Opinion Dynamics in Social Networks: A Local Interaction Game with Stubborn Agents

Javad Ghaderi and R. Srikant
Department of ECE and Coordinated Science Lab., University of Illinois at Urbana-Champaign

***Abstract*—The process by which new ideas, innovations, and behaviors spread through a large social network can be thought of as a networked interaction game: Each agent obtains information from certain number of agents in his friendship neighborhood, and adapts his idea or behavior to increase his benefit. In this paper, we are interested in how opinions, about a certain topic, form in social networks. We model opinions as continuous scalars ranging from 0 to 1 with $1(0)$ representing extremely positive(negative) opinion. Each agent has an initial opinion and incurs some cost depending on the opinions of his neighbors, his initial opinion, and his stubbornness about his initial opinion. Agents iteratively update their opinions based on their own initial opinions and observing the opinions of their neighbors. The iterative update of an agent can be viewed as a myopic cost-minimization response (i.e., the so-called best response) to the others' actions. We study whether an equilibrium can emerge as a result of such local interactions and how such equilibrium possibly depends on the network structure, initial opinions of the agents, and the location of stubborn agents and the extent of their stubbornness. We also study the convergence speed to such equilibrium and characterize the convergence time as a function of aforementioned factors. We also discuss the implications of such results in a few well-known graphs including small-world graphs.**

I. Introduction

Rapid expansion of online social networks, such as friendships and information networks, in recent years has raised an interesting question: how do opinions form in a social network? The opinion of each person is influenced by many factors such as his friends, news, political views, profession, etc. Understanding such interactions and predicting how specific opinions spread throughout social networks has triggered vast research by economists, sociologist, psychologies, physicists, etc.

We consider a social network consisting of n agents. We model the social network as a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ where agents are the vertices and edges indicate pairwise acquaintances. We model opinions as continuous scalars ranging from 0 to 1 with 1(0) representing extremely positive(negative) opinion. For example, such scalers could represent people opinions about the economic situation of the country, ranging from 0 to 1, with an opinion 1 corresponding to perfect satisfaction with the current economy and 0 representing an extremely negative view towards the economy. Agents have some private initial opinions and iteratively update their opinions based on their own initial opinions and observing the opinions of their neighbors. We study whether an equilibrium can emerge as a result of such local interactions and how such equilibrium possibly depends on the graph structure and initial opinions of the agents. In the interaction model, we also incorporate *stubbornness* of agents with respect to their initial opinions and investigate the dependency of the equilibrium on such stubborn agents. Characterizing the convergence rate to the equilibrium as a function of graph structure, location of stubborn agents and their levels of stubbornness is another goal of the current paper.

There has been an interesting line of research trying to explain emergence of new phenomenon, such as spread of innovations and new technologies, based on local interactions among agents, e.g., [6], [8], [9], [23]. Roughly speaking, a *coordination game* is played among the agents in which adopting a common strategy has a higher payoff. Agents behave according to a noisy version of the best-response dynamics which drives the system to a particular equilibrium in which all agents take the same action.

There is a rich and still growing literature on social learning using a Bayesian perspective where individuals observe the actions of others and update their beliefs iteratively about an underlying state variable, e.g., [11], [12], [13]. There is also opinion dynamics based on non-Bayesian models, e.g., those in [1], [2], [3], [10], [14]. In particular, [14] investigates a model in which agents meet and adopt the average of their pre-meeting opinions. As reported in [14], it is significantly more difficult to analyze social networks with several forceful agents that do not change their opinions and requires a different mathematical approach. Our model is closely related to the non-Bayesian framework, this keeps the computations tractable and can characterize the equilibrium in presence of agents that are biased towards their initial opinions (the so-called partially stubborn agents in our paper) or do not change their opinions at all (the so-called fully stubborn agents in our paper). The recent work [15] studies opinion dynamics based on the so-called voter model where each agent holds a binary 0-1 opinion and at each time a randomly chosen agent adopts the opinion of one of his neighbors, and there are also stubborn agents that do not change their states. Under such model, [15] shows that the opinions converge in distribution and characterizes the first and the second moments of this distribution. In addition, our paper is also related to consensus problems in which the question of interest is whether beliefs (some scalar numbers) held by different agents will converge to a common value, e.g., [16], [17], [18], [19].

This research was supported by AFOSR MURI grant FA 9550-10-1-0573 and ARO MURI grant W911NF-12-1-0385.

We use the following basic notations. x^T denotes the transpose of vector x. Given two functions f and g, $f = O(g)$ if $\sup_n |f(n)/g(n)| < \infty$. $f = \Omega(g)$ if $g = O(f)$. If both $f = O(g)$ and $f = \Omega(g)$, then $f = \Theta(g)$.

Due to page constraints, we have to omit proofs of the results and some of the details but they are provided in the complete version of the paper [31].

II. Model and Definitions

Consider a social network with n agents, denoted by a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ where agents are the vertices and edges indicate the pairs of agents that have interactions. For each agent i, define its neighborhood ∂_i as the set of agents that node i interacts with, i.e., $\partial_i := \{j : (i,j) \in \mathcal{E}\}$. Each agent i has an initial opinion $x_i(0) \in [0,1]$. Let $x(0) := [x_1(0) \cdots x_n(0)]^T$ denote the vector of initial opinions. We assume each agent i has a cost function of the form

$$J_i(x_i, x_{\partial_i}) = \frac{1}{2}\sum_{j\in\partial_i}(x_i - x_j)^2 + \frac{1}{2}K_i(x_i - x_i(0))^2, \qquad (1)$$

that he tries to minimize where $K_i \geq 0$ measures the *stubbornness* of agent i regarding his initial opinion[1]. When none of the agents are stubborn, correspondingly K_i's are all zero, the above formulation defines a *coordination game* with continuous payoffs because any vector of opinions $x = [x_1 \cdots x_n]^T$ with $x_1 = x_2 = \cdots = x_n$ is a *Nash equilibrium*. Here, we consider a synchronous version of the game between the agents. At each time, every agent observes the opinions of his neighbors and updates his opinion based on these observations and also his own initial opinion in order to minimize his cost function. It is easy to check that, for every agent i, the best-response strategy is

$$x_i(t+1) = \frac{1}{d_i + K_i}\sum_{j\in\partial_i} x_j(t) + \frac{K_i}{d_i + K_i}x_i(0), \qquad (2)$$

where $d_i = |\partial_i|$ is the degree of node i in graph $\mathcal{G}$. Define a matrix $A_{n\times n}$ such that $A_{ij} = \frac{1}{d_i + K_i}$ for $(i,j) \in \mathcal{E}$ and zero otherwise. Also define a diagonal matrix $B_{n\times n}$ with $B_{ii} = \frac{K_i}{d_i + K_i}$ for $1 \leq i \leq n$. Thus, in the matrix form, the best-response dynamics are given by

$$x(t+1) = Ax(t) + Bx(0). \qquad (3)$$

Iterating (3) shows that the vector of opinions at each time $t \geq 0$ is

$$x(t) = A^t x(0) + \sum_{s=0}^{t-1} A^s B x(0). \qquad (4)$$

In the rest of the paper, we investigate the existence of equilibrium, $x(\infty) := \lim_{t\to\infty} x(t)$, under the dynamics (3) in different social networks, with or without stubborn agents. We also characterize the convergence time of the dynamics, i.e., the amount of time that it takes for the agents' opinions to get close to the equilibrium. The equilibrium behavior is relevant only if the convergence time is reasonable [9].

Without loss of generality, assume that $\mathcal{G}$ is a connected graph (otherwise, we can consider opinion dynamics separately over each connected subgraph).

III. No Stubborn Agents

Convergence issues in the case of no stubborn agents is a special case of consensus, and has been well studied. Here, we briefly review this work to put our later results in context.

When there are no stubborn agents, A can be viewed as a transition probability matrix of a random walk over $\mathcal{G}$ (see [31] for all details). Then, it is not difficult to show that, in non-bipartite social graphs, the best-response dynamics will converge to the following unique equilibrium

$$x_i(\infty) = \frac{1}{2|\mathcal{E}|}\sum_{j=1}^{n} d_j x_j(0); \text{ for all } i \in \mathcal{V}, \qquad (5)$$

i.e., agents will reach a *consensus* in equilibrium where the impact of each agent on the equilibrium consensus is directly proportional to its degree. Also it can be shown that the error, $e(t) := x(t) - x(\infty)$, goes to zero geometrically, in some appropriate norm, at a rate equal to the *second largest eigenvalue modulus* of A. In the case that the social graph is bipartite, e.g., a ring graph with an even number of agents n, the best-response dynamics might not converge to an equilibrium. In practice, not everyone completely ignores his own previous opinion and might be slightly biased by his old opinion. Hence, we can consider a noisy version of the best-response dynamics as follows

$$x_i^{(\epsilon)}(t+1) = (1-\epsilon)\Big(\frac{1}{d_i}\sum_{j\in\partial_i} x_j(t)\Big) + \epsilon x_i(t), \qquad (6)$$

for some *self-confidence* $\epsilon > 0$[2]. Then, it can be shown that the noisy best-response dynamics, in bipartite or non-bipartite social graphs, will converge to the equilibrium (5) *independently of* ϵ. Moreover, the convergence time is $\Theta\Big(\frac{1}{1-\lambda_2(A)}\Big)$, as n grows, where $\lambda_2(A)$ is the second largest eigenvalue of A.

Example 1. *In this example, we make a comparison between two extreme interaction scenarios, namely, ring graph and complete graph with n nodes. The complete graph represents situation when all agents can communicate with each other with no constrains, while in the ring graph, each agent can only communicate with his two nearest neighbors. In both cases the (noisy) best-response dynamics converge to the average of initial opinions. It is easy to see that $\lambda_2(A)$ is $\frac{1}{n-1}$ for the complete graph and $\cos(\frac{2\pi}{n}) \approx 1 - \frac{\pi}{n^2}$ for the ring. Hence, while both of the graphs have the same equilibrium, convergence in the complete graph is much faster than convergence in the ring, in fact, $O(1)$ vs. $O(n^2)$.*

[1] Although we have considered uniform weights for the neighbors, the results in the paper hold under a more general setting when each agent puts a weight w_{ij} for his neighbor j.

[2] Here, we assume all agents have the same self-confidence but the argument can be adapted for different self-confidences as well.

There is a rich literature on approximating $\lambda_2(A)$ of the random walk A over different types of graphs, e.g., see Chapter 2 of [5] for a survey. Intuitively, the convergence time is deteriorated by the highly connected component of the graph which is loosely connected to the rest of the network (captured by the notion of *conductance*, or the edge *isoperimetric* function of the graph). Thus, we do not proceed in this direction further and in the next section, we study the more interesting case of social networks with stubborn agents.

IV. Impact of Stubborn Agents

A. Existence and characterization of equilibrium

Consider a connected social network $\mathcal{G}(\mathcal{V},\mathcal{E})$ in which at least one of the agents is stubborn, i.e., $K_i > 0$ for at least one $i \in \mathcal{V}$. Then A is an irreducible *sub-stochastic* matrix with the row-sum of at least one row less than one. Let $\rho_1(A) := \max_i |\lambda_i(A)|$ denote the spectral radius of A. It is well-known that $\rho_1(A)$ of a sub-stochastic matrix A is less than one, and hence, $\lim_{t\to\infty} A^t = 0$. Therefore, by *Perron-Ferobenius* theorem, the largest eigenvalue should be positive, real $1 > \lambda_1 > 0$ and $\rho_1(A) = \lambda_1$. Hence, in this case, based on (4), the equilibrium exists and is equal to

$$x(\infty) := \lim_{t\to\infty} x(t) = \sum_{s=0}^{\infty} A^s B x(0) = (I-A)^{-1} B x(0). \quad (7)$$

Therefore, since $B_{ii} = 0$ for all non-stubborn agents i, the initial opinions of non-stubborn agents will vanish eventually and have no effect on the equilibrium (7).

The matrix form (7) does not give any insight on how the equilibrium depends on the graph structure and the stubborn agents. Next, we describe the equilibrium in terms of explicit quantities that depend on the graph structure, location of stubborn agents and their levels of stubbornness.

Let $\mathcal{S} \subseteq \mathcal{V}$ be the set of stubborn agents and $|\mathcal{S}| \geq 1$. Any agent i in $\mathcal{S}$ is either *fully stubborn*, meaning its corresponding $K_i = \infty$, or it is *partially stubborn*, meaning $0 < K_i < \infty$. Hence, $\mathcal{S} = \mathcal{S}_F \cup \mathcal{S}_P$ where $\mathcal{S}_F$ is the set of fully stubborn agents and $\mathcal{S}_P$ is the set of partially stubborn agents. Next, we construct a *weighted graph* $\hat{\mathcal{G}}(\hat{\mathcal{V}},\hat{\mathcal{E}})$ as follows. Assign weight 1 to all the edges of $\mathcal{G}$. Connect a new vertex u_i to each $i \in \mathcal{S}_P$ and assign a weight K_i to the corresponding edge. Let $\hat{\mathcal{V}} := \mathcal{V} \cup \{u_i : i \in \mathcal{S}_P\}$ and $\hat{\mathcal{E}} := \mathcal{E} \cup \{(i,u_i) : i \in \mathcal{S}_P\}$. Also let w_{ij} denote the weight of edge $(i,j) \in \hat{\mathcal{E}}$. Then $\hat{\mathcal{G}}(\hat{\mathcal{V}},\hat{\mathcal{E}})$ is a weighted graph with weights $w_{ij} = 1$ for all $(i,j) \in \mathcal{E}$ (the edges of $\mathcal{G}$) and $w_{iu_i} = K_i$ for all $i \in \mathcal{S}_P$. Let $u(\mathcal{S}_P) = \{u_i : i \in \mathcal{S}_P\}$. Define $w_i := \sum_{j:(i,j)\in\hat{\mathcal{E}}} w_{ij}$ as the weighted degree of vertex $i \in \hat{\mathcal{V}}$. It should be clear that

$$w_i = \begin{cases} d_i + K_i & \text{for } i \in \mathcal{S}_P, \\ d_i & \text{for } i \in \mathcal{V}\backslash\mathcal{S}_P, \\ K_j & \text{for } i = u_j, j \in \mathcal{S}_P. \end{cases} \quad (8)$$

Consider the random walk $Y(t)$ over $\hat{\mathcal{G}}$ where the probability of transition from vertex i to vertex j is $P_{ij} = \frac{w_{ij}}{w_i}$. Assume the walk starts from some initial vertex $Y(0) = i \in \hat{\mathcal{V}}$. For any $j \in \hat{\mathcal{V}}$, define $\tau_j := \inf\{t \geq 0 : Y(t) = j\}$, as the first hitting time to vertex j. Also define $\tau := \bigwedge_{j\in\mathcal{S}_F\cup u(\mathcal{S}_P)} \tau_j$ as the first time that the random walk hits any of the vertices in $\mathcal{S}_F \cup u(\mathcal{S}_P)$. The following Lemma characterizes the equilibrium.

Lemma 1. *The best-response dynamics converge to a unique equilibrium where the opinion of each agent is a convex combination of the initial opinions of the stubborn agents. Based on the random walk over the graph $\hat{\mathcal{G}}$,*

$$x_i(\infty) = \sum_{j\in\mathcal{S}_P} \mathbb{P}_i(\tau = \tau_{u_j}) x_j(0) + \sum_{j\in\mathcal{S}_F} \mathbb{P}_i(\tau = \tau_j) x_j(0), \quad (9)$$

for all $1 \leq i \leq n$, where $\mathbb{P}_i(\tau = \tau_k)$, $k \in \mathcal{S}_F \cup u(\mathcal{S}_P)$, is the probability that the random walk hits vertex k first, among vertices in $\mathcal{S}_F \cup u(\mathcal{S}_P)$, given the random walk starts from vertex i.

Note that $\lim_{K_i\to\infty} \mathbb{P}_i(\tau = \tau_{u_i}) = 1$ for any partially stubborn agent $i \in \mathcal{S}_P$. This intuitively makes sense because as an agent i becomes more stubborn, his opinion will get closer to his own opinion and behaves similarly to a fully stubborn agent. It should be clear that when there is only one stubborn agent or there are multiple stubborn agents with identical initial opinions, eventually the opinion of every agent will converge to the same opinion as the initial opinion of the stubborn agents.

Alternatively, the equilibrium has an interesting electrical network interpretation as follows.

Lemma 2. *Consider $\mathcal{G}$ as an electrical network where the conductance of each edge is 1 and each stubborn agent i is a voltage source of $x_i(0)$ volts with an internal conductance K_i. Fully stubborn agents are ideal voltage sources with infinite internal conductance (zero internal resistance). Then, under the best-response dynamics, the opinion of each agent at equilibrium is just its voltage in the electrical network.*

We illustrate the use of the above lemma through the following example.

Example 2. *Consider a one-dimensional social graph, where agents are located on integers $1 \leq i \leq n$. Assume nodes 1 and n are stubborn with initial opinions $x_1(0)$ and $x_n(0)$, and stubbornness parameters $K_1 > 0$ and $K_n > 0$. Using the electrical network model, the current is the same over all edges and equal to $I = (x_1(0) - x_n(0))(\frac{1}{K_1} + \frac{1}{K_n} + n - 1)^{-1}$, and thus the voltage of each node i is $v_i = x_1(0) - I(\frac{1}{K_1} + i - 1)$, for $1 \leq i \leq n$. Hence,*

$$x_i(\infty) = (1-\alpha_i) x_1(0) + \alpha_i x_n(0),$$

where $\alpha_i := \frac{K_1^{-1}+i-1}{K_1^{-1}+K_n^{-1}+n-1}$.

B. Convergence time

Let $e(t) := x(t) - x(\infty)$ be the error vector. Trivially $e_i(t) = 0$ for all fully stubborn agents $i \in \mathcal{S}_F$, so we focus on $\tilde{e}(t) := [e_i(t) : i \in \mathcal{V}\backslash\mathcal{S}_F]^T$. For any real vector space $\mathbb{R}^m$, $m > 0$, define a convenient scalar product $\langle z, y\rangle_\pi := \sum_{i=1}^m z_i y_i \pi_i$ and its corresponding norm $\|z\|_\pi := \left(\sum_{i=1}^m z_i^2 \pi_i\right)^{1/2}$. The following lemma states that convergence to the equilibrium (7) is geometric with a rate equal to the largest eigenvalue of A.

Lemma 3. *Let $\tilde{\pi} = [\frac{w_i}{Z} : i \in \mathcal{V}\backslash\mathcal{S}_F]^T$, for w_i as in (8), and Z be the normalizing constant such that $\sum_{i\in\mathcal{V}\backslash\mathcal{S}_F} \tilde{\pi}_i = 1$. Then,*

$$\|\tilde{e}(t)\|_{\tilde{\pi}} \leq (\lambda_A)^t \|\tilde{e}(0)\|_{\tilde{\pi}}, \tag{10}$$

where λ_A is the largest eigenvalue of A.[3]

Defining the convergence time as $\tau(\nu) := \inf\{t \geq 0 : \|\tilde{e}(t)\|_{\tilde{\pi}} \leq \nu\}$, for some fixed $0 < \nu \ll 1$, shows that $\tau(\nu) = \Theta\left(\frac{1}{1-\lambda_A}\right)$ as n grows. Let $T := \frac{1}{1-\lambda_A}$. With a little abuse of terminology, we also call T the convergence time.

The exact characterization of λ_A in social networks with very large number of users and many stubborn agents is difficult, hence, we will derive appropriate upper-bounds and lower-bounds that depend on the graph structure, the location of stubborn agents and their levels of stubbornness.

Consider the weighted graph $\hat{\mathcal{G}}(\hat{\mathcal{V}}, \hat{\mathcal{E}})$. A *path* γ_{ij} from a vertex i to another vertex j in $\hat{\mathcal{G}}$ is a collection of *oriented* edges $\{(i, i_1), (i_1, i_2), \cdots, (i_m, j)\}$ that connects i to j and does not intersect itself. For any vertex $i \in \mathcal{V}\backslash\mathcal{S}_F$, consider a path γ_i from i to the set $\mathcal{S}_F \cup u(\mathcal{S}_P)$, i.e., $\gamma_i = \gamma_{ij}$ for some $j \in \mathcal{S}_F \cup u(\mathcal{S}_P)$.

Proceeding along the lines of Diaconis-Stroock [21], we get the following bound that yields an upper-bound on the convergence time.

Lemma 4. *Consider the weighted graph $\hat{\mathcal{G}}$. Given a set of paths $\{\gamma_i : i \in \mathcal{V}\backslash\mathcal{S}_F\}$, from $\mathcal{V}\backslash\mathcal{S}_F$ to $\mathcal{S}_F \cup u(\mathcal{S}_P)$, let $|\gamma_i|_w := \sum_{(s,t)\in\gamma_i} \frac{1}{w_{st}}$. Then, $T \leq 2\max_{(x,y)\in\hat{\mathcal{E}}} \xi(x, y)$, where, for each edge $(x, y) \in \hat{\mathcal{E}}$,*

$$\xi(x, y) := \sum_{i:\gamma_i\ni(x,y)} w_i|\gamma_i|_w.$$

It is also possible to proceed along the lines of Sinclair [22]. This gives a different bound stated in the following lemma.

Lemma 5. *Consider the weighted graph $\hat{\mathcal{G}}(\hat{\mathcal{V}}, \hat{\mathcal{E}})$. Given a set of paths $\{\gamma_i : i \in \mathcal{V}\backslash\mathcal{S}_F\}$ from $\mathcal{V}\backslash\mathcal{S}_F$ to $\mathcal{S}_F \cup u(\mathcal{S}_P)$, we have $T \leq 2\max_{(x,y)\in\hat{\mathcal{E}}} \eta(x, y)$, where, for each edge $(x, y) \in \hat{\mathcal{E}}$,*

$$\eta(x, y) := \frac{1}{w_{xy}} \sum_{i:\gamma_i\ni(x,y)} w_i|\gamma_i|.$$

Intuitively, both $\xi(x, y)$ and $\eta(x, y)$ are measures of *congestion* over the edge (x, y) due to paths that pass through (x, y). See [31] for examples of applications of the above bounds in complete and ring graphs and performance comparison with numerical values.

An upper bound on $1 - \lambda_A$, and thus a lower-bound on the convergence time T, is given by the following lemma.

Lemma 6. *Consider the weighted graph $\hat{\mathcal{G}}(\hat{\mathcal{V}}, \hat{\mathcal{E}})$, then*

$$1 - \lambda_A \leq \min_{U\subseteq\mathcal{V}\backslash\mathcal{S}_F} \psi(U; \hat{\mathcal{G}}), \tag{11}$$

where $\psi(U; \hat{\mathcal{G}}) := \frac{\sum_{i\in U, j\notin U} w_{ij}}{\sum_{i\in U} w_i}$. The minimum is achieved for some connected subgraph with vertex set U.

It is worth emphasizing that the above bounds are quite general and hold for social networks with any finite size and any set of stubborn agents. Next, to gain more insight into factors dominating the convergence speed, we consider the special class of *shortest paths* in social networks with large number of agents.

C. Canonical bounds via shortest paths

Let $\gamma = \{\gamma_i : i \in \mathcal{V}\backslash\mathcal{S}_F\}$ be the set of *shortest paths* from vertices $\mathcal{V}\backslash\mathcal{S}_F$ to a the set $\mathcal{S}_F \cup u(\mathcal{S}_P)$, so, in fact, for each $i \in \mathcal{V}\backslash\mathcal{S}_F$, $\gamma_i = \gamma_{ij}$ for some $j \in \mathcal{S}_F \cup u(\mathcal{S}_P)$. Let $\Gamma_j \subseteq \mathcal{V}\backslash\mathcal{S}_F$ be the set of nodes connected to $j \in \mathcal{S}_F \cup u(\mathcal{S}_P)$ via the shortest paths. We use $|\gamma| := \max_{i\in\mathcal{V}\backslash\mathcal{S}_F} |\gamma_i|$ to denote the maximum length of any shortest path and $|\Gamma| := \max_{j\in\mathcal{S}_F\cup u(\mathcal{S}_P)} |\Gamma_j|$ to denote the maximum number of nodes connected to any node in $\mathcal{S}_F \cup u(\mathcal{S}_P)$ via shortest paths.

Using Lemma 5, for each partially stubborn agent $j \in \mathcal{S}_P$,

$$\eta(j, u_j) = \frac{1}{K_j}(K_j + d_j + \sum_{i\in\Gamma_j} d_i|\gamma_i|) \leq 1 + \frac{\hat{d} + |\gamma||\Gamma|\tilde{d}}{K_{min}},$$

where $\tilde{d} := \max_{i\in\mathcal{V}\backslash\mathcal{S}} d_i$ is the maximum degree of non-stubborn agents, $\hat{d} := \max_{i\in\mathcal{S}} d_i$ is the maximum degree of stubborn agents, and $K_{min} := \min_{j\in\mathcal{S}_P} K_j$ is the minimum stubbornness. Hence, the congestion is dominated by some edge (j, u_j), $j \in \mathcal{S}_P$, only if the stubbornness K_j is sufficiently small. All the paths that pass through an internal edge $(x, y) \in \mathcal{E}$ are connected to the same $j \in \mathcal{S}_F \cup u(\mathcal{S}_P)$, or equivalently to the same stubborn agent. So for each $(x, y) \in \mathcal{E}$, $\eta(x, y) = \sum_{i:\gamma_i\ni(x,y)} d_i|\gamma_i| \leq |\gamma|B\tilde{d}$, where

$$B := \max_{(x,y)\in\mathcal{E}} |\{i : \gamma_i \ni (x, y)\}|, \tag{12}$$

is the *bottleneck constant*, i.e., the maximum number of shortest paths that pass through any edge of the social network $\mathcal{G}$. It is clear that $|\Gamma|/\hat{d} \leq B \leq |\Gamma|$ because B is at least equal to the number of paths that pass through an edge directly connected to a stubborn agent. Therefore, for $K_{min} \leq K^* := \frac{\hat{d}+|\gamma||\Gamma|\tilde{d}}{|\gamma|B\tilde{d}-1}$, η is dominated by congestion over some edge (j, u_j), $j \in \mathcal{S}_P$, and in this regime

$$T \leq 2\left(1 + \frac{\hat{d} + |\gamma||\Gamma|\tilde{d}}{K_{min}}\right). \tag{13}$$

For $K_{min} > K^*$, η is dominated by an edge of the social network which is the bottleneck, and in this regime

$$T \leq 2|\gamma|B\tilde{d}. \tag{14}$$

Dependence on $|\gamma|$, in both regimes, intuitively makes sense as it represents the minimum time required to reach any node in the network from stubborn agents. It is worth pointing out that adding more fully stubborn agents, with not necessarily equal initial opinions, or increasing the stubbornness of the agents makes the convergence faster.

[3] In Euclidian norm, $\|e(t)\|_2 \leq (\lambda_A)^t \sqrt{\frac{w_{max}}{w_{min}}}\|e(0)\|_2$, where $w_{max} := \max_{i\in\mathcal{V}\backslash\mathcal{S}_F} w_i$ and $w_{min} := \min_{i\in\mathcal{V}\backslash\mathcal{S}_F} w_i$.

D. Scaling laws in large social networks

For any social network, we can consider two cases: (i) All the stubborn agents are partially stubborn (ii) At least, one of the agents is fully stubborn. In both cases, the upper-bound on the convergence time is given by (13) and (14) depending on the levels of stubbornness of partially stubborn agents. In case (ii), if all the stubborn agents are fully stubborn, the upper-bound is given by (14).

To find a simple lower-bound, we consider the set U in (11) to include all the nodes $\mathcal{V}\backslash\mathcal{S}_F$. This gives the following lower-bound

$$T \geq \frac{\sum_{j\in\mathcal{S}_P} K_j + 2|\mathcal{E}| - \sum_{j\in\mathcal{S}_F} d_j}{\sum_{j\in\mathcal{S}_P} K_j + \sum_{j\in\mathcal{S}_F} d_j}. \quad (15)$$

In investigating the scaling laws, the scaling of the number of stubborn agents and their levels of stubbornness with n could play an important role. Here, we study scaling laws in graphs with a fixed number of stubborn agents and with fixed levels of stubbornness, as the total number of agents n in the network grows. Then, in any connected graph $\mathcal{G}$, based on (15), the smallest possible convergence time is $T = \Omega(|\mathcal{E}|)$ in the case (i) which could be as small as $\Omega(n)$, and $T = \Omega(\frac{|\mathcal{E}|}{\sum_{j\in\mathcal{S}_F} d_j})$ in the case (ii) which could be as small as $\Omega(1)$. It is possible to combine the upperbounds (13) and (14) to obtain the following (looser) upper-bound that holds for social networks with any fixed number of (partially/fully) stubborn agents and fixed levels of stubbornness

$$T = O(n\delta d_{max}). \quad (16)$$

Here, d_{max} is the maximum degree of the social graph (could possibly depend on n) and δ is the *diameter* of the graph.

Fastest convergence: It should be intuitively clear that a *star graph* $\mathcal{G}$, where a stubborn agent directly connected to $n-1$ non-stubborn agents with no edges between the non-stubborn agents, should have the fastest convergence. In fact, it is easy to check that $K^* = \Theta(n)$, hence, if the stubborn agent is partially stubborn (case (i)), by (13), $T = \Theta(n)$, and if the stubborn agent is fully stubborn (case (ii)), by (14), $T = \Theta(1)$, both achieving the smallest possible lower-bounds.

Complete graph and Ring graph: In the complete graph, with a fixed number of stubborn agents, $\tilde{d} = \hat{d} = n-1$, $|\Gamma| = \Theta(n)$, $|\gamma| = 2$, $B = 1$, and $K^* = \Theta(n)$. Hence, if at least one of the agents is partially stubborn, by (13) and (15), $T = \Theta(n^2)$. If all the stubborn agents are fully stubborn, by (14) and (15), $T = \Theta(n)$.

In the ring network, $\tilde{d} = \hat{d} = 2$, $|\Gamma| = \Theta(n)$, $|\gamma| = \Theta(n)$, $B = \Theta(n)$, and $K^* = \Theta(1)$. Thus $T = O(n^2)$ and $\Omega(n)$ in both cases (i) and (ii).

None of the graphs always has a faster convergence than the other one. For example, in the case of one stubborn agent with a fixed K_1, and n large enough (larger than a constant depending on the value of K_1), the ring network has a faster convergence than the complete graph, while for any fixed n, and K_1 large enough, the complete graph has a faster convergence than the ring.

Erdos-Renyi random graphs, Expander graphs, and Trees: Using the well-known results on the maximum degree and the diameter of such graphs, we can find appropriate bounds on the convergence time of such graphs based on (15) and (16) (see [31] for all the results). Such graph models do not capture many spatial and structural aspects of social networks and, hence, are not a realistic model of social networks [24]. Motivated by the *small-world phenomenon* observed by Milgram [25], Strogatz-Watts [26] and Kleinberg [27] proposed models that illustrate how graphs with spatial structure can have small diameters, thus, providing more realistic models of social networks. Next, we consider a two-dimensional model, proposed in [24], for simplicity but results are extendable to the higher dimensional graphs as well.

Small-world graphs: Start with a social network as a grid $\sqrt{n} \times \sqrt{n}$ of n nodes. Hence, nodes i and j are neighbors if their l_1 distance $\|i-j\| = |x_i - x_j| + |y_i - y_j|$ is equal to 1. It follows from (16) that, in presence of a fixed number of stubborn agents in a bounded degree graph, $T = O(n\delta)$, and in the grid, $\delta = 2\sqrt{n}$ obviously, hence $T = O(n\sqrt{n})$. Note that changing the location of the stubborn agents can change the convergence time only by a constant and does not change the order.

Now assume that each node creates q shortcuts to other nodes in the network. A node i chooses another node j as the destination of the shortcut with probability $\frac{\|i-j\|^{-\alpha}}{\sum_{k\neq i}\|i-k\|^{-\alpha}}$, for some parameter $\alpha > 0$. Parameter α determines the distribution of the shortcuts as large values of α produce mostly local shortcuts and small values of α increase the chance of long-range shortcuts. In particular, $q = 1$ and $\alpha = 0$ recovers the Strogatz-Watts model where the shortcuts are selected uniformly at random. It is shown in [28] that for $\alpha < 2$, the graph is an expander with high probability. Hence, using the inequality between the diameter and the spectral gap [29], its diameter is of the order of $O(\log n)$ with high probability. It can be shown that, under the small-world network model, the maximum degree $d_{max} = O(\log n)$ with high probability. Hence, putting everything together, using the upper-bound (16), we get $T = O(n \log^2 n)$. This differs from the smallest possible convergence time in case (i) by a factor of $\log^2 n$ but far from $\Omega(1)$ in case (ii).

V. Discussion

We viewed opinion dynamics as a local interaction game over a social network. When there are no stubborn agents, the best-response dynamics converge to a common opinion in which the impact of the initial opinion of each agent is proportional to its degree. In the presence of stubborn agents, agents do not reach a consensus but the dynamics converge to an equilibrium in which the opinion of each agent is a convex combination of the initial opinions of the stubborn agents. The coefficients of such convex combination are related to appropriately defined hitting probabilities of the random walk over the social network's graph. An alternative interpretation is based on an electrical network model of the social network where, at equilibrium, the opinion of each agent is simply its voltage in the electrical network.

The bounds on the convergence time in the paper can be interpreted in terms of location and stubbornness levels of stubborn agents, and graph properties such as diameter, degrees, and the so-called bottleneck constant (12). The bounds provide relatively tight orders for the convergence time in the case of a fixed number of partially stubborn agents (case (i)) but there might be a gap between the lower-bound and the upper-bound when some of the stubborn agents are fully stubborn (case (ii)). Tightening the bounds in case (ii) remains as a future work.

At this point, we discuss the implication of our results in applications where limited advertising budget is to be used to convince a few agents to adopt, for example, a certain opinion about a product/topic. The goal is the optimal selection of such agents to trigger a faster spread of the advertised opinion throughout the social network. This, in turn, implies that, over a finite time, more agents will be biased towards the advertised opinion. Similar *leader selection* problems have been discussed in [10], [7], [15] where the goal is to select a set of M agents with fixed states to optimize the system performance. The common methodology is to show that the objective is a sub-modular set function and use the sub-modular optimization framework to produce a greedy procedure where agents are added according to a greedy sequence. Although the greedy algorithm is useful, it does not answer the question in its simplest form $M = 1$, and may involve the inversion of typically large matrices [10] in social networks with very large number of users.

Using the simple bound (14), the question is reduced to where to place a fixed number of fully-stubborn agents in order to minimize $|\gamma|B\tilde{d}$. Recall that $|\gamma|$ is the maximum length of shortest paths from non-stubborn agents to stubborn agents, B is the bottleneck constant defined in (12), and $\tilde{d}$ is the maximum degree of non-stubborn agents. Since empirical graphs of social networks exhibit small-world network characteristics, $|\gamma|$ is already very small (it is less than the diameter of the graph which is already a $\log n$ quantity), and hence, $B\tilde{d}$ is the dominating factor. It is believed that degree distribution in many networks, such as social networks, WWW induced graph, etc, follows a power-law distribution [30]. Heuristically, when there are a few very high degree nodes and most of the nodes are of low degrees, selecting the high degree nodes, as stubborn agents, reduces $\tilde{d}$ dramatically and also reduces B because many agents (the neighbors of the stubborn agents) are now directly connected to the stubborn agents. On the other hand, selecting the possibly low degree nodes of bottleneck edges as stubborn agents reduces B but this reduction is at most by a factor equal to low degrees of such nodes. Therefore, in general, the high degree nodes seem to be good candidates for placement of stubborn agents. It will be certainly interesting to establish the validity of such a heuristic more rigorously.

References

[1] M. H. DeGroot, Reaching a consensus, *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118-121, 1974.
[2] G. Ellison and D. Fundenberg, Rules of thumb for social learning. *Journal of Political Economy*, vol. 110, no. 1, pp. 93-126, 1995.
[3] R. Hegselmann, U. Krause, Opinion dynamics and bounded confidence models, analysis, and simulations, *Journal of Artifical Societies and Social Simulation (JASSS)*, vol. 5, no. 3, 2002
[4] P. Bremaud, Markov chains, Gibbs fields, Monte Carlo simulation, and queues, *Springer-Verlag*, New York 1999, 2nd edition, 2001.
[5] D. Shah, Gossip Algorithms, *Foundations and Trends in Networking*, nol. 3, no. 1, 2008.
[6] M. Kandori, H. J. Mailath and F. Rob, Learning, mutation, and long run equilibria in games, *Econometrica*, vol. 61, pp. 29-56, 1993.
[7] A. Clark and R. Poovendran, A Submodular optimization framework for leader selection in linear multi-agent systems, *IEEE Conference on Decision and Control*, Orlando, December 12-15, 2011.
[8] J. Kleinberg, Cascading behavior in networks: algorithmic and economic issues, *Algorithmic Game Theory*, Cambridge University Press, 2007.
[9] G. Ellison, Learning, local interactions, and coordination, Econometrica, vol. 61, pp. 1047-1071, 1993.
[10] V. S. Borkar, J. Nair, and N. Sanketh, Manufacturing consent, *Allerton Conference on Communication, Control, and Computing*, 2010.
[11] S. Bikchandani, D. Hirshleifer, and I. Welch, A theory of fads, fashion, custom, and cultural change as information cascades, *Journal of Political Economy*, vol. 100, pp. 992-1026, 1992.
[12] A. Banerjee and D. Fudenberg, Word-of-mouth learning, *Games and Economic Behavior*, vol. 46, pp. 1-22, 2004.
[13] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, Bayesian learning in social networks, *The Review of Economic Studies*, 2011.
[14] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, Spread of (mis)information in social networks, *Games and Economic Behavior*, vol. 70, no. 2, pp. 194-227, 2010.
[15] E. Yildiz, D. Acemoglu, A. Ozdaglar, A. Saberi, and A. Scaglione, Discrete Opinion Dynamics with Stubborn Agents, submitted, 2011.
[16] J. N. Tsitsiklis, Problems in decentralized decision making and computation, *Ph.D. Thesis, Department of EECS, MIT*, 1984.
[17] J. N. Tsitsiklis, D. P. Bertsekas and M. Athans, Distributed asynchronous deterministic and stochastic gradient optimization algorithms, *IEEE Transactions on Automatic Control*, vol. 31, no. 9, pp. 803-812, 1986.
[18] A. Olshevsky and J. N. Tsitsiklis, Convergence speed in distributed consensus and averaging, *SIAM Journal on Control and Optimization*, 2008.
[19] A. Jadbabaie, J. Lin, and S. Morse, Coordination of groups of mobile autonomous agents using nearest neighbor rules, *IEEE Transactions on Automatic Control* vol. 48, no. 6, 988-1001, 2003.
[20] R. A. Horn and C. R. Johnson, Matrix analysis. *Cambridge University Press*, 1985.
[21] P. Diaconis and D. Stroock, Geometric bounds for eigenvalues of Markov chains. *Annals of Applied Probability*, vol. 1, pp. 36-61, 1991.
[22] A. Sinclair, Improved bounds for mixing rates of Markov chains and multicommodity flow, *Combinatorics, Probability and Computing*, vol. 1, pp. 351-370, 1992.
[23] A. Montarani, A. Saberi, Convergence to equilibrium in local interaction games, *FOCS Conference*, 2009.
[24] M. Draief and L. Massoulie, Epidemics and rumours in complex networks, *Cambridge University Press*, 2010.
[25] S. Milgram, The small word problem, *Psychology Today*, vol. 1, no. 1, pp. 61-67, 1967.
[26] D. Watts, Small worlds: The dynamics of networks between order and randomness, *Princeton University Press*, 1999.
[27] J. Kleinberg, The small-world phenomenon: an algorithmic perspective, *ACM Symposium on Theory of Computing*, pp. 163-170, 2000.
[28] A. Flaxman, Expansion and lack thereof in randomly perturbed graphs, *Internet Mathematics*, vol. 4, no. 2, pp. 131-147, 2007.
[29] N. Alon and V. D. Milman, λ_1, isoperimetric inequalities for graphs, and superconcentrators, *Journal of Combin. Theory Ser. B*, vol. 38 no. 1, pp. 73-88, 1985.
[30] M. E. J. Newman, Power laws, Pareto distributions and Zipf's law, available online at http://arxiv.org/pdf/cond-mat/0412004.pdf.
[31] J. Ghaderi and R. Srikant, Opinion dynamics in social networks: a local interaction game with stubborn agents, http://arxiv.org/abs/1208.5076.