Opinion fluctuations and disagreement in social networks

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

Giacomo Como
Department of Automatic Control, Lund University, BOX 118, 22100, Lund, Sweden, giacomo.como@control.lth.se, http://www.control.lth.se/Staff/GiacomoComo/

Fabio Fagnani
Dipartimento di Scienze Matematiche, Politecnico di Torino, corso Stati Uniti 24, 10129, Torino, Italy, fabio.fagnani@polito.it, http://calvino.polito.it/ fagnani/indexeng.html

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

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

The structure of the graph describing the social network and the location of the stubborn agents within it shape the opinion dynamics. The expected belief vector is proved to evolve according to an ordinary differential equation coinciding with the Kolmogorov backward equation of a continuous-time Markov chain on the graph with absorbing states corresponding to the stubborn agents, and hence to converge to a harmonic vector, with every regular agent's value being the weighted average of its neighbors' values, and boundary conditions corresponding to the stubborn agents' beliefs. Expected cross-products of the agents' beliefs allow for a similar characterization in terms of coupled Markov chains on the graph describing the social network.

We prove that, in large-scale societies which are *highly fluid*, meaning that the product of the mixing time of the Markov chain on the graph describing the social network and the relative size of the linkages to stubborn agents vanishes as the population size grows large, a condition of *homogeneous influence* emerges, whereby the stationary beliefs' marginal distributions of most of the regular agents have approximately equal first and second moment.

Key words: opinion dynamics, multi-agent systems, social networks, persistent disagreement, opinion fluctuations, social influence.
MSC2000 subject classification: Primary: 91D30, 60K35; Secondary: 93A15
OR/MS subject classification: Primary: Games/group decisions: stochastic; Secondary: Markov processes, Random walk
History:

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

Existing models of communication and learning, based on Bayesian or non-Bayesian updating mechanisms, typically lead to consensus provided that communication takes place over a strongly connected network (e.g., Smith and Sorensen [46], Banerjee and Fudenberg [7], Acemoglu, Dahleh, Lobel and Ozdaglar [2], Bala and Goyal [6], Gale and Kariv [25], DeMarzo, Vayanos and Zwiebel [19], Golub and Jackson [26], Acemoglu, Ozdaglar and ParandehGheibi [3], Acemoglu, Bimpikis and Ozdaglar [1]), and are thus unable to explain persistent disagreements. One notable exception is provided by models that incorporate a form of *homophily* mechanism in communication, whereby individuals are more likely to exchange opinions or communicate with others that have similar beliefs, and fail to interact with agents whose beliefs differ from theirs by more than some given confidence threshold. This mechanism was first proposed by Axelrod [5] in the discrete opinion dynamics setting, and then by Krause [30], and Deffuant and Weisbuch [18], in the continuous opinion dynamics framework. Such belief dynamics typically lead to the emergence of different asymptotic opinion clusters (see, e.g., [34, 10, 14]), but fail to explain persistent opinion fluctuations in the society, as well as the role of influential agents in the opinion formation process. In fact, the latter phenomena have been empirically observed and reported in the social science literature, see, e.g., the stream of work originated with Kramer's paper [29] documenting large swings in voting behavior within short periods, and the sizable literature in social psychology (e.g., Cohen [13]) documenting changes in political beliefs as a result of parties or other influential organizations.

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

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

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

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

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

Our main contribution partly stems from novel applications of several techniques of applied probability in the study of opinion dynamics. In particular, convergence in law and ergodicity of the agents' beliefs is established by first rewriting the dynamics in the form of an iterated affine function system and then proving almost sure convergence of the time-reversed process [20]. On the other hand, our estimates of the behavior of the expected values and variances of the stationary beliefs in large-scale highly fluid networks are based on techniques from the theory of Markov chains and mixing times [4, 31], as well as on results in modern random graph theory [21].

In addition to the aforementioned works on learning and opinion dynamics, this paper is related to some of the literature in the statistical physics of social dynamics: see [11] and references therein for an overview of such research line. More specifically, our model is closely related to some work by Mobilia and co-authors [36, 37, 38], who study a variation of the discrete opinion dynamics model, also called the *voter model*, with inhomogeneities, there referred to as *zealots*: such zealots are agents which tend to favor one opinion in [36, 37], or are in fact equivalent to our stubborn agents in [38]. These works generally present analytical results for some regular graphical structures (such as regular lattices [36, 37], or complete graphs [38]), and are then complemented by numerical simulations. In contrast, we prove convergence in distribution and characterize the properties of the limiting distribution for general finite graphs. Even though our model involves continuous belief dynamics, we will also show that the voter model with zealots of [38] 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 [48], Tsitsiklis, Bertsekas and Athans [49], Jadbabaie, Lin and Morse [28], Olfati-Saber and Murray [42], Olshevsky and Tsitsiklis [43], Fagnani and Zampieri [24], Nedić and Ozdaglar [39]). In consensus problems, the focus is on whether the beliefs or the values held by different units (which might correspond to individuals, sensors, or distributed processors) converge to a common value. Our analysis here does not focus on limiting consensus of values, but in contrast, characterizes the stationary fluctuations in values.

The rest of this paper is organized as follows: In Section 2, we introduce our model of interaction between the agents, describing the resulting evolution of individual beliefs, and we discuss two special cases, in which the arguments simplify particularly, and some fundamental features of the general case are highlighted. Section 3 presents convergence results on the evolution of agent beliefs over time, for a given social network: the beliefs are shown to converge in distribution, and to be an ergodic process, while in general they do not converge almost surely. Section 4 presents a characterization of the first and second moments of the stationary beliefs in terms of the hitting probabilities of two coupled Markov chains on the graph describing the social network. Section 5 presents explicit computations of the expected stationary beliefs and variances for some special network topologies. Section 6 provides bounds on the level of dispersion of the first two moments of the stationary beliefs: it is shown that, in highly fluid networks, most of the agents have almost the same stationary expected belief and variance. Section 7 presents some concluding remarks.

Basic Notation and Terminology We will typically label the entries of vectors by elements of finite alphabets, rather than non-negative integers, hence $\mathbb{R}^{\mathcal{I}}$ will stand for the set of vectors with entries labeled by elements of the finite alphabet $\mathcal{I}$. An index denoted by a lower-case letter will implicitly be assumed to run over the finite alphabet denoted by the corresponding calligraphic upper-case letter (e.g. $\sum_i$ will stand for $\sum_{i \in \mathcal{I}}$). For any finite set $\mathcal{J}$, we use the notation $\mathbb{1}_{\mathcal{J}}$ to denote the indicator function over the set $\mathcal{J}$, i.e., $\mathbb{1}_{\mathcal{J}}(j)$ is equal to 1 if $j \in \mathcal{J}$, and equal to 0 otherwise. For a matrix $M \in \mathbb{R}^{\mathcal{I} \times \mathcal{J}}$, $M^{\mathrm{T}} \in \mathbb{R}^{\mathcal{J} \times \mathcal{I}}$ will stand for its transpose, $||M||$ for its 2-norm. For a probability distribution μ over a finite set $\mathcal{I}$, and a subset $\mathcal{J} \subseteq \mathcal{I}$ we will write $\mu(\mathcal{J}) := \sum_j \mu_j$. If ν is another probability distribution on $\mathcal{I}$, we will use the notation $||\mu - \nu||_{TV} := \frac{1}{2} \sum_i |\mu_i - \nu_i| = \sup \{|\mu(\mathcal{J}) - \nu(\mathcal{J})| : \mathcal{J} \subseteq \mathcal{I}\}$, for the total variation distance between μ and ν. The probability law (or distribution) of a random variable Z will be denoted by $\mathcal{L}(Z)$. Continuous-time Markov chains on a finite set $\mathcal{V}$ will be characterized by their transition rate matrix $M \in \mathbb{R}^{\mathcal{V} \times \mathcal{V}}$, which has zero row sums, and whose non-diagonal elements are nonnegative and correspond to the rates at which the chain jumps from a state to another (see [41, Ch.s 2-3]). If $V(t)$ and $V'(t)$ are Markov chains on $\mathcal{V}$, defined on the same probability space, we will use the notation $\mathbb{P}_v(\cdot)$, and $\mathbb{P}_{vv'}(\cdot)$, for the conditional probability measures given the events $V(0) = v$, and, respectively, $(V(0), V'(0)) = (v, v')$. Similarly, for 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, $a_n \le K b_n$ for all sufficiently large n, $a_n = \Theta(b_n)$ if $b_n = O(a_n)$, $a_n = o(b_n)$ if $\lim_n a_n / b_n = 0$.

2. Belief evolution model We consider a finite population $\mathcal{V}$ of interacting agents, of possibly very large size $n := |\mathcal{V}|$. The connectivity among the agents is described by a simple directed graph $\overrightarrow{\mathcal{G}} = (\mathcal{V}, \overrightarrow{\mathcal{E}})$, whose node set is identified with the agent population, and where $\overrightarrow{\mathcal{E}} \subseteq \mathcal{V} \times \mathcal{V} \setminus \mathcal{D}$, with $\mathcal{D} := \{(v, v) : v \in \mathcal{V}\}$, stands for the set of directed links among the agents.[1]

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

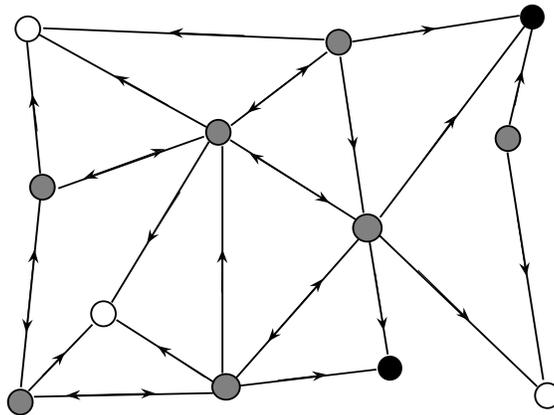[1] Notice that we don't allows for parallel links or loops.

FIGURE 1. A social network with seven regular agents (colored in grey), and five stubborn agents (colored in white, and black, respectively). The presence of a directed link (v, v') indicates that agent v is influenced by the opinion of agent v'. Therefore, links are only incoming to the stubborn agents, while links between pairs of regular agents may be uni- or bi-directional.

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:

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

In contrast, the beliefs of the regular agents are updated as follows. To every directed link 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 link (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 $(v, v') \notin \overrightarrow{\mathcal{E}}$ (hence, including loops $v' = v$). For every regular agent $a \in \mathcal{A}$, let $\mathcal{S}_a \subseteq \mathcal{S}$ be the subset of stubborn agents which are reachable from a by a directed path in $\overrightarrow{\mathcal{G}}$. We refer to $\mathcal{S}_a$ as the set of stubborn agents *influencing* a. For every stubborn agent $s \in \mathcal{S}$, $\mathcal{A}_s := \{a : s \in \mathcal{S}_a\} \subseteq \mathcal{A}$ will stand for the set of regular agents *influenced* by s.

The tuple $\mathcal{N} = \left(\overrightarrow{\mathcal{G}}, \{\theta_e\}, \{r_e\}\right)$ contains the entire information about patterns of interaction among the agents, and will be referred to as the *social network*. Together with an assignment of a probability law for the initial belief vector, $\mathcal{L}(X(0))$, the social network designates a *society*. Throughout the paper, we make the following assumptions regarding the underlying social network.

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

---

[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.

Assumption 1 may be easily removed. If there are some regular agents which are not influenced by any stubborn agent, then there is no link 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 [24] (see Example 3.5 therein), showing that, with probability one, a consensus on a random belief is achieved on every such communicating class.

We denote the *total meeting rate of agent* $v \in \mathcal{V}$ by r_v, i.e., $r_v := \sum_{v'} r_{vv'}$, and the *total meeting rate of all agents* by r, i.e., $r := \sum_v r_v$. We use $N(t)$ to denote the *total number of agent meetings (or link 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., $T_{(k)} := \inf\{t \geq 0 : N(t) \geq k\}$.

To a given social network, we associate the matrix $Q \in \mathbb{R}^{\mathcal{V} \times \mathcal{V}}$, with entries

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

In the rest of the paper, we will often consider a continuous-time Markov chain $V(t)$ on $\mathcal{V}$ with transition rate matrix Q.

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 a connected multigraph,[3] 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 $a \in \mathcal{A}$, $v \in \mathcal{V} \setminus \{a\}$, and $\{a, v\} \in \mathcal{E}$, i.e., $\overrightarrow{\mathcal{G}}$ is the directed graph obtained by making all links in $\mathcal{E}$ bidirectional except links between a regular and a stubborn agent, which are unidirectional (pointing from the regular agent to the stubborn agent). For $v, w \in \mathcal{V}$, let $\kappa_{v,w}$ denote the multiplicity of the link $\{v, w\}$ in $\mathcal{E}$ (each self-loop contributing as 2), and let $d_v = \sum_w \kappa_{v,w}$ be the degree of node v in $\mathcal{G}$. (In particular, $\kappa_{a,v} = \mathbb{1}_{\mathcal{E}}(\{a, v\})$ if $\mathcal{G}$ is a simple graph, i.e., it does not contain neither loops nor parallel links.) Let the trust parameter be constant, i.e., $\theta_{av} = \theta \in (0, 1]$ for all $(a, v) \in \overrightarrow{\mathcal{E}}$. Define

$$r_{av} = d_a^{-1} \kappa_{a,v} \mathbb{1}_{\mathcal{V} \setminus \{a\}}(v), \qquad a \in \mathcal{A}, \ v \in \mathcal{V}. \tag{3}$$

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

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

Observe that connectedness of $\mathcal{G}$ implies that Assumption 1 holds. Finally, notice that nothing prevents the multigraph $\mathcal{G}$ from having (possibly parallel) links between two nodes both in $\mathcal{S}$. However, such links do not have any correspondence in the directed graph $\overrightarrow{\mathcal{G}}$, and in fact they are irrelevant for the belief dynamics, since stubborn agents do not update their beliefs.

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

2.1. Single regular agent Consider a society consisting of a single regular agent, i.e., $\mathcal{A} = \{a\}$, and two stubborn agents, $\mathcal{S} = \{s_0, s_1\}$ (see Fig. 2(a)). Assume that $r_{as_0} = r_{as_1} = 1/2$, $\theta_{as_0} = \theta_{as_1} = 1/2$, $x_{s_0} = 0$, $x_{s_1} = 1$, and $X_a(0) = 0$. Then, one has for all $t \geq 0$,

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

---

[3] I.e., $\mathcal{E}$ is a multi-set of unordered pairs of elements of $\mathcal{V}$. This allows for the possibility of considering parallel links and loops.
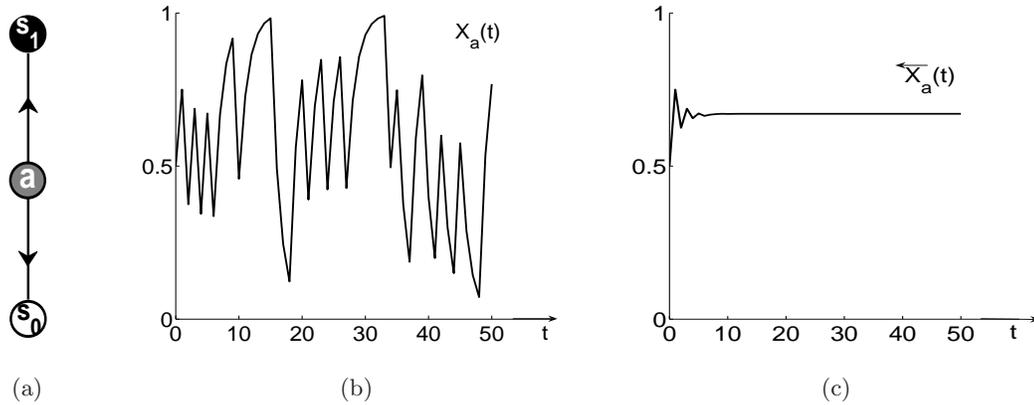
FIGURE 2. Typical sample-path behavior of the belief of the regular agent in the simple social network topology depicted in (a). In (b) the actual belief process $X_a(t)$, fluctuating ergodically on the interval $[0,1]$; in (c), the time-reversed process, rapidly converging to a stationary random belief X_a.

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

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

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

On the other hand, observe that, since $\sum_{k>n} 2^{-k}|B(k)| \leq 2^{-n}$, the series $X_a := \sum_{k \geq 1} 2^{-k}B(k)$ is sample-wise converging. It follows that, as t grows large, the time-reversed process

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

converges to X_a, with probability one, and, a fortiori, in distribution. Notice that, for all positive integer k, the binary k-tuples $\{B(1), \ldots, B(k)\}$ and $\{B(k), \ldots, B(1)\}$ are uniformly distributed over $\{0,1\}^k$, and independent from the Poisson arrival process $N(t)$. It follows 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. [45, 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

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

converges in law to the stationary belief

$$X_a := \theta(1-\theta)^{-1} \sum_{k \geq 1} (1-\theta)^k B(k). \tag{4}$$

As explained in [20, Section 2.6], for every value of θ in $(1/2,1)$, the probability law of X_a is singular, and in fact supported on a Cantor set. In contrast, for almost all values of $\theta \in (0,1/2)$,
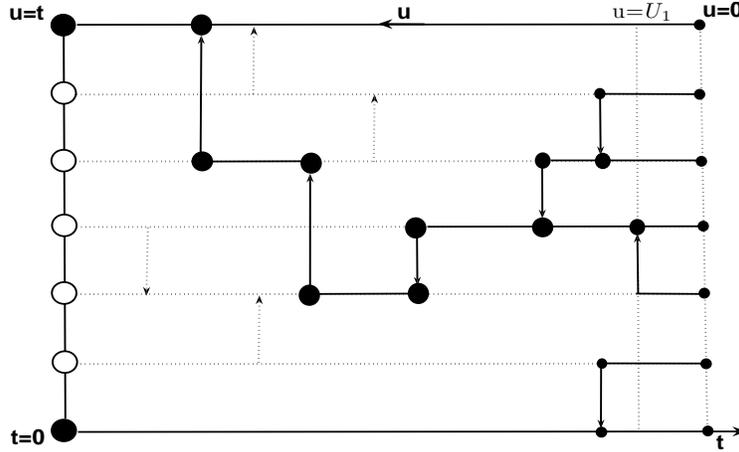
FIGURE 3. Duality between the voter model with zealots and the coalescing Markov chains with absorbing states. The network topology is a line with five regular agents and two stubborn agents placed in the two extremities. The time index for the opinion dynamics, t, runs from left to right, whereas the time index for the coalescing Markov chains 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 Markov chains 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.

the probability law of X_a is absolutely continuous with respect to Lebesgue's measure.[4] In the extreme case $\theta = 1$, it is not hard to see that $X_a(t) = B(N(t))$ converges in distribution to a random variable X_a with Bernoulli(1/2) distribution. On the other hand, observe that, regardless of the fine structure of the probability law of the stationary belief X_a, i.e., on whether it is absolutely continuous or singular, its moments can be characterized for all values of $\theta \in (0, 1]$. In fact, it follows from (4) that the expected value of X_a is given by

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

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

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

Observe that the expected value of X_a is independent from θ, while its variance increases from 0 to a maximum of $1/2$ as θ is increased from 0 to 1.

2.2. Voter model with zealots We now consider the special case when the social network topology $\overrightarrow{\mathcal{G}}$ is arbitrary, and $\theta_{av} = 1$ for all $(a, v) \in \mathcal{E}$. In this case, whenever a link $(a, v) \in \overrightarrow{\mathcal{E}}$ 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 [12], and Holley and Liggett [27]. It has been extensively studied in the framework of interacting particle systems [32, 33]. 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

---

[4] See [44]. In fact, explicit counterexamples of values of $\theta \in (0, 1/2)$ for which the asymptotic measure is singular are known. For example, Erdös [22, 23] showed that, if $\theta = (3 - \sqrt{5})/2$, then the probability law of X_a is singular.

considered, e.g., in [15, 17], [4, Ch. 14], and [21, Ch. 6.9]: in this case, consensus is achieved in some finite random time, whose distribution depends on the graph topology only.

In some recent work [38] a variant with one or more stubborn agents (there referred to as *zealots*) has been proposed and analyzed on the complete graph. We wish to emphasize that such voter model with zealots can be recovered as a special case of our model, and hence our general results, to be proven in the next sections, apply to it as well. However, we briefly discuss this special case here, since proofs are much more intuitive, and allow one to anticipate some of the general results.

The main tool in the analysis of the voter model is the *dual process*, which runs backward in time and allows one to identify the source of the opinion of each agent at any time instant. Specifically, let us focus on the belief of a regular agent a at time $t > 0$. Then, in order to trace $X_a(t)$, one has to look at the last meeting instance of agent a that occurred no later than time t. If such a meeting instance occurred at some time $t - U_1 \in [0, t]$ and the agent met was $v \in \mathcal{V}$, then the belief of agent a at time t coincides with the one of agent v at time $t - U_1$, i.e., $X_a(t) = X_v(t - U_1)$. The next step is to look at the last meeting instance of agent v occurred no later than time $t - U_1$; if such an instance occurred at time $t - U_2 \in [0, t - U_1]$, 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 continuous-time Markov chain $V_a(u)$ with state space $\mathcal{V}$, which starts at $V_a(0) = a$ and stays put there until time U_1, when it jumps to node v and stays put there in the time interval $[U_1, U_2)$, then jumps at time U_2 to node w, and so on. It is not hard to see that, thanks to the fact that the meeting instances are independent Poisson processes, the Markov chain $V_a(u)$ has transition rate matrix Q. In particular, it halts when it hits some state $s \in \mathcal{S}$. This shows that $\mathcal{L}(X_a(t)) = \mathcal{L}(X_{V_a(t)}(0))$. More generally, if one is interested in the joint probability distribution of the belief vector $X(t)$, then one needs to consider $n - |\mathcal{S}|$ continuous-time Markov chains, $\{V_a(t) : a \in \mathcal{A}\}$ each one starting from a different node in $\mathcal{A}$ (specifically, $V_a(0) = a$ for all $a \in \mathcal{A}$), and run simultaneously on $\mathcal{V}$ (see Figure 3). These Markov chains move independently with transition rate matrix Q, until the first time that they either meet, or they hit the set $\mathcal{S}$: in the former case, they stick together and continue moving on $\mathcal{V}$ as a single particle, with transition rate matrix Q; in the second case, they halt. This process is known as the *coalescing Markov chains process* with absorbing set $\mathcal{S}$. Then, one gets that

$$\mathcal{L}(\{X_a(t) : a \in \mathcal{A}\}) = \mathcal{L}(\{X_{V_a(t)}(0) : a \in \mathcal{A}\}). \tag{5}$$

Equation (5) establishes a *duality* between the voter model with zealots and the coalescing Markov chains process with absorbing states. In particular, Assumption 1 implies that, with probability one, each $V_a(u)$ will hit the set $\mathcal{S}$ in some finite random time $T_{\mathcal{S}}^a$, so that in particular the vector $\{V_a(u) : a \in \mathcal{A}\}$ converges in distribution, as u grows large, to an $\mathcal{S}^{\mathcal{A}}$-valued random vector $\{V_a(T_{\mathcal{S}}^a) : a \in \mathcal{A}\}$. It then follows from (5) that $X(t)$ converges in distribution to a stationary belief vector X whose entries are given by $X_s = x_s$ for every stubborn agent $s \in \mathcal{S}$, and $X_a = x_{V_a(T_{\mathcal{S}}^a)}$ for every regular agent $a \in \mathcal{A}$.

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

As we will see below, such behavior is rather general. In our model of social network with at least two stubborn agents having non-coincident constant beliefs, the regular agent beliefs fail to
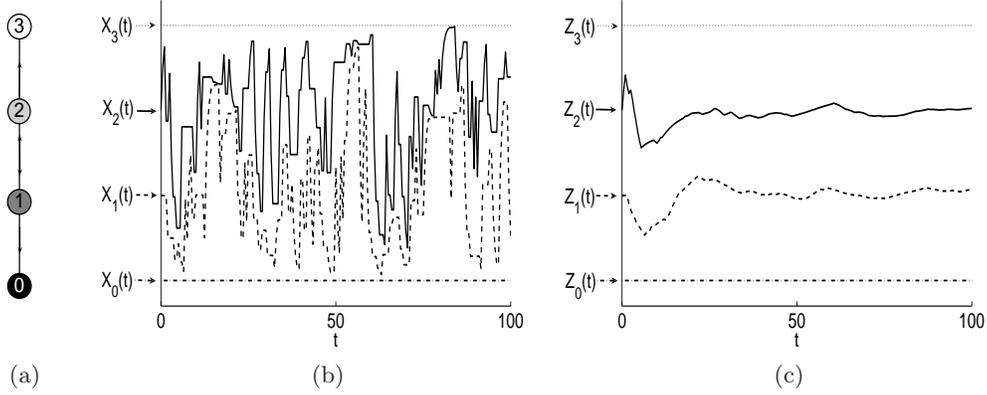
FIGURE 4. Typical sample-path behavior of the beliefs, and their ergodic averages for a social network with population size $n = 4$. The topology is a line graph, displayed in (a). The stubborn agents correspond to the two extremes of the line, $\mathcal{S} = \{0, 3\}$, while their constant opinions 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 $Z_v(t) := t^{-1} \int_0^t X_v(u) \mathrm{d}u$.

converge almost surely: we have seen this in the special cases of Section 2.1, while a general result in this sense will be stated as Theorem 2. On the other hand, we will prove that, regardless of the initial regular agents' beliefs, the belief vector $X(t)$ is convergent in distribution to a random stationary belief vector X (see Theorem 1), and in fact it is an ergodic process (see Corollary 1).

In order to prove Theorem 1, we will rewrite $X(t)$ in the form of an iterated affine function system [20]. Then, we will consider the so-called time-reversed belief process. This is a stochastic process whose marginal probability distribution, at any time $t \geq 0$, coincides with the one of the actual belief process, $X(t)$. In contrast to $X(t)$, the time-reversed belief process is in general not Markov, whereas it can be shown to converge to a random stationary belief vector with probability one. From this, we recover convergence in distribution of the actual belief vector $X(t)$.

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

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

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

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

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

$$\overrightarrow{A}(k, l) := A(l) A(l-1) \ldots A(k+1) A(k) \,, \qquad 1 \leq k \leq l \,, \tag{6}$$

with the convention that $\overrightarrow{A}(k, l) = I_{\mathcal{A}}$ for $k > l$. Then, at the time $T_{(l)}$ of the l-th belief update, one has

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

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

$$Y(t) = \overrightarrow{A}(1, N(t)) Y(0) + \sum_{1 \leq k \leq N(t)} \overrightarrow{A}(k+1, N(t)) B(k) \,, \tag{7}$$

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

$$\overleftarrow{Y}(t) := \overleftarrow{A}(1, N(t))Y(0) + \sum_{1 \le k \le N(t)} \overleftarrow{A}(1, k-1)B(k), \tag{8}$$

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 a fundamental observation (cf. [20]):

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 (7) and (8), 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.

LEMMA 2. *Let Assumption 1 hold. Then, for every value of the stubborn agents' beliefs $\{x_s\} \in \mathbb{R}^{\mathcal{S}}$, there exists an $\mathbb{R}^{\mathcal{A}}$-valued random variable Y, such that,*

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

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

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

$$\mathbb{E}[A_{aa'}(k)] = \frac{Q_{aa'}}{r}, \qquad \mathbb{E}[A_{aa}(k)] = 1 + \frac{Q_{aa}}{r}, \qquad \mathbb{E}[B_a(k)] = \frac{1}{r}\sum_s Q_{as}x_s,$$

for all $a \ne a' \in \mathcal{A}$. In particular, $\mathbb{E}[A(k)]$ is a substochastic matrix. It follows from Perron-Frobenius' theory that the spectrum of $\mathbb{E}[A(k)]$ is contained in the disk centered in 0 of radius ρ, where $\rho \in [0, +\infty)$ is its largest in module eigenvalue, with corresponding left eigenvector y with nonnegative entries. Moreover, Assumption 1 implies that, for all nonempty subsets $\mathcal{J} \subseteq \mathcal{A}$, there exists some $j \in \mathcal{J}$ and $v \in \mathcal{V} \setminus \mathcal{J}$ such that $(j, v) \in \mathcal{E}$ (otherwise $\mathcal{S}_j = \emptyset$ for all $j \in \mathcal{J}$). Therefore $\sum_a \mathbb{E}[A_{ja}] \le 1 - r^{-1}Q_{jv} < 1$. Choosing $\mathcal{J}$ as the support of the eigenvector y gives $\rho \sum_a y_a = \sum_a (\mathbb{E}[A(k)]^\mathrm{T}y)_a = \sum_j y_j \sum_a \mathbb{E}[A_{ja}(k)] < \sum_j y_j$, so that $\rho < 1$. Then, using the Jordan canonical decomposition, one can show that

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

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

$$\mathbb{E}\left[\left\|\overleftarrow{A}(1, k)\right\|_1\right] = \mathbb{E}\left[\max_{a'}\sum_a \overleftarrow{A}_{aa'}(1, k)\right] \le \sum_{a, a'} \mathbb{E}\left[\overleftarrow{A}_{a, a'}(1, k)\right] \le n\left\|\mathbb{E}\left[\overleftarrow{A}(1, k)\right]\right\|_\infty \le Cnk^{n-1}\rho^k. \tag{9}$$

Now, fix some $v \in (\rho, 1)$. From the independence of $\overleftarrow{A}(1, k-1)$ and $B(k)$ it follows that, for all $k \geq 1$,

$$
\begin{aligned}
\mathbb{P}\left(\left\|\overleftarrow{A}(1, k-1)B(k)\right\|_1 \geq v^{k-1}\right) &\leq v^{-k+1}\mathbb{E}\left[\|\overleftarrow{A}(1, k-1)B(k)\|_1\right] \\
&\leq v^{-k+1}\mathbb{E}\left[\|\overleftarrow{A}(1, k-1)\|_1\|B(k)\|_1\right] \\
&= v^{-k+1}\mathbb{E}\left[\|\overleftarrow{A}(1, k-1)\|_1\right]\mathbb{E}\left[\|B(k)\|_1\right] \\
&\leq \beta_{k-1},
\end{aligned}
$$

where $\beta_k := Cnk^{n-1}(\rho/v)^k\mathbb{E}[\|B(1)\|_1]$. Since $\sum_{k \geq 0}\beta_k < \infty$, the above bound and the Borel-Cantelli lemma imply that, with probability one, $\|\overleftarrow{A}(1, k-1)B(k)\|_1 < v^{k-1}$ for all but finitely many values of $k \geq 1$. Hence, almost surely, the series

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

is absolutely convergent. An analogous argument shows almost sure convergence of $\overleftarrow{A}(1, k)Y(0)$ to 0, as k grows large. Since, with probability one, $N(t)$ goes to infinity as t grows large, one has that

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

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

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

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

Proof. It follows from Lemma 2 $\overleftarrow{Y}(t)$ converges to Y with probability one, and a fortiori in distribution. By Lemma 1, $\overleftarrow{Y}(t)$ and $Y(t)$ are identically distributed. Therefore, $Y(t)$ converges to Y in distribution, and the first part of the claim follows by defining $X_a = Y_a$ for all $a \in \mathcal{A}$, and $X_s = x_s$ for all $s \in \mathcal{S}$. For the second part of the claim, it is sufficient to observe that the distribution of $Y = \sum_{k \geq 1}\overleftarrow{A}(1, k-1)B(k)$ is the same as the one of $Y' := A(0)Y + B(0)$, where $A(0)$, and $B(0)$, are independent copies of $A(1)$, and $B(1)$, respectively. ∎

Motivated by Theorem 1, for any agent $v \in \mathcal{V}$, we refer to the random variable X_v as the *stationary belief of agent v*. Using standard ergodic theorems for Markov processes, an immediate implication of Theorem 1 is the following corollary, which shows that time averages of continuous functions of agent beliefs with bounded expectation 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. *Let Assumption 1 hold. Then, for every value of the stubborn agents' beliefs* $\{x_s\} \in \mathbb{R}^{\mathcal{S}}$, *with probability one,*

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

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

Proof. Let $Y(t)$ and Y be the projections of the belief vector at time $t \geq 0$, and of the stationary belief vector A, respectively, to the regular agents set $\mathcal{A}$. Let $\tilde{Y}(0)$ be an $\mathbb{R}^{\mathcal{A}}$-valued random vector, independent from $Y(0)$ and such that $\mathcal{L}(\tilde{Y}(0)) = \mathcal{L}(Y)$. Let $Y(t)$ be as in (7), and

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

where $N(t)$ is the total number of agents' meetings up to time t, and $\overrightarrow{A}(k,l)$ is defined as in (6). Then, $\tilde{Y}(t) - Y(t) = \overrightarrow{A}(1, N(t))(\tilde{Y}(0) - Y(0))$. Arguing as in the proof of Lemma 2, one shows that $\lim_{t\to\infty} ||\tilde{Y}(t) - Y(t)|| = 0$, with probability one. Now, for $t > 0$, let the vectors $\tilde{X}(t)$ and $X(t)$ be defined by $\tilde{X}_a(t) = \tilde{Y}_a(t)$, $X_a(t) = Y_a(t)$ for $a \in \mathcal{A}$, and $\tilde{X}_s(t) = X_s(t) = x_s$ for $s \in \mathcal{S}$, and observe that, with probability one, $\sup_{t\geq 0} |X(t)| \leq \max_v |X_v(0)| < \infty$, $\sup_{t\geq 0} |\tilde{X}(t)| \leq \max_v |X_v| < \infty$. Then, for every continuous $\varphi : \mathbb{R}^{\mathcal{V}} \to \mathbb{R}$, one has that

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

with probability one. On the other hand, stationarity of the process $\tilde{X}(t)$ allows one to apply the ergodic theorem (see, e.g., [47, Theorem 6.2.12]), showing that, if $\mathbb{E}[\varphi(X)]$ exists and is finite, then

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

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

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

with probability one. ∎

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

THEOREM 2. *Let Assumption 1 hold, and let* $a \in \mathcal{A}$ *be such that* $|\mathcal{X}_a| \geq 2$. *Then, the stationary belief* X_a *is a non-degenerate random variable. Furthermore,* $\mathbb{P}(X_a \neq X_v) > 0$ *for all* $v \in \mathcal{V} \setminus \{a\}$.

Proof. With no loss of generality, since the distribution of the stationary belief vector X does not depend on the probability law of the initial beliefs of the regular agents, we can assume that such a law is the stationary one, i.e., that $\mathcal{L}(X(0)) = \mathcal{L}(X)$. Then, Theorem 1 implies that $\mathcal{L}(X(t)) = \mathcal{L}(X)$ for all $t \geq 0$.

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

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

$$X_v(T^{av}_{(k)}) = x_a\,, \qquad \forall n \geq 1\,. \tag{10}$$

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

$$\mathbb{P}(\text{first link activated after } T^* \text{ is } (a, v) | \mathcal{F}_{T^*}) > 0 \qquad \text{on} \qquad \{T^* < \infty\}\,, \tag{11}$$

which would contradict (10). Then, necessarily $T^* = \infty$, and hence $X_v = X_a$, with probability one.

On the other hand, assume that $\mathbb{P}(X_a = X_v) = 1$ for some $v \in \mathcal{V} \setminus \{a\}$. Then, with probability one $X_a(t) = X_v(t)$ for all rational $t \geq 0$. Since, as proved above, with probability one $X_a(t)$ is not constant in t, both $X_a(t)$ and $X_v(t)$ should jump simultaneously. However, the probability of this to occur is zero since link activations are independent Poisson processes. Therefore, necessarily $\mathbb{P}(X_a = X_v) < 1$. ∎

Even though, by Theorem 1, the belief of any agent always converges in distribution, Theorem 2 shows that, if a regular agent a is influenced by stubborn agents with different beliefs, then her stationary belief X_a is non-degenerate. By Corollary 1, this implies that, with probability one, her belief $X_a(t)$ keeps on fluctuating and does not stabilize on a limit. Similarly, Theorem 2 and Corollary 1 imply that, if a regular agents is influenced by stubborn agents with different beliefs, then, with probability one, her belief will not achieve a consensus asymptotically with any other agent in the society.

4. Expected beliefs and belief crossproducts   In this section, we provide a characterization of the expected beliefs and belief crossproducts of the agents. In particular we will provide explicit characterizations of the stationary expected beliefs and belief crossproducts in terms of hitting probabilities of a pair of coupled Markov chains on $\overrightarrow{\mathcal{G}} = (\mathcal{V}, \overrightarrow{\mathcal{E}})$.[5] Specifically, we consider a coupling $(V(t), V'(t))$ of continuous-time Markov chains with state space $\mathcal{V}$, such that both $V(t)$

---

[5] Note that the set of states for such Markov chain corresponds to the set of agents, therefore we use the terms "state" and "agent" interchangeably in the sequel.

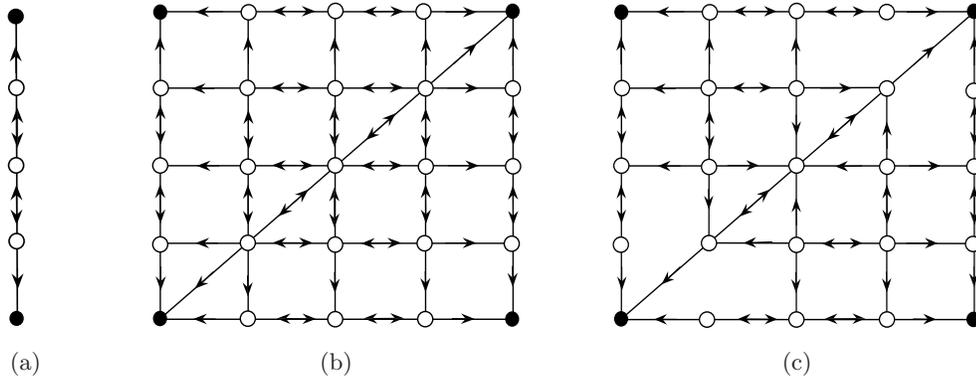(a)   (b)   (c)

FIGURE 5. In (a), a network topology consisting of a line with three regular agents and two stubborn agents placed in the extremes. In (b), the directed graph representing the possible state transitions of the corresponding coupled Markov chains pair $(V(t), V'(t))$ when $\theta_e \in (0,1)$ for all $e \in \vec{\mathcal{E}}$. Such coupled chains pair has 25 states, four of which are absorbing states. The components of $(V(t), V'(t))$ jump 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 state transition graph is reported for the Markov chains pair $(V(t), V'(t))$ in the extreme case when $\theta_e = 1$ for all $e \in \vec{\mathcal{E}}$. In this case, the coupled Markov chains 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 links from the diagonal set.

and $V'(t)$ have transition rate matrix Q, as defined in (2). The pair $(V(t), V'(t))$ is a Markov chain on the state space $\mathcal{V} \times \mathcal{V}$ with transition rate matrix K whose entries are given by

$$
K_{(v,v')(w,w')} := \begin{cases}
Q_{vw} & \text{if } v \neq v', \ v \neq w, \ v' = w' \\
Q_{v'w'} & \text{if } v \neq v', \ v = w, \ v' \neq w' \\
0 & \text{if } v \neq v', \ v \neq w, \ v' \neq w' \\
Q_{vv} + Q_{v'v'} & \text{if } v \neq v', \ v = w, \ v' = w' \\
\theta_{vw} Q_{vw} & \text{if } v = v', \ w = w', \ v \neq w \\
(1 - \theta_{vw}) Q_{vw} & \text{if } v = v', \ v \neq w, \ v' = w' \\
(1 - \theta_{vw'}) Q_{vw'} & \text{if } v = v', \ v = w, \ v' \neq w' \\
0 & \text{if } v = v', \ w \neq w', \ w \neq v, \ w' \neq v' \\
2Q_{vv} + \sum_{v'' \neq v} \theta_{vv''} Q_{vv''} & \text{if } v = v', \ v = w, \ v' = w'.
\end{cases}
\tag{12}
$$

The first four lines of (12) state that, conditioned on $(V(t), V'(t))$ being on a pair of non-coincident nodes (v, v'), each of the two components, $V(t)$ (respectively, $V'(t)$), jumps to a neighbor node w, with transition rate Q_{vw} (respectively, to a neighbor node w' with transition rate $Q_{v'w'}$), whereas the probability that both components jump at the same time is zero. On the other hand, the last five lines of (12) state that, once the two components have met, i.e., conditioned on $V(t) = V'(t) = v$, they have some chance to stick together and jump as a single particle to a neighbor node w, with rate $\theta_{vw} Q_{vw}$, while each of the components $V(t)$ (respectively, $V'(t)$) has still some chance to jump alone to a neighbor node w with rate $(1 - \theta_{vw}) Q_{vw}$ (resp., to w' with rate $(1 - \theta_{vw'}) Q_{vw'}$). In the extreme case when $\theta_{vw} = 1$ for all (v, w) in $\mathcal{E}$, the sixth and seventh line of the righthand side of (12) equal 0, and in fact one recovers the expression for the transition rates of two coalescing Markov chains: once $V(t)$ and $V'(t)$ have met, they stick together and move as a single particle, never separating from each other. See Figure 5 for a visualization of the possible state transitions of $(V(t), V'(t))$.

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

$$\gamma_w^v(t) := \mathbb{P}_v(V(t) = w) = \mathbb{P}_v(V'(t) = w), \qquad \eta_{ww'}^{vv'}(t) = \mathbb{P}_{vv'}(V(t) = w, V'(t) = w'), \qquad (13)$$

the marginal and joint transition probabilities of the two Markov chains at time t. It is a standard fact (see, e.g., [41, Theorem 2.8.3]) that such transition probabilities satisfy the *Kolmogorov backward equations*

$$\frac{\mathrm{d}}{\mathrm{d}t}\gamma_w^v(t) = \sum_{\tilde{v}} Q_{v\tilde{v}}\gamma_w^{\tilde{v}}(t), \qquad \frac{\mathrm{d}}{\mathrm{d}t}\eta_{ww'}^{vv'}(t) = \sum_{\tilde{v},\tilde{v}'} K_{(v,v')(\tilde{v},\tilde{v}')}\eta_{ww'}^{\tilde{v}\tilde{v}'}(t), \qquad v, v', w, w' \in \mathcal{V}, \qquad (14)$$

with initial condition

$$\gamma_w^v(0) = \begin{cases} 1 & \text{if } v = w \\ 0 & \text{if } v \neq w, \end{cases} \qquad \eta_{ww'}^{vv'}(0) = \begin{cases} 1 & \text{if } (v, v') = (w, w') \\ 0 & \text{if } (v, v') \neq (w, w'). \end{cases} \qquad (15)$$

The next simple result provides a fundamental link between the belief evolution process introduced in Section 2 and the coupled Markov chains, by showing that the expected values and expected cross-products of the agents' beliefs satisfy the same linear system (14) of ordinary differential equations as the transition probabilities of $(V(t), V'(t))$.

LEMMA 3. *For all $v, v' \in \mathcal{V}$, and $t \geq 0$, it holds*

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[X_v(t)] = \sum_w Q_{vw}\mathbb{E}[X_w(t)], \qquad \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[X_v(t)X_{v'}(t)] = \sum_{w,w'} K_{(v,v')(ww')}\mathbb{E}[X_w(t)X_{w'}(t)], \qquad (16)$$

so that

$$\mathbb{E}[X_v(t)] = \sum_w \gamma_w^v \mathbb{E}[X_w(0)], \qquad \mathbb{E}[X_v(t)X_{v'}(t)] = \sum_{w,w'} \eta_{ww'}^{vv'}\mathbb{E}[X_w(0)X_{w'}(0)]. \qquad (17)$$

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

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

Then, the above and the fact that the Poisson clock has rate r imply the left-most equation in (16).

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

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

as well as

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

Then, the two identities above and the fact that the Poisson clock has rate r imply the right-most equation in (16). It follows from (14), (15), and (16) that $\sum_w \gamma_w^v(t)\mathbb{E}[X_w(0)]$ and $\mathbb{E}[X_v(t)]$ satisfy the same linear system of differential equations, and the same holds true for $\mathbb{E}[X_v(t)X_{v'}(t)]$ and $\sum_{w,w'} \eta_{ww'}^{vv'}(t)\mathbb{E}[X_w(0)X_{w'}(0)]$. This readily implies (17). ∎

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

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

Observe that Assumption 1 implies that both $T_{\mathcal{S}}$ and $T_{\mathcal{S}}'$ are finite with probability one for every initial distribution of the pair $(V(0), V'(0))$. Hence, for all $v, v' \in \mathcal{V}$, we can define the *hitting probability distributions* γ^v over $\mathcal{S}$, and $\eta^{vv'}$ over $\mathcal{S}^2$, whose entries are respectively given by

$$
\begin{aligned}
\gamma_s^v &:= \mathbb{P}_v(V(T_{\mathcal{S}}) = s), \qquad s \in \mathcal{S}, \\
\eta_{ss'}^{vv'} &:= \mathbb{P}_{vv'}(V(T_{\mathcal{S}}) = s, V'(T_{\mathcal{S}}') = s'), \qquad s, s' \in \mathcal{S}.
\end{aligned}
\tag{18}
$$

Then, we have the following:

THEOREM 3. *Let Assumption 1 hold. Then, for every value of the stubborn agents' beliefs $\{x_s\} \in \mathbb{R}^{\mathcal{S}}$,*

$$
\mathbb{E}[X_v] = \sum_s \gamma_s^v x_s, \qquad \mathbb{E}[X_vX_{v'}] = \sum_{s,s'} \eta_{ss'}^{vv'} x_s x_{s'}, \qquad v, v' \in \mathcal{V}.
\tag{19}
$$

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

$$
\sum_v Q_{av}\mathbb{E}[X_v] = 0, \qquad \mathbb{E}[X_s] = x_s, \qquad \forall a \in \mathcal{A}, \quad \forall s \in \mathcal{S},
\tag{20}
$$

$$
\sum_{w,w'} K_{(a,a'),(w,w')}\mathbb{E}[X_wX_{w'}] = 0, \qquad \mathbb{E}[X_vX_{v'}] = \mathbb{E}[X_v]\mathbb{E}[X_{v'}], \qquad \forall a, a' \in \mathcal{A}, \ \forall(v,v') \in \mathcal{V}^2 \setminus \mathcal{A}^2.
\tag{21}
$$

Proof. Assumption 1 implies that $\lim_{t\to\infty} \gamma_s^v(t) = \gamma_s^v$ for every $s \in \mathcal{S}$, and $\lim_{t\to\infty} \gamma_a^v(t) = 0$ for every $a \in \mathcal{A}$. Therefore, (17) implies that

$$
\lim_{t\to\infty} \mathbb{E}[X_v(t)] = \sum_s \gamma_s^v x_s, \qquad \forall v \in \mathcal{V}.
\tag{22}
$$

Now, if the initial belief distribution $\mathcal{L}(X(0))$ coincides with the stationary one $\mathcal{L}(X)$, one has that $\mathcal{L}(X(t)) = \mathcal{L}(X)$ for all $t \geq 0$, so that in particular $\mathbb{E}[X_v(t)] = \mathbb{E}[X_v]$, and hence $\lim_{t \to \infty} \mathbb{E}[X_v(t)] = \mathbb{E}[X_v]$ for all v. Substituting in the righthand side of (22), this proves the leftmost identity in (19). The rightmost identity in (19) follows from an analogous argument.

In order to prove the second part of the claim, observe that the expected stationary beliefs and belief cross-products necessarily satisfy (20) and (21), since, by Lemma 3, they evolve according to the autonomous differential equations (16), and are convergent by the arguments above. On the other hand, uniqueness of the solutions of (20) and (21) follows from [4, Ch. 2, Lemma 27]. ∎

REMARK 1. Since the stationary beliefs X_v take values in the interval $[\min_s x_s, \max_s x_s]$, one has that both $\mathbb{E}[X_v]$ and $\mathbb{E}[X_v X_{v'}]$ exist and are finite. Hence, Corollary 1 implies that the asymptotic empirical averages of the agents' beliefs and their cross-products, i.e., of the almost surely constant limits

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

coincide with the expected stationary beliefs and belief crossproducts, i.e., $\mathbb{E}[X_v]$ and $\mathbb{E}[X_v X_{v'}]$, respectively, independently of the distribution of initial regular agents' beliefs.

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. Hence, in particular, $\mathbb{E}[X_a] = x_*$, and $\mathrm{Var}[X_a] = 0$. This can be thought of as a sort of complement to Theorem 2.

5. Explicit computations of stationary expected beliefs and variances We present now a few examples of explicit computations of the stationary expected beliefs and variances for social networks obtained using the construction in Example 1, starting from a simple undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Recall that, in this case, $Q_{av} = \theta/d_a$ for all $a \in \mathcal{A}$, and $v \in \mathcal{V}$ such that $\{a, v\} \in \mathcal{E}$. It then follows from Theorem 3 that the expected stationary beliefs can be characterized as the unique vectors in $\mathbb{R}^{\mathcal{V}}$ satisfying

$$\mathbb{E}[X_a] = \frac{1}{d_a} \sum_{v : \{v,a\} \in \mathcal{E}} \mathbb{E}[X_v], \qquad \mathbb{E}[X_s] = x_s, \qquad \forall a \in \mathcal{A}, \ \forall s \in \mathcal{S}. \tag{23}$$

Moreover, in the special case when $\theta = 1$, the second moments of the stationary beliefs are the unique solutions of

$$\mathbb{E}[X_a^2] = \frac{1}{d_a} \sum_{v : \{v,a\} \in \mathcal{E}} \mathbb{E}[X_v^2], \qquad \mathbb{E}[X_s^2] = x_s^2, \qquad \forall a \in \mathcal{A}, \ \forall s \in \mathcal{S}. \tag{24}$$

EXAMPLE 2. (**Tree**) Let us consider the case when $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a tree and the stubborn agent set $\mathcal{S}$ consists of only two elements, s_0 and s_1, with beliefs x_0, and x_1, respectively. For $v, w \in \mathcal{V}$, let $d(v, w)$ denote their distance, i.e., the length of the shortest path connecting them in $\mathcal{G}$. Let $m := d(s_0, s_1)$, and $\mathcal{W} = \{s_0 = w_0, w_1, \ldots, w_{m-1}, w_m = s_1\}$, where $\{w_{i-1}, w_i\} \in \mathcal{E}$ for all $1 \leq i \leq m$, be the unique path connecting s_0 to s_1 in $\mathcal{G}$. Then, we can partition the rest of the node set as $\mathcal{V} \setminus \mathcal{W} = \bigcup_{0 \leq i \leq m} \mathcal{V}_w$, where $\mathcal{V}_i$ is the set of nodes $v \in \mathcal{V} \setminus \mathcal{W}$ such that the unique paths from v to s_0 and s_1 both pass through w_i. Since the set of neighbors of every $v \in \mathcal{V}_i$ is contained in $\mathcal{V}_i \cup \{w_i\}$, (23) implies that

$$\mathbb{E}[X_v] = \mathbb{E}[X_{w_i}], \qquad \forall v \in \mathcal{V}_i, \ 1 \leq i \leq m. \tag{25}$$

Hence, one is left with determining the values of $\mathbb{E}[X_{w_i}]$, for $0 \leq i \leq m$. Observe that clearly

$$\mathbb{E}[X_{w_0}] = x_{s_0}, \qquad \mathbb{E}[X_{w_m}] = x_{s_1}. \tag{26}$$
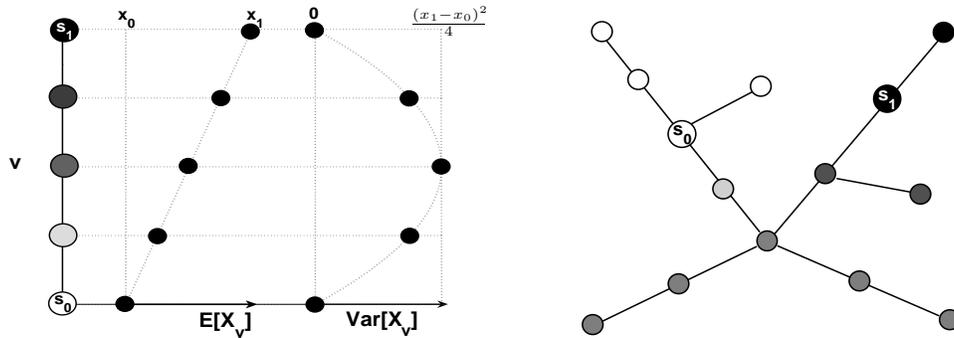
FIGURE 6. In the left-most figure, expected stationary beliefs and variances (the latter valid for the special case when $\theta = 1$) in a social network with a line graph topology with $n = 5$, and stubborn agents positioned in the two extremities. The expected stationary beliefs are linear interpolations of the two stubborn agents' beliefs, while their variances follow a parabolic profile with maximum in the central agent, and zero variance for the two stubborn agents s_0, and s_1. In the right-most figure, expected stationary beliefs in a social network with a tree-like topology, represented by different levels of gray. The solution is obtained by linearly interpolating between the two stubborn agents' beliefs, x_0 (white), and x_1 (black), on the vertices lying on the path between s_0 and s_1, and then extended by putting it constant on each of the connected components of the subgraph obtained by removing the links of such path.

On the other hand, for all $0 < i < m$, the neighborhood of w_i consists of w_{i-1}, w_{i+1}, and possibly some elements of $\mathcal{V}_i$. Then, (23) and (25) imply that

$$\mathbb{E}[X_{w_i}] = \frac{1}{2}\mathbb{E}[X_{w_{i-1}}] + \frac{1}{2}\mathbb{E}[X_{w_{i+1}}], \qquad 0 < i < m. \tag{27}$$

Now, observe that, since

$$\frac{i}{m}x_{s_1} + \frac{m-i}{m}x_{s_0} = \frac{1}{2}(i-1)x_{s_1} + \frac{1}{2}(m-i+1)x_{s_0} + \frac{1}{2}(i+1)x_{s_1} + \frac{1}{2}(m-i-1)x_{s_1},$$

then the unique solution of (26) and (27) is given by

$$\mathbb{E}[X_{w_i}] = \frac{i}{m}x_{s_1} + \frac{m-i}{m}x_{s_0}, \qquad 0 \le i \le m. \tag{28}$$

Upon observing that $d(v, s_j) = d(w_i, s_j)$ for all $w \in \mathcal{V}_i$, $0 \le i \le m$, and $j = 0, 1$, we may rewrite (25) and (28) as

$$\mathbb{E}[X_v] = \begin{cases} x_i & \text{if } |\mathcal{S}_v| = \{x_i\}, \ i = 0, 1, \\[2mm] \dfrac{d(v, s_0)x_1 + d(v, s_1)x_0}{d(v, s_0) + d(v, s_1)}, & \text{if } |\mathcal{S}_v| = 2. \end{cases} \tag{29}$$

In other words, the stationary expected beliefs are linear interpolations of the beliefs of the stubborn agents. A totally analogous argument shows that, if the confidence parameter satisfies $\theta = 1$, then (24) is satisfied by

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

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

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

The two equations above show that the belief of each regular agent keeps on fluctuating ergodically around a value which depends on the relative distance of the agent from the two stubborn agents.
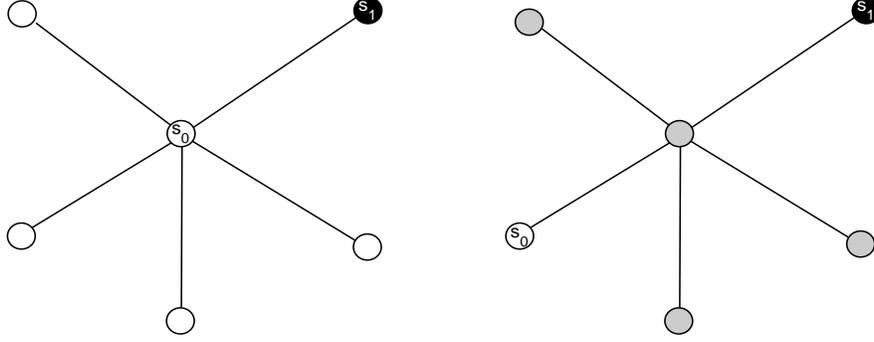
FIGURE 7. Two social network with a special case of tree-like topology, known as star graph, and two stubborn agents. In social network depicted in left-most figure one of the stubborn agents, s_0, occupies the center, while the other one, s_1, occupies one of the leaves. There, all regular agents' stationary beliefs coincide with the belief x_0 of s_0, represented in white. In social network depicted in right-most figure, none of the stubborn agents, occupy the center. There, all regular agents' stationary beliefs coincide with the arithmetic average (represented in gray) of x_0 (white), and x_1 (black).

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

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

EXAMPLE 3. (**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 link $\{a_0, a_1\}$ 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, (23) is satisfied by

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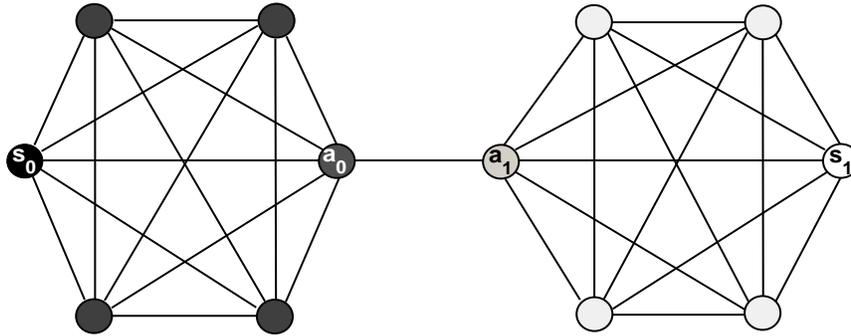$$

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

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

Let the stubborn agent set consist of only two elements: $\mathcal{S} := \{s_0, s_1\}$. Then the following formula holds (see [4, Ch. 2, Corollary 10]):

$$\gamma_{s_0}^v = \mathbb{P}_v(T_{s_1} < T_{s_0}) = \frac{E_{vs_0} - E_{vs_1} + E_{s_1 s_0}}{E_{s_0 s_1} + E_{s_1 s_0}} \tag{30}$$

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

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

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

$$E_{vw} = n^{-1}(Z_{ww} - Z_{vw}). \tag{31}$$

(See, e.g., [4, Ch. 2, Lemma 12].)

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. (See, e.g., [9, Chapter 15].) Morever, there is an orthonormal basis of eigenvectors for P good for any Θ: if $l = (l_1, \ldots l_d) \in \mathcal{V}$ define $\Upsilon_l \in \mathbb{R}^{\mathcal{V}}$ by

$$\Upsilon_l(k) = m^{-d/2} \exp\left(\frac{2\pi i}{m} l \cdot k\right), \quad k = (k_1, \ldots, k_d) \in \mathcal{V},$$
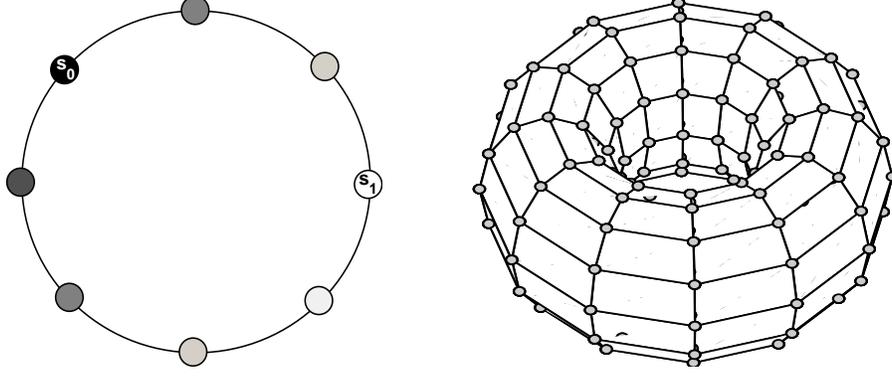
FIGURE 9. Two social networks with cycle and 2-dimensional toroidal topology, respectively.

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

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

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

$$Z_{vw} = m^{-d} \sum_{l \neq 0} \frac{\exp\left(\frac{2\pi i}{m} l \cdot (v - w)\right)}{1 - \frac{1}{|\Theta^+|} \sum_{k \in \Theta^+} \cos\left(\frac{2\pi}{m} l \cdot k\right)}, \qquad v, w \in \mathcal{V}. \tag{32}$$

From (30), (31), and the fact that $E_{s_0 s_1} = E_{s_1 s_0}$ by symmetry, one obtains

$$\gamma_{s_1}^a = \frac{1}{2} + \frac{m^{-d} \sum_{l \neq 0} \dfrac{\exp\left(\frac{2\pi i}{m} l \cdot (a - s_1)\right) - \exp\left(\frac{2\pi i}{m} l \cdot (a - s_0)\right)}{1 - \frac{1}{|\Theta^+|} \sum_{k \in \Theta^+} \cos\left(\frac{2\pi}{m} l \cdot k\right)}}{2 m^{-d} \sum_{l \neq 0} \dfrac{1 - \cos\left(\frac{2\pi}{m} l \cdot (s_0 - s_1)\right)}{1 - \frac{1}{|\Theta^+|} \sum_{k \in \Theta^+} \cos\left(\frac{2\pi}{m} l \cdot k\right)}}, \qquad a \in \mathcal{A}. \tag{33}$$

The expected stationary beliefs can then be computed using Theorem 3 and (33).

6. Homogeneous influence in highly fluid social networks   In this section, we present estimates for the expected stationary beliefs and belief variances as a function of the underlying social network. First, we will introduce the notion of fluidity of a social network, a quantity which depends only on the geometry of the network and on the size of the stubborn agent set. Then, we will prove that, when the social network is highly fluid, the influence of the stubborn agents on the rest of the society is homogeneous, meaning that most of the regular agents have approximately the same stationary expected belief and (in the special case when $\theta_e = 1$ for all $e \in \overrightarrow{\mathcal{E}}$) belief variance.

6.1. Network fluidity and homogeneous influence   Recall that Theorem 3 allows one to express the first moment of the stationary beliefs in terms of the hitting probability distributions γ^v on the stubborn agent set $\mathcal{S}$ of the continuous-time Markov chain $V(t)$ with state space $\mathcal{V}$ and transition rate matrix Q. It is a simple but key observation that such hitting probabilities only depend on the restriction of Q to $\mathcal{A} \times \mathcal{V}$ (the other rows affecting only the behavior of $V(t)$ after hitting $\mathcal{S}$), and they do not change if any row of Q is multiplied by some positive scalar (this multiplication having the only affect of speeding up or slowing down $V(t)$ without changing the jump probabilities). Formally, we have the following

LEMMA 4. *Let $P \in \mathbb{R}^{\mathcal{V} \times \mathcal{V}}$ be stochastic and such that*

$$P_{av} = \alpha_a Q_{av}, \qquad \forall a \in \mathcal{A}, v \in \mathcal{V}, a \neq v, \tag{34}$$

for some $\alpha_a > 0$. Let $W(k)$, for $k = 0, 1, \ldots$, be a discrete-time Markov chain with transition probability matrix P. Then, the hitting probability distributions of the continuous-time Markov chain $V(t)$ with transition rate matrix Q coincide with those of $W(k)$, i.e.,

$$\gamma_s^v = \mathbb{P}_v(W(U_{\mathcal{S}}) = s), \qquad \forall v \in \mathcal{V}, \, s \in \mathcal{S},$$

where $U_{\mathcal{S}} := \min\{k \geq 0 : W(k) \in \mathcal{S}\}$.

Lemma 4 allows one to consider the discrete-time chain $W(k)$ with any transition probability matrix P satisfying (34), in order to compute the hitting probability distributions γ^v. In fact, it is convenient to consider stochastic matrices $P \in \mathbb{R}^{\mathcal{V} \times \mathcal{V}}$ that, in addition to satisfying (34), are irreducible and aperiodic. Let us denote the set of all such matrices by $\mathcal{P}$. Observe that, provided that Assumption 1 holds, the set $\mathcal{P}$ is non-empty, since it can be easily checked to contain, e.g., the matrix $P \in \mathbb{R}^{\mathcal{V} \times \mathcal{V}}$ with entries $P_{sv} = 1/n$, $P_{aa} = 0$, $P_{aw} = -Q_{aw}/Q_{aa}$ for all $s \in \mathcal{S}$, $v \in \mathcal{V}$, $a \in \mathcal{A}$, and $w \in \mathcal{V} \setminus \{a\}$. Observe that, for every $P \in \mathcal{P}$, irreducibility implies the existence of a unique invariant probability measure, and, together with aperiodicity, convergence of the time-k distribution

$$p^v(k) = \{p_w^v(k) : w \in \mathcal{V}\}, \qquad p_w^v(k) := \mathbb{P}_v(W(k) = w), \tag{35}$$

irrespectively of the initial state $v \in \mathcal{V}$. We introduce the following notation.

DEFINITION 1. Given a social network satisfying Assumption 1, and $P \in \mathcal{P}$, let $\pi = P'\pi$ be the unique invariant probability measure of P. Let

$$\pi(\mathcal{S}) := \sum_s \pi_s, \qquad \pi_* := \min_v \pi_v,$$

be the size of the stubborn agents set and, respectively, the minimum weight of an agent, as measured by π. Moreover, let

$$\tau := \inf\left\{k \geq 0 : \max_v ||p^v(k) - \pi||_{TV} \leq \frac{1}{2e}\right\}, \tag{36}$$

where $p^v(k)$ are as in (35), denote the (variational distance) mixing time of the discrete-time Markov chain $W(k)$ with state space $\mathcal{V}$ and transition probability matrix P.

For the canonical construction of a social network introduced in Example 1, the quantities above have a more explicit characterization which allows for a geometric interpretation.

EXAMPLE 5. Let us consider the canonical construction of a social network from a given connected multigraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, outlined in Example 1. Define $\tilde{P} \in \mathbb{R}^{\mathcal{V} \times \mathcal{V}}$ by putting $\tilde{P}_{vw} = \kappa_{v,w}/d_v$, where $\kappa_{v,w}$ is the multiplicity of the link $\{v, w\}$ in $\mathcal{E}$, and $d_v = \sum_w \kappa_{v,w}$ is the degree of node v in $\mathcal{G}$. Then, put $P = (I + \tilde{P})/2$, where I stands for the identity matrix on $\mathcal{V}$. In fact, P defined as above is known in the literature as the transition matrix of the simple lazy random walk on $\mathcal{G}$ [31, page 9]. Observe that (34) is satisfied, connectedness of $\mathcal{G}$ implies irreducibility of $\tilde{P}$, and hence of P, while aperiodicity of P (not necessarily of $\tilde{P}$) is immediate. Hence, $P \in \mathcal{P}$. Moreover, the invariant measure π (of both P and $\tilde{P}$) is given by

$$\pi_v = d_v/(n\overline{d}), \qquad \forall v \in \mathcal{V},$$

where $\overline{d} := n^{-1} \sum_v d_v$ is the average degree of $\mathcal{G}$. Observe that, in this construction,

$$\pi(\mathcal{S}) = \left(\sum_v d_v\right)^{-1} \sum_s d_s \tag{37}$$

is the fraction of the total degree of the stubborn agents set $\mathcal{S}$ and the total degree of the whole agent set $\mathcal{V}$ in $\mathcal{G}$. Moreover, the mixing time can be bounded in terms of the *conductance* of $\mathcal{G}$, defined as

$$\Phi := \min\left\{\phi(\mathcal{W}) : \mathcal{W} \subseteq \mathcal{V}, \ 0 < \sum_w d_w \leq n\overline{d}/2\right\}, \tag{38}$$

where

$$\phi(\mathcal{W}) := \frac{\sum_{w \in \mathcal{W}} \sum_{v \in \mathcal{V} \setminus \mathcal{W}} \mathbb{1}_{\mathcal{E}}(\{v, w\})}{\sum_{w \in \mathcal{W}} d_w}$$

is the bottleneck ratio of the set $\mathcal{W} \subseteq \mathcal{V}$, i.e., the ratio between the number of links connecting $\mathcal{W}$ to the rest of the node set and its total degree. In particular, standard results imply that

$$\frac{1}{4\Phi} \leq \tau \leq \frac{2}{\Phi^2} \log \frac{2e}{\pi_*}. \tag{39}$$

(See, e.g., [31, Theorem 7.3] for the lower bound, and combine [31, Theorem 12.3] and [21, Theorem 6.2.1] in order to get the upper bound.)

We conclude this example by observing that different choices of the multigraph $\mathcal{G}$ may result in the same social network, hence in particular in the same directed graph $\overrightarrow{\mathcal{G}}$. For example, $\mathcal{G}$ could have links between pairs of nodes both belonging to $\mathcal{S}$. While such links are irrelevant for the belief dynamics (they have no correspondence in the directed graph $\overrightarrow{\mathcal{G}}$), they do affect the stochastic matrix P, hence the quantities $\pi(\mathcal{S})$ and τ. In particular, while removing such links from $\mathcal{G}$ clearly has the effect of decreasing $\pi(\mathcal{S})$, it also has the potential of increasing the mixing time τ (since less connected graphs tend to have larger mixing time). Hence, while the stationary belief distribution does not depend on the presence of links connecting pairs of stubborn agents in $\mathcal{G}$, the estimates of the stationary beliefs' moments derived in this section could potentially benefit from considering these links.

We are now ready to introduce the notion of *fluidity* of a social network.

DEFINITION 2.   Let the social network satisfy Assumption 1. For every $P \in \mathcal{P}$ let

$$\psi(P) := \frac{n\pi_*}{\tau\pi(\mathcal{S})\log\left(e^2/(\tau\pi(\mathcal{S}))\right)}. \tag{40}$$

The *fluidity* of the social network is

$$\Psi := \sup\left\{\psi(P) : P \in \mathcal{P}\right\}. \tag{41}$$

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

Our estimates will show that for large-scale highly fluid social networks, the first two moments of the stationary beliefs of most of the regular agents in the population can be approximated by those of a *weighted-mean belief* Z, supported in the finite set $\mathcal{X} := \{x_s : s \in \mathcal{S}\}$, and given by

$$\mathbb{P}(Z = z) = \sum_s \overline{\gamma}_s \mathbb{1}_{\{z\}}(x_s), \qquad z \in \mathcal{X}, \qquad \overline{\gamma}_s := \sum_v \pi_v \gamma_s^v, \qquad s \in \mathcal{S}. \tag{42}$$

We refer to the probability distribution $\{\overline{\gamma}_s : s \in \mathcal{S}\}$ as the *stationary stubborn agent distribution*. Observe that $\overline{\gamma}_s = \mathbb{P}_\pi(W(U_\mathcal{S}) = s)$ coincides with the probability that the Markov chain $W(k)$, started from the stationary distribution π, hits the stubborn agent s before any other stubborn agent $s' \in \mathcal{S}$. In fact, as we will clarify below, one may interpret $\overline{\gamma}_s$ as a relative measure of the influence of the stubborn agent s on the society compared to the rest of the stubborn agents $s' \in \mathcal{S}$.

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

$$\mathbb{E}[Z] := \sum_s \overline{\gamma}_s x_s, \qquad \sigma_Z^2 := \sum_s \overline{\gamma}_s (x_s - \mathbb{E}[Z])^2. \tag{43}$$

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

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

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

$$\Delta_* := \max\{x_s - x_{s'} : s, s' \in \mathcal{S}\}. \tag{44}$$

The next theorem presents the main result of this section.

THEOREM 4. *Let Assumption 1 hold. Then, for all $\varepsilon > 0$,*

$$\frac{1}{n}\left|\left\{v : \left|\mathbb{E}[X_v] - \mathbb{E}[Z]\right| > \Delta_*\varepsilon\right\}\right| \le \frac{1}{\varepsilon\Psi}. \tag{45}$$

Furthermore, if the trust parameters satisfy $\theta_{av} = 1$ for all $(a, v) \in \overrightarrow{\mathcal{E}}$, then

$$\frac{1}{n}\left|\left\{v : \left|\sigma_v^2 - \sigma_Z^2\right| > \Delta_*^2\varepsilon\right\}\right| \le \frac{1}{\varepsilon\Psi}. \tag{46}$$

This theorem implies that in large-scale highly fluid social networks, as the population size n grows large, the stationary expected beliefs and variances of the regular agents concentrate around fixed values corresponding to the expected weighted-mean belief $\mathbb{E}[Z]$, and, respectively, its variance σ_Z^2 (see Figures 10 and 11). We refer to this phenomenon as *homogeneous influence* of the stubborn agents on the rest of the society—meaning that their influence on most of the agents in the society is approximately the same. Indeed, it amounts to homogeneous first and second moment of the agents' stationary beliefs. This shows that in highly fluid social networks, most of the regular agents are affected by the stubborn agents in approximately the same way.

Observe that, provided that $n\pi_*$ remains bounded from below by a positive constant, as we will prove to be the case in all the considered examples, a social network is highly fluid when the stationary measure $\pi(\mathcal{S})$ of the stubborn agents set vanishes fast enough to compensate for the possible growth of the mixing time τ, as the network size n grows large. Hence, intuitively, Theorem 4 states that, if the set $\mathcal{S}$ and the mixing time τ are both small enough, 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, so that their beliefs' empirical averages and variances will converge to values very close to each other. Theorem 4 is proved in Section 6.3. Its proof relies on the characterization of the expected stationary beliefs and variances in terms of the hitting probabilities γ_s^v. The definition of highly fluid network implies that the (expected) time it takes a Markov chain to hit the set $\mathcal{S}$, when started from most of the nodes, is much larger than the mixing time τ. Hence, before hitting $\mathcal{S}$, the chain looses memory of where it started from, and approaches $\mathcal{S}$ almost as if started from the stationary distribution π.

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

Before proving Theorem 4 in Section 6.3, we present some examples of highly fluid social networks in Section 6.2.
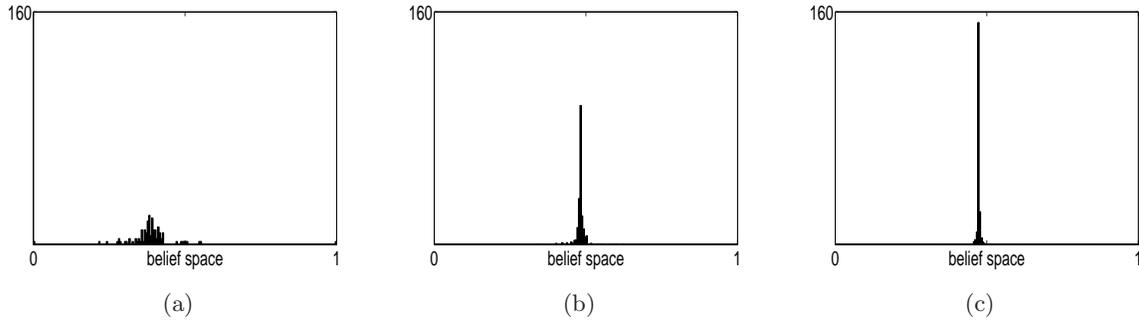
FIGURE 10. Homogeneous influence in Erdös-Rényi graphs of increasing sizes. The population sizes are $n = 100$ in (a), $n = 500$ in (b), and $n = 2000$ in (c), while $p = 2n^{-1} \log n$. In each case, there are two stubborn agents randomly selected from the node set $\mathcal{V}$ with uniform probability and holding beliefs equal to 0, and 1, respectively. The figures show the empirical density of the expected stationary beliefs, for typical realizations of such graphs. As predicted by Theorem 4, such empirical density tends to concentrate around a single value as the population size grows large.

6.2. Examples of large-scale social networks We now present some examples of families 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 Examples 1 and 5, starting from an undirected graph $\mathcal{G}$.

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 3. It is not hard to see that the minimum in the right-hand side of (38) is achieved by $\mathcal{W} = \mathcal{V}_0$, so that the conductance satisfies

$$\Phi = \left(\frac{n}{2} \left(\frac{n}{2} - 1 \right) + 1 \right)^{-1} \leq \frac{4}{(n+1)^2} \,.$$

It then follows from (39) that $\tau \geq (4\Phi)^{-1} \geq (n+1)^2/16$. Since $d_v \geq n/2 - 1$ for all v, it follows that the barbell-like network is never highly fluid provided that $|\mathcal{S}| \geq 1$. In fact, we have already seen in Example 3 that the expected stationary beliefs polarize in this case, so that the influence of the stubborn agents on the rest of the society is not homogeneous.

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 4. Since this is a regular graph, one has $\pi_* n = 1$, and $\pi(\mathcal{S}) = |\mathcal{S}|/n$. Moreover, it is well known that (see, e.g., [31, Theorem 5.5]) $\tau \leq C_d n^{2/d}$, for some constant C_d depending on the dimension d only. Then, $\tau \pi(\mathcal{S}) \leq C_d |\mathcal{S}| n^{2/d-1}$. For $d \geq 3$, this implies that the social network with toroidal topology is highly fluid, and hence homogeneous influence holds, provided that $|\mathcal{S}| = o(n^{1-2/d})$.

In contrast, for $d \leq 2$, our arguments do not allow one to prove high fluidity of the social network. In fact, using the explicit calculations of Example 2, one can see that the stubborn agents' influence is not homogeneous in the case $d = 1$, since the expected stationary beliefs do not concentrate. On the other hand, in the case $d = 2$, we conjecture that, using the explicit expression (33) and Fourier analysis, one should be able to show that the condition $|\mathcal{S}| = o(n^{1/2})$ would be sufficient for homogeneous influence. In fact, a more general conjecture is that $|\mathcal{S}| = o(n^{1-1/d})$ should suffice for homogeneous influence, when $d \geq 2$. Proving this conjecture would require an analysis finer than the one developed in this section, possibly based on discrete Fourier transform techniques. The motivation behind our conjecture comes from thinking of a limit continuous model, which can be informally summarized as follows. First, recall that the expected stationary beliefs vector solves the Laplace equation on $\mathcal{G}$ with boundary conditions assigned on the stubborn agent set $\mathcal{S}$. Now,

consider the Laplace equation on a d-dimensional manifold with boundary conditions on a certain subset. Then, in order for the problem to be well-posed, such subset should have dimension $d-1$. Similarly, one should need $|\mathcal{S}| = \Theta(n^{1-1/d}) = \Theta(m^{d-1})$ in order to guarantee that the expected stationary beliefs vector is not almost constant in the limit of large n.

We now present four examples of random graph sequences which have been the object of extensive research. Following a common terminology, we say that some property of such graphs holds with high probability, if the probability that it holds approaches one in the limit of large population size n.

EXAMPLE 8. (**Connected Erdös-Rényi**) Consider the Erdös-Rényi 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 a link with probability p, independently from the others. We focus on the regime $p = cn^{-1}\log n$, with $c > 1$, where the Erdös-Rényi graph is known to be connected with high probability [21, Thm. 2.8.2]. In this regime, results by Cooper and Frieze [16] 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 [21, Lemma 6.5.2]. In particular, it follows that, with high probability, $(\pi_* n)^{-1} \leq 4/\delta$. Hence, using (37), one finds that the resulting social network is highly fluid, provided that $|\mathcal{S}| = o(n/\log n)$, as n grows large. Figure 10 shows the empirical density of the expected stationary beliefs for typical realizations of Erdös-Rényi graphs of increasing size $n = 100, 500, 2000$, and constant stubborn agents number $|\mathcal{S}| = 2$.

EXAMPLE 9. (**Fixed degree distribution**) Consider a random graph $\mathcal{G} = \mathcal{FD}(n, \lambda)$, generated as follows. Fix $\mathcal{V}$ with $|\mathcal{V}| \geq 2$, and let $\{d_v : v \in \mathcal{V}\}$ be a family of independent and identically distributed random variables with $\mathbb{P}(d_v = k) = \lambda_k$, for $k \in \mathbb{N}$. Assume that $\lambda_1 = \lambda_2 = 0$, that $\lambda_{2k} > 0$ for some $k \geq 2$, and that the first two moments $\sum_k \lambda_k k$, and $\sum_k \lambda_k k^2$ are finite. Then, let $\mathcal{G} = \mathcal{FD}(n, \lambda)$ be the multigraph of vertex set $\mathcal{V}$ generated by conditioning on the event $E_n := \{\sum_v d_v$ is even$\}$ (whose probability converges either to $1/2$ or to 1 as n grows large) and matching the vertices uniformly at random given their degree. (See [21, Ch. 3] for details on this construction) Then, results in [21, Ch. 6.3] show that the mixing time of the lazy random walk on $\mathcal{G}$ satisfies $\tau = O(\log n)$ with high probability. Therefore, using (37), one finds that the resulting social network is highly fluid with high probability provided that $\sum_s d_s = o(n/\log n)$.

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 [21, Ch. 4] and consider the random multigraph $\mathcal{G} = \mathcal{PA}(n, m)$ with n vertices, generated by starting with two vertices connected by m parallel links, 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 [21, Th. 4.1.4], the degree distribution converges in probability to the power law $\lambda_k = 2m(m+1)/k(k+1)(k+2)$, and the graph is connected with high probability [21, 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. [35] (see also [21, Th. 6.4.2]) imply that the mixing time of the lazy random walk satisfies $\tau = O(\log n)$, with high probability. Therefore, thanks to (37), the resulting social network is highly fluid with high probability if $\sum_s d_s = o(n/\log n)$.

EXAMPLE 11. (**Watts & Strogatz's small world**) Watts and Strogatz [50], and then Newman and Watts [40] 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
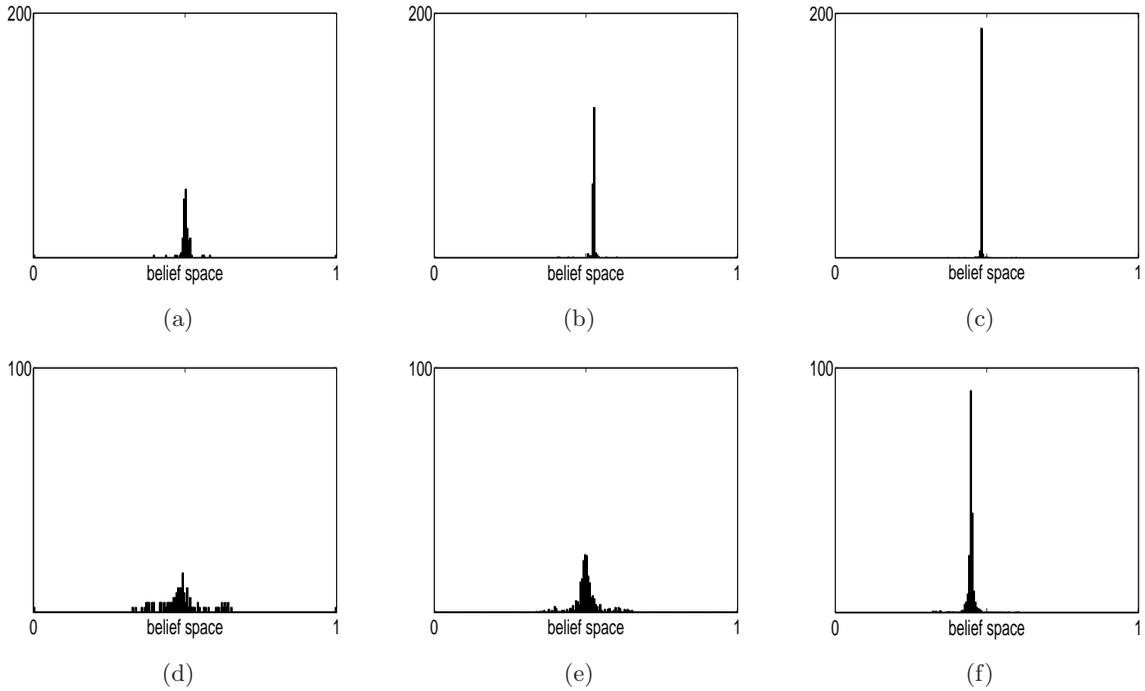
FIGURE 11. Homogeneous influence in preferential attachment networks of increasing sizes. The figures show the empirical density of the expected stationary beliefs, for typical realizations of such graphs. The population sizes are $n = 100$ in (a) and (d), $n = 500$ in (b) and (e), and $n = 2000$ in (c) and (f), while $m = 4$. In each case, there are two stubborn agents holding beliefs equal to 0, and 1, respectively. In (a), (b), and (c), the stubborn agents are chosen to coincide with the two latest attached nodes, and therefore tend to have the lowest degree. In contrast, in (d), (e), and (f), the stubborn agents are chosen to be the two initial nodes, and therefore tend to have the highest degrees. As predicted by Theorem 4, the empirical density of the expected stationary beliefs tends to concentrate around a single value as the population size grows large. The rate at which this concentration occurs is faster in the top three figures, where $\sum_s d_s$ is smaller, and slower in the bottom three figures, where $\sum_s d_s$ is larger.

case, the average degree remains bounded with high probability as n grows large, while results by Durrett [21, Th. 6.6.1] show that the mixing time $\tau = O(\log^3 n)$. This, and (37) imply that the network is highly fluid with high probability provided that $\sum_s d_s = o(n/\log^3 n)$.

6.3. Proof of Theorem 4 In order to prove Theorem 4, we will obtain estimates on the hitting probability distributions γ^v.

The following result 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 5. *Let the social network satisfy Assumption 1 and $P \in \mathcal{P}$. Then, for all $P \in \mathcal{P}$, and $k \geq 0$,*

$$||\gamma^v - \overline{\gamma}||_{TV} \leq \mathbb{P}_v(U_{\mathcal{S}} < k) + \exp(-\lfloor k/\tau \rfloor), \qquad v \in \mathcal{V}, \tag{47}$$

where τ is the mixing time of the discrete-time Markov chain $W(k)$ with transition probability matrix P (cf. (36)), and $U_{\mathcal{S}} := \min\{k \geq 0 : W(k) \in \mathcal{S}\}$ is the hitting time of such chain on $\mathcal{S}$.

Proof. Notice that (47) is trivial when $k = 0$ or $v \in \mathcal{S}$. For $k \geq 1$ and $a \in \mathcal{A}$, one can reason as follows. Let $\tilde{U}_{\mathcal{S}} := \inf\{k' \geq k : W(k') \in \mathcal{S}\}$. Thanks to Lemma 4, one has that the distributions of $W(U_{\mathcal{S}})$ and $W(\tilde{U}_{\mathcal{S}})$, conditioned on $W(0) = a$, are given by γ^a, and $\sum_v p_v^a(k)\gamma^v$, respectively. Using the identity

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

(see, e.g., [31, Prop. 4.5]), and observing that the event $\{U_{\mathcal{S}} \geq k\}$ implies $\{W(U_{\mathcal{S}}) = W(\tilde{U}_{\mathcal{S}})\}$, one gets that

$$
\begin{aligned}
\left\|\gamma^a - \sum_v p_v^a(k)\gamma^v\right\|_{TV} &= \frac{1}{2} \sup_{f \in [-1,1]^{\mathcal{V}}} \left\{ \mathbb{E}_a\left[f(W(U_{\mathcal{S}})) - f(W(\tilde{U}_{\mathcal{S}})) \right] \right\} \\
&= \frac{1}{2} \sup_{f \in [-1,1]^{\mathcal{V}}} \left\{ \mathbb{E}_a\left[\mathbb{1}_{\{U_{\mathcal{S}}<k\}} \left(f(W(U_{\mathcal{S}})) - f(W(\tilde{U}_{\mathcal{S}})) \right) \right] \right\} \\
&\leq \mathbb{P}_a\left(U_{\mathcal{S}} < k\right).
\end{aligned}
$$

On the other hand, since the Markov kernel $v \mapsto \gamma^v$ is contractive in total variation distance, one has that

$$
\left\|\sum_v p_v^a(k)\gamma^v - \overline{\gamma}\right\|_{TV} = \left\|\sum_v (p_v^a(k) - \pi_v)\gamma^v\right\|_{TV} \leq \|p^a(k) - \pi\|_{TV}.
$$

Finally, submultiplicativity of the maximal total variation distance from the stationary distribution (see, e.g., [31, Lemma 4.12]) implies that

$$
\|p^a(k) - \pi\|_{TV} \leq \exp(-\lfloor k/\tau \rfloor).
$$

By applying the triangle inequality and the three bounds above, one gets that

$$
\|\gamma^a - \overline{\gamma}\|_{TV} \leq \left\|\gamma^a - \sum_v \gamma_v^a(t)\gamma^v\right\|_{TV} + \left\|\sum_v \gamma_v^a(t)\gamma^v - \overline{\gamma}\right\|_{TV} \leq \mathbb{P}_a\left(U_{\mathcal{S}} < k\right) + \exp(-\lfloor k/\tau \rfloor),
$$

thus proving the claim. ∎

Lemma 6, stated below, is the main technical result of this section.

LEMMA 6. *Consider a social network satisfying Assumption 1. Then,*

$$
\frac{1}{n}\left|\{v \in \mathcal{V} : \|\gamma^v - \overline{\gamma}\|_{TV} \geq \varepsilon\}\right| \leq \frac{1}{\Psi\varepsilon},
$$

for every $\varepsilon > 0$.

Proof. Fix an arbitrary $P \in \mathcal{P}$, and let $\pi = P'\pi$ be its invariant measure. Let $W(k)$ be a discrete-time Markov chain with transition probability matrix P. For every nonnegative integer k, stationarity of π and the union bound yield

$$
\mathbb{P}_\pi(U_{\mathcal{S}} < k) = \mathbb{P}_\pi\left(\bigcup_{0 \leq j < k} \{W(j) \in \mathcal{S}\}\right) \leq \sum_{0 \leq j < k} \mathbb{P}_\pi(W(j) \in \mathcal{S}) = k\pi(\mathcal{S}), \tag{48}
$$

Combining (48) with Lemma 5, one gets that

$$
\sum_v \pi_v \|\gamma^v - \overline{\gamma}\|_{TV} \leq \mathbb{P}_\pi(U_{\mathcal{S}} < k) + \exp(-\lfloor k/\tau \rfloor) \leq k\pi(\mathcal{S}) + \exp(-\lfloor k/\tau \rfloor).
$$

Choosing $k = \tau\lfloor \log(e/(\tau\pi(\mathcal{S}))) \rfloor$ yields

$$
\frac{1}{n}\sum_v \|\gamma^v - \overline{\gamma}\|_{TV} \leq \sum_v \frac{\pi_v}{n\pi_*}\|\gamma^v - \overline{\gamma}\|_{TV} \leq \frac{1}{n\pi_*}\left(\tau\pi(\mathcal{S})\log\frac{e}{\tau\pi(\mathcal{S})} + \exp\left(-\log\frac{e}{\tau\pi(\mathcal{S})} + 1\right)\right) = \frac{1}{\psi(P)}.
$$

Then,

$$
\frac{\varepsilon}{n}\left|\{v \in \mathcal{V} : \|\gamma^v - \overline{\gamma}\|_{TV} \geq \varepsilon\}\right| \leq \frac{1}{n}\sum_v \|\gamma^v - \overline{\gamma}\|_{TV} \leq \frac{1}{\psi(P)}.
$$

The claim now follows from the arbitrarinessss of $P \in \mathcal{P}$. ∎

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 $|y_s| \le \Delta_*/2$. Then, it follows from Theorem 3 that

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

so that (45) immediately follows from Lemma 6.

On the other hand, in order to show (46), first recall that, if $\theta_e = 1$ for all $e \in \overrightarrow{\mathcal{E}}$, then Eq. (12) provides the transition rates of coalescing Markov chains. In particular, if $V(0) = V'(0)$, then $V(T_{\mathcal{S}}) = V'(T'_{\mathcal{S}})$, so that $\eta_{ss'}^{vv} = \gamma_s^v$ if $s = s'$, and $\eta_{ss'}^{vv} = 0$ otherwise. Then, it follows from Theorem 3 that

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

$$= \sum_{s,s'} \eta_{ss'}^{vv} x_s x_{s'} - \left(\sum_s \gamma_s^v x_s \right)^2$$

$$= \sum_s \gamma_s^v x_s^2 - \left(\sum_s \gamma_s^v x_s \right)^2$$

$$= \frac{1}{2} \sum_s \sum_{s'} \gamma_s^v \gamma_{s'}^v (x_s - x_{s'})^2 \,.$$

Similarly,

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

so that

$$|\sigma_v^2 - \sigma_Z^2| \le \frac{1}{2} \sum_{s,s'} |\gamma_s^v \gamma_{s'}^v - \overline{\gamma}_s \overline{\gamma}_{s'}| (x_s - x_{s'})^2$$

$$\le \frac{1}{2} \sum_{s,s'} |\gamma_s^v \gamma_{s'}^v - \overline{\gamma}_s \overline{\gamma}_{s'}| \Delta_*^2$$

$$\le \frac{1}{2} \sum_{s,s'} \left(\gamma_s^v |\gamma_{s'}^v - \overline{\gamma}_{s'}| + \overline{\gamma}_{s'} |\gamma_s^v - \overline{\gamma}_s| \right) \Delta_*^2$$

$$= ||\gamma^v - \overline{\gamma}||_{TV} \Delta_*^2 \,.$$

Now, (46) follows again from a direct application of Lemma 6. ■

REMARK 3. For a stochastic matrix $P \in \mathcal{P}$ which is reversible, i.e., such that $\pi_v P_{vw} = \pi_w P_{wv}$ for all $v, w \in \mathcal{V}$, (observe that this additional property is enjoyed by the matrix P considered in Example 5 for the canonical construction of a social network from an undirected graph $\mathcal{G}$) one can potentially obtain tighter estimates on the homogeneity of the agents' influence. In fact, one could use the results on the approximate exponentiality of hitting times (i.e., the property that the distribution of $T_{\mathcal{S}}/\mathbb{E}_\pi[T_{\mathcal{S}}]$ is close to a rate-1 exponential distribution, see, e.g., [4, Ch. 3.5]) in order to show that, for a continuous-time Markov chain with transition rate matrix $P - I$, one has $\mathbb{P}_\pi(T_{\mathcal{S}} \ge t) \le (t + \tau)/\mathbb{E}_\pi[T_{\mathcal{S}}]$ for all $t \ge 0$. Using this bound in place of (48), arguments analogous to those developed in this section imply that $\tau/\mathbb{E}_\pi[T_{\mathcal{S}}] = o(1)$ is a sufficient condition for homogeneous influence. Observe that, using Markov's inequality and (48) with $k = \lfloor 1/(2\pi(\mathcal{S})) \rfloor$, gives

$$\mathbb{E}_\pi[T_{\mathcal{S}}] = \mathbb{E}_\pi[U_{\mathcal{S}}] \ge \left\lfloor \frac{1}{2\pi(\mathcal{S})} \right\rfloor \mathbb{P}_\pi \left(U_{\mathcal{S}} \ge \left\lfloor \frac{1}{2\pi(\mathcal{S})} \right\rfloor \right) \ge \left\lfloor \frac{1}{2\pi(\mathcal{S})} \right\rfloor \left(1 - \pi(\mathcal{S}) \left\lfloor \frac{1}{2\pi(\mathcal{S})} \right\rfloor \right) \ge \frac{1 - 4\pi^2(\mathcal{S})}{4\pi(\mathcal{S})} \,.$$

Hence, this argument would provide potentially a weaker sufficient condition for homogenous influence in situations where $\pi(\mathcal{S}) = o(1/\mathbb{E}_\pi[T_{\mathcal{S}}])$.

7. Conclusion   In this paper, we have studied a possible mechanism explaining persistent disagreement and opinion fluctuations in social networks. We have considered an inhomogeneous stochastic gossip model of continuous opinion dynamics, whereby some stubborn agents in the network never change their opinions. We have shown that the presence of these stubborn agents leads to persistent fluctuations and disagreements among the rest of the society: the beliefs of regular agents do not converge almost surely, and keep on fluctuating in an ergodic fashion. A duality argument allows for characterizing expected stationary beliefs in terms of the hitting probabilities of a Markov chain on the graph describing the social network, while the correlation between the stationary beliefs of any pair of regular agents can be characterized in terms of the hitting probabilities of a pair of coupled Markov chains. We have shown that in highly fluid social networks, whose associated Markov chains have mixing times which are sufficiently smaller than the inverse of the stubborn agents' set size, the vectors of the stationary expected beliefs and variances are almost constant, so that the stubborn agents have homogeneous influence on the rest of the society. We wish to emphasize that homogeneous influence in highly fluid societies needs not imply approximate consensus among the agents, whose beliefs may well fluctuate in an almost uncorrelated way. A deeper understanding of this topic is ongoing work.

References
[1] D. Acemoglu, K. Bimpikis, and A. Ozdaglar, *Dynamics of information exchange in endogenous social networks*, available at http://econ-www.mit.edu/files/6606, 2011.

[2] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, *Bayesian learning in social networks*, Review of Economic Studies **78** (2011), no. 4, 1201–1236.

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

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

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

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

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

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

[9] E. Behrends, *Introduction to markov chains with special emphasis on rapid mixing*, Vieweg Verlag, 2000.

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

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

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

[13] G. Cohen, *Party over policy: The dominating impact of group influence on political beliefs*, Journal of Personality and Social Psychology **85** (2003), no. 5, 808–822.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

[29] G. H. Kramer, *Short-term fluctuations in U.S. voting behavior, 1896-1964*, American Political Science Review **65** (1971), no. 1, 131–143.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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