

Strategic Seeding of Rival Opinions

Samuel D. Johnson¹(✉), Jemin George², and Raissa M. D’Souza³

¹ HRL Laboratories, LLC, Malibu, CA 90265, USA

`sdjohnson@hrl.com`

² United States Army Research Laboratory, Adelphi, MD 20783, USA

`jemin.george.civ@mail.mil`

³ University of California, Davis, CA 95616, USA

`raissa@cse.ucdavis.edu`

Abstract. We present a network influence game that models players strategically seeding the opinions of nodes embedded in a social network. A social learning dynamic, whereby nodes repeatedly update their opinions to resemble those of their neighbors, spreads the seeded opinions through the network. After a fixed period of time, the dynamic halts and each player’s utility is determined by the relative strength of the opinions held by each node in the network *vis-à-vis* the other players. We show that the existence of a pure Nash equilibrium cannot be guaranteed in general. However, if the dynamics are allowed to progress for a sufficient amount of time so that a consensus among all of the nodes is obtained, then the existence of a pure Nash equilibrium can be guaranteed. The computational complexity of finding a pure strategy best response is shown to be NP-complete, but can be efficiently approximated to within a $(1 - 1/e)$ factor of optimal by a simple greedy algorithm.

Keywords: Social networks · Opinion dynamics · Game theory · Nash equilibrium · Computational complexity · Approximation algorithm

1 Introduction

Opinions are shaped by the information individuals obtain through their social connections. These opinions inform our career decisions, political views, and purchasing behaviors, which can ultimately spread to affect an entire society. The *influence maximization problem* [14] asks: Given the ability to seed a small number of individuals (nodes) in a social network to adopt a desired behavior (*e.g.*, to purchase a product, support a political candidate, contract an infection, etc.), which nodes should be selected so that the behavior subsequently spreads to a maximum fraction of the entire population?

Strategic extensions to the influence maximization problem that involve two or more players representing competing, substitutable behaviors (products, opinions, infections, etc.) where each player selects a set of seed nodes so as to maximize the fraction of the population that eventually adopts their represented

This research was conducted while S.D. Johnson was a graduate student at the University of California, Davis.

behavior, have subsequently been studied; see, for example, [1, 4, 5, 8, 11, 17]. A common feature among most of these models is that the diffusion dynamics they use involve nodes making binary decisions to determine whether to fully adopt one of the diffused opinions. Furthermore, these models often employ progressive SI-style dynamics in which nodes that do not have an opinion must first make an irreversible decision to fully commit themselves to a single opinion in order to become a participant in the subsequent propagation of their adopted opinion. Although such dynamics meet the desiderata of many applications with practical importance, we believe them to be unnatural for the diffusion of opinions through social networks because of the requirement that opinions are only allowed to pass through nodes that have already committed irreversibly to a particular preference.

In the current work, we present the *network influence game* (NIG) to model the strategic seeding of opinions in a social network by $m \geq 2$ players and employs a dynamic that is arguably more appropriate than the SI-style dynamics used elsewhere. In our model, each node maintains a vector of opinions toward the m players, and in each discrete time step, a node's opinion is updated to reflect a weighted averaging of their previous opinion and the opinions of their neighbors. These dynamics proceed in accordance with the *consensus dynamic* [6], and terminate after a period of T time steps. The consensus dynamic has played a prominent role in the study of opinion dynamics in social networks – see, for example, the surveys by Jackson [12, 13]. In our competitive model, players' utilities are defined to be a function of each node's relative opinions toward the players upon the conclusion of the dynamic process.

The NIG model is formally presented in Sect. 2, followed by a discussion on the nature of influence pertaining to it in Sect. 2.1. Section 3 contains our results on the existence of pure Nash equilibrium strategies. We show that existence can be guaranteed if the dynamics run until a consensus opinion is reached, but cannot be guaranteed otherwise. In Sect. 4 we show that computing a pure strategy best response is NP-complete, yet can be efficiently approximated to within a factor of $(1 - 1/e)$ of optimal. Section 5 concludes the paper with a discussion and suggested topics for future research.

2 Model

The *network influence game* NIG is specified by a tuple $\langle M, G, \Delta_C, T, \mathbf{b}, \pi \rangle$ with a player set, $M = \{1, \dots, m\}$; a network, $G = (V, E)$, represented by a weighted digraph with $|V| = n$ nodes; the opinion dynamic, Δ_C , that determines how influence spreads between adjacent nodes in the network; the length of time, $T \in \mathbb{N}$, that the diffusion dynamic is allowed to proceed; a profile, $\mathbf{b} = (b_1, \dots, b_m)$, of integer seed budgets $b_i > 0$ for each player $i \in M$; and a utility function, $\pi_i(\cdot)$, that aggregates the nodes' opinions upon the conclusion of the diffusion dynamics at time T into a non-negative payoff for each player $i \in M$.

We assume that the graph G is strongly connected. Furthermore, we stipulate that for all edges $(u, v) \in E$, the edge weight $w(u, v) > 0$ and, for all nodes $v \in V$, the sum of all incoming edge weights equals one.

A (pure) strategy for player i is a subset $s_i \subseteq V$ of at most b_i seed nodes. A strategy profile $s = (s_1, \dots, s_m)$ specifies the seed nodes chosen by all m players, and is used to set the initial conditions for the opinion dynamics, the result of which determines the players' utilities.

The NIG's opinion dynamics, Δ_C , is based on DeGroot's *consensus dynamic* [6]. For this dynamic, each node $v \in V$ maintains a length- m opinion vector $\mathbf{x}^v = (x_1^v, \dots, x_m^v)$ with the entry $0 \leq x_i^v \leq 1$ representing v 's opinion toward player i . If $x_i^v > x_j^v$, then this is to be interpreted as v holding a more favorable opinion toward player i than toward player j . We require that the sum of a node's opinions is at most one; $|\mathbf{x}^v|_1 = \sum_{i \in M} x_i^v \leq 1$. Since we need to refer to the evolution of a node's opinion over time, we will use $\mathbf{x}^v(t)$ to denote v 's vector of opinions at time t , with entry $x_i^v(t)$ representing v 's opinion toward player i at time t .

The diffusion process is initialized by s at time $t = 0$. This involves each player i implanting a "seed opinion" \mathbf{y}^i into each of the nodes included in s_i . The seed opinions $\mathbf{y}^i = (y_1^i, \dots, y_m^i)$ are defined as $y_i^i = 1$ and $y_j^i = 0$ for all $j \neq i$, meaning that \mathbf{y}^i specifies a high (maximum) opinion toward player i and a low (minimum) opinion toward all other players j . Let $M_v(s) = \{i \mid v \in s_i\} \subseteq M$ denote the subset players that include a given node v in their strategy, and define $m_v(s) = |M_v(s)|$. Finally, let V_s denote the subset of nodes that are designated as seed nodes by at least one player. We initialize the opinions of the nodes $v \in V$ towards each player $i \in M$

$$x_i^v(0) = \begin{cases} \frac{1}{m_v(s)} \sum_{j \in M_v} y_j^i & \text{if } v \in V_s \\ \varepsilon & \text{if } v \in V \setminus V_s, \end{cases} \quad (1)$$

where $0 < \varepsilon \ll 1/m$ is a small constant. Equation (1) specifies that each seed node $v \in V_s$ is initialized to the average seed opinion of the players that include v in their strategies. Otherwise, for a node $v \notin V_s$, initialization involves assigning a ε opinion value toward every player $i \in M$.

The consensus dynamic, Δ_C , proceeds in discrete time steps $t = 1, 2, \dots, T$ and specifies that the opinion of a node v at time t is a weighted average of its prior opinion and the opinions of its neighbors at time $t - 1$. Specifically,

$$x_i^v(t) = (1 - \alpha) \cdot x_i^v(t - 1) + \alpha \sum_{u \in N^+(v)} w(u, v) \cdot x_i^u(t - 1), \quad \forall i \in M, \quad (2)$$

where $N^+(v) = \{u \mid (u, v) \in E\}$ is the set of v 's incoming neighbors in G and $0 < \alpha < 1$ is a model parameter.

This dynamic can be expressed more succinctly in matrix form. Let A be the weighted adjacency matrix of G with entries $a_{ij} = w(i, j)$, and define $\Gamma = (1 - \alpha)I + \alpha A^\top$ to be the $n \times n$ influence matrix¹ with entries γ_{vu} conveying the amount of direct influence that the opinions of node u shape those of node v from one time step to the next. By construction, Γ is an aperiodic stochastic matrix. Let $\mathbf{x}_i(t)$ denote the length- n vector containing entries for each node's opinion

¹ Γ is sometimes referred to as a *listening structure* [7] or *interaction matrix* [10].

toward player i at time t . Using Γ and $\mathbf{x}_i(\cdot)$, we can rewrite the dynamics in Eq. (2) as $\mathbf{x}_i(t) = \Gamma \mathbf{x}_i(t-1)$. In particular, the opinions toward player i after T time steps is simply $\mathbf{x}_i(T) = \Gamma^T \mathbf{x}_i(0)$.²

Upon the termination of the opinion dynamics after T steps, each node v is left holding an opinion vector $\mathbf{x}^v(T)$. The utility for player $i \in M$ is defined to be the average relative opinion held by the population toward i ,

$$\pi_i(s) = \frac{1}{n} \sum_{v \in V} \frac{x_i^v(T)}{|\mathbf{x}^v(T)|_1}. \quad (3)$$

Notice that, (3) implies that for any strategy profile s , we have $\sum_{i \in M} \pi_i(s) = 1$, so the NIG is a constant-sum game.

2.1 Influence

A player’s best response strategy in the NIG involves selecting an “influential” subset of seed nodes so that, upon the termination of the dynamics, the average relative opinion held by the nodes toward the player is maximized. The precise character of *influence* in this context deserves some examination.

It is well-known that the DeGroot consensus dynamic converges so that a common opinion is shared by every node in the network is guaranteed as $T \rightarrow \infty$ when the matrix Γ is aperiodic and stochastic. The consensus obtained is described by a weighted sum of the nodes’ initial opinions (at time $t = 0$), with the weights given by the entries in the eigenvector of Γ corresponding to the eigenvalue 1. Hence, for sufficiently large T , a node’s influence is directly related to their eigenvector centrality. However, if T is not large enough to obtain consensus, then a node’s eigenvector centrality no longer corresponds to the influence they exert on the average opinions upon the termination of the dynamics.

The importance of T in characterizing the influence that a node exerts on a diffusion process was recently identified in an empirical study by Banerjee *et al.* [2] (see also [3]), which led them to define a quantity called *diffusion centrality*. Although their definition corresponds to a different diffusion dynamic than the one we consider in this paper, we can still offer a definition that is qualitatively similar to theirs but tailored to the consensus dynamic.

Let $\delta[j]$ denote the n -dimensional column vector consisting of a one in row j and zeros everywhere else. The diffusion centrality of a node v is defined to be the vector $\mathbf{c}^v = \Gamma^T \delta[v]$ whose entries c_u^v describe the fraction of node u ’s opinion at time T that is due to node v ’s initial opinion at time $t = 0$. As $T \rightarrow \infty$, the convergence of the consensus dynamic ensures that $|c_u^v - c_{u'}^v| \rightarrow 0$ for all $u, u' \in V$; let c^v denote this uniform amount of influence that v ’s brings to bear upon the final opinion of every node in V . The value c^v is precisely the v ’th entry of the unique left eigenvector \mathbf{c} of Γ corresponding to the eigenvalue 1, and the sum of the entries in \mathbf{c} equal 1.

² We use Γ^T to denote the matrix Γ raised to the T th power. For matrix transposition, we use the notation Γ^\top .

3 On the Existence of Pure Nash Equilibrium

This section presents our results on the existence of pure Nash equilibrium strategies. Recall that a pure strategy profile $s = (s_i, s_{-i})$ is a Nash equilibrium if, for every player $i \in M$ and every possible pure strategy s'_i for that player, we have $\pi_i(s_i, s_{-i}) \geq \pi_i(s'_i, s_{-i})$. Throughout this paper, all of the results regarding Nash equilibrium will be with respect to pure strategies, and all unqualified mentions of *strategy* should be understood to refer to a *pure strategy*; all mentions of *Nash equilibrium* refer to *pure Nash equilibrium*.

3.1 At Consensus

In this section we establish the existence of pure strategy Nash equilibria for NIGs in which T is sufficiently large to ensure that the opinions reach a consensus. In the consensus regime, all nodes share the same final opinion vector, $\mathbf{x}(T) = (x_1(T), \dots, x_m(T))$, and the utility for each player $i \in M$ is simply

$$\pi_i(s) = \frac{x_i(T)}{|\mathbf{x}(T)|_1}. \quad (4)$$

With the profile of weights $\mathbf{c} = (c^{v_1}, c^{v_2}, \dots, c^{v_n})$, where c^v denotes the weight of node v 's contribution to this consensus, we can express the consensus opinion toward player i as $x_i(T) = \sum_{v \in V} c^v x_i^v(0)$. Notice that the consensus opinion $x_i(T)$ for player i is a countably additive function of player i 's strategy, s_i . This implies that in order for a deviation from s_i to s'_i to increase the consensus opinion toward player i , then any deviation from s_i to s''_i where $s''_i = (s_i \setminus \{v\}) \cup \{u\}$, $v \in s_i \setminus s'_i$, and $u \in s'_i \setminus s_i$ will also increase the consensus opinion toward player i . In other words, if i can increase their share of the consensus opinion by swapping k nodes in their strategy, then they can also gain from swapping only a single node.

A key feature to the consensus case is that, since all nodes will converge to the same opinion, there is no longer any need for players to make strategic trade-offs involving increasing their opinion share among one particular subset of nodes at the cost of decreasing their opinion share elsewhere. This is reflected by the fact that the utility function in Eq. (3) reduces to Eq. (4) in the consensus regime (*i.e.*, when T is sufficiently large).

Proposition 1. *Every NIG in which T is sufficiently large so as to ensure that a consensus opinion $\mathbf{x}^v(T)$ is reached among every node $v \in V$ has at least one pure Nash equilibrium.*

Proof (Sketch). Assume that the m players are ordered so that $b_1 \geq b_2 \geq \dots \geq b_m$ and the nodes $V = \{v_1, v_2, \dots, v_n\}$ are ordered so that $c^{v_1} \geq c^{v_2} \geq \dots \geq c^{v_n}$. Set player 1's seed strategy s_1 such that it maximizes $x_1(T)$. The strategies for players $i = 2, 3, \dots, m$ will be built sequentially, so that s_i is a best response to the profile $s^{i-1} = (s_1, \dots, s_{i-1})$. We then argue that, given $s^i = (s_1, \dots, s_{i-1}, s_i)$, for every player $j < i$, s_j is a best response to $s^i_{-j} = (s_1, \dots, s_{j-1}, s_{j+1}, \dots, s_i)$.

Here, we will give the proof for $m = 2$ players; the proof for $m \geq 2$ players follows from an inductive argument that is based on similar reasoning.

Set $s_1 = \{v_1, \dots, v_{b_1}\}$, and let s_2 be a best response to s_1 that maximizes the value of the consensus opinion toward player 2, $x_2(T)$. We have two cases to consider: (i) $s_1 \cap s_2 = \emptyset$, and (ii) $s_1 \cap s_2 \neq \emptyset$.

The first case is trivial: if s_2 does not contain any of the nodes in s_1 , then player 1 enjoys exclusive access to b_1 of the most influential seed nodes.

For the second case, let $r = s_1 \cap s_2$ be the set of seed nodes chosen by player 2 that are also in s_1 . Suppose, toward a contradiction, that player 1 can strictly benefit from changing to a strategy $s'_1 = (s_1 \setminus \{u\}) \cup \{v\}$ for some nodes $u \in r$ and $v \in V \setminus (s_1 \cup s_2)$; *i.e.*, player 1 swaps out a shared node u for exclusive access to another node v . By swapping out u for v , player 1 may increase the consensus opinion toward them self, but they would also be increasing the consensus opinion toward player 2 by relinquishing their share of the influence weight c^u . But, by virtue of player 2's inclusion of u instead of v in their own best response, it must be the case that the relative opinion toward player 1 – and, thus, player 1's utility – would not improve by swapping u for v . Therefore, $\pi_1(s'_1, s_2) \leq \pi_1(s_1, s_2)$, contradicting the claim that s'_1 can earn player 1 a strictly higher utility than s_1 . \square

3.2 The General Case

In this section, we show that the existence of a pure Nash equilibrium is *not* guaranteed in NIGs when the dynamic does not reach a consensus opinion. We prove the non-existence for the symmetric setting, in which every player shares the same seed budget; a proof for the asymmetric setting can be found in the full version of this paper.

Proposition 2. *For any $m \geq 2$ and symmetric seed budget b , there exist NIGs that do not admit a pure strategy Nash equilibrium.*

Proof. A construction that does not admit a pure Nash equilibrium is as follows: Create $\mu = m(b + 1) + 1$ nodes $v_0, v_1, \dots, v_{\mu-1}$, and add directed edges from each v_i to v_{i+k} for $k = 1, \dots, b$. Next, for each v_i , add two additional “petal nodes” $v_{i,l}$ and $v_{i,r}$, and add the following four links: $(v_i, v_{i,l})$, $(v_i, v_{i,r})$, $(v_{i,l}, v_{i,r})$, and $(v_{i,r}, v_i)$. Define the weight of each edge (u, v) to be $w(u, v) = \frac{1}{\lfloor N^\mp(v) \rfloor}$. Note that the sum of each node's incoming edge weights equals one, and that the resulting graph is strongly connected. (See Fig. 1 for an example.)

Let G be a graph that implements the above construction, and set the parameters $\alpha = 1/2$ and $T = b$. Because a player's optimal strategy will never involve seeding a petal node, we restrict our attention to strategies comprising only central nodes. Let V_C denote these μ central nodes. Since $|V_C| = \mu = m(b + 1) + 1$ and the total seed budget among all m players is bm , it will always be the case that a player's best response will not include nodes that are already included in another player's strategy, and exactly $m + 1$ of the central nodes will not be




Fig. 1. Example of the graph construction in the proof of Proposition 2 with $m = 3$ players and budgets $b = 2$. The dark nodes are the “central nodes” and the white ones are the “petal nodes”. This network does not admit a pure Nash equilibrium.

included as seed nodes in any players’ strategies. Let $U \subset V_C$ denote these $m + 1$ unseeded central nodes. By design, the graph is constructed so that players prefer selecting seed nodes with successors in U . Of the $|U| = m + 1$ nodes, at most m of them will have no predecessors that are also in U ; and there will always be at least one node $u^* \in U$ that does have a predecessor in U . Let u' denote u^* ’s predecessor in U . A player i with a seed node $v \in s_i$ that is a successor of u' can always improve their utility by changing to a strategy $s'_i = \{u'\} \cup (s_i \setminus \{v\})$. Since the existence of (at least one) such u^* is guaranteed, then there will always be a player i that can change strategies for an increase in utility. Hence, a Nash equilibrium cannot be obtained. \square

4 Computational Properties of Best Response

This section considers the computational problem of finding, for a given player i and strategy profile s_{-i} , a strategy s_i that maximizes i ’s utility, $\pi_i(s_i, s_{-i})$. Such a strategy s_i is called a *best response* to the profile s_{-i} of strategies belonging to every other player $j \neq i$.

Our first result in this section establishes the computational complexity of the best response problem.

Proposition 3. *Finding a best response strategy for the NIG is NP-complete.*

Proof (Sketch). Hardness follows by reduction from SET COVER and completeness is due to the fact that the utility function can be computed in polynomial time. \square

Next, we turn the problem of finding an approximate best response. It will be useful to adopt the following definition of the utility function, which is equivalent to Eq. (3):

$$\pi_i(s) = \frac{1}{n} \sum_{v \in V} g_v^i(s) \quad (5)$$

where

$$g_v^i(s) = \frac{f_v^i(s)}{f_v(s)}, \quad (6)$$

$$f_v^i(s) = \sum_{u \in s_i} \frac{c_v^u}{m_u(s)} + \sum_{u \in V \setminus s} \varepsilon c_v^u, \quad (7)$$

and

$$f_v(s) = \sum_{j \in M} f_v^j(s).$$

Here, we employ the notation $m_u(s)$ to denote the cardinality of the set $M_u(s) = \{j \mid u \in s_j\} \subseteq M$. The quantities c_v^u , which measure the amount of influence that node u exerts on node v after T time steps, was defined in Sect. 2.1.

Our main result for this section establishes that the utility function $\pi_i(\cdot)$ is submodular,³ from which it follows from the classic result by Nemhauser, Wolsey, and Fisher [16] that a $(1 - 1/e) \approx 0.6321$ approximation can be computed using a greedy algorithm.

Proposition 4. *The utility function $\pi_i(\cdot)$ is monotonic and submodular.*

Proof (Sketch). The submodularity and monotonicity follow immediately from establishing that (7) is increasing and (6) is submodular and the fact that since $\pi_i(\cdot)$ is a function that is defined by a linear combination of submodular functions (cf., Eq. (5)), then $\pi_i(\cdot)$ is itself submodular. \square

5 Discussion and Future Work

This paper presented a model for the strategic seeding of opinion dynamics using the simple, well-studied DeGroot consensus dynamic. We established that the existence of pure Nash equilibria cannot be guaranteed if the dynamic is not allowed to run to consensus. The amount of time required for the dynamic to reach to consensus is known to be slower in networks with modular (homophilic) connectivity patterns [9]. This implies that in practice, strategic opinion seeding on real-world social networks, which often exhibit modular structures, should not assume that there will be enough time for the population to coalesce around a shared, consensus opinion; and, crucially, the individuals that appear to be attractive seeds in the steady state regime when a consensus is reached (those with high eigenvector centrality) may not be the best choice if the dynamics halt in the transient regime.

Our findings in Sect. 4 on the computational problem of finding best response strategies are in alignment with similar competitive influence models that

³ A set function $f : \Omega \rightarrow \mathbb{R}$ is *submodular* if, for every $X \subseteq Y \subset \Omega$ and element $x \in \Omega \setminus Y$, we have $f(X \cup \{x\}) - f(X) \geq f(Y \cup \{x\}) - f(Y)$.

employ opinion dynamics that are more complex and less amenable to analytical tractability than the simplistic DeGroot consensus dynamic we use. For example, in a competitive extension of the probabilistic *independent cascade* model of Kempe *et al.* [14], Bharathi *et al.* [4] establish a $(1 - 1/e)$ approximation guarantee using monotone and submodularity arguments that extend those used in [14] for the “single-player” setting (see also Mossel and Roch [15]). However, some models have approximation guarantees that can be significantly worse than $(1 - 1/e)$. For example, Borodin *et al.* [5] show that competitive extensions of the *threshold* diffusion model do not share the $(1 - 1/e)$ approximation guarantee established in [14] for the “single player” optimization setting. They do, however, offer a variant of the threshold model that does admit a $(1 - 1/e)$ approximation guarantee. Similar $(1 - 1/e)$ approximation guarantees are established more recently in competitive influence maximization models by Goyal *et al.* [11] and Fotakis *et al.* [8].

Our analysis in Sect. 3 highlights the importance of the length of the diffusion process, T , in guaranteeing the existence pure strategy Nash equilibria. The identification of conditions that are sufficient to guarantee the existence of equilibria for small, non-consensus reaching values of T is an interesting open problem. Related to this is an intriguing extension to the model that would allow players to not only choose *which* nodes to seed, but also *when* to seed them. In our preliminary investigations into this extension, we have observed in simulations that some graphs contain nodes whose influence “peaks” at a greater magnitude in the dynamic’s transient regime than in the steady state. We believe that such an extension would also more closely model many real-world applications, such as political contests and advertising campaigns, where timing can be an important consideration.

Acknowledgements. The authors gratefully acknowledge support from the US Army Research Office MURI Award No. W911NF-13-1-0340 and Cooperative Agreement No. W911NF-09-2-0053.

References

1. Alon, N., Feldman, M., Procaccia, A.D., Tennenholtz, M.: A note on competitive diffusion through social networks. *Inf. Process. Lett.* **110**(6), 221–225 (2010)
2. Banerjee, A., Chandrasekhar, A.G., Dufflo, E., Jackson, M.O.: The diffusion of microfinance. *Science* **341**(6144), 363–371 (2013)
3. Banerjee, A.V., Chandrasekhar, A., Dufflo, E., Jackson, M.O.: Gossip: Identifying central individuals in a social network. Working Paper, August 2014
4. Bharathi, S., Kempe, D., Salek, M.: Competitive influence maximization in social networks. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 306–311. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-77105-0_31](https://doi.org/10.1007/978-3-540-77105-0_31)
5. Borodin, A., Filmus, Y., Oren, J.: Threshold models for competitive influence in social networks. In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 539–550. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-17572-5_48](https://doi.org/10.1007/978-3-642-17572-5_48)
6. DeGroot, M.H.: Reaching a consensus. *J. Am. Stat. Assoc.* **69**(345), 118–121 (1974)

7. DeMarzo, P.M., Vayanos, D., Zweibel, J.: Persuasion bias, social influence, and unidimensional opinions. *Q. J. Econ.* **118**(3), 909–968 (2003)
8. Fotakis, D., Lykouris, T., Markakis, E., Obratzsova, S.: Influence maximization in switching-selection threshold models. In: Lavi, R. (ed.) *SAGT 2014*. LNCS, vol. 8768, pp. 122–133. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-44803-8_11](https://doi.org/10.1007/978-3-662-44803-8_11)
9. Golub, B., Jackson, M.O.: How homophily affects the speed of contagion, best response and learning dynamics. Working Paper, September 2010
10. Golub, B., Jackson, M.O.: Naïve learning in social networks and the wisdom of crowds. *Am. Econ. J. Microeconomics* **2**(1), 112–149 (2010)
11. Goyal, S., Heidari, H., Kearns, M.: Competitive contagion in networks. *Games and Economic Behavior* (2014) (in press)
12. Jackson, M.O.: *Social and Economic Networks*. Princeton University Press, Princeton (2008)
13. Jackson, M.O.: An overview of social networks and economic applications. In: Benhabib, J., Bisin, A., Jackson, M.O. (eds.) *Handbook of Social Economics*, chap. 12, pp. 511–585. North Holland, San Diego (2011)
14. Kempe, D., Kleinberg, J., Tardos, É.: Maximizing the spread of influence through a social network. In: *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2003*, pp. 137–146. ACM, New York (2003)
15. Mossel, E., Roch, S.: Submodularity of influence in social networks: from local to global. *SIAM J. Comput.* **39**(6), 2176–2188 (2010)
16. Nemhauser, G.L., Wolsey, L.A., Fisher, M.L.: An analysis of approximations for maximizing submodular set functions—i. *Math. Program.* **14**(1), 265–294 (1978)
17. Tzoumas, V., Amanatidis, C., Markakis, E.: A game-theoretic analysis of a competitive diffusion process over social networks. In: Goldberg, P.W. (ed.) *WINE 2012*. LNCS, vol. 7695, pp. 1–14. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-35311-6_1](https://doi.org/10.1007/978-3-642-35311-6_1)