

Scalable Equilibrium Computation in Multi-agent Influence Games on Networks

Fotini Christia,¹ Michael Curry,² Constantinos Daskalakis,¹ Erik Demaine,¹ John P. Dickerson,²
MohammadTaghi Hajiaghayi,² Adam Hesterberg,¹ Marina Knittel,² Aidan Milliff¹

¹ Massachusetts Institute of Technology

² University of Maryland, College Park

cfotini@mit.edu, curry@cs.umd.edu, costis@csail.mit.edu, edemaine@mit.edu, john@cs.umd.edu, hajiaghayi@gmail.com,
acheesterberg@gmail.com, mknittel@cs.umd.edu, milliff@mit.edu

Abstract

We provide a polynomial-time, scalable algorithm for equilibrium computation in multi-agent influence games on networks, extending work of Bindel, Kleinberg, and Oren (2015) from the single-agent to the multi-agent setting. In games of influence, agents have limited advertising budget to influence the initial predisposition of nodes in some network towards their products, but the eventual decisions of the nodes are determined by the stationary state of DeGroot opinion dynamics on the network, which takes over after the seeding (Ahmadinejad et al. 2014, 2015). In multi-agent systems, how should agents spend their budgets to seed the network to maximize their utility in anticipation of other advertising agents and the network dynamics? We show that Nash equilibria of this game are pure and (under weak assumptions) unique, and can be computed in polynomial time; we test our model by computing equilibria using mirror descent for the two-agent case on random graphs.

1 Introduction

Opinion formation and spreading dynamics over networks are important and well-researched topics in sociology, economics, probability, and a variety of other applied and theoretical fields. They are related to a variety of other problems on networks, such as the spreading of technologies or epidemics, and are becoming increasingly important from a computational standpoint as much of economic activity, advertising, and political influencing takes place on online social networks. There is an array of models for how opinions spread over networks, but the basic premise is simple: at the beginning of time, nodes in the network have some initial beliefs, reflecting political opinions, affiliation to groups, technologies that they use, or proclivity to buy different products; from time zero onward, network spreading dynamics determine how nodes express and update their beliefs as a function of their own state and the expressed beliefs of the other nodes in the network. A long line of work investigates how different types of belief spaces, state spaces, expressed beliefs, update rules, and, importantly, the structure of the network influence the opinion spreading dynamics and its final outcomes; see, e.g. (DeGroot 1974; Friedkin and Johnsen 1990; Ellison 1993; Gale and Kariv 2003;

Kuran and Sandholm 2008; Hagenbach and Koessler 2010; Jackson 2010; Montanari and Saberi 2010; Acemoglu et al. 2011; Acemoglu and Ozdaglar 2011; Bhawalkar, Gollapudi, and Munagala 2013; Bindel, Kleinberg, and Oren 2015), but note that this literature is vast and we cannot do it justice here.

Clearly, the outcome of the opinion spreading dynamics, be it a limiting state, a stationary distribution, or a transient measure over beliefs or expressed beliefs, is intimately related to the “seeding” of the network’s nodes with initial beliefs. This naturally suggests an algorithmic challenge. Given a network, an objective function assigning payoffs to different configurations of beliefs on the network’s nodes, and a seeding budget, how should this budget be allocated to influence the initial beliefs of the nodes so as to maximize the objective function in the outcome of the opinion spreading dynamics? This problem has been studied intensively in the literature, leading to efficient algorithms attaining (approximately) optimal solutions under various instantiations of our problem; see e.g. (Kempe, Kleinberg, and Tardos 2003; Chen, Yuan, and Zhang 2010; Gionis, Terzi, and Tsaparas 2013; Ahmadinejad et al. 2014, 2015).

Of course, many settings involve multiple entities competing for influence in a network. These could be different brands selling competing products. They could be different political parties competing for voters. They could be different social network influencers competing for followers. Examples abound, suggesting a natural extension of the afore-described optimization problem to the following game-theoretic challenges: How should a strategic agent optimally spend their budget to seed the network in order to maximize their final influence, in anticipation of not only how their seeding will affect the opinion spreading dynamics but also in anticipation of the strategies of other agents who are simultaneously seeding the network and may have competing objectives? What are the resulting equilibria? And how do the optimal strategies depend on the structure of the network, which nodes are deemed most valuable by the agents’ objectives, and what the initial beliefs of the nodes are?

These are the challenges that we target in this paper. We study multi-agent influence games over networks, using a natural extension, to multidimensional beliefs, of Friedkin and Johnsen’s extension (Friedkin and Johnsen 1990) of

the celebrated DeGroot model (DeGroot 1974) of belief-spreading dynamics, discussed in Section 3. In this model, there is an underlying (possibly directed) weighted network $G = (V, E, w : E \rightarrow \mathbb{R}_{\geq 0})$, whose nodes $i \in V$ are endowed with internal beliefs $\mathbf{s}_{i,\cdot} = (s_{i,k})_{k \in K} \in \Delta^K$, where K is the set of agents and $\Delta^K = \{x \in \mathbb{R}_{\geq 0}^K \mid \sum_{k \in K} x_k = 1\}$. Vector $\mathbf{s}_{i,\cdot}$ corresponds to (potentially fractional) affinities of node i to the agents in K in the beginning of time. Keeping these internal beliefs fixed, the nodes update their external beliefs $(z_{i,k}^t)_{i,k}$ according to the following rule, at every $t \in \mathbb{N}$:

$$\forall i, k : z_{i,k}^t \leftarrow \frac{s_{i,k} + \sum_{(j,i) \in E} w_{j,i} z_{j,k}^{t-1}}{1 + \sum_{(j,i) \in E} w_{j,i}}, \quad (1)$$

with the following initialization $z_{i,k}^0 = s_{i,k}$, for all i, k . In particular, given the seeding $(s_{i,k})_{i,k}$, the update dynamics defined by (1) decomposes into one Friedkin-Johnsen process for each k . This process satisfies that for all nodes i and times t : $\mathbf{z}_{i,\cdot}^t \in \Delta^K$, i.e. across all times the fractional affinities of each node to the different agents in K are non-negative and sum to 1.

Given the above setup, we now describe how agents in K compete for influence. We assume that each agent $k \in K$ has a payoff function $f_k : [0, 1]^V \rightarrow \mathbb{R}$, mapping external beliefs $\mathbf{z}_{\cdot,k} \equiv (z_{i,k})_i$ to a scalar payoff, and a budget B_k . They can spread B_k over nodes in V , spending $b_{i,k}$ on every node $i \in V$. The total amount spent on a node by all agents alters its internal belief as follows:

$$s_{i,k} \leftarrow \frac{s_{i,k} + b_{i,k}}{1 + \sum_{\ell \in K} b_{i,\ell}}. \quad (2)$$

We observe that (2) guarantees that $\mathbf{s}_{i,\cdot}$ remains in Δ^K after the update.

So now we have a multi-agent game. The agents have to decide how to spend their budgets over nodes in anticipation of how the other agents will spend their budgets, how the overall spending on the nodes will influence their internal beliefs, and how this, in turn, will influence the limiting external beliefs $(z_{i,k}^\infty)_{i,k}$ evolved according to (1). The goal of agent k is to maximize $f_k(\mathbf{z}_{\cdot,k}^\infty)$, where $\mathbf{z}_{\cdot,k}^\infty \equiv (z_{i,k}^\infty)_i$.

Our goal is to compute optimal strategies and equilibria in the afore-described family of games. The challenge is two-fold: (i) The strategy space of each agent k is high-dimensional and continuous, namely $S_k = \{(b_{i,k})_i \mid b_{i,k} \geq 0, \forall i, \text{ and } \sum_i b_{i,k} \leq B_k\}$, corresponding to all possible ways to split budget B_k among the nodes in V ; thus, randomized strategies correspond to distributions over S_k . (ii) Even for a given collection of deterministic strategies, we only have implicit access to the resulting payoffs; indeed, we need to determine how these affect the internal beliefs, compute the resulting limiting external beliefs, and plug those into the functions $(f_k)_k$ to see how much payoff every agent derives. Our main contribution is the following:

Theorem 1.1. *If the payoff functions f_k are non-decreasing and convex, and $\sum_k f_k(\mathbf{z}_{\cdot,k}) = 1$ for all $\mathbf{z} = (z_{i,k})_{i,k}$, i.e. the game is constant-sum, then there exists a pure Nash equilibrium, which is unique if additionally, for all k , f_k is*

continuous, strictly increasing and the internal beliefs prior to budget allocation satisfy $s_{i,k} \neq 0$ for all i . Under non-decreasing and convex payoff functions f_k , if the agents play the game repeatedly and use no-regret learning dynamics to update their strategies (meaning that the difference in average utility obtained by chosen strategies and that obtained by the best fixed in hindsight strategy vanishes over time), then the average trajectory of the dynamics converges to a pure Nash equilibrium. If, additionally, every f_k is a linear function, then a pure Nash equilibrium can be computed in polynomial-time using convex programming.

We note that the constraints placed by our theorem on the payoff functions are natural constraints. Indeed, a non-decreasing f_k simply ensures that increased final influence of agent k on the network does not decrease k 's utility, while convex, constant-sum payoffs encompass a broad range of natural games. The constant sum of 1 is not important, of course, as Nash equilibria are invariant under scaling and shifting the payoffs. The fact that the game is constant-sum captures that the agents are competing for fixed resources in a closed system. It is well-known that non-constant-sum games are intractable (Daskalakis, Goldberg, and Papadimitriou 2008; Chen, Deng, and Teng 2009), so it is natural to restrict our attention to constant-sum games in order to obtain polynomial-time, scalable algorithms.

Our main technical insight is that, in our setting, the payoff of each agent is concave in their own strategy and convex in the strategies of the other agents. Combining this insight with the fact that the game is constant-sum allows us to write a convex program to compute a Nash equilibrium using ideas from (Daskalakis and Papadimitriou 2009; Cai and Daskalakis 2011; Cai et al. 2016). Finally, we show that our game is a socially-concave game, which implies that no-regret learning converges to equilibrium (Even-Dar, Mansour, and Nadav 2009).

We apply these insights to compute optimal strategies in a two-player case of the multi-agent game. We show that in the two-player case, where the game becomes a convex-concave saddle-point problem, equilibrium strategies can be computed using the mirror-descent algorithm (Bubeck 2014; Nemirovsky and Yudin 1983). We implement and validate this approach on families of random graphs.

Additional Related Work. We have already provided some historical and modern references on opinion spreading and learning dynamics on networks, as well as relevant references on single-agent influence maximization. Work on the multi-agent front has been sparser, due to the significant difficulty of computing equilibria in this high-dimensional setting. We discuss this work here. (Alon et al. 2010) initiated the multiagent network diffusion problem where agents directly influence nodes. This was followed by a number of additional works (Garimella et al. 2017; Maehara, Yabe, and Kawarabayashi 2015; Tzoumas, Amanatidis, and Markakis 2012), however (Alon et al. 2010)'s initial diffusion model is the most comparable to this work. They introduce a graph-based multi-agent game that is strikingly similar to (yet notably distinct from) our model. In their model, each agent

corresponds to a color, and they can select one node each to assign to their color. From there, the colors diffuse in the following manner: an uncolored vertex is colored k if all of its neighbors are either colored k or uncolored and at least one is colored k , and an uncolored vertex is eliminated if at least two of its neighbors have different colors. They showed that a Nash equilibrium exists and can be found efficiently on all graphs with diameter at most two, however there exists a graph with diameter three which yields no Nash equilibrium. (Alon et al. 2010)’s model can almost be thought of as a simplified and discretized version of our multi-agent extension of the DeGroot model: internal beliefs all start as zero, color assignments correspond to external beliefs and may only have binary values, and nodes with multiple influences are simply ignored. Conversely, our model can be thought of as a more continuous extension of the (Alon et al. 2010) model, except for subtle differences in the diffusion process.

2 Preliminaries & our multi-agent influence games

Notation. For convenience, we may denote a matrix $(a_{i,k})_{i,k}$ by a_{\cdot} , or by boldface \mathbf{a} . For a given i , we will denote by $\mathbf{a}_{i\cdot}$ the i -th row of \mathbf{a} , and, for a given k , we will denote by $\mathbf{a}_{\cdot k}$ its k -th column. As is common in game theory, if b_1, \dots, b_K are strategies of K players in a game, we denote by b_{-i} the strategies of all players but player i , and by $(b_i; b_{-i})$ the strategies of all players.

Single-agent games of influence

A Single-Agent Influence Process. As a basis for our research, we use (Friedkin and Johnsen 1990)’s extension of (DeGroot 1974)’s model for influence networks. In this model, we consider the spread of a single belief across a graph of nodes via a Friedkin-Johnsen process. This model is formally defined as follows:

Definition 2.1. A *single-agent Friedkin-Johnsen influence spreading process* takes as input: a (possibly directed) graph $G = (V, E, w : E \rightarrow \mathbb{R}_{\geq 0})$ of nodes; and for every node $i \in V$, an “internal belief” $s_i \in [0, 1]$. Then for every node $i \in V$, it sets the node’s initial “external belief” to $z_i^0 = s_i$, and updates the nodes’ external beliefs over a series of steps $t \in \mathbb{N}$ as follows:

$$z_i^t \leftarrow \frac{s_i + \sum_j w_{j,i} z_j^{t-1}}{1 + \sum_j w_{j,i}}.$$

Notably, (Bindel, Kleinberg, and Oren 2015) showed that this process converges: there exists some vector of limiting external beliefs $\mathbf{z}^\infty = \lim_{t \rightarrow \infty} (z_i^t)_i$.

A useful tool for analyzing and solving problems in this model is called the *influence matrix*, which is defined using a Laplacian L of the input graph G . To start, define the matrix $M_{i,i} = 1 + \sum_{j \in N(i)} w_{j,i}$, $M_{i,j} = -w_{j,i}$ for $j \in N(i)$, and $M_{i,j} = 0$ for $j \notin N(i)$, where $N(i)$ are the in-neighbors of i . In other words, $M = L + I$, where L is a Laplacian of the graph. Let $A = M^{-1T}$. (Ahmadinejad

et al. 2014) showed that the vector \mathbf{z}^∞ of equilibrium external beliefs and the vector \mathbf{s} of internal beliefs satisfies: $\mathbf{z}^{\infty T} = \mathbf{s}^T \cdot A$. Therefore, A_{ij} measures to what extent the internal belief s_i impacts the equilibrium external belief z_j^∞ . Thus, we call A the influence matrix.

Lemma 2.1. In the single-agent Friedkin-Johnsen network process, the influence matrix A is strictly increasing on $[0, 1]^{|V|}$ (i.e. for vectors $\mathbf{s}, \mathbf{s}' \in \Delta^{|V|}$ such that $s_i \geq s'_i$, for all i , and $s_j > s'_j$, for some j , it holds that $(\mathbf{s}^T A)_i \geq (\mathbf{s}'^T A)_i$, for all i , and $(\mathbf{s}^T A)_j > (\mathbf{s}'^T A)_j$).

Proof. Suppose that $\mathbf{s}, \mathbf{s}' \in [0, 1]^{|V|}$ such that $s_i \geq s'_i$ for all i and $s_j > s'_j$ for some j . By definition, \mathbf{s} and \mathbf{s}' are both valid internal belief vectors for the Friedkin-Johnsen process. It follows that, with those internal belief vectors, the process has limiting external belief vectors $\mathbf{z}^{\infty T} = \mathbf{s}^T A$ and $\mathbf{z}'^{\infty T} = \mathbf{s}'^T A$ respectively. Hence, to show the desired property, it is sufficient to show that the property is satisfied, at all timesteps t , by \mathbf{z}^t and \mathbf{z}'^t , which are the external beliefs maintained by the Friedkin-Johnsen process with internal beliefs \mathbf{s} and \mathbf{s}' respectively. That is, we will show that, for all t , $z_i^t \geq z_i'^t$, for all i , and $z_j^t > z_j'^t$. We will use induction. The property is clearly true at $t = 0$ since $\mathbf{z}^0 = \mathbf{s}$ and $\mathbf{z}'^0 = \mathbf{s}'$. Note that, by the update formula of the Friedkin-Johnsen process, z_i^t is (strictly) monotone increasing in s_i and monotone non-decreasing with respect to all z_ℓ^{t-1} . Since, by our induction hypothesis, $z_\ell^{t-1} \geq z_\ell'^{t-1}$ for all ℓ , it follows that: $z_i^t \geq z_i'^t$, for all i , and $z_j^t > z_j'^t$, because $s_j > s'_j$. \square

We will later show that the influence matrix and its properties extend to our multi-agent model.

A Single-Agent Influence Game. So far, we have only discussed a process wherein the involved actors are nodes, and their actions define the belief diffusion process. This paper, however, is more concerned with games of influence. (Ahmadinejad et al. 2015) introduce and analyze a number of optimization problems, wherein a single agent optimizes how to seed the network to spread an idea. Specifically, they consider an instance of Definition 2.1 and add an external agent who chooses an influencing strategy $(b_i)_i$, which impacts the nodes’ internal beliefs as $s_i \leftarrow s_i + b_i$ before the game starts. The goal of the agent is to optimize some objective. In many problems they consider, the objective as some function over the external beliefs of the nodes at equilibrium, and the influencing strategy is constrained to some budget B where $\sum_i b_i \leq B$. This idea will serve as a single-agent basis for our multi-agent optimization problem in the next section.

Multi-agent games of influence

The Multi-Agent Influence Process. Our multi-agent games of influence are defined in terms of a natural multi-agent analogue of the Friedkin-Johnsen influence spreading process of Definition 2.1. In this process, the nodes of a network have internal and external beliefs over a set of agents K .

Definition 2.2. A multi-agent Friedkin-Johnsen influence spreading process takes as input:

1. a (possibly directed) graph $G = (V, E, w : E \rightarrow \mathbb{R}_{\geq 0})$ of nodes;
2. a set of agents K ; and
3. for every node $i \in V$ an “internal belief vector” $\mathbf{s}_i \equiv (s_{i,k})_k \in \Delta^K$, where $\Delta^K = \{x \in \mathbb{R}_{\geq 0}^K \mid \sum_{k \in K} x_k = 1\}$; in particular $s_{i,k} \in [0, 1]$ is the internal belief of node i in agent k .

Every node $i \in V$ initializes their “external belief vector” as follows $\mathbf{z}_{i,\cdot}^0 = \mathbf{s}_{i,\cdot}$, and over a sequence of steps the nodes update their external belief vectors according to the Friedkin-Johnsen process:

$$z_{i,k}^t \leftarrow \frac{s_{i,k} + \sum_j w_{j,i} z_{j,k}^{t-1}}{1 + \sum_j w_{j,i}}.$$

A nice property of the multi-agent Friedkin-Johnsen process is that external belief vectors remain in the simplex in the course of the dynamics, as stated in the following simple lemma.

Lemma 2.2. For each i , at any timestep t , $\mathbf{z}_i^t \equiv (z_{i,k}^t)_k \in \Delta^K$.

Proof. To prove this, we simply need to show that, for all t : $\sum_k z_{i,k}^t \leq 1$ for all i , as all external beliefs remain non-negative in the course of the dynamics, because of the non-negativity of the weights. We show this by induction. The property is clearly true before influence spreading starts, i.e. $t = 0$, since for each i , $\sum_k z_{i,k}^0 = \sum_k s_{i,k} \leq 1$. We now show the inductive step. At time t , and for some node i :

$$\begin{aligned} \sum_k z_{i,k}^t &= \sum_k \frac{s_{i,k} + \sum_j w_{j,i} z_{j,k}^{t-1}}{1 + \sum_j w_{j,i}} \\ &= \frac{\sum_k s_{i,k} + \sum_j (\sum_k z_{j,k}^{t-1}) w_{j,i}}{1 + \sum_j w_{j,i}} \\ &\leq \frac{1 + \sum_j w_{j,i}}{1 + \sum_j w_{j,i}} = 1, \end{aligned}$$

where we used the inductive hypothesis (that for all j : $\sum_k z_{j,k}^{t-1} \leq 1$) to get the inequality. So the claimed property is preserved at each step. \square

It is easy to see that, once the internal beliefs and external beliefs are initialized, the multi-agent Friedkin-Johnsen influence spreading process decomposes into $|K|$ independent copies of the single-agent Friedkin-Johnsen influence spreading process. It thus follows from the properties of the single-agent process, discussed in Section 2, that for all $k \in K$ there is a limiting vector of external beliefs in agent k , $\mathbf{z}_{i,\cdot}^\infty = \lim_{t \rightarrow \infty} \mathbf{z}_{i,\cdot}^t$, which is linked to the vector of internal beliefs in agent k as follows: $\mathbf{z}_{i,\cdot}^\infty = \mathbf{s}_{i,\cdot}^T A$, where A is the influence matrix defined in Section 2.

The Multi-Agent Influence Game. We define multi-agent influence games by allowing the agents in K to modify the

internal beliefs of the nodes of the network, before the network influence process begins, in order to achieve some goal. What is left to define is: (1) how agents modify the internal beliefs of the nodes and with what limitations, and (2) what an agent’s metric for success is.

(1) We assume that each agent $k \in K$ is given some budget $B_k > 0$ and is allowed to allocate some portion $b_{i,k} \geq 0$ of that budget to influence each $i \in V$, subject to $\sum_i b_{i,k} \leq B_k$. Recall that in our single-agent influence game of Section 2, we assumed that the impact to the internal belief of a node of the agent’s budget allocation to that node was additive (i.e. $s_i \leftarrow s_i + b_i$). However, in a multi-agent game, the impact of any one agent’s budget allocation on some node i will be reduced as other agents allocate budget to i . To keep internal beliefs normalized, we assume that the impact of budget allocation to the nodes’ internal beliefs is determined by the following update rule:

$$\forall i, k : z_{i,k}^0(\mathbf{b}) = s_{i,k}^{\text{post}} \leftarrow \frac{s_{i,k}^{\text{pre}} + b_{i,k}}{1 + \sum_{\ell \in K} b_{i,\ell}}.$$

Notice, in particular, that if the internal belief vector of node i lies in Δ^K before the agents allocate budget to that node, then the above update guarantees that the internal belief vector will still lie in Δ^K post budget allocation as well.

(2) After agents allocate their budgets, and the internal belief vectors of the nodes are updated, a multi-agent Friedkin-Johnsen process is executed, resulting in some limiting beliefs $\mathbf{z}_{i,\cdot}^\infty$ for each k . The goal of each agent k is to optimize $f_k(\mathbf{z}_{i,\cdot}^\infty)$, given some function $f_k : [0, 1]^V \rightarrow \mathbb{R}$.

We note here that a single-player version of this game does not reduce to the optimization problems defined by (Ahmadinejad et al. 2015), although it is conceptually similar and the influence-spreading dynamics are the same.

3 Scalable solution methods for multi-agent games of influence

The most natural questions to address in competitive diffusion models are whether or not a Nash equilibrium exists and, if it exists, whether it is unique and whether it can be found efficiently. In this section, we address these by discussing and proving Theorem 1.1. The results of this section are three-fold. First, we show that if each f_k is non-decreasing and convex and the game is constant sum (formally: $\sum_k f_k(\mathbf{z}_{i,\cdot}) = 1$), then a Nash equilibrium must exist, indeed a pure Nash equilibrium must exist. Second, we show that for these games, agents can employ no-regret learning dynamics to update their strategies to converge to a pure Nash equilibrium. Finally, we formulate a convex program whose purpose is two-fold: first, we use it to establish the uniqueness of pure Nash equilibrium when the f_k ’s are strictly increasing, convex and continuous and prior to budget allocation all nodes of the network allocate non-trivial internal belief to each agent; second, we use it to show that a pure Nash equilibrium can be computed efficiently using convex programming, in the natural class of non-decreasing linear f_k ’s. We prove Theorem 1.1, after we establish two auxiliary lemmas:

Lemma 3.1. *In the setting of Theorem 1.1, if for some agent k the function f_k is non-decreasing and convex then the utility of agent k is a convex function of the strategy of each other agent k' , as well as the joint strategy vector of all other agents. If f_k is strictly increasing and convex and $s_{i,k}^{\text{pre}} \neq 0$ for all i , then the utility of agent k is strictly convex in the strategy of each other k' , as well as the joint strategy vector of all other agents.*

Proof. First, we note that the initial belief $z_{i,k}^0 = \frac{s_{i,k}^{\text{pre}} + b_{i,k}}{1 + \sum_{\ell} b_{i,\ell}}$ of node i in agent k is convex in $b_{i,k'}$ for $k' \neq k$, since as a function of $b_{i,k'}$ it takes the form $\frac{a}{a' + b_{i,k'}}$ for some non-negative a and positive a' . It is strictly convex if $s_{i,k}^{\text{pre}} \neq 0$.

Hence, each coordinate of the vector of initial beliefs $\mathbf{z}_{i,k}^0(\mathbf{b})$ is convex in $\mathbf{b}_{\cdot,k'}$. In particular, if we use $(b'_{i,k'}; b_{i,-k'})$ to represent the vector $\mathbf{b}_{i,\cdot}$ with $b_{i,k'}$ replaced by $b'_{i,k'}$, then $\forall \mathbf{b}_{\cdot,-k'}, \forall \lambda \in [0, 1], \forall i, \forall \mathbf{b}_{\cdot,k'} \forall \mathbf{b}'_{\cdot,k'}$,

$$\begin{aligned} z_{i,k}^0(\lambda b_{i,k'} + (1-\lambda)b'_{i,k'}; b_{i,-k'}) \\ \leq \lambda z_{i,k}^0(\mathbf{b}_{i,\cdot}) + (1-\lambda)z_{i,k}^0(b'_{i,k'}; b_{i,-k'}), \end{aligned}$$

and the inequality is strict if $s_{i,k}^{\text{pre}} \neq 0$ and $b_{i,k'} \neq b'_{i,k'}$.

By Lemma 2.1, namely because A is a strictly increasing matrix, it holds that $\forall \mathbf{b}_{\cdot,-k'}, \forall \lambda \in [0, 1], \forall i, \forall \mathbf{b}_{\cdot,k'} \forall \mathbf{b}'_{\cdot,k'}$:

$$\begin{aligned} (A^T \mathbf{z}_{i,k}^0(\lambda \mathbf{b}_{\cdot,k'} + (1-\lambda)\mathbf{b}'_{\cdot,k'}; \mathbf{b}_{\cdot,-k'}))_i \\ \leq (A^T(\lambda \mathbf{z}_{i,k}^0(\mathbf{b}) + (1-\lambda)\mathbf{z}_{i,k}^0(\mathbf{b}'_{\cdot,k'}; \mathbf{b}_{\cdot,-k'})))_i, \end{aligned}$$

and the inequality is strict if $s_{i,k}^{\text{pre}} \neq 0$ and $b_{i,k'} \neq b'_{i,k'}$.

Since the utility function f_k is increasing,

$$\begin{aligned} f_k(A^T \mathbf{z}_{i,k}^0(\lambda \mathbf{b}_{\cdot,k'} + (1-\lambda)\mathbf{b}'_{\cdot,k'}; \mathbf{b}_{\cdot,-k'})) \\ \leq f_k(\lambda A^T \mathbf{z}_{i,k}^0(\mathbf{b}) + (1-\lambda)A^T \mathbf{z}_{i,k}^0(\mathbf{b}'_{\cdot,k'}; \mathbf{b}_{\cdot,-k'})), \end{aligned}$$

and the inequality is strict if f_k is strictly increasing, $s_{i,k}^{\text{pre}} \neq 0$, for all i , and $\mathbf{b}_{\cdot,k'} \neq \mathbf{b}'_{\cdot,k'}$.

Finally, if we apply the fact that f_k is a convex function of its inputs to the right hand side, we get that

$$\begin{aligned} f_k(A^T \mathbf{z}_{i,k}^0(\lambda \mathbf{b}_{\cdot,k'} + (1-\lambda)\mathbf{b}'_{\cdot,k'}; \mathbf{b}_{\cdot,-k'})) \\ \leq \lambda f_k(A^T \mathbf{z}_{i,k}^0(\mathbf{b})) + (1-\lambda)f_k(A^T \mathbf{z}_{i,k}^0(\mathbf{b}'_{\cdot,k'}; \mathbf{b}_{\cdot,-k'})), \end{aligned}$$

and the inequality is strict if f_k is strictly increasing, $s_{i,k}^{\text{pre}} \neq 0$, for all i , and $\mathbf{b}_{\cdot,k'} \neq \mathbf{b}'_{\cdot,k'}$.

That is, the utility of agent k is a convex function of the strategy of k' , as claimed, and it is strictly convex if f_k is strictly increasing and $s_{i,k}^{\text{pre}} \neq 0$, for all i . Since the strategies of the other agents appear additively in the utility of agent k , we can use the same proof to show convexity/strict convexity of the utility of k to the joint strategy vector of all other agents. \square

Lemma 3.2. *In the setting of Theorem 1.1, if for some agent k the function f_k is non-decreasing and linear, then the utility of agent k is a concave function of the strategy of agent k .*

Proof. First, we note that the initial belief $z_{i,k}^0 = \frac{s_{i,k}^{\text{pre}} + b_{i,k}}{1 + \sum_{\ell} b_{i,\ell}}$ of node i in agent k is concave in $b_{i,k}$, since as a function of $b_{i,k}$ it takes the form $1 - \frac{a}{a' + b_{i,k}}$ for $a \geq 0$ and $a' > 0$.

Hence each coordinate of the vector $\mathbf{z}_{i,k}^0(\mathbf{b})$ of initial beliefs in k is concave in $\mathbf{b}_{\cdot,k}$. Namely, for all $\lambda \in [0, 1]$, for any pair of strategies $\mathbf{b}_{\cdot,k}, \mathbf{b}'_{\cdot,k}$ of agent k , and any strategies $\mathbf{b}_{\cdot,-k}$ of the other agents, we have that for all i :

$$\begin{aligned} z_{i,k}^0(\lambda \mathbf{b}_{\cdot,k} + (1-\lambda)\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k}) \\ \geq \lambda z_{i,k}^0(\mathbf{b}) + (1-\lambda)z_{i,k}^0(\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k}). \end{aligned}$$

By Lemma 2.1, since A is a monotone increasing matrix, all coordinates of the vector of beliefs in k at $t = \infty$, which takes the form $A^T \mathbf{z}_{i,k}^0(\mathbf{b})$, are also concave in $\mathbf{b}_{\cdot,k}$, i.e. for all i :

$$\begin{aligned} (A^T(\mathbf{z}_{i,k}^0(\lambda \mathbf{b}_{\cdot,k} + (1-\lambda)\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k})))_i \\ \geq (\lambda A^T \mathbf{z}_{i,k}^0(\mathbf{b}) + (1-\lambda)A^T \mathbf{z}_{i,k}^0(\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k}))_i. \end{aligned}$$

Since f_k is non-decreasing,

$$\begin{aligned} f_k(A^T(\mathbf{z}_{i,k}^0(\lambda \mathbf{b}_{\cdot,k} + (1-\lambda)\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k}))) \\ \geq f_k(\lambda A^T \mathbf{z}_{i,k}^0(\mathbf{b}) + (1-\lambda)A^T \mathbf{z}_{i,k}^0(\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k})). \end{aligned}$$

Finally, if we apply the fact that f_k is a linear function of its inputs to the right hand side, we get that the utility of agent k is a concave function of the strategy of k , as claimed. \square

Proof. (of Theorem 1.1)

First, we prove that if the f_k 's are nondecreasing convex functions and the game is constant-sum, then the game is *socially concave*, that is, that

1. there exists a strict convex combination of the agents' utility functions which is a concave function of the agents' strategies; and
2. the utility of each agent k is a convex function of the strategy of each other agent k' .

Lemma 3.1 proves the second condition. For the first condition, we use that the game is constant-sum, i.e. $\sum_k f_k(\mathbf{z}_{\cdot,k}) = 1$, and 1 is a concave function.

As shown by Even-Dar, Mansour, and Nadav (Even-Dar, Mansour, and Nadav 2009), a socially concave game has a pure Nash equilibrium, and if agents play the game repeatedly and use no-regret learning dynamics to update their strategies, then the average trajectory of the dynamics converges to a pure Nash equilibrium.

Next, we state a convex program to which the pure Nash equilibria are optimal solutions. We will use the convex program to achieve two goals:

- (a) argue that there is a unique pure Nash equilibrium if, for all k , the utility function f_k is strictly increasing, convex and continuous, and $s_{i,k}^{\text{pre}} \neq 0$, for all i ;
- (b) argue that the convex program can be solved in polynomial time to compute a pure Nash equilibrium, if each utility function is non-decreasing and linear.

The variables of our convex program are $\mathbf{b} = (\mathbf{b}_{\cdot,k})_k$, the collection of agent strategies, and $\omega = (\omega_k)_k$, a collection of scalars. For intuition ω_k will be an upper bound to the best response payoff of agent k to the strategies of the other agents. At the optimum of the convex program, it will be equal to the payoff of agent k . Our convex program is the following:

$$\begin{aligned} & \min_{\mathbf{b}, \omega} \sum_k \omega_k \text{ subject to:} \\ & 1. \forall i, \forall k: b_{i,k} \geq 0 \\ & 2. \forall k: \|\mathbf{b}_{\cdot,k}\|_1 \leq B_k \\ & 3. \forall k, \forall \mathbf{b}'_{\cdot,k} \geq 0 \text{ s.t. } \|\mathbf{b}'_{\cdot,k}\|_1 \leq B_k: \\ & \quad \omega_k \geq f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k})). \end{aligned}$$

We argue that this is indeed a convex program: the first two conditions clearly define convex spaces; for the third, we need to check that if $\lambda \in [0, 1]$ and we have two points in a space defined by a constraint of the third type, say $\omega_k \geq f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k}))$ and $\omega'_k \geq f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b}''_{\cdot,k}; \mathbf{b}_{\cdot,-k}))$ for some k and $\mathbf{b}'_{\cdot,k}$, then $\lambda\omega_k + (1 - \lambda)\omega'_k \geq f_k(\mathbf{z}_{\cdot,k}^\infty(\lambda\mathbf{b}'_{\cdot,k} + (1 - \lambda)\mathbf{b}''_{\cdot,k}; \mathbf{b}_{\cdot,-k}))$. But by Lemma 3.1 (notice that the proof of the lemma implies that the utility of agent k is jointly convex in the strategies of the other agents) it follows that: $f_k(\mathbf{z}_{\cdot,k}^\infty(\lambda\mathbf{b}'_{\cdot,k} + (1 - \lambda)\mathbf{b}''_{\cdot,k}; \mathbf{b}_{\cdot,-k})) \leq \lambda f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k})) + (1 - \lambda)f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b}''_{\cdot,k}; \mathbf{b}_{\cdot,-k})) \leq \lambda\omega_k + (1 - \lambda)\omega'_k$, as desired.

Now let us argue that if (\mathbf{b}, ω) is an optimum of the program then \mathbf{b} is a pure Nash equilibrium. We first argue that the optimum of the program has objective value at most 1. Indeed, as we have argued a pure Nash equilibrium exists. Say that \mathbf{b} is a pure Nash equilibrium. This means that, under strategies \mathbf{b} , no agent can increase their utility by changing strategies, i.e. for all k , and for all valid strategies $\mathbf{b}'_{\cdot,k}$: $f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b})) \geq f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k}))$. Thus setting $\omega_k = f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b}))$ attains an objective value of $\sum_k \omega_k = \sum_k f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b})) = 1$.

Now let us pick an optimal solution (\mathbf{b}, ω) to our convex program. If \mathbf{b} is not a Nash equilibrium, then some agent k can increase its utility by changing strategies. Hence, there is some k and some valid strategy $\mathbf{b}'_{\cdot,k}$ such that $f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b})) < f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k})) \leq \omega_k$. For all other k' , $f_{k'}(\mathbf{z}_{\cdot,k'}^\infty(\mathbf{b})) \leq \omega_{k'}$. Adding these up gives $1 = \sum_k f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b})) < \sum_k \omega_k$, so this cannot be a solution to the program, as its objective value is larger than 1. Thus, any solution to the program must be a pure Nash equilibrium, as claimed.

So we have written a convex program, and have argued that its optimal solutions are pure Nash equilibria. Next we want to argue properties (a) and (b), as promised above.

(a) Uniqueness: Suppose that, for all k , f_k is strictly increasing and $s_{i,k}^{\text{pre}} \neq 0$, for all i . We have already shown that an optimum (\mathbf{b}, ω) of the convex program gives a pure Nash equilibrium (if we drop the ω part) and a pure Nash equilibrium \mathbf{b} can be associated with some ω so that (\mathbf{b}, ω) is an optimum of the convex program. Thus, to argue uniqueness of pure Nash equilibrium, it is sufficient to show that the convex program does not have two optimal solutions (\mathbf{b}, ω) and (\mathbf{b}'', ω'') such that $\mathbf{b} \neq \mathbf{b}''$. Towards proving a contradiction, suppose that it does.

By convexity of the program, $(\mathbf{b}''', \omega''') := (\frac{1}{2}\mathbf{b} + \frac{1}{2}\mathbf{b}'', \frac{1}{2}\omega + \frac{1}{2}\omega'')$ must also be an optimal solution. Pick some k^* such that $\mathbf{b}_{\cdot,-k^*} \neq \mathbf{b}''_{\cdot,-k^*}$. Because f_{k^*} is strictly increasing and $s_{i,k^*}^{\text{pre}} \neq 0$, for all i , it follows from the proof of Lemma 3.1 that the utility of agent k^* is jointly strictly convex in the strategies of the other agents. Thus,

$$\forall \mathbf{b}'_{\cdot,k^*} \geq 0 \text{ s.t. } \|\mathbf{b}'_{\cdot,k^*}\|_1 \leq B_{k^*}: \quad (3)$$

$$\begin{aligned} & f_{k^*}(\mathbf{z}_{\cdot,k^*}^\infty(\mathbf{b}'_{\cdot,k^*}; \mathbf{b}''_{\cdot,-k^*})) \\ & \equiv f_{k^*}\left(\mathbf{z}_{\cdot,k^*}^\infty\left(\mathbf{b}'_{\cdot,k^*}; \frac{1}{2}\mathbf{b}_{\cdot,-k^*} + \frac{1}{2}\mathbf{b}''_{\cdot,-k^*}\right)\right) \\ & < \frac{1}{2}f_{k^*}(\mathbf{z}_{\cdot,k^*}^\infty(\mathbf{b}'_{\cdot,k^*}; \mathbf{b}_{\cdot,-k^*})) + \frac{1}{2}f_{k^*}(\mathbf{z}_{\cdot,k^*}^\infty(\mathbf{b}'_{\cdot,k^*}; \mathbf{b}''_{\cdot,-k^*})) \\ & \leq \frac{1}{2}\omega_{k^*} + \frac{1}{2}\omega''_{k^*} = \omega'''_{k^*}, \end{aligned} \quad (4)$$

where the first (strict) inequality follows from the strict convexity of the utility of agent k^* in the strategies of the other agents and the fact that $\mathbf{b}_{\cdot,-k^*} \neq \mathbf{b}''_{\cdot,-k^*}$, and the second inequality follows from the fact that (\mathbf{b}, ω) and (\mathbf{b}'', ω'') are feasible solutions of the convex program.

Observe, however, that (4) contradicts the optimality of $(\mathbf{b}''', \omega''')$, since the fact that these inequalities are strict implies that ω'''_{k^*} can be pushed lower to improve the objective function without violating these inequalities. (Formally, we use compactness and continuity of f_{k^*} to argue that at least one of these inequalities should have been tight at an optimal solution.)

(b) Computational Efficiency: To argue that the convex program can be solved in polynomial time when the f_k 's are linear non-decreasing, it suffices to show that there exists a separation oracle that, given ω and \mathbf{b} , either finds a violated constraint or verifies that none exists. Checking that $\forall i \forall k, b_{i,k} \geq 0$ and $\forall k, \|\mathbf{b}_{\cdot,k}\|_1 \leq B_k$ is trivial. So, fix k , and we need to check whether $\forall \mathbf{b}'_{\cdot,k}, \omega_k \geq f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k}))$. Since by assumption f_k is linear, it follows by Lemma 3.2 that $f_k(\mathbf{z}_{\cdot,k}^\infty(\mathbf{b}'_{\cdot,k}; \mathbf{b}_{\cdot,-k}))$ is concave in $\mathbf{b}'_{\cdot,k}$, so we can check the constraint by maximizing w.r.t. $\mathbf{b}'_{\cdot,k}$. \square

To conclude, we showed that for a broad range of multi-agent influence games, a pure Nash equilibrium exists and can be found using no-regret learning. Additionally, for the natural set of games with linear payoffs, a Nash equilibrium can be found efficiently through convex programming.

4 Case study: Two-agent games of influence

As a case study, we apply the competitive diffusion model in a setting where there are two agents trying to spread influence. Our game then becomes a concave-convex max-min problem, and we can use standard algorithms to approximate the saddle point.

Solution via mirror descent

In the two-agent case, because our game is constant-sum, we can simply assume that the first agent is trying to maximize their utility while the second agent is trying to minimize it.

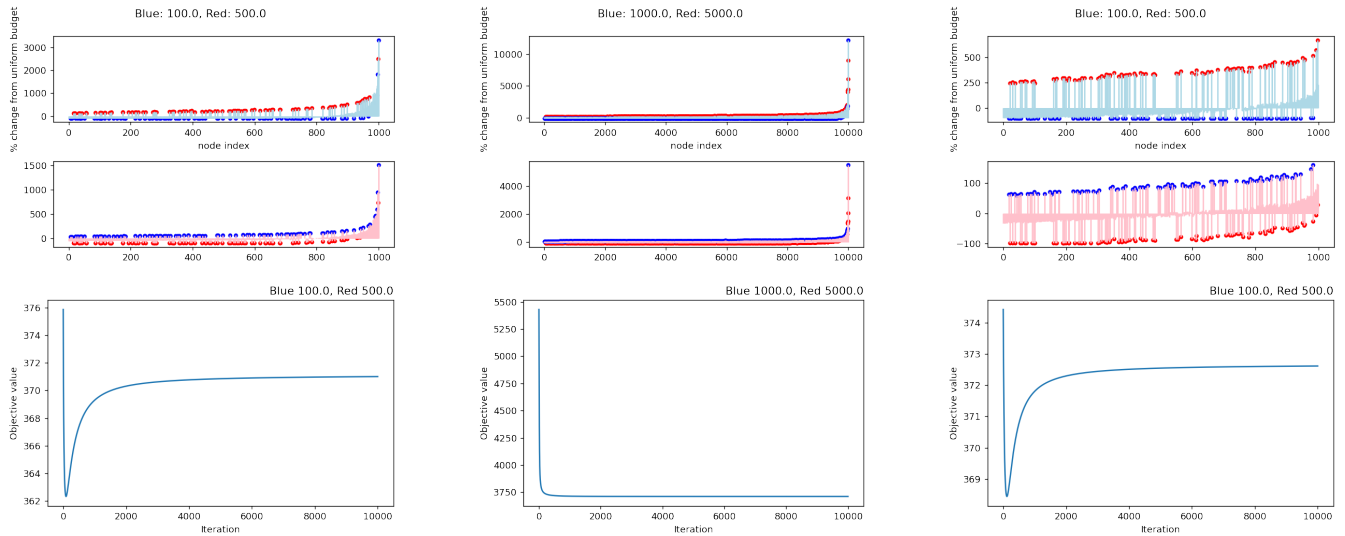


Figure 1: Top plots show percent change (from uniform strategy) of budgets at equilibrium; nodes are sorted by their component of the second eigenvector of the adjacency matrix, reflecting graph structure. The bottom panel illustrates the convergence rate of the equilibrium computation algorithm. Red and blue nodes are highlighted with dots. From left to right: a Barabasi-Albert graph of 1000 nodes and degree parameter 3, a Barabasi-Albert graph with 10000 nodes and degree parameter 3, and a Watts-Strogatz graph with 1000 nodes, initial degree 3, and rewiring probability 0.2.

The two-agent game is therefore a saddle-point problem (after rescaling budgets to lie on the simplex):

$$\max_{\mathbf{b}_{\cdot,1} \in \Delta^n} \min_{\mathbf{b}_{\cdot,2} \in \Delta^n} \left(\frac{s_{i,1} + B_1 b_{i,1}}{1 + B_1 b_{i,1} + B_2 b_{i,2}} \right)_{i \in V}^T \mathbf{Ac}$$

where \mathbf{c} is the vector representation of the payoff function (i.e. $f(\mathbf{z}) = \mathbf{z}^T \mathbf{c}$). By Lemma 3.2, the objective of this problem is concave in $\mathbf{b}_{\cdot,1}$; by Lemma 3.1 it is convex in $\mathbf{b}_{\cdot,2}$.

The gradients of the objective function are:

$$\nabla_{\mathbf{b}_{\cdot,1}} f(\mathbf{b}_{\cdot,1}, \mathbf{b}_{\cdot,2}) = B_1 \left(\frac{1 + B_2 b_{i,2} - s_{i,1}}{(1 + B_1 b_{i,1} + B_2 b_{i,2})^2} \right)_{i \in V} \odot \mathbf{Ac}$$

$$\nabla_{\mathbf{b}_{\cdot,2}} f(\mathbf{b}_{\cdot,1}, \mathbf{b}_{\cdot,2}) = -B_2 \left(\frac{B_1 b_{i,1} + s_{i,1}}{(1 + B_1 b_{i,1} + B_2 b_{i,2})^2} \right)_{i \in V} \odot \mathbf{Ac}$$

Then it becomes straightforward to use the classic mirror descent approach (Nemirovsky and Yudin 1983; Bubeck 2014) with mirror map $\Phi(x) = \sum_i x_i \log x_i$, $\nabla \Phi(x) = 1 + \log x$, which is equivalent to exponentiated gradient descent (Kivinen and Warmuth 1997) or multiplicative weight updates (Freund and Schapire 1997).

Results on random graphs

We experiment with computing equilibrium strategies on a number of random graphs from the Barabasi-Albert and Watts-Strogatz families (as implemented in NetworkX) (Watts and Strogatz 1998; Barabási and Albert 1999), with incoming edge weights set to $1/(\deg(v) + 1)$. We treat blue and red as max and min agents respectively. Blue nodes have initial opinions s which are 0.9 favorable to blue, red nodes 0.1 favorable to blue, and all others 0.5 favorable to blue.

We run all experiments for 10,000 iterations with a step size of 3×10^{-4} .

Figure 1 shows equilibrium budget allocations for several random graphs, where the blue budget (the constraint B on strategies b) is 10% of the $|V|$ nodes and the red budget is 50%. We randomly assign 10% of nodes to be favorable to red or blue. The solutions converge quite quickly to equilibrium, and the chosen budgets reflect graph structure to some extent as well as initial opinions of the nodes. Plots from additional experiments on a range of parameters are included in the technical appendix.

5 Conclusions & future research

We provide a model of multi-agent games of influence on networks, which admits scalable, polynomial-time computation of Nash equilibria, unlike many other models of multi-agent influence maximization. Additionally, our game is socially concave, implying that these Nash equilibria are pure, and we also show that under very weak additional conditions, they are unique. By instantiating the game as a saddle point problem, we test our algorithm for the two-player case on a number of random graphs.

Our game is closely related to well-studied models of influence spreading on social networks, and it has many computational advantages over other models of competitive diffusion. Work remains to be done, however, on seeing how real-world social phenomena may be similar to or different from the behaviors of nodes and players under our model. In particular, because our model results in a game that is socially concave, mirror-descent-based approaches such as the one we presented in Section 4 could potentially be applicable to general multi-agent setting in practice; as motivated

by our general multi-agent theory, this setting is applicable to a broad swath of real-world settings.

Ethical implications of the work

Diffusion processes in real-world networks are myriad, including: contagion or disease spread based on an underlying sociobehavioral and geospatial network, the spread of malware in computer networks, and the spread of (mis)information in social networks. In turn, one or more agents may be interested in manipulating that process. As with many research directions, we see both potential positive and negative impacts of the work.

Take the spread of (mis)information in social networks such as Twitter, Facebook, and Sina Weibo. Here, a user may share a news article—a form of expressing a belief publicly—that they find interesting—a function of their internal beliefs. That news article may influence the internal beliefs of the direct connections of the initial user; those influenced users then may choose to share or not share to their local neighborhood—again, a form of expressing their beliefs publicly—and the process continues. It may be that one or more agents (e.g., state actors, advertisers) would benefit from greater spread of particular types of associated news. Other actors including news organizations, competing advertisers, or recently even the platforms themselves may be interested in *limiting* the diffusion of that same news, or maximizing the diffusion of a correction or opposing evidence. Our work could potentially lead to greater ability for an agent to *push* a particular message, and could simultaneously lead to a greater ability for an agent to *defend against* that same message. Herein lies the potential for both positive and negative impacts.

Additionally, advertising or competitive influence diffusion is a feature of many political and military conflicts. As with other applications described above, the social benefit or disadvantage associated with our model depends on the content of the belief or influence it is being used to spread. Some agents in conflict seek to influence opinions by providing public services or building infrastructure. Others spread propaganda messages intended to incite participation in violence or influence public approval of violence.

The ongoing COVID-19 public health crisis presents a third important situation where various agents are interested in influencing nodes to adopt certain beliefs—and more importantly, behaviors that follow from those beliefs—about the crisis. Governments and others trying to spread accurate public health information and scientifically sound behavioral guidelines always operate with limited budgets and without the ability to directly “advertise” important information to every node. At the same time, we see in the ongoing crisis that other agents (with similarly limited budgets) are interested in propagating inaccurate information and unfounded behavioral guidelines across the same networks. For networks with different initial conditions with respect to trust in/influence of government, different resource levels for relevant agents, and different underlying structures, our model has the potential to create positive impacts by helping agents identify advertising strategies that propagate accurate

information and good behavioral influences as far as possible given limited budgets. Of course, information about optimal advertising distribution could also help those interested in spreading bad behavioral influences.

These are three real-world instantiations of the same underlying mathematical model, all of which could potentially be captured by our work. Any action taken at the agent level in each of these scenarios would be a morally-laden decision. Because our model facilitates efficient allocation of influencing resources in competitive settings, our work in theory benefits agents who use it, and disadvantages whoever they are competing against. If the weaker agent in some hypothetical competition uses our model and the stronger agent does not, our work has the effect of reducing power disparities. If the stronger agent uses the model, our work has the effect of amplifying power disparities.

Because our work might increase power disparities in competitive settings where only the “stronger” agent uses the model, it is important to point out the steps we have taken to make our work accessible and implementable for agents with relatively lower resources. First, our work can be implemented using only open source tools, and relies mainly on well-known Python libraries like the NumPy stack and NetworkX. Second, equilibrium strategies can be computed without specialized hardware—the experiments in this paper were run on a laptop. The major remaining impediment to equitable access to our work is access to the raw inputs: appropriately structured network data for which equilibrium strategies can be computed.

Acknowledgments

All authors acknowledge funding support from DARPA SI3-CMD #S4761. Dickerson and Curry were supported in part by NSF CAREER Award IIS-1846237, NSF Award CCF-1852352, NSF D-ISN Award #2039862, NIST MSE Award #20126334, DoD WHS Award #HQ003420F0035, NIH R01 Award NLM-013039-01, DARPA GARD Award #HR00112020007, and a Google Faculty Research Award. Daskalakis was supported by NSF Awards IIS-1741137, CCF-1617730 and CCF-1901292, by a Simons Investigator Award, and by the DOE PhILMs project (No. DE-AC05-76RL01830). Hajiaghayi and Knittel were supported in part by Guggenheim Fellowship, NSF Awards CCF-SPX 1822738, IIS-BIGDATA 1546108, and a Northrop Grumman Faculty Award. The authors acknowledge the MIT SuperCloud and Lincoln Laboratory Supercomputing Center for providing HPC, database, and consultation resources that have contributed to the research results reported within this paper/report.

References

- Acemoglu, D.; Dahleh, M. A.; Lobel, I.; and Ozdaglar, A. 2011. Bayesian learning in social networks. *The Review of Economic Studies* 78(4): 1201–1236.
- Acemoglu, D.; and Ozdaglar, A. 2011. Opinion dynamics and learning in social networks. *Dynamic Games and Applications* 1(1): 3–49.

- Ahmadinejad, A.; Dehghani, S.; Hajiaghayi, M.; Mahini, H.; Seddighin, S.; and Yazdanbod, S. 2014. How effectively can we form opinions? In *the 23rd International World Wide Web Conference (WWW)*.
- Ahmadinejad, A.; Dehghani, S.; Hajiaghayi, M.; Mahini, H.; Seddighin, S.; and Yazdanbod, S. 2015. Forming external behaviors by leveraging internal opinions. In *IEEE Conference on Computer Communications (INFOCOM)*, 1849–1857.
- Alon, N.; Feldman, M.; Procaccia, A. D.; and Tennenholtz, M. 2010. A note on competitive diffusion through social networks. *Information Processing Letters* 110(6): 221–225.
- Barabási, A.-L.; and Albert, R. 1999. Emergence of scaling in random networks. *science* 286(5439): 509–512.
- Bhawalkar, K.; Gollapudi, S.; and Munagala, K. 2013. Co-evolutionary opinion formation games. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, 41–50.
- Bindel, D.; Kleinberg, J.; and Oren, S. 2015. How bad is forming your own opinion? *Games and Economic Behavior* 92: 248–265.
- Bubeck, S. 2014. Theory of convex optimization for machine learning. *arXiv preprint arXiv:1405.4980* 15.
- Cai, Y.; Candogan, O.; Daskalakis, C.; and Papadimitriou, C. 2016. Zero-sum polymatrix games: A generalization of minmax. *Mathematics of Operations Research* 41(2): 648–655.
- Cai, Y.; and Daskalakis, C. 2011. On minmax theorems for multiplayer games. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete algorithms*, 217–234. SIAM.
- Chen, W.; Yuan, Y.; and Zhang, L. 2010. Scalable influence maximization in social networks under the linear threshold model. In *2010 IEEE international conference on data mining*, 88–97. IEEE.
- Chen, X.; Deng, X.; and Teng, S.-H. 2009. Settling the Complexity of Computing Two-Player Nash Equilibria. *Journal of the ACM*.
- Daskalakis, C.; Goldberg, P.; and Papadimitriou, C. 2008. The Complexity of Computing a Nash Equilibrium. *SIAM Journal on Computing*.
- Daskalakis, C.; and Papadimitriou, C. H. 2009. On a network generalization of the minmax theorem. In *International Colloquium on Automata, Languages, and Programming*, 423–434. Springer.
- DeGroot, M. H. 1974. Reaching a consensus. *Journal of the American Statistical Association* 69(345): 118–121.
- Ellison, G. 1993. Learning, local interaction, and coordination. *Econometrica: Journal of the Econometric Society* 1047–1071.
- Even-Dar, E.; Mansour, Y.; and Nadav, U. 2009. On the convergence of regret minimization dynamics in concave games. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, 523–532.
- Freund, Y.; and Schapire, R. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences* 55(1): 119–139.
- Friedkin, N. E.; and Johnsen, E. C. 1990. Social influence and opinions. *Journal of Mathematical Sociology* 15(3-4): 193–206.
- Gale, D.; and Kariv, S. 2003. Bayesian learning in social networks. *Games and Economic Behavior* 45(2): 329–346.
- Garimella, K.; Gionis, A.; Parotsidis, N.; and Tatti, N. 2017. Balancing information exposure in social networks. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, 4663–4671.
- Gionis, A.; Terzi, E.; and Tsaparas, P. 2013. Opinion maximization in social networks. In *Proceedings of the 2013 SIAM International Conference on Data Mining*, 387–395. SIAM.
- Hagenbach, J.; and Koessler, F. 2010. Strategic communication networks. *The Review of Economic Studies* 77(3): 1072–1099.
- Jackson, M. O. 2010. *Social and economic networks*. Princeton university press.
- Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 137–146.
- Kivinen, J.; and Warmuth, M. K. 1997. Exponentiated gradient versus gradient descent for linear predictors. *information and computation* 132(1): 1–63.
- Kuran, T.; and Sandholm, W. H. 2008. Cultural integration and its discontents. *The Review of Economic Studies* 75(1): 201–228.
- Maehara, T.; Yabe, A.; and Kawarabayashi, K. 2015. Budget Allocation Problem with Multiple Advertisers: A Game Theoretic View. In Bach, F. R.; and Blei, D. M., eds., *International Conference on Machine Learning (ICML)*, 428–437.
- Montanari, A.; and Saberi, A. 2010. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences* 107(47): 20196–20201.
- Nemirovsky, A. S.; and Yudin, D. B. 1983. Problem complexity and method efficiency in optimization. .
- Tzoumas, V.; Amanatidis, C.; and Markakis, E. 2012. A Game-Theoretic Analysis of a Competitive Diffusion Process over Social Networks. In Goldberg, P. W., ed., *International Workshop On Internet And Network Economics (WINE)*, 1–14. Springer.
- Watts, D. J.; and Strogatz, S. H. 1998. Collective dynamics of ‘small-world’ networks. *nature* 393(6684): 440–442.