

Competition of Influencers: A Model for Maximizing Online Social Impact

Original

Competition of Influencers: A Model for Maximizing Online Social Impact / Galante, Franco; Garetto, Michele; Leonardi, Emilio. - ELETTRONICO. - (2024), pp. 1-11. (Intervento presentato al convegno WEBSCI 2024 tenutosi a Stuttgart (D) nel May 2024).

Availability:

This version is available at: 11583/2985740 since: 2024-03-01T14:09:35Z

Publisher:

ACM

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

ACM postprint/Author's Accepted Manuscript, con Copyr. autore

(Article begins on next page)

Competition of Influencers: A Model for Maximizing Online Social Impact

Franco Galante
Politecnico di Torino
Torino, Italy
franco.galante@polito.it

Michele Garetto
Università degli Studi di Torino
Torino, Italy
michele.garetto@unito.it

Emilio Leonardi
Politecnico di Torino
Torino, Italy
emilio.leonardi@polito.it

ABSTRACT

The landscape of human interaction has undergone a profound transformation since the advent of Online Social Networks (OSNs). Not only are they changing interpersonal dynamics, but they are also redefining the way businesses, political figures, and media organizations engage with the broader population. In today's digital landscape, OSNs have spawned a new class of social media influencers who play a crucial role in shaping opinion. These influencers actively compete within social media to seize users' attention. Through these targeted efforts, influencers seek to captivate users and build a loyal and engaged fan base, solidifying their position as an authoritative voice in the digital world. In this work, we develop a game-theoretic model for the interactions between users and influencers, where the latter compete to maximize their impact on the population's opinions. The goal of this work is twofold: first, we formalize the problem of maximizing social media impact and study the structure of the optimal solution. Then, taking inspiration from the optimal strategy, we develop a game with two opposing players trying to maximize their influence on users' opinions, for which we characterize the Nash equilibria in pure strategy. The model allows us to evaluate the impact of influencer differences and user population characteristics. In addition, we study the effect of the speed at which user popularity evolves in such a competitive environment. The proposed model proves valuable for brand competition, marketing campaigns, and the ever-evolving political arena.

CCS CONCEPTS

• **Applied computing** → *Sociology*; • **Information systems** → *Social advertising*; **Social networking sites**; **Social networks**; • **Theory of computation** → **Social networks**; • **Computing methodologies** → *Simulation evaluation*.

KEYWORDS

Online social networks, game theory, competition, Nash equilibria, social impact maximization, opinion dynamics

ACM Reference Format:

Franco Galante, Michele Garetto, and Emilio Leonardi. 2024. Competition of Influencers: A Model for Maximizing Online Social Impact. In *ACM Web*

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Websci '24, May 21–24, 2024, Stuttgart, Germany

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0334-8/24/05.

<https://doi.org/10.1145/3614419.3644031>

Science Conference (Websci '24), May 21–24, 2024, Stuttgart, Germany. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3614419.3644031>

1 INTRODUCTION

The proliferation of online social networks (OSNs) has impacted various components of modern living such as the way we communicate, consume information, and navigate daily hurdles. It touches a variety of dimensions of everyday life and has not spared the way people form their opinions and make their consumption decisions. With billions of users around the world and an ever-expanding range of platforms and features, online social networks can produce alluring opportunities both for companies and individuals (e.g., influencers). A defining characteristic of online social networks (OSNs) is the remarkable asymmetry among users. On the one hand, there is a small number of highly visible and influential users, commonly known as *influencers* or *opinion leaders*. On the other hand, the vast majority of users are “*regular*” users and typically have a much more modest *following*. We consider an opinion model that explicitly separates these two classes of users [9] and considers the interactions between the two¹. Another distinctive feature is the content filtering performed by the platform on the posts suggested to the users.

In this paper, our goal is to represent the competition between different influencers over an OSN and investigate the best possible strategies they can adopt to maximize their own utility function by extending the opinion model for OSN proposed in [9]. As a first step in our investigation, we address the fundamental question of how an influencer, taken in isolation, should best proceed to increase her influence on a particular group of users. An influencer's strategy is a sequence of posts that convey an opinion, and that she publishes on her social media profile. The “best” strategy is the one that brings the greatest benefit, however, it is defined. Although an individual's opinion generally does not frequently fluctuate, an influencer's stance can be influenced by both external and personal factors. For example, it is common for influencers to retract certain positions due to the pressure of public opinion. Another example, influencers may change their collaborative partnerships and promote other brands' products. Similarly, politicians often adjust their positions based on the opinions of their electorate on certain issues. Therefore, assuming an influencer's opinion can span the entire opinion space, we show that the *greedy* strategy which maximizes influence at every post emission is not always *optimal*. Our experiments hint that it is optimal to group the user base, which has diverse initial views on the topic, around a common viewpoint, and then move the group towards the desired opinion.

¹In this work, to avoid the confusion which would arise using the gender-neutral “they”, we will, arbitrarily, use she/her for an *influencer*, and he/his for a *regular* user.

Following the analysis of this scenario, we then consider the competition of two influencers in a population. We aim to characterize the resulting strategic competition, which can be mathematically formulated as a zero-sum game. Our goal is to identify the Nash equilibria that represent the optimal outcome in the context of the game's rules and strategies. Even if considering only a duopolistic situation may appear restrictive, it encompasses many real-world situations, such as political opposition: democrats vs liberals, brand competition, or software rivalry, e.g. Linux vs Windows.

The article is structured as follows. Section 2 discusses relevant work in the area of opinion dynamics and strategic competition. Our novel *social impact maximization* problem is formalized in Section 3, along with the underlying model of opinion dynamics. Section 4 is concerned with the *optimal* solution and a comparison with the *greedy* one. Then, Section 5 models influencers' competition by defining a two-player game for which we study the Nash equilibria in pure strategy. We extend the model and consider arbitrary user distributions thanks to a discretization procedure in Section 6. Section 7 concludes the article. The source code for this work can be found on GitHub².

2 RELATED WORK

Our work can be considered at the intersection between the literature on opinion dynamics [3] and the more classic economic literature on competition [18]. Widely used models of opinion dynamics include linear models proposed by DeGroot [6] and Friedkin–Johnsen [8], and non-linear models by Deffuant–Weisbuch [5] and Hegselman–Krause [20]. We refer the reader to [19] for a recent review on different models of opinion propagation, and to [1] for results specific to non-linear (bounded confidence) dynamics.

Influence maximization is a different problem, slightly related to ours [13]. It aims at finding a (small) set of k agents so that to maximize the adoption of a certain product. This problem differs largely from ours in that it considers an explicit network structure and uses a simpler underlying (cascade) opinion model whose states are binary. Many extensions appeared, [4] considers a *negativity* bias, and [15] non only competitive behavior. In the following, we limit ourselves to mentioning works more similar in spirit to ours, i.e., those dealing with non-linear (e.g., bounded confidence) dynamics, and focused on opinion manipulation by a restricted set of strategic agents aiming at maximizing their impact on a population of users.

A bounded confidence model of opinion dynamics on a fixed network structure comprising both influencers and followers is proposed and analyzed by simulation in [2]. In [12] the impact of charismatic leaders is taken into account in bounded confidence dynamics as a constant exogenous signal. Interestingly, they discovered that higher signals may have less effect in attracting other agents. Opinion manipulation through (possibly time-varying) exogenous inputs is analyzed in [17] for an Eulerian (i.e., by considering a probability distribution of agents) bounded-confidence system. In [23], the authors consider a continuous-time bounded confidence model with a single leader, showing that it is possible to control the leader velocity to ensure final consensus at a prescribed opinion value.

We mention that optimal control approaches have also been proposed to optimize opinion manipulation. For example, [11] takes this direction considering the case of a single influencer. In [14], a set of (coordinated) controllers is optimally placed to minimize the convergence time of the system to a final stable state in the influence graph.

To the best of our knowledge, there are only a few game-theoretical attempts to study the problem of opinion formation/manipulation. Papers [7, 16, 21, 22] consider a scenario in which two competing marketers play a resource allocation game, whose goal is to establish how many resources to allocate to each potential customer in the network. Customers dynamics obey a voter model in [16], to a consensus model over a graph in [21, 22], and to a Friedkin–Johnsen model in [7]. A Nash control formulation is proposed in [10] for the case of two opposing influencers. Opinions spread on an underlying graph in an epidemic manner.

Differently from previous papers, we consider non-linear opinion dynamics where the interaction between influencers and users is mediated by a social media platform. Influencers operate on a fixed budget, and have to decide what opinion to express within each of N posts.

3 FORMULATION OF THE PROBLEM

We first discuss the underlying opinion model that describes how popular individuals (i.e., influencers) interact with the pool of *regular* users through an online platform. The opinion model is a simplified version of the model presented in [9]. We will use the full version of this model in the framework discussed in Section 6.1. Then, we introduce the *social impact maximization* problem for an influencer taken in isolation.

3.1 The opinion model

The foundational element at the core of our approach is the model that determines how social media users interact with each other. The social media audience can be divided into two macro-categories: the popular and influential, commonly referred to as *influencers*, and the “other” users, who represent the vast majority of OSN users, whom we refer to as *regular* users. We are interested in the dynamics between users and influencers and consider only this type of interaction. Indeed, we use a simplified version of the Communication Asymmetry opinion model [9], considering a one-dimensional opinion space $\mathcal{X} = [0, 1]$ and a simple characterization of the influencers.

Let us denote by \mathcal{I} (indexed by i) the set of influencers and by \mathcal{U} (indexed by u) the set of regular users. These social network users at a certain time instant $n \in \mathbb{N}_0^+$ hold opinions $x_n^{(i)} \in \mathcal{X}$ and $x_n^{(u)} \in \mathcal{X}$, respectively. We assume that at each discrete time instant n influencer $i \in \mathcal{I}$ publishes a post over the social media network conveying her *expressed* opinion. In this paper, as opposed to [9], we consider influencers with time-varying opinions, i.e., the influencers generate a sequence of posts $\{x_n^{(i)}\}$, $n \in [0, \dots, N]$, where N represents a considered (finite) time horizon. An influencer will adapt her expressed point of view strategically in order to maximize her impact on the population (see Section 3.2).

²Link: <https://github.com/Franco-Galante/onine-social-impact-maximization>

A given post is not guaranteed to reach a regular user $u \in \mathcal{U}$, due to the *filtering* effect of the social media platform. This represents the *content personalization* performed by most platforms (e.g., the ranking of posts to be shown on the user's timeline). We assume a post effectively reaches a regular user with a probability $\psi(|x^{(i)} - x^{(u)}|)$ which depends on the distance between the influencer's *expressed* opinion and the user's opinion, modeling a homophilic behavior, whereby individuals are more likely to interact with others who share similar beliefs.

A user holding opinion $x_n^{(u)}$ at instant n updates his opinion according to the following rule:

$$x_{n+1}^{(u)} = \begin{cases} \alpha z^{(u)} + \beta x_n^{(u)} + \gamma x_n^{(i)} & \text{if } \Psi(|x_n^{(u)} - x_n^{(i)}|) = 1 \\ \frac{\alpha}{\alpha+\beta} z^{(u)} + \frac{\beta}{\alpha+\beta} x_n^{(u)} & \text{otherwise} \end{cases} \quad (1)$$

where α, β, γ are fixed parameters in $[0, 1]$ such that $\alpha + \beta + \gamma = 1$. $\Psi(\psi)$ is a Bernoulli random variable with parameter ψ and determines whether a post will be visible to a given user or not. Note that when a user receives a new post from the influencer, i.e., when $\Psi = 1$ (first row of Eq. (1)), he moves to a new position in the opinion space, which is a convex combination of three contributions: i) his prejudice $z^{(u)}$, i.e., the preconceived opinion about a certain matter³; ii) his current opinion $x_n^{(u)}$; and iii) the influencer's opinion $x_n^{(i)}$ expressed in the post.

Otherwise (second row of Eq. (1)), the influencer's contribution is not present, either because the post has not been proposed to the user or he has not been influenced (e.g., not liked it), so a renormalization of the weights is required. This case in Eq. (1) models a process of *self-thinking*, namely that users who are not reached by an influencer's post gradually return to their prejudice $z^{(u)}$.

The above updating rule is simple but does not allow a direct interpretation of the parameters. Noting that the free parameters in the first line of Eq. (1) are only two, it is possible to rewrite the update rule:

$$x_{n+1}^{(u)} = \begin{cases} (1 - \beta) \left[\delta z^{(u)} + (1 - \delta) x_n^{(i)} \right] + \beta x_n^{(u)} & \text{if } \Psi(\psi) = 1 \\ \frac{\delta}{\delta + \frac{\beta}{1-\beta}} z^{(u)} + \frac{\frac{\beta}{1-\beta}}{\delta + \frac{\beta}{1-\beta}} x_n^{(u)} & \text{otherwise} \end{cases} \quad (2)$$

where $\delta, \beta \in [0, 1]$ have a direct interpretation as the *inertia* (β) of the user, i.e., the weight the users give to their current opinion, and the *degree of stubbornness* (δ), i.e., the weight on the user's preconceived opinion.

REMARK 1. (*Large population*) In the large population limit (i.e., when $|\mathcal{U}| \rightarrow \infty$) fluctuations of aggregate random variables around their average smooth out. Therefore macroscopic dynamics tend to become deterministic.

Thanks to the *large-population* assumption, we do not have to track the microscopic interactions described in Eq. (2) but we can consider a distribution of regular users characterized by the probability density function $\mu_n(x)$, whose evolution is driven by influencers' posts emission. In particular, every time influencer i generates a new post, a fraction $\psi(|x - x^{(i)}|)$ of the population

³This quantity could also be interpreted as the *field effect* of the total population, which would represent a process of peer interaction not otherwise included in the model.

placed at x will be hit by the influencer's messages while the remaining fraction of users will not be reached by it. This assumption greatly simplifies the analysis, in particular in Section 6.2 where we consider a more complicated function ψ , while it is not strictly necessary for the rest of the work (see Remark 2). In our framework, the assumption is not restrictive, as our focus lies in the mean-field effects observed across a large population of individuals.

3.2 Online social impact maximization for an influencer in isolation

We are interested in determining the influencer posting pattern that maximizes her online social impact on a population of regular users. Even considering the case where a single influencer seeks to maximize her impact over a finite time horizon, assuming there are no other influencers, is insightful and complicated enough to be worth exploring. We will then use the observations gathered in this simplified setting to develop our game of online competition.

Our novel *social impact maximization* problem, for the case of a single influencer, can be formulated as follows. Recall that we consider a fixed time horizon N , where the influencer has to choose the temporal sequence of opinions $\{x_n^{(i)}\}_{n=1}^N$ to convey through her posts in order to attract regular users towards a desired *target* opinion x^T in the opinion space. This value can represent, for example, the *true* opinion of the influencer regarding a certain topic or a certain consumption behavior to be installed in the population. Regular users obey the dynamics in Eq. (1).

We assume that the influencer knows how users would react to her posts, i.e., the parameters of Eq. (1), and in particular the shape of $\psi(\cdot)$ as a function of the opinion distance $d \triangleq |x_n^{(u)} - x_n^{(i)}|$, which dictates whether her posts are received by users in the first place (platform filtering). Moreover, we assume that the influencers know the initial distribution of users $\mu_0(x)$ (e.g., through polls, surveys, reviews and other forms of users' feedback).⁴

The benefit an influencer obtains from a particular distribution of users' opinions may vary. This variability can be captured by an arbitrary function $f(\cdot)$ that provides the influencer's benefit from a generic user at a given distance from the target opinion. Thus, in its greatest generality, the problem can be formulated as follows:

$$\begin{aligned} \max_{\{x_n^{(i)}\}_{n=1}^N} \mathbb{E}_x \left[f(|x_N^{(u)} - x^T|) \right] &= \int f(|x - x^T|) d\mu_N(x) \quad (3) \\ \text{s.t. dynamics in (2)} \end{aligned}$$

where $\mu_N(x)$ is the *final* distribution reached by users over the opinion space at time N . Note however that the maximization in Eq. (4) is over the entire sequence of N posts generated by the influencer. The influencer *benefit function* $f(\cdot)$ is reasonably a non-increasing function of the distance from the target opinion.

The formulation in Eq. (4) leads to a complex optimization problem, given the generality of the initial user distribution, the probabilistic movement of users (in the *finite population* case), the arbitrary choice of the influencer's opinion at each step, and the arbitrariness of function $f(\cdot)$. Therefore, we will now make a series

⁴It would be possible to incorporate a *noisy* version of such distribution into the model, along the lines of what happens in politics, where polls provide a noisy measure of the true distribution of public opinion. We leave this possibility for future work.

of simplifications that eventually lead to a problem that is solvable in polynomial time with N .

First, we assume that $f(\cdot)$ is a linearly decreasing function of the distance from the target point x^T , i.e., $f = -|x_N^{(u)} - x^T|$. While this simplification does not effectively reduce the complexity, it allows us to get a reasonable case study that does not require us to specify details of the shape function $f(\cdot)$. Therefore, by linearity of the mean and as $\max -g(\cdot) = \min g(\cdot)$, the problem becomes:

$$\min_{\{x_n^{(i)}\}_{n=1}^N} \mathbb{E}_x \left[|x_N^{(u)} - x^T| \right] = \int |x - x^T| d\mu_N(x) \quad (4)$$

s.t. dynamics in (2)

This reformulation also corresponds to a simpler interpretation: the influencer aims to bring the overall opinion of the population of regular users as close as possible to her target opinion.

Our main simplification assumes a binary (0-1) behavior for the event related to whether a post reaches a user at a certain distance d from the opinion expressed in the post. This is achieved, for example, by the following natural choice for the function ψ :

$$\psi(d = |x_n^{(u)} - x_n^{(i)}|; w) = \begin{cases} 0 & \text{if } d > w \\ 1 & \text{if } d \leq w \end{cases} \quad (5)$$

where w is a fixed width parameter. This means that a post conveying the opinion of the influencer $x_n^{(i)}$ is deterministically read at time n by users who have an opinion that is at most w away from it. This formulation of the model essentially implements the well known concept of *bounded confidence* (see e.g. [20] and [5]), where here we consider a single entity (the influencer) publishing posts to attract regular users. We emphasize that bounded confidence dynamics have proven quite difficult to analyze, so most of the known results have been obtained through Monte Carlo simulations.

REMARK 2. *Due to the assumption of a 0-1 $\psi(\cdot)$ function in Eq. (5), the model becomes deterministic even in the case of $|\mathcal{U}| < +\infty$ (finite population). Indeed, a user either moves or does not move, depending solely on his distance from the influencer's expressed opinion.*

Even with this simplifying assumption, the optimization problem remains significantly challenging. Indeed, it can be formulated as a Markov Decision Process (MDP), in which the number of states of the underlying Markov Chains (MCs), i.e., the MCs obtained by fixing the influencer sequence of posts, is combinatorially exponential with the number of users.

To gain initial insights into the problem and as a useful benchmark, we consider the *greedy* solution that, at each time step, selects the influencer's opinion $x_n^{(i)}$ that produces the best instantaneous improvement of the objective function. More formally, let $\Delta x_n^{(u)} \triangleq x_{n+1}^{(u)} - x_n^{(u)}$ be the *opinion shift* of a user u holding opinion $x_n^{(u)}$ at time n :

$$\Delta x_n^{(u)} = \begin{cases} (1 - \beta) \left[\delta z^{(u)} + (1 - \delta) x_n^{(i)} - x_n^{(u)} \right] & \text{if } \Psi(\psi) = 1 \\ \frac{\delta}{\delta + 1 - \beta} \left(z^{(u)} - x_n^{(u)} \right) & \text{if } \Psi(\psi) = 0 \end{cases}$$

For simplicity of exposition, but without loss of generality, consider the case in which all users are initially to the left of target point x^T , so that positive values of $\Delta x_n^{(u)}$ translate into equivalent

improvements of the objective function, whereas negative values translate into equivalent utility losses.

Then, given the users' distribution $\mu_n(x)$, the *greedy* algorithm selects, at each step n , the influencer opinion $x_n^{(i)}$ maximizing the overall users' opinion shifts:

$$x_n^{(i)} = \arg \max_{x_n^{(i)}} \mathbb{E} \left[\Delta x_n^{(u)} \right] = \arg \max_{x_n^{(i)}} \int \Delta x_n^{(u)} d\mu_n(x) \quad (6)$$

CLAIM 1. *The greedy strategy is not always optimal.*

As might have been expected, the above *greedy* strategy is, in general, suboptimal. We will demonstrate this in the next section in a simple but representative scenario. The reason for the suboptimality lies in the fact that the greedy algorithm does not “look into the future,” ruling out solutions that initially reduce the overall utility but, in the long run, lead to a better final configuration of users in the opinion space. Understanding when the *greedy* strategy may be suboptimal and by how much is of great interest both theoretically and practically. In cases where the *greedy* strategy does not lead to an optimal outcome, it is essential to apply strategies that sacrifice short-term gains in favor of long-term benefits.

4 THE OPTIMAL STRATEGY

The optimal strategy can be computed (numerically) under the simplifying assumptions introduced in the previous section by resorting to a discretization of the opinion space and the user distribution. In particular, let us assume that both the users' opinion and the influencer's opinion expressed in each post can only take values in a discrete set of cardinality B . In practice, we divide the opinion space $\mathcal{X} = [0, 1]$ into bins of constant width, and assume that only the mid-point of each bin is a feasible opinion value for the influencer's expressed opinion. For the sake of simplicity, we take as target points $x^{(T)} \in \{0, 1\}$ for the influencers operating in the system: these points can be interpreted as two opposing political views or as two different brands offering the same product to customers. In the case of a single influencer, we assume $x^T = 1$ (the case $x^T = 0$ is completely symmetrical).

Let \mathcal{B} be the set of feasible opinion values, and j be the index running on it. We will also discretize the prejudice of users, assuming that it belongs to a finite set $\mathcal{Z} \subseteq \mathcal{B}$ of prejudice values, of cardinality Z . Note that the population is indeed described by two distributions, the time-varying opinion and the static prejudice.

At last, we assume that the distribution $\mu_n(x)$ of users over the opinion space can be well approximated by considering the users belonging to a finite set \mathcal{M} of “groups”, of cardinality M , indexed by m . Users of a given group m share the same (time-varying) opinion $x_n^{(m)} \in \mathcal{B}$, and the same (static) prejudice value $z^{(m)} \in \mathcal{Z}$. Groups represent discrete “masses” of users moving together as a single unit, that cannot split into smaller sub-units over time. This is guaranteed by the assumption stated in Eq. (5), which leads to deterministic opinion movements. Note that without the 0-1 assumption for ψ , the position of a group could not be used to identify the state of the system and that as long as $0 < \psi < 1$, there would be an exponential growth of group subdivisions over time.

It follows that the system state at time n can be fully specified by the vector $\{x_n^{(m)}\}_m$, and that there are B^M possible system states.




Figure 1: A simple illustration with only four states ($B = 2$) to demonstrate the method used to determine the optimal solution.

At each time step, influencer i has to choose an opinion $x^{(i)} \in \mathcal{B}$ to convey in her n -th post, and we can separately evaluate the effect of this post on each group. This can be efficiently done by exploiting the Trellis-like structure of system dynamics described next.

4.1 Trellis-like structure for optimal strategy

Under the simplifying assumptions introduced before, the optimal solution can be computed in polynomial time for any n but exploiting the trellis-like structure sketched in Figure 1 for the toy case of $M = 2$ groups, $m \in \{0, 1\}$, and $B = 2$ opinion values, $x_n^{(m)} \in \{a, b\}$, leading to 4 possible system states. Let for short $S_n^{(s)} \in \mathcal{S}$ denote the possible system states at time n . From each state $S_n^{(s)} \in \mathcal{S}$, it is possible to evaluate all reachable states and the *action* $x^{(i)}$ that leads to the transition. This allows us to define a transition matrix T , shown on the left in Figure 1. Note that there may be multiple $x^{(i)}$ leading to the same target state, and for our purposes these transitions are equivalent. Indeed, we are interested only in the final *best* state, and not in finding all particular sequences of traversed states leading to it.

With the transition matrix in hand, it is possible to unfold the process over time, starting from a given initial state S_0^1 (see right part of Figure 1). The resulting trellis-like structure allows us to account for all paths starting from the initial state and efficiently compute the one to the final *best* state.

Without loss of generality, we can consider the case where $x^T = 1$ and the objective of our social impact maximization problem becomes $\max_{\{x^{(i)}\}_1^N} \mathbb{E}[x_N]$. Therefore, the *best* state is the one that leads to the highest average opinion in the population for $n = N$. In principle, the path leading to this state is not unique, and thus we consider a path as *optimal* if it leads to the *best* state in the shortest possible time.

4.2 The case of two user groups

Now that we have a method for deriving the *optimal* solution to the influence maximization problem, we can prove the correctness of Claim 1. For this, we have considered a very simple system that is computationally tractable and contains enough features to be of interest. We consider only two regular users' groups whose

Table 1: Scenarios of the greedy vs optimal comparison

#	$z^{(0)}$	$x_0^{(0)}$	$z^{(1)}$	$x_0^{(1)}$	α	β	N	w
1	0.5	0.005	0.5	0.995	0.05	0.3	50	0.1
2	0.005	0.005	0.995	0.995	0.05	0.3	50	0.1

prejudice and initial opinion are $z^{(m)}, x_0^{(m)}$, $m \in \{0, 1\}$ respectively, and a large number B of possible opinion values, leading to B^2 possible system states.

The discrepancy between greedy and optimal strategies can be evaluated for a variety of parameters. However, for the sake of compactness, we limit ourselves to two representative scenarios whose parameters are summarized in Table 1.

There are two ways in which the *optimal* solution can be better. It is either faster, i.e., it reaches the *best* state with a smaller number of posts (first scenario, Fig. 2a), or it leads to a higher value of $\mathbb{E}[x_N]$, i.e., a better *best* final state (second scenario Fig. 2d).

As a general rule of thumb, we have derived the following empirical rule from the optimal numerical solution of the system:

REMARK 3. *The optimal strategy for the online impact maximization problem is to first bring the (two) user groups close together, and then gradually persuade them towards the target opinion x^T .*

Note that as an effect of our choice of the platform filtering function ψ , group u is affected by the influencer's expressed opinion only when $|x^{(u)} - x^{(i)}| < w$. This means that a group whose opinion $x^{(u)}$ is too far from the influencer's expressed opinion $x^{(i)}$ is not affected by the post. Recall, however, that a group can still change its current opinion in the absence of a post's influence, as it is also attracted to its prejudice. In the next Section, we will go into more detail about what happens in the two-group scenario and explain the underlying system dynamics. We will also briefly discuss the effects of the parameters.

4.3 Numerical experiments: parameters impact

The first scenario considers two user groups with equal prejudice of 0.5 (representing a moderate position toward a certain topic




Figure 2: (2a) Scenario # 1 in Table 1, (2b) greedy, (2c) optimal. (2d) Scenario # 2 in Table 1. (2e) greedy (2f) optimal.

or indifference between the two extreme choices), and extreme initial opinions around 0 and 1, respectively. The second scenario instead assumes that the prejudices of the two groups coincide with their initial opinions, which are again set at the extremes of the opinion domain. It is interesting to note that the two strategies lead to different results in both cases, albeit in different ways. This becomes clear by looking at figures 2b-2c (for the first scenario) and figures 2e-2f (for the second scenario), showing detailed locations of the two user groups, as well as the influencer opinion, at each time step n , separately for the greedy and optimal strategy.

In the first scenario, the greedy strategy (plot 2b) leads the influencer to focus on the first, closest group $x^{(u=1)}$ and ignore the second, distant group $x^{(u=0)}$, which, in the absence of influencer stimulus, starts to gradually shift towards its prejudice $z^{(0)} = 0.5$.

Only when the second group is close enough, the greedy influencer finds it temporarily convenient to jump close to the second group, and immediately after go back to $x^{(i)} = 1$ to bring the first group close to the extreme of the opinion domain. This results in the erratic behavior of the orange trajectory in plot 2b.

These “hectic” moves make the greedy strategy inefficient, and in fact, the greedy strategy is largely outperformed by the optimal strategy (plot 2c) in terms of the number of steps (posts) to reach the *best* state. Note that, on the contrary, the optimal strategy focuses on the second distant group, which at some point (around $n = 10$) merges with the first group. The coalesced groups are then efficiently moved together towards the best state (see plot 2c).

In the second scenario, the greedy strategy (plot 2e) never allows the first group $x^{(u=1)}$ to step away from its (already taken) radicalized opinion at $x^{(u=1)} = 0.995$. Note that since $w = 0.1$ the

first group would move away from its initial position whenever the influencer conveys an opinion $0.895 \leq x^{(i)} < 0.995$

On the contrary, the optimal solution (plot 2f) accepts to temporarily worsen the opinion of the closer group, and by so doing it is able to pull the second group up to a better (closer to the target point) final position. This requires, at some point, a non-greedy step (see the non-monotonicity⁵ in plot 2d).

The difference between the two strategies increases when w takes small values and gradually diminishes as w increases. For $w > 0.5$, the greedy and optimal strategies coincide in virtually all cases, as larger w values lead to a more significant influence over the population. The extreme scenario is the one in which the entire population is always reached by a post. The optimal strategy in this case is to publish the target opinion exclusively, as there is no advantage in taking other viewpoints. Finally, it is worth mentioning that the weight coefficients α and β in the opinion update determine both the inertia of the system and the maximum achievable opinion value of each user group. Slow dynamics (high values of $\alpha + \beta$) are more challenging to study numerically because of the smaller opinion shifts Δx produced, which require a denser discretization of the opinion space (larger number of bins B)

Results (in terms of greedy vs optimal performance) are consistent despite the choice of these parameters, considering that the final opinion value depends very weakly on β (which primarily impacts the convergence time), while α essentially determines the best target opinion that *any* strategy can eventually achieve in the long run.

⁵Sometimes we observe non-monotonic behavior on the greedy trajectory, which can occur when *any* greedy step produces a negative increment of the objective function.

5 THE GAME: COMPETING FOR INFLUENCE

The above framework is interesting in that it shows how to optimally target a population of users, and it provides evidence that it is worth accepting short-term losses to achieve long-term gains. However, it completely disregards competition from other influencers on the social media platform. Inspired by the fact that in the *online impact maximization* problem an influencer has an incentive to first take a more moderate position in order to then exert influence over a larger user base, we will develop a game that embodies this idea. Indeed, Remark 3 suggests that it is better for the influencer to first group users, i.e. to express an opinion that allows it to reach “distant” users even, potentially, at the expense of de-radicalizing users close to its target opinion, and then gradually draw them towards its target. This consideration allows us to propose a fixed structure for an influencer’s strategy that greatly simplifies the problem and makes it tractable.

We define a game for the duopolistic competition of influencers aiming at maximizing their online social impact. Indeed, we consider a set $\mathcal{I} = \{0, 1\}$ of two *players* (influencers), each of which has a *target* opinion $x^{(i;T)}$: the opinion around which the influencer wants to attract users, recall Eq. (4). We will assume, without loss of generality, to have players having diametrically opposing views: $x^{(i=0;T)} = 0$ and $x^{(i=1;T)} = 1$, representing, for example, two opposing political parties.

The players/influencers are characterized by their willingness to deviate from their target opinion $x^{(i;T)}$ when trying to reach new users. This aspect is modeled by the parameter $\delta^{(i)}$, which allows us to define $x^{(i;E)} = x^{(i;T)} - \text{sgn}(x^{(i;T)} - 0.5) \cdot \delta^{(i)}$, which we will refer to as the *exploratory*⁶ opinion. We thus assume that the influencers will only assume one of the two opinions $x^{(i;T)}, x^{(i;E)}$. Moreover, the “exploration” phase will always precede the “targeting” phase, the situation is sketched in Figure 3.




Figure 3: Two-phase strategy $\{x^{(i)}\}_{i=1}^N$ structure.

Our assumptions may appear rather restrictive, however, they are motivated in part by the Remark 3 and by the necessity of reducing the space of possible *actions* for each player for tractability purposes. As a matter of fact, the set of *actions* becomes $\mathcal{A}^{(i)} = \{0, 1, \dots, N\}$, and the action $a^{(i)}$ determines how long the exploration phase lasts for influencer i , defining a particular $\{x^{(i)}\}$ sequence. Recall that knowing (the system’s parameters) and the influencer post sequence, the movement of the regular users is described by Eq. (2). Note that, once we fix the sequences of posts emitted by the two influencers (and knowing the system’s parameters), the movement of the regular users can be deterministically predicted by opinion dynamics in Eq. (2), by superposition of the effects.

⁶This variable represents the opinion that the influencer expresses while trying to approach a group of users who are far apart in their opinions and whom it would otherwise not reach due to the filter function ψ .

The last element to be specified to characterize our game is the *payoff* (or *utility*) function, which in our setting is a function $u_i : \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{R}$ and specifies the preferences of the players over the outcomes of the game, given the strategy of the other player(s). In our case, the payoff function corresponds to the objective function in Eq. (4), where the opinion configuration depends on the combined *actions* of the two influencers: $x_N^{(u)}(a^{(0)}, a^{(1)})$.

The game is a *simultaneous* game, i.e. both players choose their strategies at the same time and then stick to their choice for the entire duration of the game. It is also a game with *complete information*, i.e. both players know perfectly the rules of the game, i.e., they know the set of actions playable by the other player and the effect that such actions exert on the population of users.

5.1 The two-groups scenario

We first consider a simplified but illustrative scenario, similar to the one in Section 4.2, for which it is possible to provide an exact procedure to determine the Nash Equilibria (NE) under the assumption that only pure strategies may be adopted.




Figure 4: Schematic representation of the simplified scenario.

The restriction to two user groups is not strictly necessary, and we could consider a population with M groups, each with a (static) prejudice $z^{(m)} \in \mathcal{Z}$. Indeed, the complexity of the procedure to determine the NE would only scale linearly with M , as will become clearer later. However, we decided to limit the number of groups to two for continuity with Section 4.2 and ease of interpretation. We (first) consider a 0-1 ψ function to ensure identical reactions from the bulk of users belonging to the same group. So doing, we avoid the exponential growth of the user groups over time generated by their splitting. In Section 6 the case of a general user distribution and ψ function is discussed. Each group is characterized by $z^{(m)} \in \mathcal{Z}$ (assuming $x_0^{(m)} = z^{(m)}$) and its “proportion” defined as $\rho^{(m)} \triangleq \frac{\int d\mu_N(x) \psi(z^{(m)})}{\int d\mu_N(x)}$ with respect to the overall population. Lastly, we consider influencers whose target opinions lay at the extreme of the opinion domain: $x^{(i;T)} \in \{0, 1\}$. Fig. 4 sketches the setting. We first look into some straightforward solutions to the problem:

PROPOSITION 1. *If $|x^{(i;T)} - z^{(m)}| \leq w, \forall m \in \mathcal{M}, \forall i \in \mathcal{I}$ then $(0, 0)$ is a Nash equilibrium (NE) of the game.*

PROOF. Without loss of generality, we can consider i' such that $x^{(i';T)} = 1$ (recall $x^{(i;T)} \in \{0, 1\}$). Given that $x_0^{(m)} = z^{(m)}, \forall m \in$

\mathcal{M} and $|1 - x_0^{(m)}| \leq w, \forall m \in \mathcal{M}$, then $a^{(i')} = 0$ is a *dominant* strategy for i' , i.e., $\arg \max_{a^{(i')}} \mathbb{E}_x[x_N^{(u)}(\{a^{(i)}\}_i)] = 0, \forall a^{(i)}, i \in \mathcal{I} \setminus i'$. Indeed, the total opinion increment $\Delta x(\{x_n^{(i)}\}_n) \triangleq \Delta x^{(0)} + \Delta x^{(1)}$ for a general $x_n^{(i)}$ sequence is weakly smaller than $\Delta x(\{x^{(i;T)}\}_n)$, for any strategy of the other players. Every player would play this dominant strategy, so $(0, 0)$ is a NE. \square

It is clear that such situations are not of interest. Note also that in the exploration phase, a player favors the other player in some way (i.e. the closer user group moves further away from the player). A player would only do this if there is a possibility of reaching a larger user base, which is otherwise hindered by filtering.

We make a structural assumption that is in no way restrictive but has the dual goal of avoiding trivial solutions (see proposition 1) and ensures that the influencer does not deviate too much from her target opinion. This captures the fact that the influencer needs to still keep in contact with (and somehow maintain under control) the more radical individuals (those closer to $x^{(i;T)}$).

ASSUMPTION 1. *a) Only the closest group in opinion m is reachable ($\psi(|x^{(i;T)} - x^{(m)}|) = 1 \neq 0$) in the targeting phase while b) both (m, M) are reachable in the exploring phase.*

5.2 Exact solution for the two-groups scenario

Since the set of possible actions \mathcal{A} is finite and as $|\mathcal{I}| = 2$, it is possible to find the Nash equilibria for the two-players game or to determine when no Nash equilibrium exists.

THEOREM 1. *The procedure in Algorithm 1 identifies all the Nash equilibria in pure strategy of the game, if any exist.*

PROOF. The rationale for the algorithm is that in a Nash equilibrium, each player plays a *best response* $\mathcal{BR}_i(a^{(i)}, a^{(-i)})$ to the other player(s). In a two-player game, it is possible to define a matrix \mathbf{P} in which we have as rows the strategies of player $i = 0$ and as columns those of player $i = 1$. Each element p_{ij} is defined as (u_0, u_1) . This matrix is well-defined because the action space \mathcal{A} has finite cardinality, \mathbf{P} is $N + 1 \times N + 1$. For each column k (strategy of player $i = 1$), we compute $\arg \max_{a^{(0)}} u_0(a^{(0)}, a^{(1)} = k)$, which is the *best response* of player $i = 0$ to the action k of player $i = 1$. We do the same over the rows and then consider the elements (i, j) , if any, for which the procedure identified an $\arg \max$ over both the rows and the columns. These are the Nash equilibria in pure strategy, as both players play their *best response* to each other's strategy. \square

5.3 Characterization of the Nash equilibria as a function of the population characteristics

Figure 5 shows the Nash equilibria computed with the method in Algorithm 1, as a function of the user *inertia* β and the degree of stubbornness δ and under the Assumption 1. The experimental setting includes two equal user groups with $z^{(0)} = 0.25, z^{(1)} = 0.75$, the influencers are identical ($\delta^{(0)} = \delta^{(1)}$) and the considered time horizon $N = 5$, with a rectangular ψ function with width $w = 0.7$. We see that the less the population can be influenced, the more time the influencers spend on exploration. This can be explained by the fact that the closer an influencer is to a group of

Algorithm 1: Determine Nash equilibria

Data: players \mathcal{I} , groups of users \mathcal{M} , ψ function, actions \mathcal{A}
 // Compute the payoff matrix \mathbf{P}
foreach $(a^{(0)}, a^{(1)}) \in \mathcal{A} \times \mathcal{A}$ **do**
 $v \leftarrow \arg \min_i a^{(i)}$
 $V \leftarrow \mathcal{I} \setminus m$
 for $n \in \{1, \dots, N\}$ **do**
 if $n < m$ **then**
 $x_{n+1}^{(m)} = \text{UpdateBelief}(x_n^{(v;E)}, x_n^{(V;E)}), \forall m \in \mathcal{M}$
 else if $m < n < M$ **then**
 $x_{n+1}^{(m)} = \text{UpdateBelief}(x_n^{(v;T)}, x_n^{(V;E)}), \forall m \in \mathcal{M}$
 else
 $x_{n+1}^{(m)} = \text{UpdateBelief}(x_n^{(v;T)}, x_n^{(V;T)}), \forall m \in \mathcal{M}$
 $\mathbf{P}[a^{(0)}, a^{(1)}] \leftarrow (1 - \mathbb{E}_x[x_N^{(u)}], \mathbb{E}_x[x_N^{(u)}])$
 // Compute Best Responses and Nash equilibria
 $\mathcal{BR}_0(a^{(1)}) \leftarrow \arg \max_{a^{(0)}} \mathbf{P}[\cdot, a^{(1)}]$
 $\mathcal{BR}_1(a^{(0)}) \leftarrow \arg \max_{a^{(1)}} \mathbf{P}[a^{(0)}, \cdot]$
 $NE = \{(a^{(0)}, a^{(1)}) : a^{(0)} \in \mathcal{BR}_0(a^{(1)}), a^{(1)} \in \mathcal{BR}_1(a^{(0)})\}$

users, the more influence she can exert. This is also related to the fact that $(0, 0)$ is not a NE for the game. Since influencers in the exploration phase have a higher influence and can reach more users, any influencer who knows that the *competing* influencer can switch to a *milder* opinion would also carry out an exploration phase. In such a symmetric setting it is reasonable to expect symmetric equilibria (as in Fig. 5), and this is summarized in Proposition 2:

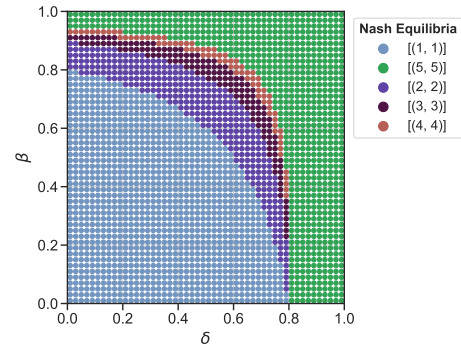


Figure 5: Nash equilibria in a symmetric scenario as a function of the population characteristics β, δ . The Nash Equilibria are a list (of one element) of tuples of the form $(a^{(0)}, a^{(1)})$.

PROPOSITION 2. *In the two-groups scenario, in the symmetric case ($\delta^{(i)} = \delta \forall i, z^{(0)} = 1 - z^{(1)}, \rho^{(0)} = \rho^{(1)}$) the actions of the two players in the Nash equilibrium are the same if the NE is unique.*

PROOF. It is rather straightforward to see that, in this situation, the best response of player 0 is $BR_0 = \arg \max_i u(\mu_n(x) | a^{(1)})$, and

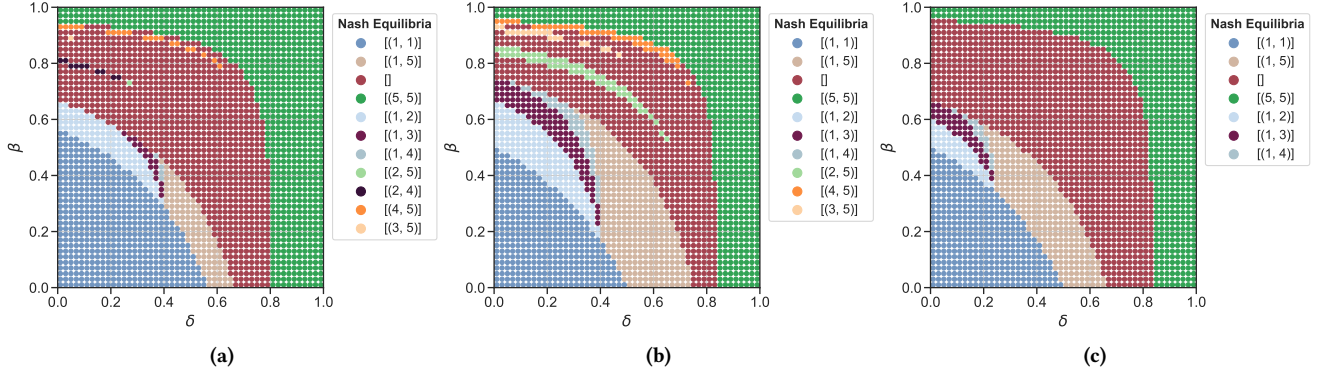


Figure 6: Nash equilibria as a function of β and δ for different shifts $\delta^{(i)}$ from the target opinion $x^{(i,T)}$, we organize the shifts for the two influencers as $(\delta^{(0)}, \delta^{(1)})$ and we consider $(0.2, 0.2)$ in (6a), $(0.1, 0.2)$ in (6b), and $(0.2, 0.3)$ in (6c).

it is equal to the best response of player 1 due to symmetry. Thus, if the NE is unique, the two players play the same strategy. \square

Under Assumption 1, the equilibria of the game are non-trivial, i.e., different from $(0, 0)$, and there is for the players a strategic incentive to compete.

So far we discussed a fully symmetric situation but, it is interesting to consider also an asymmetric distribution, i.e., $z^{(0)} = 0.15$, $z^{(1)} = 0.75$, in Figure 6 we depict the Nash equilibria. This setup implies that player 0 has an initial advantage in terms of user distribution. The structure of the above Nash equilibria suggests that the one in a disadvantageous position (here player 1) has a strategic incentive to prolong the exploration phases in order to approach the average collective opinion value of the population and thus also reach the most distant group and try to persuade it.

6 TOWARDS A MORE REALISTIC SCENARIO

Online social networks are characterized by pronounced asymmetries in the interaction between entities. Furthermore, in [9], the authors model the closed feedback between regular users and influencers. User feedback determines the *popularity* of influencers, which in turn is tied to their ability to reach users, i.e., their *visibility* over the platform. In this section, we use the full specification of the model in [9] in our framework.

6.1 Closed-loop opinion model

The users are subject to the dynamics described in Section 3.1 with the addition of the closed loop between regular users and influencers. For simplicity, we already present the model in its deterministic form, see Remark 1. The user *feedback* can be directly derived by the value of ψ , as it describes when a user is reached by a certain post and the probability of he moving in the direction of the influencer’s opinion because he likes it (positive feedback). Therefore we can define the total feedback provided by the users to a post of an influencer as $T_n^{(i)} = \int \psi_n^{(i)} d\mu_N(x)$. The feedback allows us to define the popularity $p^{(i)}$ update:

$$p^{(i)}(t+1) = p^{(i)}(t) + T_n^{(i)} \quad (7)$$

To close the loop, it is necessary to make ψ a function of popularity, i.e. the more popular an influence is in the OSN, the more users it can reach. To this end, consider the factorization of the function ψ into two contributions ($\psi = \omega \cdot \theta$), where the first factor models the *homophily* of interactions across OSNs together with the filtering of the platform, and the second describes the degree to which a user likes a post⁷. Considering θ a decreasing function of opinion distance also models the fact that users with very different opinions are less likely to be convinced. So even if an influencer is very popular, she will find it difficult to convince users who are distant in opinion.

6.2 Approximate solution of the game for arbitrary ψ and user distribution

The closed-loop scenario is more realistic but clearly more complicated. The procedure in Algorithm 1 cannot be applied directly because a group of users m cannot be tracked perfectly. At each time instant n , a certain group “splits” into a subgroup that is influenced and updates its opinion (with probability $\psi(d, p^{(i)})$), and the complementary group, which does not move. The number of subgroups grows exponentially as $2^N M$. To avoid the exponential explosion, one can discretize the opinion space into \mathcal{B} bins and keep the proportion of users from the groups in \mathcal{M} in each of the bins. With this simplification, it is possible to apply Algorithm 1 and consider any ψ function and also any (discrete) distribution of users. The ψ function in the product form we used in our experimental setting is:

$$\psi_n^{(i)}(d, p^{(i)}) = \omega(d, p^{(i)}) \theta(d) = e^{\nu \cdot \frac{d^2}{p^{(i)} / \sum_i p^{(i)}}} (1-d) \quad (8)$$

only the first factor depends on the relative popularity of the given influencer, and the parameter ν controls to what extent the visibility of an influencer decays with the distance.

From now on, we consider two possible scenarios, the “slow-dynamics”, in which popularity evolves considerably slower compared to opinions, and the “fast-dynamics”, in which popularity evolves quicker and has a greater impact on the visibility over the

⁷Only the product $\omega\theta$ is relevant for the total feedback, since a user must be reached ($\Omega = 1$) to give positive feedback ($\Theta = 1$) in the general stochastic model.

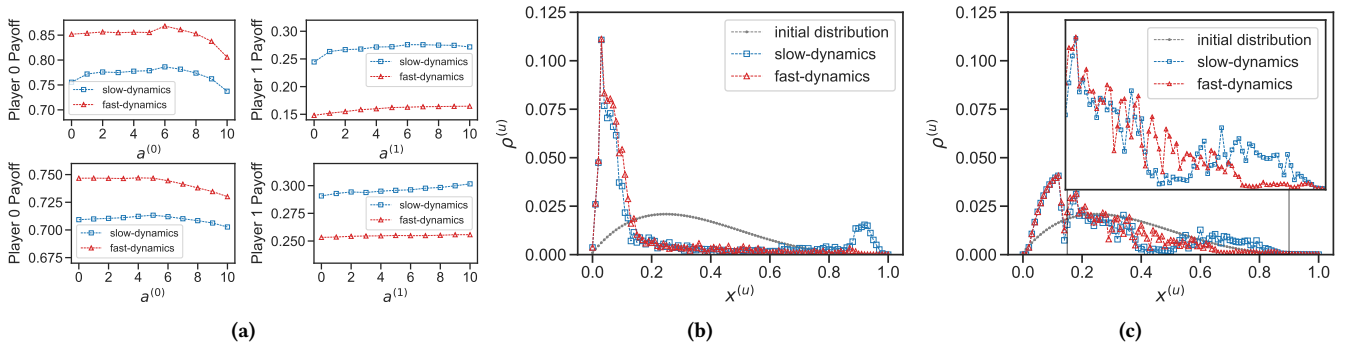


Figure 7: (a) Player’s payoffs considering the other influencer as stubborn with $x_n^{(i;T)}$, $\forall n$ for Scenario 1 (top) and Scenario 2 (bottom). Final opinion distribution at NE (b) Scenario 1 (c) Scenario 2, with different popularity growth speeds.

OSN. It is immediate to do so by appropriately normalizing the total feedback update $T_n^{(i)}$ in Eq. (7), choosing a large enough constant.

6.3 The effects of strategic behavior and popularity evolution

We developed a game for which the solution is limited to a very simplified structure. We will briefly discuss whether this simplified strategic behavior leads to a sizeable advantage for the players who behave strategically. To do this, we consider that one of the players is *stubborn*, i.e., she only posts her target opinion $x^{(i;T)}$. We can do this by looking at the first column and the first row of the payoff matrix P , which correspond respectively to a stubborn player 1 and a stubborn player 2.

We consider a Beta-shaped prejudice distribution (see grey line in Fig 7b-7c), skewed towards $x = 0$ and whose characterizing parameters $a = 2$ and $b = 4$. Similarly to previous sections, we assume that the initial opinion coincides with the users’ prejudice. The two players have $x^{(i;T)}$ of 0 and 1 respectively and both have $\delta^{(i)} = 0.1$. We consider two possible choices for the underlying opinion update weights, i.e., $\alpha = 0.1, \beta = 0.5$ (Scenario 1) and $\alpha = 0.1, \beta = 0.8$ (Scenario 2). In this setting we slightly modified the dynamical behavior so that the popularity can evolve macroscopically. For each action, we consider $N = 10$ posts are emitted carrying either the target opinion $x^{(i;T)}$ or the exploratory one $x^{(i;E)}$.

The results are shown in Figure 7a and reveal two aspects: first, the player who has a *structural advantage*, i.e., is favored by the initial distribution, benefits from the fast-dynamic settings. Second, the two players would exhibit somewhat opposite behavior, in that the favored player is harmed by long exploration phases, while the other receives better payoffs with long exploration phases. This supports the claim that the one in a disadvantaged position would benefit from being more “aggressive” and compromising her target opinion to get closer to the majority opinion in the population. This is also supported by the fact that in our experiments, when looking at the Nash equilibria, we have that: $a_{NE}^{(0)} \leq a_{NE}^{(1)}$.

Finally, from Figure 7b-7c, it is clear that the rapid evolution of popularity erases competition, as the *advantaged* influencer is able to become more visible on the platform (see Eq. (8)) than the other, and eventually attracts all users. Note the bimodal nature of

the final distribution in the slow-dynamics setting (blue curve) and the flat nature of the distribution in the fast-dynamics setting (red curve).

7 CONCLUSION

We formalized the problem of maximizing online social impact over an online social network by considering a model that incorporates communication asymmetry and platform filtering in OSNs. In the case of a single influencer over the network, we characterized the optimal strategy, highlighting that the greedy strategy does not provide the same benefits. Inspired by this experiment, we developed a competitive game that we extended for the case of closed-loop interactions (post-feedback-popularity-filtering). We showed that a disadvantaged influencer should compromise her target opinion more to get closer to the mass of the population to have a greater degree of persuasiveness towards them. We found that when popularity develops rapidly, an influencer (the advantaged one) tends to monopolize attention and attract virtually all users to her opinion.

REFERENCES

- [1] Carmela Bernardo, Claudio Altafini, Anton Proskurnikov, and Francesco Vasca. 2024. Bounded confidence opinion dynamics: A survey. *Automatica* 159 (2024), 111302. <https://doi.org/10.1016/j.automatica.2023.111302>
- [2] Heather Z. Brooks and Mason A. Porter. 2020. A model for the influence of media on the ideology of content in online social networks. *Phys. Rev. Res.* 2 (Apr 2020), 023041. Issue 2. <https://doi.org/10.1103/PhysRevResearch.2.023041>
- [3] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. 2009. Statistical physics of social dynamics. *Reviews of Modern Physics* 81, 2 (may 2009), 591–646. <https://doi.org/10.1103/revmodphys.81.591>
- [4] Wei Chen, Alex Collins, Rachel Cummings, Te Ke, Zhenming Liu, David Rincon, Xiaorui Sun, Yajun Wang, Wei Wei, and Yifei Yuan. 2011. *Influence Maximization in Social Networks When Negative Opinions May Emerge and Propagate*. 379–390. <https://doi.org/10.1137/1.9781611972818.33> arXiv:<https://arxiv.org/abs/10.1137/1.9781611972818.33>
- [5] Guillaume Deffuant, David Neau, Frederic Amblard, and Gérard Weisbuch. 2000. Mixing beliefs among interacting agents. *Advances in Complex Systems* 03, 01n04 (Jan. 2000), 87–98. <https://doi.org/10.1142/s0219525900000078>
- [6] Morris H. Degroot. 1974. Reaching a Consensus. *J. Amer. Statist. Assoc.* 69 (1974), 118–121. <https://api.semanticscholar.org/CorpusID:120077293>
- [7] Swapnil Dhamal, Walid Ben-Ameur, Tijani Chahed, and Eitan Altman. 2019. Optimal investment strategies for competing camps in a social network: A broad framework. *IEEE Transactions on Network Science and Engineering* 6 (10 2019), 628–645. Issue 4. <https://doi.org/10.1109/TNSE.2018.2864575>
- [8] Noah Friedkin and Eugene Johnsen. 1999. Social Influence Networks and Opinion Change. *Advances in Group Processes* 16 (01 1999).
- [9] Franco Galante, Luca Vassio, Michele Garetto, and Emilio Leonardi. 2023. Modeling communication asymmetry and content personalization in online social

- networks. *Online Social Networks and Media* 37-38 (2023), 100269. <https://doi.org/10.1016/j.osnem.2023.100269>
- [10] Mohak Goyal and D. Manjunath. 2020. Opinion Control Competition in a Social Network. In *2020 International Conference on COMMunication Systems & NETWORKS (COMSNETS)*. 306–313. <https://doi.org/10.1109/COMSNETS48256.2020.9027368>
- [11] Rainer Hegselmann, Stefan König, Sascha Kurz, Christoph Niemann, and Jörg Rambau. 2015. Optimal Opinion Control: The Campaign Problem. *Journal of Artificial Societies and Social Simulation* 18, 3 (2015), 1–18. <https://ideas.repec.org/a/jas/jasssj/2014-124-3.html>
- [12] Rainer Hegselmann and Ulrich Krause. 2015. Opinion dynamics under the influence of radical groups, charismatic leaders, and other constant signals: A simple unifying model. *Networks and Heterogeneous Media* 10 (07 2015), 477–509. <https://doi.org/10.3934/nhm.2015.10.477>
- [13] David Kempe, Jon Kleinberg, and Éva Tardos. 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 (Washington, D.C.) (KDD '03)*. Association for Computing Machinery, New York, NY, USA, 137–146. <https://doi.org/10.1145/956750.956769>
- [14] Sascha Kurz. 2014. Optimal control of the freezing time in the Hegselmann–Krause dynamics. *Journal of Difference Equations and Applications* 21 (2014), 633 – 648. <https://api.semanticscholar.org/CorpusID:6635240>
- [15] Wei Lu, Wei Chen, and Laks V. S. Lakshmanan. 2015. From competition to complementarity: comparative influence diffusion and maximization. *Proc. VLDB Endow.* 9, 2 (oct 2015), 60–71. <https://doi.org/10.14778/2850578.2850581>
- [16] Antonia Masucci and Alonso Silva. 2014. Strategic Resource Allocation for Competitive Influence in Social Networks. *2014 52nd Annual Allerton Conference on Communication, Control, and Computing, Allerton 2014* (02 2014). <https://doi.org/10.1109/ALLERTON.2014.7028557>
- [17] Anahita Mirtabatabaei, Peng Jia, and Francesco Bullo. 2014. Eulerian Opinion Dynamics with Bounded Confidence and Exogenous Inputs. *SIAM Journal on Applied Dynamical Systems* 13, 1 (2014), 425–446. <https://doi.org/10.1137/130934040>
- [18] K. Sridhar Moorthy. 1985. Using Game Theory to Model Competition. *Journal of Marketing Research* 22, 3 (1985), 262–282. <http://www.jstor.org/stable/3151424>
- [19] Hossein Noorazar. 2020. Recent advances in opinion propagation dynamics: a 2020 survey. *The European Physical Journal Plus* 135 (06 2020). <https://doi.org/10.1140/epjp/s13360-020-00541-2>
- [20] Hegselmann Rainer and Ulrich Krause. 2002. Opinion Dynamics and Bounded Confidence: Models, Analysis and Simulation. *Journal of Artificial Societies and Social Simulation* 5, 3 (2002).
- [21] Vineeth Varma, Morărescu Irinel-Constantin, Samson Lasaulce, and Samuel Martin. 2017. Opinion Dynamics Aware Marketing Strategies in Duopolies. In *IEEE 56th Annual Conference on Decision and Control (CDC)*. <https://doi.org/10.1109/CDC.2017.8264227>
- [22] Vineeth S. Varma, Samson Lasaulce, Julien Mounthanyvong, and Irinel-Constantin Morărescu. 2019. Allocating Marketing Resources Over Social Networks: A Long-Term Analysis. *IEEE Control Systems Letters* 3, 4 (2019), 1002–1007. <https://doi.org/10.1109/LCSYS.2019.2919959>
- [23] Suttida Wongkaew, Marco Caponigro, and Alfio Borzi. 2015. On the control through leadership of the Hegselmann–Krause opinion formation model. *Mathematical Models and Methods in Applied Sciences* 25, 03 (2015), 565–585. <https://doi.org/10.1142/S0218202515400060>

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009