tilde{V}'(\tilde{T}'_{\mathcal{S}}) = V'(T'_{\mathcal{S}})$. Now, if $V(T_{\mathcal{S}}) = s$, and $V'(T'_{\mathcal{S}}) = s'$, for some $s \neq s'$, then necessarily $V(t)$ and $V'(t)$ have not coalesced before hitting $\mathcal{B}$, i.e., $T_{\mathcal{C}} = T_{\mathcal{B}}$, and hence $\tilde{V}(\tilde{T}_{\mathcal{S}}) = V(T_{\mathcal{S}})$, and $\tilde{V}'(\tilde{T}'_{\mathcal{S}}) = V'(T'_{\mathcal{S}})$. Let $\zeta_a^{vv'} := \mathbb{P}_{vv'}(V(T_{\mathcal{D}}) = V'(T_{\mathcal{D}}) = a|\, T_{\mathcal{D}} < T_{\mathcal{B}})$, and

$p_{vv'} := \mathbb{P}_{v,v'}(T_{\mathcal{D}} < T_{\mathcal{B}})$. It follows that, if $s \neq s'$, then

$$
\begin{aligned}
\gamma^v_s \gamma^{v'}_{s'} - \eta^{vv'}_{ss'} &= \mathbb{P}_{vv'}\left(\tilde{V}(\tilde{T}_{\mathcal{S}}) = s, \tilde{V}'(\tilde{T}'_{\mathcal{S}}) = s'\right) - \mathbb{P}_{vv'}\left(V(T_{\mathcal{S}}) = s, V'(T'_{\mathcal{S}}) = s'\right) \\
&= \mathbb{P}_{vv'}\left(\tilde{V}(\tilde{T}_{\mathcal{S}}) = s, \tilde{V}'(\tilde{T}'_{\mathcal{S}}) = s'\right) \\
&\quad -\mathbb{P}_{vv'}\left(\tilde{V}(\tilde{T}_{\mathcal{S}}) = V(T_{\mathcal{S}}) = s, \tilde{V}'(\tilde{T}'_{\mathcal{S}}) = V'(T'_{\mathcal{S}}) = s'\right) \\
&= \mathbb{P}_{vv'}\left(T_{\mathcal{C}} = T_{\mathcal{D}}, \tilde{V}(\tilde{T}_{\mathcal{S}}) = s, \tilde{V}'(T'_{\mathcal{S}}) = s', (V(T_{\mathcal{S}}), V'(T'_{\mathcal{S}})) \neq (s, s')\right) \\
&= \mathbb{P}_{vv'}\left(T_{\mathcal{C}} = T_{\mathcal{D}}, \tilde{V}(\tilde{T}_{\mathcal{S}}) = s, \tilde{V}'(T'_{\mathcal{S}}) = s'\right) \\
&= \textstyle\sum_a p_{vv'} \zeta^{vv'}_a \gamma^a_s \gamma^a_{s'} \,.
\end{aligned}
$$

On the other hand,

$$
\begin{aligned}
\eta^{vv'}_{ss} - \gamma^v_s \gamma^{v'}_s &= \mathbb{P}_{vv'}\left(V(T_{\mathcal{S}}) = V'(T'_{\mathcal{S}}) = s\right) - \mathbb{P}_{vv'}\left(\tilde{V}(\tilde{T}_{\mathcal{S}}) = \tilde{V}'(\tilde{T}'_{\mathcal{S}}) = s\right) \\
&= \mathbb{P}_{vv'}\left(V(T_{\mathcal{S}}) = V'(T'_{\mathcal{S}}) = s, T_{\mathcal{C}} = T_{\mathcal{D}}\right) \\
&\quad -\mathbb{P}_{vv'}\left(\tilde{V}(\tilde{T}_{\mathcal{S}}) = \tilde{V}'(\tilde{T}'_{\mathcal{S}}) = s, T_{\mathcal{C}} = T_{\mathcal{D}}\right) \\
&\quad +\mathbb{P}_{vv'}\left(V(T_{\mathcal{S}}) = V'(T'_{\mathcal{S}}) = s, T_{\mathcal{C}} = T_{\mathcal{B}}\right) \\
&\quad -\mathbb{P}_{vv'}\left(\tilde{V}(\tilde{T}_{\mathcal{S}}) = \tilde{V}'(\tilde{T}'_{\mathcal{S}}) = s, T_{\mathcal{C}} = T_{\mathcal{B}}\right) \\
&= \mathbb{P}_{vv'}\left(T_{\mathcal{C}} = T_{\mathcal{D}}, V(T_{\mathcal{S}}) = V'(T'_{\mathcal{S}}) = s, (\tilde{V}(\tilde{T}_{\mathcal{S}}), \tilde{V}'(\tilde{T}'_{\mathcal{S}})) \neq (s, s)\right) \\
&= \textstyle\sum_a p_{vv'} \zeta^{vv'}_a \left(\gamma^a_s - \gamma^a_s \gamma^a_s\right) \,.
\end{aligned}
$$

It follows that

$$
\begin{aligned}
\mathrm{Cov}\left(X_v, X_{v'}\right) &= \textstyle\sum_{s,s'} \left(\eta^{vv'}_{ss'} - \gamma^v_s \gamma^{v'}_{s'}\right) x_s x_{s'} \\
&= p_{vv'} \textstyle\sum_a \zeta^{vv'}_a \left(\sum_s \gamma^i_s x^2_s - \sum_{s,s'} \gamma^a_s \gamma^a_{s'} x_s x_{s'}\right) \\
&= p_{vv'} \textstyle\sum_a \zeta^{vv'}_a \sigma^2_a \\
&\leq p_{vv'} \textstyle\sum_a \zeta^{vv'}_a \frac{1}{2} \Delta^2_* \\
&\leq p_{vv'} \tfrac{1}{2} \Delta^2_* \,.
\end{aligned}
$$

Finally, one has that

$$
\begin{aligned}
\sigma^2_{\overline{X}} &= n^{-2} \textstyle\sum_{v,v'} \mathrm{Cov}\left(X_v, X_{v'}\right) \\
&\leq \tfrac{1}{2} \Delta^2_* \textstyle\sum_{v,v'} \pi_v \pi_{v'} \dfrac{1}{n^2 \pi_v \pi_{v'}} p_{vv'} \\
&\leq \tfrac{1}{2} \Delta^2_* (n\pi_*)^{-2} \textstyle\sum_{v,v'} \pi_v \pi_{v'} p_{vv'} \\
&= \tfrac{1}{2} \Delta^2_* (n\pi_*)^{-2} p_{\mathcal{D}} \,,
\end{aligned}
$$

which completes the proof. ∎

The probability $p_{\mathcal{D}}$ can in turn be upper-bounded using the approximate exponentiality of the hitting times.

LEMMA 9. *For coalescing random walks with a reversible irreducible transition probability matrix P,*

$$p_{\mathcal{D}} \leq \frac{\mathbb{E}_\pi[T_{\mathcal{S}}]}{\mathbb{E}_\pi[T_{\mathcal{D}}]} \log \frac{e\mathbb{E}_\pi[T_{\mathcal{D}}]}{\mathbb{E}_\pi[T_{\mathcal{S}}]} + \frac{\tau_2}{\mathbb{E}_\pi[T_{\mathcal{D}}]} + \frac{\tau_2}{\mathbb{E}_\pi[T_{\mathcal{S}}]} \,.$$

PROOF. For every $t \geq 0$, one has that

$$\begin{aligned}
p_{\mathcal{D}} &= \mathbb{P}_\pi(T_{\mathcal{D}} < T_{\mathcal{B}}) \\
&\leq \mathbb{P}_\pi(\{T_{\mathcal{D}} \leq t\} \cup \{T_{\mathcal{B}} > t\}) \\
&\leq \mathbb{P}_\pi(T_{\mathcal{D}} \leq t) + \mathbb{P}_\pi(T_{\mathcal{B}} > t) \\
&\leq \mathbb{P}_\pi(T_{\mathcal{D}} \leq t) + \mathbb{P}_\pi(T_{\mathcal{S}} > t) \\
&\leq 1 - e^{-t/\mathbb{E}_\pi[T_{\mathcal{D}}]} + e^{-t/\mathbb{E}_\pi[T_{\mathcal{S}}]} + \frac{\tau_2}{\mathbb{E}_\pi[T_{\mathcal{D}}]} + \frac{\tau_2}{\mathbb{E}_\pi[T_{\mathcal{S}}]} \\
&\leq \frac{t}{\mathbb{E}_\pi[T_{\mathcal{D}}]} + e^{-t/\mathbb{E}_\pi[T_{\mathcal{S}}]} + \frac{\tau_2}{\mathbb{E}_\pi[T_{\mathcal{D}}]} + \frac{\tau_2}{\mathbb{E}_\pi[T_{\mathcal{S}}]} \,,
\end{aligned}$$

the forth inequality following from, Proposition 2, the last one from the inequality $e^{-1} \leq 1-x$. With the optimal choice $t = \mathbb{E}_\pi[T_{\mathcal{S}}] \log \left(\mathbb{E}_\pi[T_{\mathcal{D}}]/\mathbb{E}_\pi[T_{\mathcal{S}}]\right)$, the foregoing gives the claim. ∎

In order to apply Lemma 9, one needs an upper bound on the ratio $\mathbb{E}_\pi[T_{\mathcal{S}}]/\mathbb{E}_\pi[T_{\mathcal{D}}]$. In the absence of more precise information about these expected hitting times, one can estimate $\mathbb{E}_\pi[T_{\mathcal{D}}]$ from below using Proposition 3. On the other hand the following general upper bound on $\mathbb{E}_\pi[T_{\mathcal{S}}]$ can be applied.

PROPOSITION 4. *[3, Ch. 3, Prop. 21] Let $V(t)$ be a continuous-time reversible random walk on $\mathcal{V}$ with irreducible transition probability matrix P, and stationary distribution π. Then,*

$$\mathbb{E}_\pi[T_{\mathcal{S}}] \leq \frac{1-\pi(\mathcal{S})}{\pi(\mathcal{S})}\tau_2$$

Combining Lemmas 10 and 9 with Propositions 3 and 4, one obtains the following:

LEMMA 10. *Consider a social network satisfying Assumptions 1-4. Assume that $\pi(\mathcal{S}) \leq 1/4$, and $8\pi(\mathcal{D})\tau_2 \leq \pi(\mathcal{S})$. Then,*

$$\sigma^2_{\overline{X}} \leq \frac{\Delta_*^2}{2(n\pi_*)^2}\left(\frac{8\pi(\mathcal{D})\tau_2}{\pi(\mathcal{S})}\log\frac{e\pi(\mathcal{S})}{8\pi(\mathcal{D})\tau_2} + 8\tau_2\pi(\mathcal{S}) + 8\tau_2\pi(\mathcal{D})\right)$$

PROOF. From Proposition 3, and the assumption $\pi(\mathcal{S}) \leq 1/4$, one gets

$$(\mathbb{E}_\pi[T_\mathcal{S}])^{-1} \leq \frac{2\pi(\mathcal{S})}{1-3\pi(\mathcal{S})} \leq 8\pi(\mathcal{S})\,. \tag{50}$$

Moreover, one has $\pi(\mathcal{D}) \leq \pi(\mathcal{S})/(8\tau_2) \leq \pi(\mathcal{S})/8 \leq 32$, and, arguing as above, $(\mathbb{E}_\pi[T_\mathcal{D}])^{-1} \leq 8\pi(\mathcal{D})$. From this inequality, and Proposition 4, one finds that

$$\frac{\mathbb{E}_\pi[T_\mathcal{S}]}{\mathbb{E}_\pi[T_\mathcal{D}]} \leq \frac{8\tau_2\pi(\mathcal{D})}{\pi(\mathcal{S})} \leq 1\,. \tag{51}$$

Then, Lemmas 10 and 9, together with (50) and (51), imply that

$$\sigma^2_{\overline{X}} \leq \frac{\Delta_*^2}{2(n\pi_*)^2}p_\mathcal{D} \leq \frac{\Delta_*^2}{2(n\pi_*)^2}\left(\frac{8\pi(\mathcal{D})\tau_2}{\pi(\mathcal{S})}\log\frac{e\pi(\mathcal{S})}{8\pi(\mathcal{D})\tau_2} + 8\tau_2\pi(\mathcal{S}) + 8\tau_2\pi(\mathcal{D})\right),$$

thus proving the claim. ∎

Now, it is easily seen, using Lemma 10, that $\sigma^2_{\overline{X}} = o(1)$ in the large population limit of a family of highly fluid social networks with a significant presence of stubborn agents. From this, and Lemma 7, one gets that $\Delta^2 = \sigma^2_Z + o(1)$, and Theorem 5 follows.

8. Conclusion.

In this paper, we have studied a possible mechanism explaining persistent disagreement and opinion fluctuations in social networks. We have considered a stochastic gossip model of continuous opinion dynamics, combined with the assumption that there are some stubborn agents in the network who never change their opinions. We have shown that the presence of these stubborn agents leads to persistent oscillations and disagreements among the rest of the society: the beliefs of regular agents do not converge almost surely, and keep on oscillating according to an ergodic distribution. First and second moments of the ergodic beliefs distribution can be characterized in terms of the hitting probabilities of a random walk on the network, while the correlation between the ergodic beliefs of any pair of regular agents can be characterized in terms of the hitting probabilities of a pair of coupled random walks. We have shown that in highly fluid, reversible social networks, whose associated random walks have mixing times which are sufficiently smaller than the inverse of the stubborn agents' set size, the vectors of the expected ergodic beliefs and of the ergodic variances are almost constant, so that the stubborn agents have approximately the same influence on the society. Finally, we have also shown that in highly fluid social networks in which there is a significant presence of stubborn agents, the variance of the ergodic aggregate belief of the system vanishes in the limit of large population size, and the ergodic distribution of the agents beliefs approaches an approximately chaotic condition. This implies that, if the influence of any of the stubborn agents' opinions does not dominate the influence of the rest, then the mean square disagreement does not vanish in the large population size. We conjecture that, in highly fluid social networks without a significant presence of stubborn agents, i.e., with $\pi(\mathcal{S})$ and $\pi(\mathcal{D})$ of the same asymptotic order, an intermediate condition between approximate consensus and chaotic ergodic belief distribution should emerge in the large population limit.

REFERENCES

[1] Acemoglu, D., Dahleh, M., Lobel, I. and Ozdaglar, A. (2008). Bayesian Learning in Social Networks. Preprint.

[2] Acemoglu, D., Ozdaglar, A. and ParandehGheibi, A. (2009). Spread of (Mis)Information in Social Networks. LIDS report 2812, forthcoming in Games and Economic Behavior.

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

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

[5] Axelrod, R. (1997). The dissemination of culture. *Journal of Conflict Resolution* **42** 203-226.

[6] Bala, V. and Goyal, S. (1998). Learning from Neighbours. *Review of Economic Studies* **65** 595–621.

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

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

[9] Blondel, V. D., Hendickx, J. and Tsitsiklis, J. N. (2009). On Krause's consensus formation model with state-dependent connectivity. *IEEE Transactions on Automatic Control* **54** 2586–2597.

[10] Breiman, L. (1960). The Strong Law of Large Numbers for a Class of Markov Chains. *The Annals of Mathematical Statistics* **31** 801-803.

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

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

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

[14] Cooper, C. and Frieze, A. (2006). The cover time of sparse random graphs. *Random Structures and Algorithms* **30** 1-16.

[15] Cox, J. T. (1989). Coalescing random walks and voter model consensus times on the thorus in $\mathbb{Z}_d$. *The Annals of Probability* **17** 1333-1366.

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

[17] DeMarzo, P. M., Vayanos, D. and Zwiebel, J. (2003). Persuasion Bias, Social Influence, and Unidimensional Opinions. *The Quarterly Journal of Economics* **118** 909–968.

[18] Diaconis, P. and Friedman, D. (1999). Iterated Random Functions. *SIAM Review* **41** 45-76.

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

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

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

[22] Fagnani, F. and Zampieri, S. (2008). Randomized consensus algorithms over large scale networks. *IEEE Journal on Selected Areas of Communications* **26** 634–649.

[23] Gale, D. and Kariv, S. (2003). Bayesian Learning in Social Networks. *Games and Economic Behavior* **45** 329–346.

[24] Golub, B. and Jackson, M. O. (2010). Naive Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics* **2** 112–149.

[25] Holley, R. A. and Liggett, T. M. (1975). Ergodic Theorems for Weakly Interacting Infinite Systems and The Voter Model. *The Annals of Probability* **3** 643-663.

[26] Jadbabaie, A., Lin, J. and Morse, S. (2003). Coordination of Groups of Mobile Autonomous Agents using Nearest Neighbor Rules. *IEEE Transactions on Automatic Control* **48** 988–1001.

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

[28] Levin, D. A., Perez, Y. and Wilmer, E. L. (2010). *Markov Chains and Mixing Times*. American Mathematical Society.

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

[30] Liggett, T. M. (1999). *Stochastic interacting systems: Contact, voter, and exclusion processes*. Springer Berlin.

[31] Lorenz, J. (2005). A stabilization theorem for continuous opinion dynamics. *Physica A* **355** 217–223.

[32] Mihail, M., Papadimitriou, C. and Saberi, A. (2004). On Certain Connectivity Properties of the Internet Topology. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science 2003*.

[33] Mobilia, M. (2003). Does a Single Zealot Affect an Infinite Group of Voters? *Physical Review Letters* **91** 028701.

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

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

[36] Nedić, A. and Ozdaglar, A. (2009). Distributed Subgradient Methods for Multi-Agent Optimization. *IEEE Transactions on Automatic Control* **54** 48–61.

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

[38] Olfati-Saber, R. and Murray, R. M. (2004). Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control* **49** 1520–1533.

[39] Olshevsky, A. and Tsitsiklis, J. N. (2008). Convergence Speed in Distributed Consensus and Averaging. forthcoming in SIAM Journal on Control and Optimization.

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

[41] Rao, M. M. and Swift, R. J. (2006). *Probability Theory with Applications*. Springer.

[42] Smith, L. and Sorensen, P. (2000). Pathological Outcomes of Observational Learning. *Econometrica* **68** 371–398.

[43] Tsitsiklis, J. N. (1984). Problems in Decentralized Decision Making and Computation PhD thesis, Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology.

[44] Tsitsiklis, J. N., Bertsekas, D. P. and Athans, M. (1986). Distributed Asynchronous Deterministic and Stochastic Gradient Optimization Algorithms. *IEEE Transactions on Automatic Control* **31** 803–812.

[45] Villani, C. (2009). *Optimal transport: old and new*. Springer-Verlag Berlin-Heidelberg.

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

Department of Economics,
Massachusetts Institute of Technology,
77 Massachusetts Ave, Cambridge, MA 02139, US
E-mail: daron@mit.edu

Laboratory for Information and Decision Systems,
Massachusetts Institute of Technology,
77 Massachusetts Ave, Cambridge, MA 02139, US
E-mail: giacomo@mit.edu

Dipartimento di Matematica,
Politecnico di Torino,
Corso Stati Uniti 24, 10129, Torino, Italy
E-mail: fabio.fagnani@polito.it

Laboratory for Information and Decision Systems,
Massachusetts Institute of Technology,
77 Massachusetts Ave, Cambridge, MA 02139, US
E-mail: asuman@mit.edu

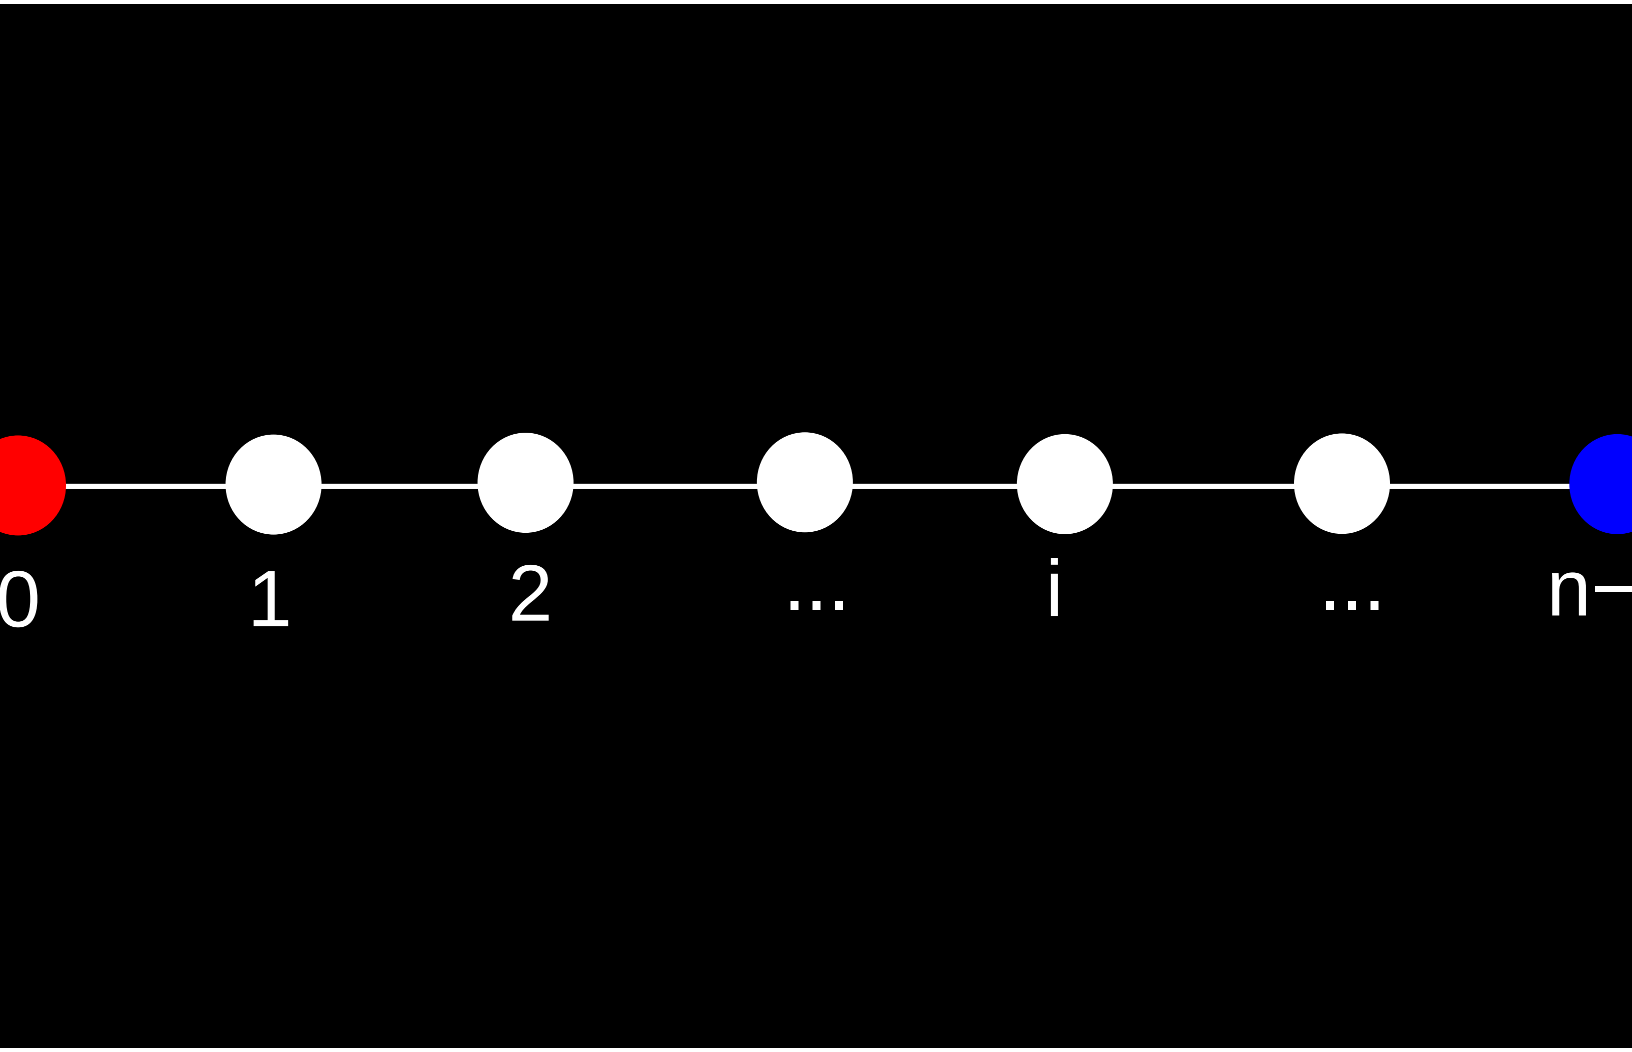
0
1
2
...
i
...
n−