



HAL
open science

Dynamic competition over social networks

Antoine Mandel, Xavier Venel

► **To cite this version:**

Antoine Mandel, Xavier Venel. Dynamic competition over social networks. *European Journal of Operational Research*, Elsevier, 2020, 280 (2), pp.597-608. 10.1016/j.ejor.2019.07.028 . halshs-02334595

HAL Id: halshs-02334595

<https://halshs.archives-ouvertes.fr/halshs-02334595>

Submitted on 20 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial| 4.0 International License

Dynamic competition over social networks

Antoine Mandel*, Xavier Venel (corresponding author) †

Abstract

We propose an analytical approach to the problem of influence maximization in a social network where two players compete by means of dynamic targeting strategies. We formulate the problem as a two-player zero-sum stochastic game. We prove the existence of the uniform value: if the players are sufficiently patient, both can guarantee the same mean-average opinion without knowing the exact length of the game. Furthermore, we put forward some elements for the characterization of equilibrium strategies. In general, players must implement a trade-off between a forward-looking perspective, according to which they aim to maximize the future spread of their opinion in the network, and a backward-looking perspective, according to which they aim to counteract their opponent's previous actions. When the influence potential of players is small, we describe an equilibrium through a one-shot game based on eigenvector centrality.

Keywords: Game theory, Social Network, Dynamic games, Targeting, Stochastic games

1 Introduction

"*Delivering the right message to the right person at the right time*" is a common dictum among influencers in politics, lobbying, and marketing. The large increase in information about individual characteristics and social interactions brought about by the development of internet and online social networks has generated tremendous interest, among both practitioners and scientists, on the problem of identifying appropriate targets to maximize influence. In marketing, targeted advertising on social networks has become a cornerstone of the industry in less than a decade. In politics, Barack Obama's campaign for the U.S. presidential election in 2008 has illustrated the tremendous potential for influence carried by social networks (Cogburn and Espinoza-Vasquez, 2011). In the academic literature, identifying the key target in a network in order to gain control, influence, or market shares has become a central focus of computer science, economics, and operations research. Most of the existing literature examines optimal/efficient strategies for a single agent (Kempe et al., 2003; Ballester et al., 2006). More recently, the competitive nature of the targeting problem has been emphasized and game-theoretic contributions have started to analyze the behavior of players competing for prominence over a network (Goyal et al., 2014; Bimpikis et al., 2016; Grabisch et al., 2017).

We note, however, that a common feature of most of the existing approaches is their focus on identification of the key target, i.e., "*the right person*". The temporal dimension, "*the right time*", hasn't yet been approached in a competitive setting. Its role is nevertheless crucial in applications. In political campaigning, as emphasized by Granato and Wong (2004), "*the relation between voters and campaign strategists is dynamic and evolves until voters' views on a candidate crystallize.*" In marketing, an important element of a firm's strategy is the sequence in which a product is offered to potential buyers (Hartline et al., 2008).

*Paris School of Economics, Centre d'économie de la Sorbonne, Université Paris 1 Panthéon-Sorbonne, 106-110 Boulevard de de l'hôpital, 75647 Paris Cedex 13, France E.mail: antoine.mandel@univ-paris1.fr

†Paris School of Economics, Centre d'économie de la Sorbonne, Université Paris 1 Panthéon-Sorbonne, 106-110 Boulevard de de l'hôpital, 75647 Paris Cedex 13, France E.mail: xavier.venel@univ-paris1.fr

The main objective of this paper is to include this temporal dimension in the analysis of the competitive targeting problem on networks. Therefore, we place ourselves in a competitive setting where two players choose a sequence of targets in a social network in order to influence the average opinion that gets formed about a given topic, e.g., a political candidate, a product, or the truth value of a piece of information. The network is made up of non-strategic agents whose opinion is represented by a real number and who update their opinions using a weighted average of their neighbors’ opinions (as in Golub and Jackson, 2010). Strategic players choose a sequence of targets among non-strategic agents. Hence, they directly influence the opinion of the targeted agent and indirectly influence the whole network through the dynamics of opinion formation.

In this setting, we are concerned with two main issues: the existence of a uniform value in the underlying (infinite-horizon) game and the characterization of optimal strategies, i.e., the determination of an optimal sequence of targets. A priori, both issues can be sensitive, on the one hand to the influence potential of the two strategic players, which measures the extent to which strategic players are able to change the opinion of the targeted agent and which may differ between players, and on the other hand to the informational structure of the game, that is, what each player observes of his opponent’s actions.

The existence of a uniform value is non-trivial in dynamic/stochastic games with compact state spaces. However, in the general case, it may not exist, as shown by (Ziliotto, 2016). The first part of this paper is devoted to this issue. We highlight the relation between the linear updating of the vector of opinions, using the network’s matrix of weights that we consider here, and the conventional stochastic game obtained when weights are interpreted as transition probabilities. We show that if the matrix is primitive, strategic influence “contracts” the vector of opinions in the network. This suffices to show that there exists a uniform value and that this value is independent of the initial distribution of opinions. Furthermore, this result holds independently of the informational structure of the game.

As for the characterization of equilibrium strategies, i.e., the determination of an optimal sequence of targets, a key observation is that players must implement a trade-off between a forward-looking perspective, according to which they aim to maximize the future spread of their opinion in the network, and a backward-looking perspective, according to which they aim to counteract their opponent’s previous actions. It is likely impossible to provide a closed-form analytical characterization of dynamic strategies implementing this trade-off in the general case. Nevertheless, we provide several series of results that highlight the relationship between network structure, the strength of the influence exerted by strategic players, and the type of strategies that ought to be used at equilibrium. We first show that, in highly symmetric networks, players must adopt a purely backward-looking perspective. They should focus on counteracting the influence of their opponent so as to keep the distribution of opinions in an acceptable state. Second, we show that in very hierarchical, hence asymmetric, networks, players must adopt a purely forward-looking perspective. Independently of the current distribution of opinions, they should target (sequences of) influential nodes in order to foster the forward diffusion of their opinion in the network. Finally, we show that, if the level of influence is small, the game can be approximated by a static game similar to the one considered in Grabisch et al. (2017). Note in particular that, if the players have equal influence and the non-strategic agents are equally influenced, they should target the node with the largest eigenvector centrality (a fact that went unnoticed in Grabisch et al. (2017)). Our results also highlight the importance of the temporal structure of strategies. For certain networks, there exist pure equilibria with dynamic targeting strategies, whereas no pure equilibria exist when a fixed target must be chosen (as in Grabisch et al., 2017). This stems from the fact that dynamic strategies are less easily counteracted than static ones.

The remainder of the paper is organized as follows. Section 2 briefly reviews the related literature. Section 3 presents a game-theoretic model of dynamic competitive targeting and proves the existence of the uniform value. Section 4 provides a partial characterization of equilibrium strategies. Section 5 compares dynamic and static strategies. Section 6 concludes. All proofs are given in the appendix.

2 Related literature

Since the seminal work in quantitative sociology of the 1970s (notably Bonacich, 1972; Freeman, 1977), the identification of central agents has been a key research question in the analysis of socio-economic networks. With the rise of digital social networks and the correlative flow of data, many applications have been identified in communication and marketing, and the problem has been approached from a more goal-oriented perspective at the interface between economics and computer science. The targeting problem, i.e., the choice of an optimal target in view of the diffusion of a product or a message, has been widely investigated from an algorithmic perspective in computer science, following on from the early contributions of Domingos and Richardson (2001) and Kempe et al. (2003, 2005). In economics, two important contributions are Ballester et al. (2006) and Banerjee et al. (2013). The first characterizes a key target through an “inter-centrality measure” that takes into account both a player’s centrality and his contribution to the centrality of others, while the second develops a model that discriminates between information passing and endorsement and derives a measure of “diffusion centrality” to characterize efficient targets.

Now, the bulk of the literature has approached the targeting problem as an optimization problem for a single influencing agent. A competitive perspective has been introduced in a few recent contributions that develop game-theoretic models in which agents choose, ex-ante, a targeting strategy to maximize their influence on the outcome of the ensuing opinion dynamics. The contributions of (Carnes et al., 2007; Bharathi et al., 2007; Goyal et al., 2014; He and Kempe, 2013) offer an algorithmic perspective on the issue while (Lever, 2010; Dubey et al., 2014; Bimpikis et al., 2016; Grabisch et al., 2017) focus on the characterization of optimal targets. In particular, Grabisch et al. (2017) is a direct predecessor of this work and puts forward “intermediacy centrality” as a minimax characterization of the key target which accounts for the fact that, in a competitive context, agents should focus on relative, rather than absolute, centrality. Dubey et al. (2014) considers a different model where strategic players can interact with several non-strategic agents at the same time: the players choose an expenditure for each non-strategic agent at a given cost. Their model is also static in the sense that the expenditures of the players are fixed at the beginning of the game.

The importance of adopting a dynamic perspective when designing influence strategies is also clearly identified in the literature. In political science, Box-Steffensmeier and Lin (1996) emphasizes the fact that *“at different stages of the campaign process candidates have different goals, and their expenditures should have different effects on the final election outcome”*. Granato and Wong (2004) provides a simulation model of the dynamics of an electoral campaign from this perspective, while Demange and Van der Straeten (2009); Aragonés et al. (2015) and Amorós and Puy (2013) investigate in a game-theoretic setting how candidates should allocate their campaign time to the different issues of concern for the electorate.

The computer science literature documents empirically both the diffusion of influence in social networks such as Twitter and the dynamic behavior of agents in view of influence maximization. With respect to the former, Raghavan et al. (2013) shows that the predictive power of models of users’ activity on Twitter is improved by taking into account the dynamics of social interactions. With respect to the latter, Lee et al. (2014) provides empirical evidence for the evolution of behavior in online social networks, while Mekouar et al. (2017) highlights the importance of the dynamic updating of content in social networks in view of revenue maximization. More specifically related to the present paper, a number of contributions in the computer science literature have considered the problem of dynamic targeting in a social network for a single influencer. Golovin and Krause (2011) considers the design of adaptive algorithms to address the targeting problem when influences spread randomly over the network. Zhuang et al. (2013) also puts forward an adaptive approach in a setting where the influencer periodically probes the network in order to adjust his influence strategy. Yadav et al. (2016) approaches the same problem, but from the more formal perspective of a partially observable Markov decision process (POMDP). This is further developed by Yu et al. (2010), which focuses on the identification of spread blockers in dynamic networks. Mekouar et al. (2015) includes both a dynamic and a competitive perspective, since it focuses on the timely delivery of content from two competing sources.

From a more economic standpoint, Hartline et al. (2008) analyses the marketing strategy of a firm

through a social network in a dynamic programming framework and emphasizes the complementarity between pricing and targeting: “*In general it is advantageous to get influential buyers to buy the item early in the sequence; it even makes sense to offer such buyers smaller prices to get them to buy the item.*” There is also a wide literature on competition in advertising. Chakrabarti and Haller (2011) provides a simple model of advertising wars. These have a long history, as documented in Beard (2010). Doraszelski and Markovich (2007) models dynamic competition among firms through advertising and analyzes its impact on industrial dynamics. Finally, Pastine and Pastine (2002) analyzes the use of advertising by two competing firms as an expectation coordination device in the purchasing decision of consumers when there are consumption externalities.

However, the present paper is to our knowledge the first to include both a dynamic and a competitive perspective on the problem of targeting in a social network. As such, it raises a number of theoretical issues. The first is the definition of an appropriate solution concept. We choose to focus on the mean-average payoff and on the uniform approach (Mertens and Neyman, 1981), which seems to us well suited to extend the results of Grabisch et al. (2017) to a dynamical framework. However, the existence of the uniform value is a complicated question in our framework, where the set of states is compact. It has been solved positively in several classes of stochastic games (see Solan, 2012; Jaśkiewicz and Nowak, 2016) or as an intermediate result for stochastic games with incomplete information (Renault, 2012; Aumann et al., 1995). However, none of these results apply here. Moreover, it has been shown recently that, even under reasonable assumptions, the uniform value might not exist (Ziliotto, 2016).

3 Model

3.1 Notation

We shall use the following notation throughout:

- In the vector space \mathbb{R}^K , where $K \in \mathbb{N}^*$, e_i denotes the i th vector of the canonical basis and $e' = (1, \dots, 1) = \sum_{i=1}^K e_i$.
- Given \mathcal{K} a finite set of cardinality K , we denote by $\Delta(\mathcal{K})$ the set of probability distributions over \mathcal{K} . It can be viewed as a subset of \mathbb{R}^K : $\{a \in \mathbb{R}_+^K, \sum_{k \in \mathcal{K}} a_k = 1\}$. We consider the $\|\cdot\|_\infty$ norm on $\Delta(\mathcal{K})$ defined by

$$\forall a \in \Delta(\mathcal{K}), \quad \|a\|_\infty = \max_{k \in \mathcal{K}} |a_k|.$$

- Given two natural numbers $i, j \in \mathbb{N}$, $\delta_{i,j}$ denotes the Kronecker symbol such that $\delta_{i,j} = 1$ if $i = j$ and zero otherwise.
- Given a matrix M of size $K \times K$, M_k denotes the k th row vector.
- Given two vectors $x, y \in \mathbb{R}^K$, $x \cdot y$ denotes their scalar product.
- Given a family of matrices $(M_i)_{i \in \mathbb{N}}$, $\prod_{i=a}^{\leftarrow b} B_i$ stands for $\prod_{j=0}^{b-a} B_{b-j}$ and $\prod_{i \in \emptyset} B_i$ stands for the identity matrix I .

3.2 Opinion game

We consider two strategic players who compete for influence on a social network through dynamic targeting strategies. The network is formed by K non-strategic agents whose social ties are represented by a row-stochastic matrix M of size $K \times K$, accounting for the fact that the network is in general directed. Each of the non-strategic agents is characterized by his opinion on an issue of concern for the two strategic players, e.g., the relative quality of two products or the proximity relative to two opposing electoral platforms. This opinion is represented by a number in $[-1, 1]$ and, by default, evolves according to the social

influence exerted on the agent, i.e., each agent updates his opinion by linearly combining the opinions of his neighbors in proportions determined by the corresponding row of the social network matrix M . This boundedly rational model of opinion formation was introduced by DeGroot (1974) and has received wide attention in the economic literature as a model of opinion formation in social networks (see e.g. Golub and Jackson, 2010, and references therein). It is particularly well suited to contrasting the myopic behavior of a “crowd” with the strategic behavior of external influencers such as lobbies, advertisers, or political campaigners (see Bimpikis et al., 2016; Grabisch et al., 2017).

In our framework, the two strategic players are referred to as Player 1 and Player -1 . They are characterized by their fixed opinions, 1 and -1 , respectively, and their levels of influence, λ and μ , respectively. They aim to bring the average opinion in the network as close as possible to their own opinion, e.g., to trigger the purchase of their product or a vote for their political platform. In each period, they simultaneously choose a target among the non-strategic agents with a view to influencing his opinion. Targeted agents then update their opinions through a convex combination of the opinions of their non-strategic neighbors and of the strategic agent(s) who targeted them. A strategic player might only partially observe the decision of the other player.

3.2.1 Structure of the game

Formally, we define an n -stage game $\Gamma_n := (n, \mathcal{K}, M, x_1, \lambda, \mu, \nu, C^1, C^{-1}, s^1, s^{-1})$ through the following parameters: a non-negative integer n representing the length of the game, a finite set \mathcal{K} representing non-strategic agents, a row-stochastic matrix M representing the local interaction between the non-strategic agents, a vector x_1 in $X := [-1, 1]^K$ representing the initial vector of opinions of the non-strategic agents, a vector ν in \mathbb{R}_+^K representing the influenceability of the non-strategic agents, a positive real number λ representing the relative influence of Player 1, a positive real number μ representing the relative influence of Player -1 , a finite set C^1 of signals and a function s^1 from $\mathcal{K} \times \mathcal{K}$ to C^1 representing the signaling function of Player 1, a finite set C^{-1} of signals and s^{-1} a function from $\mathcal{K} \times \mathcal{K}$ to C^{-1} representing the signaling function of Player -1 .¹

The game Γ_n is played in n periods between the two strategic players, Player 1 and Player -1 , as follows:

- Initially, each non strategic agent, $k \in \mathcal{K}$, holds an opinion $x_{1,k} \in [-1, 1]$ given by the vector $x_1 \in [-1, 1]^K$.
- In every period $t \geq 1$:
 - Player 1 chooses a target $i_t \in \mathcal{K}$ and Player -1 chooses a target $j_t \in \mathcal{K}$.
 - Each non-strategic agent updates his belief using a convex combination of the opinions of his non-strategic neighbors and of the strategic players who target him. For all $k \in \mathcal{K}$,

$$x_{t+1,k} = q_k(x_t, i_t, j_t) := \begin{cases} M_k \cdot x_t & \text{if } \nu_k = 0, k \neq i_t \text{ and } k \neq j_t, \\ \frac{\nu_k M_k \cdot x_t + \lambda \delta_{k,i_t} - \mu \delta_{k,j_t}}{\nu_k + \lambda \delta_{k,i_t} + \mu \delta_{k,j_t}} & \text{otherwise} \end{cases} \quad (3.1)$$

This can be summarized in matrix notation as follows. Let $A^1(i, j) = \frac{\lambda}{\nu_k + \lambda + \mu \delta_{j,k}} e_i$, $A^{-1}(i, j) = \frac{\mu}{\nu_k + \lambda \delta_{i,k} + \mu} e_j$ and $B(i, j)$ be the $K \times K$ matrix with coefficients

$$B(i, j)_{kl} := \begin{cases} M_{kl} & \text{if } \nu_k = 0, k \neq i \text{ and } k \neq j, \\ \frac{\nu_k M_{kl}}{\nu_k + \lambda \delta_{k,i} + \mu \delta_{k,j}} & \text{otherwise.} \end{cases} \quad (3.2)$$

¹In the following, we shall also use the notation $\Gamma_n(x_1)$ to emphasize the dependence on the initial vector of opinions x_1 .

Then, we obtain

$$x_{t+1} = B(i_t, j_t)x_t + A^1(i_t, j_t) - A^{-1}(i_t, j_t). \quad (3.3)$$

$A^1(i_t, j_t)$ is the influence of Player 1 in one stage and $A^{-1}(i_t, j_t)$ is the influence of Player -1, while $B(i_t, j_t)$ represents the remaining local interactions between non-strategic agents.

– Player 1 then receives the signal $s^1(i_t, j_t)$, while Player -1 receives the signal $s^{-1}(i_t, j_t)$.

- If $t < n$, a new period begins. If $t = n$, the game terminates.

Remark 3.1 *Given a sequence of n pairs of targets $(i_1, j_1, \dots, i_n, j_n)$, the vector of opinions of the non-strategic agents at stage $t \leq n$ can be written in matrix notation as*

$$x_t = \prod_{m=1}^{t-1} B(i_m, j_m)x_1 + \sum_{l=1}^{t-1} \left(\prod_{m=l+1}^{t-1} B(i_m, j_m) \right) (A^1(i_l, j_l) - A^{-1}(i_l, j_l)). \quad (3.4)$$

The aim of Player 1 is to maximize the mean-average opinion across time and across non-strategic agents, while the aim of Player -1 is to minimize the mean-average opinion. More precisely, strategies and payoffs are defined as follows.

At stage t , Player 1 chooses a target, possibly using randomization, as a function of the information he learned in the previous stages. Hence, a strategy of Player 1 is a function $\sigma : \cup_{m=1}^{+\infty} H_m^1 \rightarrow \Delta(\mathcal{K})$, with $H_m^1 := (\mathcal{K} \times C^1)^{m-1}$ (the player observes his signals and his own sequence of actions). The corresponding set of strategies is denoted by Σ . Similarly, a strategy for player -1 is a function $\tau : \cup_{m=1}^{+\infty} H_m^{-1} \rightarrow \Delta(\mathcal{K})$, with $H_m^{-1} := (\mathcal{K} \times C^{-1})^{m-1}$, and the corresponding set of strategies is denoted by \mathcal{T} . The strategy of a player is said to be pure if it does not involve randomization, i.e., if for every history, the player targets a single node with probability 1. For convenience, we will refer to the choices of the players as actions. Moreover, we introduce the set of actions $H_n^a = (\mathcal{K} \times \mathcal{K})^n$ taken in the n -stage game by both players, while $H_\infty^a = (\mathcal{K} \times \mathcal{K})^\mathbb{N}$ is the set of infinite histories of actions.

Taken with the information structure, a triple (x_1, σ, τ) naturally induces a probability distribution on $H_n := (\mathcal{K} \times \mathcal{K} \times [-1, 1]^K \times C^1 \times C^{-1})^n$ denoted by $\mathbb{P}_{x_1, \sigma, \tau}$. We denote by $\mathbb{E}_{x_1, \sigma, \tau}$ the expectation under this probability distribution. When Player 1 follows the strategy σ and Player -1 follows the strategy τ , the average opinion across time and across non-strategic agents is therefore given by

$$\gamma_n(x_1, \sigma, \tau) = \mathbb{E}_{x_1, \sigma, \tau} \left(\frac{1}{Kn} \sum_{t=1}^n \sum_{k \in \mathcal{K}} x_{t,k} \right). \quad (3.5)$$

The payoff of Player 1 is given by $\gamma_n(x_1, \sigma, \tau)$ and the payoff of Player -1 by $-\gamma_n(x_1, \sigma, \tau)$. Hence, Player 1 (resp. -1) aims to maximize (resp. minimize) the average opinion across time and across agents.

3.2.2 Interpretation of the parameters

Hence, we consider a zero-sum dynamic game with incomplete information. The parameters of the game define on the one hand an influence structure (M, ν, λ, μ) and on the other hand an information structure $(C^1, C^{-1}, s^1, s^{-1})$.

Within the influence structure, the matrix M determines the dynamics of opinions among non-strategic agents in the absence of external influence. It is row-stochastic to represent the fact that an agent forms his opinion using a weighted average of his neighbours' opinions. The vector ν measures the sensitivity of non-strategic agents to external influence: the lower the value of ν_k , the more influenceable agent k . Finally, λ and μ measure the influence potential of the two strategic players. An interesting special case occurs when each non-strategic agent puts equal weight on his neighbours and the share of influence

received from the network is proportional to the number of neighbors. In this case, ν_k is the degree of agent k in the network and all non-zero entries of M_k have value $1/\nu_k$.

The information structure determines the type and extent of information a strategic agent receives about the actions of his opponent. Two cases will be of particular interest. First, the *blind information* structure, where s^1 and s^{-1} are constant on $I \times J$. Then, the players have no information about their opponent's actions and only take into consideration their own history of play. Second, the *complete information* structure where $s^1 = s^{-1} = \text{Id}_{I \times J}$ and both players are perfectly informed about the actions of their opponent. Note that, given a vector of initial opinions, the complete information structure is equivalent to a situation where the players perfectly observe the vector of opinions at the beginning of each period.

3.3 Value of the finite horizon game

For a given n , the game Γ_n is in fact equivalent to a finite zero-sum game, i.e., a matrix game. It thus follows from the classical result of von Neumann (1928) that the game has a value $v_n(x_1)$. That is, Player 1 can choose a strategy that ensures that the time average of opinions is above $v_n(x_1)$ independently of the strategy of Player -1, and reciprocally, Player -1 can ensure that the average opinion is below $v_n(x_1)$ independently of the strategy of Player 1. The corresponding pair of strategies form a Nash equilibrium of the game. More formally, one has:

Proposition 3.2 *The n -stage game $\Gamma_n(x_1)$ has a value:*

$$v_n(x_1) := \max_{\sigma \in \Sigma} \min_{\tau \in \mathcal{T}} \gamma_n(x_1, \sigma, \tau) = \min_{\tau \in \mathcal{T}} \max_{\sigma \in \Sigma} \gamma_n(x_1, \sigma, \tau).$$

In particular, there exists σ^* , a strategy of Player 1, such that for every strategy τ of Player -1, we have

$$\gamma_n(x_1, \sigma^*, \tau) \geq v_n(x_1).$$

Such a strategy will be said to be *optimal*. While following this strategy, Player 1 can guarantee the payoff $v_n(x_1)$ against any strategy played by Player -1. Similarly a strategy τ^* of Player -1 is optimal if

$$\forall \sigma \in \Sigma, \gamma_n(x_1, \sigma, \tau^*) \leq v_n(x_1).$$

Note in particular that the value of the finite-horizon game depends on the vector of initial opinions because the strategic players can only exert a limited amount of influence in finite time.

3.4 Uniform value

Strategic agents ought to take into account the diffusion and mixing of opinions in the network, which might be arbitrarily long. In order to account for this process, we focus on games with an arbitrarily long horizon. The notion of value does not extend straightforwardly to a setting with an arbitrarily long horizon. As a matter of fact, many notions of value coexist in this context. However, the notion of uniform value (Mertens and Neyman, 1981) provides a well-suited solution concept in our setting. Informally, it implies that there exist strategies that can be used independently of the length of the game to guarantee a given value. Formally, the uniform value is defined as follows. We denote the sequence $(\Gamma_n)_{n \in \mathbb{N}}$ by Γ , and by $\Gamma(x_1)$ when we wish to emphasize the initial vector of opinions.

Definition 3.3 *Let v be a real number.*

- *Player 1 can guarantee v in $\Gamma(x_1)$ if, for all $\varepsilon > 0$, there exists a strategy $\sigma^* \in \Sigma$ of Player 1 such that*

$$\liminf_n \inf_{\tau \in \mathcal{T}} \gamma_n(x_1, \sigma^*, \tau) \geq v - \varepsilon.$$

We say that such a strategy σ^ guarantees $v - \varepsilon$ in $\Gamma(x_1)$.*

- Player -1 can guarantee v in $\Gamma(x_1)$ if, for all $\varepsilon > 0$, there exists a strategy $\tau^* \in \mathcal{T}$ of Player -1 such that

$$\limsup_n \sup_{\sigma \in \Sigma} \gamma_n(x_1, \sigma, \tau^*) \leq v + \varepsilon.$$

We say that such a strategy τ^* guarantees $v + \varepsilon$ in $\Gamma(x_1)$.

- If both players can guarantee v , then v is called the uniform value of the game $\Gamma(x_1)$ and denoted by $v_\infty(x_1)$. The preceding strategies are then said to be ε -optimal in the infinite game.

In our setting, the existence of the uniform value ensures that both players can guarantee that the mean-average opinion is above (resp. below) a fixed threshold. Moreover, the optimal strategy does not depend on the length of the game. This implies that a player can play almost optimally in any long enough game without knowing either the exact length or the strategy of his adversary beforehand. The existence of the uniform value $v_\infty(x_1)$ also implies that the sequence $(v_n(x_1))_{n \geq 1}$ converges to $v_\infty(x_1)$ (Mertens and Neyman, 1981).

In our framework where the state space is infinite, the existence of the uniform value is a complicated question. It was recently shown that, even under reasonable assumptions, the uniform value might not exist (Ziliotto, 2016). There are positive existence results for special classes of games (see Jaśkiewicz and Nowak, 2016, for recent survey), but they do not apply in our setting. The next section is thus devoted to the existence of the uniform value in Γ .

3.5 Existence of the uniform value

A key assumption in models of opinion diffusion is the connectivity of the graph, which ensures that an opinion can propagate between any pair of agents. While studying the emergence of consensus, DeGroot (1974) stressed the fact that connectivity is not sufficient for the emergence of consensus. For example, if we consider a cycle of 2 agents $\{\alpha, \beta\}$, where each agent is influenced only by the other agent, then the dynamic of opinions is a cycle of length 2: the agents exchange their opinions at every stage. The associated Markov chain is periodic. This implies in particular that there is no consensus if the agents disagree at the beginning. This leads DeGroot (1974) to use the following notion from the theory of Markov chains.

Definition 3.4 A positive matrix M of size $K \times K$ is primitive if there exists $m \geq 1$ such that

$$\forall k \in \mathcal{K}, \forall l \in \mathcal{K}, (M^m)_{kl} > 0.$$

By extension, we will say that the network is primitive when the matrix M associated with the network is primitive. Informally, when the network is primitive, after m periods of diffusion, every non-strategic agent $k \in \mathcal{K}$ will be taking into account the opinion initially held by every agent: the network will be mixing opinions. A typical example of a network which is not primitive is a network in which every agent is listening to only one other agent. In this setting, the opinions will be circulating across the network but never mixing. On the contrary, a sufficient condition for the network to be primitive is that it should be strongly connected and that at least one non-strategic agent should put some strictly positive weight on his own opinion. **This implies on the one hand that each agent eventually influences each other agent and on the other hand that at least one agent has some confidence in his own opinion.**

In our setting, the primitiveness of the network ensures that the opinion game admits a uniform value and, moreover, that this value is independent of the initial distribution of opinions. In particular, one has:

Theorem 3.5 Let Γ be an opinion game such that M is primitive. For every $x_1 \in X$, the opinion game $\Gamma(x_1)$ has a uniform value. Moreover, there exists $v \in \mathbb{R}$ such that, for every $x_1 \in X$,

$$v_\infty(x_1) = v.$$

A key insight brought about by Theorem 3.5 is that influencers can design robust targeting strategies that are ϵ -optimal independently of the initial distribution of opinions in the network and independently of the length of the game (provided it is long enough). In other words, influence strategies can be determined on the sole basis of the structure of the network. Furthermore, these strategies can be designed a priori, before the player obtains information about the actions of his opponent. From the point of view of opinion dynamics, Theorem 3.5 implies the convergence of the average opinion in the network at equilibrium, but not the convergence of opinions towards a consensus. The latter result is standard in a setting where agents are non-strategic and thus contribute, not necessarily purposely, to a common objective (e.g., the discovery of the truth in Golub and Jackson (2010) or distributed optimization in Nedic et al. (2010)). In contrast, in our setting, strategic agents have a similar influence on the dynamics to stubborn agents in Yildiz et al. (2013): they prevent convergence to a consensus and imply that the asymptotic distribution of opinions is independent of the initial opinions of non-strategic agents (whence the independence of the value from initial opinions in Theorem 3.5). A major difference with Yildiz et al. (2013), however, is that we also consider our “stubborn” agents to be strategic in the choice of their placement in the network. Hence the game-theoretic nature of our model.

The key step in the proof of Theorem 3.5 is to divide the game into blocks of sufficient time length to ensure that the transition function is contracting. Hence, our model shares some similarities with stochastic games with state-independent transitions (Thuijsman, 1992), where the transition does not depend on the state variable. Nevertheless, our assumption is weaker since it is only asymptotically true. From this first result, we deduce that the value in very long games is almost independent of the original vector of opinions, whence the existence of the uniform value.

Remark 3.6 *Our proof fails if the network is not primitive, so the existence of the uniform value is open in the general case. It is interesting to comment on the non-primitive case and in particular why Grabisch et al. (2017) were able to assume only that the network is strongly connected, i.e., why they did not need the stronger notion that it is primitive. In Grabisch et al. (2017), the strategic agents cannot change their target during the game. Hence, even if the original network between non-strategic agents is periodic, the new network augmented by the strategic agents will be aperiodic: after some m' stages of the game, the opinion of every non-strategic agent $k \in \mathcal{K}$ is influenced at every stage by the opinion of the strategic players. In our model, the strategic agents may change their targets during the game and therefore adopt a cyclical behavior. In this case, it is possible for the opinion of a non-strategic agent at a given stage to be independent of the influence of the strategic agents. For example, in our previous example with two agents $\{\alpha, \beta\}$ who influence each other, this is the case if the two strategic players target α at even stages and β at odd stages: the opinion of β at an even stage is never influenced by the strategic agents. Under such strategies, the transition is not contracting and our proof does not hold.*

4 Characterization of equilibrium strategies

In this section, we focus on the characterization of equilibrium strategies. As emphasized in Theorem 3.5, the uniform value of the game is independent of the initial distribution of opinions among non-strategic agents. Therefore, we will focus on the case where initial opinions are uniformly set to 0 unless otherwise specified. In particular, we let v_n stand for $v_n(0)$, unless otherwise specified. For finite games, this assumption is not without loss of generality, but we will emphasize, where relevant, how our results can be extended to the case of arbitrary initial opinions.

The problem faced by the strategic players in a dynamic context involves a number of trade-offs. At each stage, a strategic agent must first decide whether to influence an easily influenceable agent, with low ν_k , or a central agent, which might, however, have a high ν_k and be harder to influence. Second, he must decide whether to confront the other influencer by choosing the same target or to shy away by choosing another target or playing a mixed strategy (e.g., if he has a much weaker influence potential than his opponent). Both trade-offs are already present in the static targeting problem where the influenced agent

is fixed through time (Bimpikis et al., 2016; Grabisch et al., 2017). The specific issue in a dynamic context is the trade-off between adopting a forward-looking perspective, according to which players aim to maximize the future spread of their opinion in the network, and a backward-looking perspective, according to which they aim to counteract their opponent’s previous actions. It is likely impossible to provide a closed-form analytic characterization of dynamic strategies implementing this trade-off in the general case. Nevertheless, we have identified three cases that illustrate the relation between network structure, level of influence, and equilibrium strategies in the opinion game. In highly symmetric networks, players must adopt a purely backward-looking perspective. They should only focus on counteracting the influence of their opponent so as to keep the current distribution of opinions in an “acceptable” state. In very hierarchical, hence asymmetric, networks, players must adopt a purely forward-looking perspective. Independently of the current distribution of opinions, they should target (sequences of) influential nodes in order to foster the diffusion of their influence in the network. Finally, if the level of influence of both players is small, the game is similar to the static game in Grabisch et al. (2017); note in particular that, if the players have equal influence and the non-strategic agents are equally influenced, they should target the node with the largest eigenvector centrality.

4.1 Backward-looking strategies

A purely backward strategy consists in targeting the node that is most likely to be influenced by the action the opponent took during the preceding period. The implementation of such a strategy clearly requires complete information. A first class of network where this strategy can be effective is the class where each node influences a single node, that is, when there exists a permutation σ of \mathcal{K} such that the adjacency matrix of the network is of the form $M_{i,j}^\phi = \delta_{j,\phi(i)}$. This strategy is also effective (like every other strategy) in the complete network (with equal weights).

In this section, we will focus on the broader class of networks with adjacency matrix of the form $M_{i,j}^{\phi,\alpha} = \alpha\delta_{j,\phi(i)} + (1-\alpha)/K$, where ϕ is a permutation of \mathcal{K} and $\alpha \in [0, 1]$. These are networks where each node puts a weight $(1 - \alpha)$ on the average opinion and a weight α on a specific node. This class notably contains the complete network with equal weights (when $\alpha = 0$), the circle (when $\alpha = 1$ and $\phi(i) = (i + 1) \bmod K$), the network where all agents are stubborn (when $\alpha = 1$ and $\phi(i) = i$), and perturbations of these (when $\alpha < 1$). Formally, we introduce the following assumption.

Assumption 4.1 *The game Γ is such that:*

- *The network has an adjacency matrix of the form $M_{i,j}^{\phi,\alpha} = \alpha\delta_{j,\phi(i)} + (1-\alpha)/K$, where $\alpha \in [0, 1]$ and ϕ is a permutation of \mathcal{K} .*
- *The players have equal influence (i.e., $\lambda = \mu$) and there exists $\nu \in \mathbb{R}_+$ such that, for all $k, k' \in \mathcal{K}$, $\nu_k = \nu_{k'} = \nu \leq \lambda$.*

One can then show that backward-looking strategies are optimal in networks satisfying Assumption 4.1. Notice that, when $\alpha = 1$, the matrix is not primitive, and this means that Theorem 3.5, which proves the existence of the uniform value, does not apply. Nevertheless, our construction of the optimal strategies will imply the existence of the uniform value.

The condition $\lambda = \mu$ allows the computation to be tractable. Unfortunately, when $\lambda \neq \mu$ it becomes impossible to characterize optimal strategies even in our restricted class. Qualitatively, it is nevertheless clear that the player with a higher influence has a positive advantage, while the other player still has some non-negligible influence: for example, when there is only one node, his opinion converges to $(\lambda - \mu)/(\lambda + \mu)$.

Theorem 4.2 *Suppose that Assumption 4.1 holds and that the players have complete information. Let (σ^*, τ^*) be the pair of strategies defined as follows.*

- In period 1, $\sigma^*(\emptyset)$ is an arbitrary target in \mathcal{K} and, for every $t \geq 1$ and every history $h_t = (i_1, j_1, \dots, i_t, j_t)$,

$$\sigma^*(h_t) = \phi^{-1}(j_t).$$

- In period 1, $\tau^*(\emptyset)$ is an arbitrary target in \mathcal{K} and, for every $t \geq 1$ and every history $h_t = (i_1, j_1, \dots, i_t, j_t)$,

$$\tau^*(h_t) = \phi^{-1}(i_t).$$

Then, the strategy σ^* (resp. τ^*) guarantees a value of 0 to Player 1 (resp. Player -1) in every n -stage game. Therefore, the value of any n -stage game Γ_n is 0, as is the uniform value of Γ , and (σ^*, τ^*) is, in each case, a pair of optimal strategies.

Remark 4.3 The proof of Theorem 4.2 highlights the fact that the equilibrium strategies σ^* and τ^* can also be defined as the strategies for which, in every period, a player targets the non-strategic agent whose opinion is the most different from his.

Theorem 4.2 can be contrasted with the fact that there does not exist an equilibrium in pure strategies in the absence of information in networks satisfying Assumption 4.1

Proposition 4.4 Suppose Assumption 4.1 holds with $K \geq 3$, $\alpha > 0$, and $n \geq 3$. If the players have blind information, there does not exist an equilibrium in pure strategies in the game Γ_n .

The contrast between Theorem 4.2 and Proposition 4.4 highlights the role of information and dynamics in devising efficient targeting strategies. Indeed, in the class of networks satisfying Assumption 4.1 with $\alpha > 0$, there is no equilibrium in pure strategies in the absence of information, whereas there exist equilibria with pure dynamic targeting strategies when information is complete. This contrast hinges on two facts. First, the structure of the network. In regular networks satisfying Assumption 4.1, the efficient strategy for players is to counteract the influence of their opponent. Second, the set of strategies available to the players. On the one hand, with complete information, players can counteract the influence of their opponent by implementing a purely backward-looking strategy. On the other hand, in the absence of information, players choosing a pure strategy ought to commit ex-ante to a given sequence of actions (see Section 4.2 below). His opponent can then gain a major advantage over the player by best-responding to this given sequence of targets. Hence, to prevent his influence from being too easily counteracted, the player must use a mixed strategy.

To further emphasize the role of network structure, one should note that the purely backward-looking strategies defined in Theorem 4.2 are not optimal in every type of network. Example 4.5 highlights this fact.

Example 4.5 We consider the set of non-strategic agents $\mathcal{K} = \{\alpha, \beta, \gamma\}$ and the local interaction matrix M given by

$$\begin{array}{c} \alpha \quad \beta \quad \gamma \\ \alpha \quad \left(\begin{array}{ccc} 1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 \\ 1/2 & 1/2 & 0 \end{array} \right) \\ \beta \\ \gamma \end{array}$$

Informally, each non-strategic agent puts equal weight on his own opinion and that of his predecessor. We further assume that, for every $k \in \mathcal{K}$, $\nu_k = 0$ and $\lambda = \mu = 1$, and also that the initial vector of opinions is 0.

In this setting, the backward-looking strategy put forward in Theorem 4.2 is not optimal in the game of length 3. Indeed if Player -1 plays according to τ^* , we consider the following strategy of Player 1:

- In stage 1, play uniformly over \mathcal{K} ,

- In the second stage, considering permutation symmetries, there are three types of vectors of opinions to take into account:

1. If $x_2 = (1, -1, 0)$, i.e., Player -1 targeted the predecessor of the target of Player 1, then Player 1 follows σ^* . This yields, in stage 3, the vector of opinions $p_3 = (-1, 1, 1/2)$.
2. If $x_2 = (-1, 1, 0)$, i.e., Player 1 targeted the predecessor of the target of Player -1, then Player 1 targets the agent with opinion 0. This yields the new vector of opinions $(0, -1, 1)$.
3. If $x_2 = (0, 0, 0)$, i.e., they chose the same targets, Player 1 then plays, in stage 3, a successor of his action in stage 2 such that x_3 is either $(0, 0, 0)$ or $(1, -1, 0)$.

We see that the payoff is positive in all cases and strictly positive in one of them, whence the mean average opinion is strictly positive. The strategy τ^* is not optimal. One can check that the result holds true for any n -stage game with $n \geq 1$, nevertheless, the computation shows that the short term advantage that is obtained by Player 1 at stage 2 decays over time. Hence, although the strategy is not optimal, it is ε -optimal in the infinite game. It is not clear whether this holds true in the general case.

4.2 Blind games and forward-looking strategies

In the absence of information, i.e., in the blind information game, the players cannot condition their actions on the history of the game. More precisely, a pure strategy of Player 1 in a blind game of length n is a sequence of actions $(i_1, \dots, i_n) \in \mathcal{K}^n$. Such a strategy guarantees a value v_n for Player 1 if, for every strategy $(j_1, \dots, j_n) \in \mathcal{K}^n$ of Player -1, one has $\gamma((i_1, \dots, i_n), (j_1, \dots, j_n)) \geq v_n$. Hence, for an equilibrium in pure strategies to exist, Player 1 must determine a priori a sequence of actions that will guarantee him the value v_n independently of the sequence of actions that his opponent chooses: he must be forward-looking. More precisely, a pair of strategies $(\sigma^*, \tau^*) = ((i_1^*, \dots, i_n^*), (j_1^*, \dots, j_n^*)) \in \mathcal{K}^n \times \mathcal{K}^n$ is an equilibrium in pure strategies of a blind game with value v_n if one has

$$\forall \sigma \in \mathcal{K}^n, \forall \tau \in \mathcal{K}^n, \quad \gamma_n(\sigma^*, \tau) \geq v_n \geq \gamma_n(\sigma, \tau^*) \quad (4.1)$$

Note that an equilibrium in pure strategies need not necessarily exist as Equation 4.1 is very demanding: players must be able to counteract a priori any sequence of actions chosen by their opponent. Yet, if such a pure equilibrium strategy exists, it will be very robust. It can, in fact, be used in any information structure. Indeed, the opponent cannot take advantage of any additional information as the original pure strategy is a best-response against any sequence of actions. In particular, given a pure strategy σ of Player 1 for the blind-information structure, let us denote by $\sigma_{(C^1, C^{-1}, s)}$ the strategy in the game with information structure (C^1, C^{-1}, s) that plays the sequence of actions σ independently of the signal received (and respectively $\tau_{(C^1, C^{-1}, s)}$ given a strategy τ of Player -1). Let us also denote by Γ_n^b the blind game associated with an arbitrary game Γ_n by replacing its information structure by the blind information structure. Then one has:

Proposition 4.6 *Let Γ_n be a game with information structure (C^1, C^{-1}, s) . Assume that (σ, τ) is an equilibrium in pure strategies in the associated n -stage blind game Γ_n^b . Then, $(\sigma_{(C^1, C^{-1}, s)}, \tau_{(C^1, C^{-1}, s)})$ is an equilibrium in the n -stage game Γ_n . Moreover the two games have the same value².*

In order to characterize these blind equilibrium strategies, it will prove useful to measure the total influence over time generated by the actions of the players in a given period. Given a history of actions $h = ((i_1, j_1), \dots, (i_n, j_n)) \in H_n^a$, this intertemporal influence can be measured as follows:

- The actions taken in stage n of the game only have an impact from stage $n+1$ onward. Hence, they have no influence on the game of length n . In other words, the “influence potential” in period n is $y_n = 0 \in \mathbb{R}^K$.

²This fact was already noticed in a weaker form by Amir (2003) in deterministic stochastic games with complete information.

- The actions taken in stage $n - 1$ only have an impact on opinions in period n (there is no forward diffusion). Hence, the potential influence exerted through each target (not accounting for influenceability) is identical. Thus, the “influence potential” is distributed proportionally to $y_{n-1} = \frac{1}{K}e'$. Accounting for the influenceability of the target, the actual impact on opinions of the actions of the players is given by $y_{n-1}(A^1(i_{n-1}, j_{n-1}) - A^{-1}(i_{n-1}, j_{n-1}))$.
- The actions taken by the players in stage $n - 2$ of the game have an influence in periods $n - 1$ and n . The potential influence exerted directly on each of the non-strategic agents in period $n - 1$ is identical, and is thus distributed according to $\frac{1}{K}e'$. The potential influence exerted indirectly in period n depends on the forward diffusion of the target’s opinion in the network and also the forward diffusion of the influence potential of agents in the next period. It is given by $y_{n-1}B(i_{n-1}, j_{n-1})$. Overall, the influence potential in stage $n - 2$ is distributed according to $y_{n-2} = y_{n-1}B(i_{n-1}, j_{n-1}) + \frac{1}{K}e'$, and the total intertemporal impact on opinions of the actions of players in period $n - 2$ is given by $y_{n-2}(A^1(i_{n-2}, j_{n-2}) - A^{-1}(i_{n-2}, j_{n-2}))$.
- More generally, the influence potential can be defined recursively as the sum of the instantaneous and forward influences. In particular, the distribution of influence potential in stage t is given by

$$y_t = y_{t+1}B(i_{t+1}, j_{t+1}) + \frac{1}{K}e'. \quad (4.2)$$

Accordingly, the intertemporal impact on opinions of the actions of players in period t is given by $y_t(A^1(i_t, j_t) - A^{-1}(i_t, j_t))$.

- Overall, the intertemporal impact on opinions of a history of actions $h = ((i_1, j_1), \dots, (i_n, j_n))$ is given by

$$\overleftarrow{\gamma}_n(h) = \sum_{t=1}^{n-1} y_t(A^1(i_t, j_t) - A^{-1}(i_t, j_t)). \quad (4.3)$$

The intertemporal impact on opinions defined above does coincide with the mean belief across time and across agents, i.e., with the payoff. Namely, one has:

Proposition 4.7 *Assume $x_1 = 0$. Let $n \geq 1$ and $h \in H_n^a$ be a history of actions. Then,*

$$\gamma_n(h) = \overleftarrow{\gamma}_n(h).$$

Remark 4.8 *The case of arbitrary initial opinions x_1 can be accounted for by altering the definition of the influence potential y_n in period n .*

One can then build on Proposition 4.7 to give a characterization of pure equilibrium strategies in the blind game with equal influence, provided that such strategies exist. Indeed, in the case of equal influence, the game is zero-sum and symmetric, so its value is zero. In this setting, at an equilibrium in pure strategy, a player ought to target, in each period, the node where his intertemporal impact is maximal. Otherwise, his opponent would gain a strictly positive profit by targeting the maximal impact node when the player doesn’t, and mimicking his strategy otherwise. In particular, one has:

Proposition 4.9 *Assuming that a blind game Γ_n^b admits a pure Nash equilibrium (σ, τ) , then for every $t \in \{1, \dots, n - 1\}$,*

$$\sigma_n \in \left\{ k \mid \forall j \in \mathcal{K}, \frac{y_{n,k}}{\lambda + \nu_k} \geq \frac{y_{n,j}}{\lambda + \nu_j} \right\},$$

and

$$\tau_n \in \left\{ k \mid \forall j \in \mathcal{K}, \frac{y_{n,k}}{\lambda + \nu_k} \geq \frac{y_{n,j}}{\lambda + \nu_j} \right\}.$$

Hence, in an equilibrium in pure strategy of the blind game with equal influence, a player ought to adopt a purely forward-looking strategy: in each period he ought to maximize the forward propagation of his influence in the network. Proposition 4.9 also shows why there might not exist equilibria in pure strategies in the blind game. The identity of the node in which the influence potential is maximum depends, in general, on the history of play, whereas in the blind game a player commits ex-ante to a fixed sequence of targets when he chooses a pure strategy.

Nevertheless, when one can decouple the identification of the maximal influence node and the history of play, necessary conditions for the existence of a blind equilibrium in pure strategies can be put forward. One has the following proposition:

Proposition 4.10 *Assume that there exists a sequence of nodes $(\bar{i}_1, \dots, \bar{i}_n) \in \mathcal{K}^n$ such that, for all $t = 1, \dots, n-1$ and all histories $h \in H_\infty^a$, one has*

$$\bar{i}_t = \operatorname{argmax}_{k \in \mathcal{K}} \frac{y_{t,k}}{\lambda + \nu_k}. \quad (4.4)$$

Then, $(\bar{i}_1, \dots, \bar{i}_{n-1})$ is an equilibrium strategy for both players in the blind game of length n with equal influence λ .

Remark 4.11 *If there exists ν such that, for all $k \in \mathcal{K}$, $\nu_k = \nu$, Equation 4.4 can be replaced by*

$$\bar{i}_t = \operatorname{argmax}_{k \in \mathcal{K}} y_{t,k}. \quad (4.5)$$

To ease the interpretation of our results, we shall focus on this particular case in the examples given in the remainder of this section.

Proposition 4.10 can then be used to characterize some networks for which there exist blind equilibria in pure strategies. To this end, one has to relate the influence potential y_t defined above to the network characteristics. With this in mind, we introduce the following notions.

Definition 4.12 *Let a network with adjacency matrix M be given.*

- *A path of length $n \in \mathbb{N}$ is a sequence of n nodes $p = (i_1, \dots, i_n) \in \mathcal{K}^n$ such that, for all $t < n$, $m_{i_t, i_{t+1}} > 0$. The weight of such a path is defined, with a slight abuse of notation, as $m_p = \prod_{t=1}^{n-1} m_{(i_t, i_{t+1})}$. We denote by $\mathcal{P}^n(M)$ the set of paths of length n in the network M , and by $\mathcal{P}_k^n(M)$ the set of paths of length n whose end node is k .*
- *A path constraint on $\mathcal{P}^n(M)$ is an n dimensional vector $C = (C_1, \dots, C_n) \in (2^{\mathcal{K}})^n$ inducing constraints on admissible paths. **Informally, we restrict to paths that are not going through C_t in stage t .** Formally, a path (i_1, \dots, i_n) satisfies the constraint $C \in (2^{\mathcal{K}})^n$ if, for all $t < n$, $i_t \notin C_t$ (if $C_t = \emptyset$, there is actually no constraint on i_t). We denote by $\mathcal{P}_n^C(M)$ the set of paths of length n in the network M satisfying the constraint C , i.e., $\mathcal{P}_n^C(M) = \{(i_1, \dots, i_n) \in \mathcal{P}^n(M) \mid \forall t \leq n, i_t \notin C_t\}$. In the following, we focus on path constraints $C = (C_1, \dots, C_n)$ such that, for all $t \leq n$, C_t has at most 2 elements. Accordingly we denote by \mathcal{C} the set of subsets of \mathcal{K} with at most 2 elements.*
- *The n -centrality of a node $k \in \mathcal{K}$, denoted by $\zeta_k^n(M) \in \mathbb{R}_+$, is the sum of the weights of those paths of length less than $n \in \mathbb{N}$ leading to k , i.e., $\zeta_k^n(M) = \sum_{t=2}^n \sum_{p \in \mathcal{P}_k^t(M)} m_p$. In vector form, one has $\zeta^n(M) = e' \sum_{t=1}^{n-1} M^t$.*
- *Given a constraint $C \in \mathcal{C}^n$, the C -constrained n -centrality of a node $i \in \mathcal{K}$, denoted by $\zeta_i^{n,C}(M) \in \mathbb{R}_+$, is the sum of the weights of those paths of length less than $n \in \mathbb{N}$ leading to k and satisfying the constraint C , i.e., $\zeta_k^{n,C}(M) = e' \sum_{t=2}^n \sum_{p \in \mathcal{P}_k^{t,C}(M)} m_p$ (where the constraints are truncated according to path length). In vector form, one has $\zeta^{n,C}(M) = e' \sum_{t=1}^{n-1} \prod_{s=1}^t M_{-C_s}$, where, given $C_s \in \mathcal{C}$, M_{-C_s} denotes the matrix obtained from M by substituting 0 for those rows whose indices are in C_s .*

A notable class of networks for which there exist blind equilibrium in pure strategies are those where the sequence of nodes with maximal t -centrality is not altered if the network is perturbed by a C -constraint, i.e., if at most two nodes are discarded every period. We introduce the following assumption:

Assumption 4.13 *The game Γ_n is such that:*

- *The strategic players have equal influence, i.e., $\lambda = \mu$.*
- *For all $t < n$, there exists $i_{n-t}^* \in \mathcal{K}$ such that, for all $C \in \mathcal{C}^{n-t-1}$,*

$$i_{n-t}^* = \operatorname{argmax}_{k \in \mathcal{K}} \frac{\zeta_k^{n,c}(M)}{\lambda + \nu_k}.$$

Remark 4.14 *Similarly to Remark 4.11, in the case where there exists $\nu > 0$ such that $\nu_k = \nu$ for all $k \in \mathcal{K}$, the condition in the second part of Assumption 4.13 reduces to $i_{n-t}^* = \operatorname{argmax}_{k \in \mathcal{K}} \zeta_k^{n,c}(M)$ independently of $\lambda \in \mathbb{R}_+$.*

One then has:

Proposition 4.15 *Suppose Assumption 4.13 holds in the game Γ_n . Then there exists a blind equilibrium in the game Γ_n in which the equilibrium strategy of both players is given by the sequence of actions (i_1^*, \dots, i_n^*) satisfying Assumption 4.13 (and i_n^* is chosen arbitrarily).*

Hence, under Assumption 4.13, in period t , a player ought to target the node that has maximal $n-t-1$ centrality, i.e., the node for which the sum of paths of length less than $n-t-1$ is maximal. The player discards previous actions and current beliefs and aims at maximizing his influence over the course of the remaining periods. The earlier in the game it is, i.e., the smaller t is, the longer the horizon of the player, i.e., the longer the paths of the network that he takes into consideration.

In order to satisfy Assumption 4.13 and to admit a blind equilibrium in pure strategies, a network must channel the flow of influence towards specific nodes, even when the structure is perturbed by the deletion of certain nodes. Examples 4.16 and 4.18 put forward two classes of networks which satisfy the assumption in the benchmark case where non-strategic agents are equally influenceable (i.e., there exists $\nu > 0$ such that, for all $k \in \mathcal{K}$, $\nu_k = \nu$). A basic example is the star network (Example 4.16) in which a blind equilibrium strategy is to target the center of the star. More broadly tree-like structures are natural candidates for Assumption 4.13, as they strongly constrain the flow of influence. Example 4.18 illustrates this point by defining a family of trees satisfying the assumption.

Example 4.16 *The (connected) star is a network with K nodes with an adjacency matrix such that $M_1 = \frac{1}{K}e'$ and, for $i > 1$, $M_i = e_1$. The node 1 is then the center of the star. As soon as $K \geq 3$, it is straightforward to check that, for $t \geq 2$, the center of the star has maximal C -constrained t -centrality, independently of the constraint C .*

Remark 4.17 *Example 4.16 mainly emphasizes the role of the topology of the network in the determination of optimal targets. The weights of links can also play a fundamental role. To illustrate this point, consider a star with two centers with unequal incoming weights. More precisely, consider the network with $K+1$ nodes and adjacency matrix M such that $M_1 = M_2 = \frac{1}{K}e'$ (nodes 1 and 2 are the two centers of the star) and, for all $i > 2$, $M_i = (1-\epsilon)e_1 + \epsilon e_2$. It is straightforward to check that, for ϵ small enough, node 1 has maximal C -constrained t -centrality, independently of the constraint C and for all $t \geq 2$. Thus, although nodes 1 and 2 are perfectly symmetric from a topological point of view, node 1 is the only relevant target in the game because of its quantitative characteristics.*

Example 4.18 We consider rooted trees in which the root node is denoted by 1, the descendants of each node are indexed by natural numbers, and an arbitrary node is characterized by the sequence of its parents, i.e., $(1, u_1, u_2, \dots, u_t) \in \mathbb{N}^t$ denotes the node at distance t from the root which is the u_t th descendant of the node $(1, u_1, \dots, u_{t-1})$ (and so on, recursively). We then consider the tree F_T of height T , defined recursively as follows:

- The root has n_0 descendants.
- Each of the descendants of the root has $n_1 \geq n_0$ descendants, except $(1, 1)$, which has $n'_1 \geq n_1 + 3$ descendants.
- For each $t \in \mathbb{N}$, each node at a distance t from the root has $n_t \geq n_{t-1}$ descendants, except $(1, 1, \dots, 1)$, which has $n'_t \geq n_t + 3$ descendants.

As above, we denote by \mathcal{K} the set of nodes of the network. The adjacency matrix M of the tree is such that $m_{i,j} = 1$ if i is the descendant of j and $m_{i,j} = 0$ otherwise. In order to turn F_T into a connected network, we add to each node of the network (including the root) one link towards the root, i.e., we set $m_{1,k} = 1/|\mathcal{K}|$, for all $k \in \mathcal{K}$. The resulting network, denoted by \overline{F}_T , has a primitive adjacency matrix.

It is then straightforward to check that, for any constraint $C \in \mathcal{C}^T$, the node $i_1^* := (1, \dots, 1) \in \mathbb{N}^T$ has a strictly greater C -constrained 1-centrality than every other node. Indeed, its C -constrained 1-centrality is at least $n_T + 1$, while any other node has a C -constrained 1-centrality of at most n_T . Similarly, the node $i_2^* := (1, \dots, 1) \in \mathbb{N}^{T-1}$ has a C -constrained 2-centrality strictly greater than $n_T n_{T-1}$, while that of every other node is at most $n_T n_{T-1}$. More generally, one can show by recursion that, for any constraint $C \in \mathcal{C}^T$, the node $i_t^* := (1, \dots, 1) \in \mathbb{N}^{T+1-t}$ has a strictly greater C -constrained t -centrality than every other node.

According to Proposition 4.15, this implies that, for every $t \leq T$, the sequence of nodes (i_t^*, \dots, i_1^*) induces a blind equilibrium in the game Γ_t associated with the network \overline{F}_T , when strategic players have equal influence.

Any subset of nodes, except 1 and the descendants of $(1, 1)$, can be deleted from the network without affecting this result. We thereby define a family of trees with a “dominant branch” in which there exists a blind equilibrium.

4.3 Characterization in the case of small influence

Finally, we focus on the limit case where the influence exerted by the strategic players is vanishingly small. In this setting, the dynamics without external influence provides a useful benchmark: one knows that the vector of opinions converges to a consensus in which the influence of each initial opinion is proportional to the eigenvector centrality of the corresponding node (see e.g. Jackson, 2008; Golub and Jackson, 2010). If strategic players have a small influence, they shall take as given the relative importance of non-strategic agents and thus choose their targets according to their eigenvector centrality. We show below that this is indeed the case and, more broadly, provide an equivalence between dynamic games with small influence and a class of static games.

Fix a game Γ . We first define, for every positive measure y over \mathcal{K} , a one-shot game G_y where Player 1 chooses a probability distribution a over \mathcal{K} and player -1 chooses a probability distribution b over \mathcal{K} . The definition is in two steps. First, for every pair of non-strategic agents $i, j \in \mathcal{K}$, we define

$$f_y^1(i, j) = \begin{cases} y(i) \frac{\lambda}{\lambda + \nu_i} & \text{if } i \neq j, \\ y(i) \frac{\lambda}{\lambda + \mu + \nu_i} & \text{if } i = j. \end{cases}$$

and

$$f_y^{-1}(i, j) = \begin{cases} y(j) \frac{\mu}{\mu + \nu_j} & \text{if } i \neq j, \\ y(j) \frac{\mu}{\lambda + \mu + \nu_j} & \text{if } i = j. \end{cases}$$

Second, given $a \in A = \Delta(\mathcal{K})$ and $b \in B = \Delta(\mathcal{K})$, we define the payoff of Player 1 by

$$g_y^1(a, b) = \frac{\sum_{i, j \in \mathcal{K}} a(i)b(j)f_y^1(i, j)}{\sum_{i, j \in \mathcal{K}} a(i)b(j)(f_y^{-1}(i, j) + f_y^1(i, j))}$$

and the payoff of Player -1 in the opposite way by

$$g_y^{-1}(a, b) = 1 - g_y^1(a, b)$$

In line with the preceding section, y can be interpreted as a measure of the influence potential associated with each node. Accordingly, in the game G_y , players aim to maximize their expected relative influence given the influence potential y .

Remark 4.19 *Following Acemoglu et al. (2013), the influence potential y_n and the overall intertemporal impact $\overleftarrow{\gamma}_n$ defined in Section 4.2 can be interpreted in the context of an n -stage ball-catching game in which players try to catch balls that move backwards in time in the network. More precisely, we consider that a new ball is added to the network uniformly and at random in every period, and that, given a pair of actions (i, j) , Player 1 captures a ball present in node i with probability $A^1(i, j)$ and player -1 captures a ball present in node j with probability $A^{-1}(i, j)$. Conditionally on the pair of actions, the capture of different balls is supposed to be independent. Then, considering that time flows backwards from period n to period 1, given a sequence of actions $h = ((i_1, j_1), \dots, (i_n, j_n)) \in H_n^a$, one has:*

- *The initial distribution of balls in the network is given by $y_n = 0$.*
- *For every $t < n$, the expected number of balls in period t is given by $y_t = y_{t+1}B(i_{t+1}, j_{t+1}) + \frac{1}{K}e'$.*
- *The expected numbers of balls caught by Player 1 and Player -1 in period t are given by $y_t A^1(i_t, j_t)$ and $y_t A^{-1}(i_t, j_t)$, respectively.*

In this setting, $\overleftarrow{\gamma}_n(h)$, corresponds to the difference between the numbers of balls caught by Player 1 and player -1 across time.

The game G_y can then be interpreted with this terminology. Assume that a large number of balls are distributed in the network along the measure y . Then, given that Player 1 targets i and Player -1 targets $j \in \mathcal{K}$, $f_y^1(i, j)$ and $f_y^{-1}(i, j)$ are respectively the expected number of balls caught by Player 1 and by Player -1. Then $g_y^1(a, b)$ is the relative proportion of balls caught in one stage by Player 1 and $1 - g_y^1(a, b)$ is the relative proportion of balls caught by Player -1. Notice that this is not the linear extension of the proportion of balls obtained through pure actions.

Using Sion's minimax theorem Sion et al. (1958), one can show that the game G_y has a value.

Proposition 4.20 *For every positive measure $y \in \mathbb{R}_+^K$ over \mathcal{K} , the game G_y admits a value w_y^* :*

$$w_y^* = \max_{a \in A} \min_{b \in B} g_y^1(a, b) = \min_{b \in B} \max_{a \in A} g_y^1(a, b).$$

We denote by a_y^ an optimal strategy of Player 1 and by b_y^* an optimal strategy of Player -1 in this one-shot game.*

In line with the dynamic scenario without external influence, we shall focus on games in which the influence potential is proportional to the eigenvector centrality. We denote the eigenvector centrality of M by $\chi \in \Delta(\mathcal{K})$, defined as the unique invariant probability distribution associated with M , i.e., one has $\chi M = \chi$ and $\sum_{k \in \mathcal{K}} \chi(k) = 1$.

One can then show that the equilibrium of the auxiliary game G_χ induces an equilibrium in the infinite influence game Γ . In particular, if a player plays an optimal strategy in G_χ at every stage, then this strategy is ε -optimal in the infinite game. More formally, we obtain the following result.

Theorem 4.21 *Assume that for every $k \in \mathcal{K}$, $\nu_k > 0$. Let σ^* (resp. τ^*) be the strategy of Player 1 (resp. Player -1) that plays a_χ^* (resp. b_χ^*) at every stage in the opinion game Γ . There exists $\lambda_0 \in \mathbb{R}_+$ such that, for every $\lambda, \mu \leq \lambda_0$, (σ^*, τ^*) is ε -optimal in the infinite game $\Gamma(x_1)$, for every $x_1 \in X$ and every information structure. Moreover, one has:*

$$\forall x_1 \in X, \lim_{n \rightarrow +\infty} v_n(x_1) = 2w_\chi^* - 1.$$

Note that, although Theorem 4.21 requires the influence to be arbitrarily small, strategic players nevertheless have a real impact on the outcome because the game is arbitrarily long. It is also worth noting that Theorem 4.21 implies that Γ admits an equilibrium in pure strategies whenever the auxiliary game G_χ admits an equilibrium in pure strategy. As emphasized above, this is notably the case when players have equal influence and there is a unique node with maximum eigenvector centrality. One thus has the following corollary:

Corollary 4.22 *Assume that $\lambda = \mu$ and that there exists $\nu > 0$ such that, for every $k \in \mathcal{K}$, $\nu_k = \nu$. Further, assume that there exists a unique non-strategic agent $k^* \in \mathcal{K}$ such that*

$$\chi(k^*) = \max_{k \in \mathcal{K}} \chi(k).$$

Then there exists $\lambda_0 \in \mathbb{R}_+$ such that, for every $\lambda \leq \lambda_0$, it is an ε -optimal strategy in the infinite game $\Gamma(x_1)$ to target the node k^ for every $x_1 \in X$ and every information structure.*

Note that such an ‘‘eigenvector targeting strategy’’ is another instance of a purely forward-looking strategy: it is optimal independently of the actions of the other player. Strategic agents focus on the diffusion of their own opinion. This also provides evidence for the importance of eigenvector centrality as a measure of influence.

In Corollary 4.22, we focus on arbitrarily long games. One can, however, build on Proposition 4.10 to characterize equilibrium in pure strategies in ‘‘short’’ games, i.e., of a given fixed length $N \in \mathbb{N}$, with small influence. Indeed, if influence is sufficiently small and the time-horizon is bounded, strategic actions do not modify the ordering of nodes in terms of t -centrality. In this context, let us assume that there exists a unique sequence of nodes that maximize t -centrality in the following sense.

Assumption 4.23 *For all $1 \leq t \leq n$, there exists a unique $\bar{i}_{n-t} \in \mathcal{K}$ such that*

$$\bar{i}_{n-t} = \operatorname{argmax}_{k \in \mathcal{K}} c_k^{n-t-1}(M).$$

For a fixed n , this assumption holds generically, i.e., up to an arbitrarily small modification of the network weights. It suffices to guarantee the existence of a unique equilibrium in pure strategies in the game of length n with equal and small influence.

Proposition 4.24 *Suppose that Assumption 4.23 is satisfied, $x_1 = 0$, $\lambda = \mu$, and there exists $\nu > 0$ such that, for every $k \in \mathcal{K}$, $\nu_k = \nu$. Then, there exists $\lambda_0 \in \mathbb{R}_+$ such that, for every $\lambda < \lambda_0$, the sequence $(\bar{i}_1, \dots, \bar{i}_n) \in \mathcal{K}^n$ of nodes with maximal t -centrality is the only blind equilibrium of the game Γ_n (up to an arbitrary choice of \bar{i}_n).*

5 Comparison between static and dynamic strategies

The influence game considered in this paper is a dynamic extension of Grabisch et al. (2017). We consider the same model for the diffusion of opinions, but we allow strategic agents to update their targets dynamically in the network, whereas Grabisch et al. (2017) treat the case where, once chosen, the target agent is fixed. The model of Grabisch et al. (2017) can be emulated in our framework by considering static strategies in the following sense.

Definition 5.1 We say that a pure strategy τ of Player -1 is static if there exists $j^* \in \mathcal{K}$ such that, for every history $h \in \cup_{n=1}^{+\infty} H_n^{-1}$, $\tau(h) = j^*$. A strategy is static if it is a mixture of pure static strategies. *Informally, Player -1 chooses a target once and for all at the beginning of the game.*

This allows one to compare the efficiency of dynamic and static strategies. In particular, let us note that, in the class of regular networks considered in Section 4.1, there is no equilibrium in pure strategies. In fact, one has:

Proposition 5.2 Suppose Assumption 4.1 holds with $K \geq 3$, $\alpha > 0$, ϕ is such that the orbit of every element is at least 3 (i.e., for all $k \in \mathcal{K}$, $\phi(k) \neq k$ and $\phi^2(k) \neq k$), and $n \geq 3$. If the players are constrained to use a static strategy, there does not exist an equilibrium in pure strategies in the game Γ_n .

The non-existence of pure equilibria with static strategies was already noted by Grabisch et al. (2017) in the case of the circular network. The comparison between this result and Theorem 4.2 shows that there might exist equilibria in pure strategies for networks for which there is no equilibrium in pure strategies in a static setting. Updating targeting strategies through time is, in a sense, akin to mixing strategies in the static game. It provides the players with sufficient flexibility to counteract the actions of their opponents.

Another corollary of this result is that, in regular networks in the sense of Assumption 4.1, a player adopting a dynamic strategy necessarily gains an edge against an opponent using a static strategy, i.e., if Player 1 is “dynamic” while Player -1 is “static”, Player 1 can guarantee a strictly positive value (assuming both players have equal influence). Indeed, from periods 2 onward, Player 1 knows the action played by Player -1 for the rest of the game and can thus gain an advantage by best-replying as in the proof of Proposition 5.2.

n. nodes proba	0.1	0.25	0.5	n. nodes / proba	0.1	0.25	0.5	n. nodes / proba	0.1	0.25	0.5
10	18	10	11	10	12	10	9	10	13	11	7
25	14	12	10	25	10	4	5	25	3	1	0
50	10	2	4	50	2	0	0	50	1	0	0

Table 1: Number of runs (out of 20 Monte-Carlo simulations for each parameter combination) in which the genetic algorithm identifies a dynamic strategy that dominates any static strategy in games of length 10 (left), 25 (center), and 50 (right).

In order to investigate the relative performance of dynamic and static targeting strategies in a broader range of networks, we ran a series of numerical simulations. In particular, we have drawn a sample of Erdos–Renyi random graphs with linkage probability varying in $\{0.1, 0.25, 0.5\}$ and number of nodes varying in $\{10, 25, 50\}$. We investigated whether dynamic strategies perform better than static strategies in games of length 10, 25, and 50 using the following algorithm. We applied a genetic algorithm³ to see whether there was a sequence of actions of Player 1 that would guarantee a payoff strictly greater than 0 against every fixed action of player -1 . If such a sequence of actions exists, it implies that Player 1 is strictly better off using the corresponding dynamic strategy than any static strategy, i.e., there is indeed an advantage in using dynamic strategies. Table 5 gives the results of these simulations. It shows first that, in a large proportion of networks, dynamic strategies do improve upon static ones. Second, it provides insights into the influence of network structure on the relative performance of dynamic and static targeting strategies. In this respect, the share of runs in which the genetic algorithm identified an improvement upon static strategies can be interpreted as a lower bound on the number of networks for which such an improvement exists, or a measure of the computational complexity of finding such an improvement⁴. From this perspective, our results suggest that the relative performance of dynamic

³The default genetic algorithm implemented in matlab.

⁴The fact that the genetic algorithm did not always find an improvement upon static strategies does not mean that such an improvement did not exist, but rather that the algorithm hadn’t been able to identify one. Only a comprehensive exploration of the set of all dynamic strategies could allow one to make a categorical statement about non-existence, but this would rapidly lead to combinatorial explosion.

strategies decreases with the size and density of the network, as well as with the length of the game. This might be due to the increasing complexity of finding such strategies or to the emergence of more central nodes, which ought to be the sole target of players, as the size and density of the network grows.

6 Conclusion

We consider the problem of maximizing influence through the choice of a sequence of targets when two players compete for predominance over a social network. We show that the problem can be modelled by a stochastic game with compact state space. This raises the question of the existence of a uniform value, which we answer affirmatively under the assumption that the adjacency matrix of the network is primitive. Hence, strategic players can design optimal targeting strategies independently of the length of the game.

Further analysis of the structure of equilibrium strategies shows that they depend strongly on the level of information available to each player about the actions of his opponents, and hence about the state of opinions in the network. In the absence of information, there exist pure equilibria only if players can devise strategies that are efficient independently of the current state of the network, i.e., they must adopt a purely forward-looking perspective on the diffusion of their influence. When players are well informed, they can implement more flexible strategies that take into account the state of opinions in the network and implement a trade-off between a forward-looking perspective, according to which they try to maximize the future spread of their opinion in the network, and a backward-looking perspective, according to which they try to counteract their opponent's previous actions. This flexibility sustains the existence of equilibria in pure strategies in networks where there is no such equilibrium with low levels of information or in the static game à la Grabisch et al. (2017) with fixed targets.

These results emphasize the relevance of adopting a dynamic approach, in particular in view of empirical applications: sequential strategies allow one to solve the targeting problem in a wide range of networks and independently of the length of the game. It is likely impossible to obtain a complete analytical characterization of these strategies, but it seems to us that an algorithmic approach could be fruitfully applied to the problem in future work. The model could also be extended, analytically or numerically, along with a number of dimensions, such as number of strategic players, number of targets, or cost of targeting.

Acknowledgement: Mandel acknowledges the support of the EU H2020 FET project DOLFINS (grant nb. 640772). Venel acknowledges the support of the Agence National de la Recherche [ANR CIGNE ANR-15-CE38-0007-01].

References

- D. Acemoglu, G. Como, F. Fagnani, and A. Ozdaglar. Opinion fluctuations and disagreement in social networks. *Mathematics of Operations Research*, 38(1):1–27, 2013.
- R. Amir. Stochastic games in economics and related fields: an overview. In *Stochastic Games and Applications*, pages 455–470. Springer, 2003.
- P. Amorós and M. S. Puy. Issue convergence or issue divergence in a political campaign? *Public Choice*, 155(3-4):355–371, 2013.
- E. Aragones, M. Castanheira, and M. Giani. Electoral competition through issue selection. *American journal of political science*, 59(1):71–90, 2015.
- R. J. Aumann, M. Maschler, and R. E. Stearns. *Repeated games with incomplete information*. MIT press, 1995.
- C. Ballester, A. Calvo-Armengol, and Y. Zenou. Who’s who in networks. Wanted: The key player. *Econometrica*, 74(5):1403–1417, 2006.
- A. Banerjee, A. Chandrasekhar, E. Duflo, and M. Jackson. Diffusion of microfinance. *Science*, 341, DOI: 10.1126/science.1236498, 2013.
- F. Beard. Comparative advertising wars: An historical analysis of their causes and consequences. *Journal of Macromarketing*, 30(3):270–286, 2010.
- S. Bharathi, D. Kempe, and M. Salek. Competitive influence maximization in social networks. In *International Workshop on Web and Internet Economics*, pages 306–311. Springer, 2007.
- K. Bimpikis, A. Ozdaglar, and E. Yildiz. Competitive targeted advertising over networks. *Operations Research*, 64(3):705–720, 2016.
- P. Bonacich. Factoring and weighting approaches to status scores and clique identification. *Journal of Mathematical Sociology*, 2(1):113–120, 1972.
- J. M. Box-Steffensmeier and T.-M. Lin. A dynamic model of campaign spending in congressional elections. *Political Analysis*, pages 37–66, 1996.
- T. Carnes, C. Nagarajan, S. M. Wild, and A. Van Zuylen. Maximizing influence in a competitive social network: a follower’s perspective. In *Proceedings of the ninth international conference on Electronic commerce*, pages 351–360. ACM, 2007.
- S. Chakrabarti and H. Haller. An analysis of advertising wars. *The Manchester School*, 79(1):100–124, 2011.
- D. L. Cogburn and F. K. Espinoza-Vasquez. From networked nominee to networked nation: Examining the impact of web 2.0 and social media on political participation and civic engagement in the 2008 obama campaign. *Journal of Political Marketing*, 10(1-2):189–213, 2011.
- M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- G. Demange and K. Van der Straeten. A communication game on electoral platforms. 2009.
- P. Domingos and M. Richardson. Mining the network value of customers. *Seventh International Conference on Knowledge Discovery and Data Mining*, pages 57–66, 2001.
- J. L. Doob. *Stochastic processes*, volume 7. Wiley New York, 1953.

- U. Doraszelski and S. Markovich. Advertising dynamics and competitive advantage. *The RAND Journal of Economics*, 38(3):557–592, 2007.
- P. K. Dubey, R. Garg, and B. De Meyer. Competing for customers in a social network. *Journal of Dynamics and Games*, 1(3):377–409, 2014.
- L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, pages 35–41, 1977.
- D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.
- B. Golub and M. O. Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.
- S. Goyal, H. Heidari, and M. Kearns. Competitive contagion in networks. *Games and Economic Behavior*, 2014.
- M. Grabisch, A. Mandel, A. Rusinowska, and E. Tanimura. Strategic influence in social networks. *Mathematics of Operations Research*, forthcoming, 2017.
- J. Granato and M. S. Wong. Political campaign advertising dynamics. *Political Research Quarterly*, 57(3):349–361, 2004.
- J. Hartline, V. Mirrokni, and M. Sundararajan. Optimal marketing strategies over social networks. In *Proceedings of the 17th international conference on World Wide Web*, pages 189–198. ACM, 2008.
- X. He and D. Kempe. Price of anarchy for the n-player competitive cascade game with submodular activation functions. In *International Conference on Web and Internet Economics*, pages 232–248. Springer, 2013.
- M. O. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.
- A. Jaśkiewicz and A. S. Nowak. Zero-sum stochastic games. *Handbook of Dynamic Game Theory*, pages 1–64, 2016.
- D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. *Proceedings of 9th International Conference on Knowledge Discovery and Data Mining*, pages 137–146, 2003.
- D. Kempe, J. Kleinberg, and E. Tardos. Influential nodes in a diffusion model for social networks. *Proceedings of 32nd International Colloquium on Automata, Languages and Programming*, pages 1127–1138, 2005.
- A. J. Lee, F.-C. Yang, H.-C. Tsai, and Y.-Y. Lai. Discovering content-based behavioral roles in social networks. *Decision Support Systems*, 59:250–261, 2014.
- C. R. Lever. *Strategic Competitions over Networks*. Ph.D. dissertation, Stanford University, 2010.
- S. Mekouar, S. El-Hammani, K. Ibrahimi, and E.-H. Bouyahf. Optimizing diffusion time of the content through the social networks: Stochastic learning game. In *International Conference on Networked Systems*, pages 367–381. Springer, 2015.
- S. Mekouar, E.-H. Bouyahf, S. El-Hammani, and K. Ibrahimi. Access and sharing contents through the social network: A pomdp approach. *Computational Intelligence*, 33(4):680–705, 2017.
- J.-F. Mertens and A. Neyman. Stochastic games. *International Journal of Game Theory*, 10(2):53–66, 1981.

- A. Nedic, A. Ozdaglar, and P. A. Parrilo. Constrained consensus and optimization in multi-agent networks. *IEEE Transactions on Automatic Control*, 55(4):922–938, 2010.
- I. Pastine and T. Pastine. Consumption externalities, coordination, and advertising. *International Economic Review*, 43(3):919–943, 2002.
- V. Raghavan, G. Ver Steeg, A. Galstyan, and A. G. Tartakovsky. Coupled hidden markov models for user activity in social networks. In *2013 IEEE International Conference on Multimedia and Expo Workshops (ICMEW)*, pages 1–6. IEEE, 2013.
- J. Renault. The value of repeated games with an informed controller. *Mathematics of operations Research*, 37(1):154–179, 2012.
- M. Sion et al. On general minimax theorems. *Pacific Journal of mathematics*, 8(1):171–176, 1958.
- E. Solan. Stochastic games. In *Computational Complexity*, pages 3064–3074. Springer, 2012.
- F. Thuijsman. Optimality and equilibria in stochastic games. *CWI Tracts*, 82:1–107, 1992.
- J. von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1):295–320, 1928.
- A. Yadav, E. Kamar, B. Grosz, and M. Tambe. Healer: Pomdp planning for scheduling interventions among homeless youth. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1504–1506. International Foundation for Autonomous Agents and Multiagent Systems, 2016.
- E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione. Binary opinion dynamics with stubborn agents. *ACM Transactions on Economics and Computation (TEAC)*, 1(4):19, 2013.
- Y. Yu, T. Y. Berger-Wolf, J. Saia, et al. Finding spread blockers in dynamic networks. In *Advances in Social Network Mining and Analysis*, pages 55–76. Springer, 2010.
- H. Zhuang, Y. Sun, J. Tang, J. Zhang, and X. Sun. Influence maximization in dynamic social networks. In *Data Mining (ICDM), 2013 IEEE 13th International Conference on*, pages 1313–1318. IEEE, 2013.
- B. Ziliotto. Zero-sum repeated games: Counterexamples to the existence of the asymptotic value and the conjecture $\max \min = \lim v_n$. *The Annals of Probability*, 44(2):1107–1133, 2016.