A Game-Theoretic Analysis of a Competitive Diffusion Process over Social Networks*

Vasileios Tzoumas[1,3], Christos Amanatidis[2], and Evangelos Markakis[2]

[1] National Technical University of Athens, School of Electrical and Computer Engineering
[2] Athens University of Economics and Business, Department of Informatics
[3] University of Pennsylvania, Department of Electrical and Systems Engineering

Abstract. We study a game-theoretic model for the diffusion of competing products in social networks. Particularly, we consider a simultaneous non-cooperative game between competing firms that try to target customers in a social network. This triggers a competitive diffusion process, and the goal of each firm is to maximize the eventual number of adoptions of its own product. We study issues of existence, computation and performance (social inefficiency) of pure strategy Nash equilibria in these games. We mainly focus on 2-player games, and we model the diffusion process using the known linear threshold model. Nonetheless, many of our results continue to hold under a more general framework for this process.

1 Introduction

A large part of research on social networks concerns the topic of diffusion of information (e.g., ideas, behaviors, trends). Mathematical models for diffusion processes have been proposed ever since [11, 19] and also later in [9]. Given such a model, some of the earlier works focused on the following optimization problem: find a set of nodes to target so as to maximize the spread of a given product (in the absence of any competitors). This problem was initially studied by Domingos and Richardson [8], Kempe et al. [13], and subsequently by [6, 18]. Their research builds on a "word-of-mouth" approach, where the initial adopters influence some of their friends, who in turn recommend it to others, and eventually a cascade of recommendations is created. Within this framework, finding the most influential set of nodes is NP-hard, and approximation algorithms as well as heuristics have been developed for various models.

Different considerations, however, need to be made in the presence of multiple competing products in a market. In real networks, customers end up choosing a product among several alternatives. Hence, one natural approach to model this competitive process is the use of game-theoretic analysis with the players being the firms that try to market their product. The game-theoretic approaches that have been proposed along this direction mainly split into two types. The first is to view the process as a Stackelberg game, where the competitors of a product first choose their strategy, and then a

* Work supported by the project ALGONOW of the research funding program Thales (co-financed by the European Social Fund-ESF and Greek national funds). A full version of this work with all missing proofs is available at the authors' websites.

last mover needs to make a decision on the set of nodes to target [3–5, 14]. This approach essentially reduces to the algorithmic question of finding the best response for the firm that moves last. The main results that have been obtained along this direction is that, in certain cases, the algorithm of [13], for the case of a single product, can be applied in the competitive environment as well. For more models and related problems under this context see also [2]. A different approach is to capture the competition as a simultaneous game, where firms pick their initial set of nodes at the same time, and then the diffusion process follows (after first taking care of ties). This was first proposed in [1], and has also been studied very recently by [10]. The approach of [1] and [10] as a noncooperative normal-form game is the focus of our work as well.

1.1 Contributions

Our work is an attempt to further understand game theoretic aspects of viral marketing. To this end, we first define in Section 2 a general framework for a competitive diffusion process in a social network, generalizing the model of [1]. This corresponds to a class of non-cooperative games where firms target customers in order to maximize the spread of their own product. We study issues of existence, computation, and performance (social inefficiency) of pure Nash equilibria (PNE). The games we define are one-stage games, as in [1, 10], i.e., all firms spend their budget in one step, a fact that renders natural the analysis of PNE. We use as instantiations of the competitive diffusion process the well-known linear threshold model, however some of our results also hold for more general local interaction schemes.

In more detail, we mostly deal with 2-player games, as in [10], and in Section 3 we first illustrate that such games may not possess PNE, even for simple graphs. On top of that, we also prove that it is co-NP-hard to decide whether a PNE exists for a given game. We then move on to investigate conditions for the existence of PNE. In Section 4, we begin with studying the improvement paths induced by our games. We exhibit that networks with special in and out-degree distributions — e.g. power law — are not expected to be more stable than others, in the sense that all possible dynamics can be realized essentially by any graph. Motivated by all these, we then focus on sufficient conditions for the existence of PNE via generalized ordinal potential functions. We also consider ϵ-approximate generalized ordinal potentials, and we provide tight upper bounds on the existence of such approximations, as well as, polynomial time algorithms for computing approximate PNE. Finally, we study the Price of Anarchy and Stability for games with an arbitrary number of players, and we show that PNE (when they exist) can be quite inefficient. We conclude with a discussion of the effects on the payoff of a single player (or a coalition of players), as the number of competitors increases.

We view as one of the main contributions the fact that we unveil new decisive factors for the existence of PNE that are intertwined with structural characteristics of the underlying network. For example, some of the factors that play a role in our model for obtaining generalized ordinal potentials (exact, or approximate) involve *i)* the *diffusion depth* of a game (defined in Section 2 as the maximum possible duration of the diffusion process), *ii)* the *ideal spread* (defined as the maximum possible spread that a strategy can achieve) and *iii)* the *diffusion collision factor* (defined in Section 4.3 as a measure for comparing how two strategies of one player perform against a given strategy of another player). We advocate that our results motivate further empirical research on social networks for identifying a typical range of these quantities in real networks. Regarding the diffusion depth, some empirical research has already provided new insights for certain recommendation networks [15].

1.2 Related work

Our work has been largely motivated by [1], (see also the erratum [21]). To the best of our knowledge, this was the first article to consider such games over networks with the players being the firms. The diffusion process of [1] is a special case of our model, in particular, it is a linear threshold model where each firm is allowed to target only one node as a seed, and the thresholds and the weights are equal to $1/|N(v)|$ (with $N(v)$ being the neighborhood of v). We consider the general class of linear threshold models, and in some cases our results hold even for arbitrary local interaction schemes beyond threshold models. In [1] the existence of equilibria is linked to bounding the diameter of the graph. In our model we find that the diameter is not much correlated to existence. Instead we identify other parameters that influence the existence of equilibria.

Besides [1], a very recent related work is [10]. One of the major differences between [10] and our work is that they study the set of *mixed* Nash equilibria of a similar diffusion game, and focus on the Price of Anarchy, and another measure denoted as the *Budget Multiplier*. We, on the other hand, focus on *pure* Nash equilibria. Another difference is that [10] is studying stochastic processes whereas our local interaction schemes induce deterministic processes, as in [1].

Other game-theoretic approaches have also been considered for social networks. One line of work concerns models of Stackelberg games as mentioned earlier [3–5, 14]. A different approach is to consider a game where the players are the individual nodes of the network, who have a utility function depending on their own choice, and that of their neighbors, see e.g., [17, 20]. This leads to very different considerations.

2 Preliminaries

2.1 Social Networks

The underlying structure of the social network is assumed *static*, and is modeled by a fixed finite directed graph $G = (V, E)$ with no parallel edges and no self-loops. Each node $v \in V$ represents an individual within the social network, while each directed edge $(u, v) \in E$ represents that v can be influenced by u. We assume that there are two competing products (or trends, ideas, behavioral patterns) produced by two different firms $\mathcal{M} = \{1, 2\}$, and to each such product we assign a distinct color. Throughout this work, we shall use the terms product, color, and firm interchangeably. Further, each node can have at most one color, and as with most of the literature, we assume that all decisions are *final*; i.e., no node that has adopted a particular product will later alter its decision. Moreover, if a node has adopted a product, we shall refer to it as **colored**, or **infected**, otherwise we will call it a **white** node.

We denote the (in-)neighbors of a node v as $N(v) = \{u \in V | (u, v) \in E\}$, i.e., $N(v)$ is the set of nodes that can influence v. Also, we denote as d_v^{in} and d_v^{out} the in-degree and out-degree of v. The way that a node v can be influenced by $N(v)$ is usually described by a **local interaction scheme** (LIS). Hence, a local interaction scheme is essentially a function that takes as input a node v, the status of its neighbors, a product c under consideration, and possibly other characteristics of the graph, and determines if node v is eligible to adopt this product. An example of a LIS, that was initially studied for the spread of a single product, is the **linear threshold model** (LTM) [11, 19]. Under LTM, there is a weight $w_{uv} \in [0, 1]$ for every edge (u, v) such that for every node v, it holds that $\sum_{u \in N(v)} w_{uv} \leq 1$. Every node v also has a threshold value $\theta_v \in (0, 1]$. The condition that needs to hold, under LTM, so that node v can adopt a product c is

$$\sum_{u \in N(v)} \mathbb{I}_u w_{uv} \geq \theta_v,$$

where $\mathbb{I}_u$ is 1 if u has already adopted product c, and zero otherwise. Note that in a local interaction scheme, the eligibility condition may hold for more than one product at a given time (e.g., under LTM this could happen if $\theta_v < 1/2$ for some node v).

Given a local interaction scheme, and a set of competing firms, we consider the following competitive diffusion process, which evolves over discrete time steps:

The diffusion process. Initially each firm tries to infect a set of "seeds". The number of seeds for each firm may depend on its budget for advertising and marketing. We assume here that the firms have the same power so that in the beginning they can target a set of k nodes each (we think of k as being much smaller than $|V|$ but not necessarily a constant).

- At time step $t = 0$: This is the **initiation** step. In the beginning, all nodes are colored white. If a node v was targeted by a single firm c, then v adopts product c. Since each firm may pick to target an arbitrary set of k nodes, some overlaps may also occur. Thus, we assume that a tie-breaking criterion TBC1 is applied to resolve such dilemmas. This may be a global rule, or a rule that depends on each node.
- At any time step $t > 0$: We look at each remaining white node and check if it is eligible to adopt any of the products, i.e., if the adoption condition, as determined by LIS, holds. For this, we take into account *only* the neighbors of v that were infected up until time step $t - 1$, hence the order with which we examine the white nodes does not matter. During this process, a white node v may be eligible to adopt more than one product. To resolve such dilemmas a second tie-breaking criterion TBC2 should be considered. The process terminates at a time step t, when no white node is eligible to adopt any product. We allow that TBC1 may differ from TBC2, since TBC2 may depend on specific features of the diffusion process, whereas TBC1 occurs only at the initiation step.

A particular instance of a tie-breaking criterion, that we shall often use, is the rule that is also used in [12, 17], where ties are resolved in favor of the "best quality" product: all the individuals within the social network share a *common* **reputation ordering**, say $R^{\prec} \equiv 1 \succ 2$, over the products and in case of ties they decide according to $R^{\prec}$. We shall also see later that some of our results are independent of the tie-breaking rules.

Note 1. All definitions above can be generalized in a straightforward manner to an arbitrary number of m firms, i.e., $\mathcal{M} = \{1, \ldots, m\}$. In Section 3 and Section 4 we focus mostly on the *2-player* case. Section 5 deals with arbitrary m-player games as well.

Definition 1. *A **social network** $\mathcal{N}$ is defined through the tuple* $(G, \texttt{LIS}, \texttt{TBC1}, \texttt{TBC2})$.

2.2 Strategic Games induced by Diffusion Processes

A game $\Gamma = (\mathcal{N}, \mathcal{M}, k)$ is induced by a social network $\mathcal{N} = (G, \texttt{LIS}, \texttt{TBC1}, \texttt{TBC2})$ and the set of firms $\mathcal{M} = \{1, \ldots, m\}$, which we shall refer to as a **diffusion game**. In a diffusion game, all participating firms choose simultaneously a set of k seeds, which then triggers a diffusion process according to the interaction scheme and tie-breaking criteria of $\mathcal{N}$. We denote as $\mathcal{S} = \{S : |S| = k\}$ the set of available strategies, which is the same for each firm. We shall use the phrases *strategy* S and *subset* S interchangeably. A pure strategy profile is a vector $\mathbf{s} = (S_1, \ldots, S_m) \in \mathcal{S}^m$, where S_i corresponds to the strategy played by player $i \in \mathcal{M}$. Also, we set $\mathbf{s_{-i}} \equiv \{S_1, \ldots, S_{i-1}, S_{i+1}, \ldots, S_m\}$.

Given a strategy profile $\mathbf{s} \in \mathcal{S}^m$, the utility of firm $i \in \mathcal{M}$, denoted by $u_i(\mathbf{s})$, is the total number of nodes that have been colored by firm i at the end of the competitive diffusion process. We denote the associated game matrix as $\Pi(\Gamma)$. Moreover, a pure strategy profile $\mathbf{s} \in \mathcal{S}^m$ is a **pure Nash equilibrium (PNE)** of game Γ if $u_i(S_i', \mathbf{s_{-i}}) \leq u_i(\mathbf{s})$, $\forall i$ and $\forall S_i'$.

An important parameter in our games is the so-called diffusion depth defined below.

Definition 2. *The **diffusion depth** $D(\Gamma)$ of a game Γ is defined as the maximum number of time steps that the competitive diffusion process may need, where the maximum is taken over all strategy profiles* $\mathbf{s} \in \mathcal{S}^m$.

Observe that the diffusion depth can take values either lower, equal, or greater than the diameter of the underlying graph G.

Another important notion in our analysis is defined below. Consider a hypothetical scenario where only one player participates in the game. Then his payoff will not be obstructed by anybody else, and any strategy that he chooses achieves its best possible performance. This is useful for quantifying the players' utilities as we shall see later on.

Definition 3. *Assume that only one player from $\mathcal{M}$ participates in the game, and let $S \in \mathcal{S}$ be one of his strategies. We define as **ideal spread of** $\mathbf{S}$, denoted by H_S, the set of nodes that have adopted by the end of the diffusion process the product of this player under strategy S. This includes the initial seed as well, i.e., $S \subseteq H_S$.*

3 Existence: Examples and Complexity

We start with some remarks concerning the presentation. In Section 3 and Section 4 we consider mostly *2-player* games. Furthermore, our results mainly hold for the linear threshold model but some of them can be generalized to arbitrary models. Whenever in stating a theorem, we do not specify a parameter of the network, it means that it holds independent of its value (e.g. in some results we do not specify the tie-breaking criteria, or the local interaction scheme).

The games that we study do not always possess PNE and we present an example below to illustrate this. We note that this is independent of the tie-breaking criteria used. For any other choice of such criteria (deterministic or even randomized), we can construct analogous examples.

Example 1. Consider the game $(\mathscr{N}, \mathcal{M} = \{1, 2\}, k = 1)$ over the graph of Figure 1, where $\mathscr{N} = (G, \texttt{LIS} = \texttt{LTM}, \texttt{TBC1} = R^{\prec}, \texttt{TBC2} = R^{\prec})$. It is easy to generalize this to a line with an arbitrary number of nodes. We assume that all nodes have threshold 1.

The game matrix is seen in Table 1 and it is easy to check that no PNE exists.

$$n_1 \xrightarrow{1} n_2 \xrightarrow{1} n_3$$

Fig. 1: A network with underlying structure a line.

	n_1	n_2	n_3
n_1	$3, 0$	$1, 2$	$2, 1$
n_2	$2, 1$	$2, 0$	$1, 1$
n_3	$1, 2$	$1, 1$	$1, 0$

Table 1: The payoff matrix for the game of Figure 1.

The example reveals that even simple graph structures may fail to have PNE. This holds for larger values of k as well, and we have also found other examples with no PNE, where the graph G is a cycle, a clique, or belongs to certain classes of trees.

Given these examples, the next natural question is whether it is easy to decide if a given game has at least one PNE. We assume that the input to this problem is not the game matrix, which can be exponentially large, but simply the graph G and a description of the local interaction scheme. Note that for $k = O(1)$ the problem is easy, hence the challenge is for larger values of k. We establish the following hardness result.

Theorem 1. *Deciding whether a game* $((G, \texttt{LIS} = \texttt{LTM}, \texttt{TBC1} = R^{\prec}, \texttt{TBC2} = R^{\prec}), \mathcal{M}, k)$ *has a PNE is co-NP-hard and belongs to* Σ_2^p.

Remark 1. The reduction in the proof of Theorem 1 produces instances where the network is a directed acyclic graph (DAG) and the diffusion depth is $D = 3$.

The proof of Theorem 1 is based on a reduction from 3SAT. Note that we have not obtained membership in the class co-NP. This is because there seems to be no short certificate for checking that a game does not have any PNE (one would need to check all strategy profiles). It is an open problem to determine if the problem is complete for Σ_2^p. Another open problem would be to determine the complexity for games with diffusion depth $D = 1$, or $D = 2$.

4 Towards Characterizations

To understand better the issue of existence of PNE, we start with quantifying the utility functions $u_i : \mathcal{S}^2 \mapsto \mathbb{N}$, $i \in \mathcal{M} = \{1, 2\}$. For this we need to introduce some important notions. A convenient way to calculate the utility of a player under a profile s, is by utilizing the definition of H_S in Section 2, which is the ideal spread of a product if the firm was playing on its own and used S as a seed. In the presence of a competitor, the firm will lose some of the nodes that belong to H_S. The losses happen due to three reasons. First, the competitor may have managed to infect a node at an earlier time step than the step that the firm would reach that node. Second, the firm may lose nodes due to the tie-breaking criteria, if both firms are eligible to infect a node at the same time step. Finally, there may be nodes that belong to H_S, but the firm did not manage to infect enough of their neighbors so as to color them as well. These nodes either remain white, or are eventually infected by the other player. All these are captured below:

Definition 4. *Consider a game* $((G, \mathtt{LIS}, \mathtt{TBC1}, \mathtt{TBC2}), \mathcal{M}, k)$*, and a strategy profile* $\mathbf{s} = (S_1, S_2)$*. For* $i \in \{1, 2\}$*,*

i. we denote by $\alpha_i(\mathbf{s})$ *the number of nodes that belong to* H_{S_i}*, and under profile* $\mathbf{s}$*, player* i *would be eligible to color them at some time step* t *but the other player has already infected them at some earlier time step* $t' < t$ *(e.g., this may occur under the threshold model when* $\theta_v < 1/2$ *for some node* v*).*

ii. we denote by $\beta_i(\mathbf{s})$ *the number of nodes in* H_{S_i}*, such that under profile* $\mathbf{s}$*, both firms become eligible to infect them at the same time step, and due to tie-breaking rules, they get infected by the competitor of* i*.*

iii. we denote by $\gamma_i(\mathbf{s})$ *the number of nodes that belong to* H_{S_i}*, but under* $\mathbf{s}$*, firm* i *never becomes eligible to infect them (because* i *did not manage to color the right neighbors under* $\mathbf{s}$*).*

Finally, we set $\alpha_{i,max}$ (respectively $\beta_{i,max}, \gamma_{i,max}$) to be the maximum value of $\alpha_i(\mathbf{s})$ over all valid strategy profiles and also $\alpha_{max} = \max\{\alpha_{1,max}, \alpha_{2,max}\}$ (similarly for β_{max}, and γ_{max}). We refer the reader to an example in our full version for an illustration of these concepts.

When we use $R^{\prec}$ for tie-breaking, clearly $\beta_1(\mathbf{s}) = 0$. Hence for *2-player* games of the form $((G, \mathtt{LIS}, \mathtt{TBC1} = R^{\prec}, \mathtt{TBC2} = R^{\prec}), \mathcal{M}, k)$, the utility functions of the players, given a strategy profile $\mathbf{s} = (S_1, S_2) \in \mathcal{S}^2$, are

$$u_1(\mathbf{s}) = |H_{S_1}| - \alpha_1(\mathbf{s}) - \gamma_1(\mathbf{s}), \quad (1)$$

$$u_2(\mathbf{s}) = |H_{S_2}| - \alpha_2(\mathbf{s}) - \beta_2(\mathbf{s}) - \gamma_2(\mathbf{s}). \quad (2)$$

4.1 Realizability of Improvement Paths

Following Section 3, we unwind further the richness and complexity of our games motivated by the study of their *improvement paths*. We establish that one of the main structural properties of social networks, their degree distribution, does not play a role on its own to the existence of equilibria. This fact motivates the search for other important parameters related to existence, which is the topic of the next subsections.

An **improvement path** is any sequence $(\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_j, \ldots)$ of strategy profiles $\mathbf{x}_j \in \mathcal{S}^2$ such that for every j the strategy profiles $\mathbf{x}_j$ and $\mathbf{x}_{j+1}$ differ in exactly one coordinate, say the $i(j)$-th, i.e., only player $i(j)$ has switched to another strategy, and also $u_{i(j)}(\mathbf{x}_j) < u_{i(j)}(\mathbf{x}_{j+1})$, $\forall j \geq 1$. It is called a **best response path** if also

$u_{i(j)}(\mathbf{x}_{j+1}) = max_{x \in \mathcal{S}} u_{i(j)}(x, (\mathbf{x}_{j+1})_{-i(j)})$. We can also define **improvement cycles** in a similar fashion.

A well-known sufficient condition for existence of PNE is the *Finite Improvement Property* (FIP), saying that all improvement paths are finite [16]. In our case, the FIP does not hold, but in order to find conditions for the existence of PNE, one could still try to understand how do cycles occur. For example, do the cycles have some particular form? Does the degree distribution affect the formation of cycles? We obtain a negative result in this direction, showing that essentially in any given graph, any possible set of cycles may be realized, independent of its structure.

We proceed with some more definitions. Given a finite 2-player game Γ, played on a $r \times r$ matrix, let $P(\Gamma)$ denote the set of all improvement paths (including infinite ones) that are induced by the game starting from any entry in the matrix. Let P denote any possible set of *consistent* improvement paths (including cycles) that can be created on a $r \times r$ matrix. By a consistent set we mean that if, e.g., there is a path with a move from entry (i, j) of the matrix to (i, l), then there cannot be another path in P that contains a move from (i, l) to (i, j). We say that P is **realizable** if there is a game Γ such that $P = P(\Gamma)$. We show that any such set P is realizable by the family of our games, essentially by any graph. Hence *all* possible dynamics can be captured by these games.

To prove our claim, we will argue about an appropriate submatrix of the games we construct, since some strategy profiles may need to be eliminated due to domination. Particularly, we need the following form of domination.

Definition 5. *Given a 2-player game, assume that $\mathcal{S} = X \cup Y$, where $X \cap Y = \emptyset$. We say that X is a **sink** in $\mathcal{S}$, if* at least one *of the following holds:*

i. $\forall (a, b) \in Y \times (Y \cup X)$, *$\exists x \in X$ such that* $u_1(x, b) > u_1(a, b)$*, and* $\forall (a, b) \in X \times Y$, $\exists x \in X$ *such that* $u_2(a, x) > u_2(a, b)$.
ii. $\forall (a, b) \in (Y \cup X) \times Y$, *$\exists y \in X$ such that* $u_2(a, y) > u_2(a, b)$*, and* $\forall (a, b) \in Y \times X$, $\exists x \in X$ *such that* $u_1(x, b) > u_1(a, b)$.

The definition says that any improvement path that is not a cycle, starting from the Y-region of the matrix, will eventually come to the X-region.

Given a 2-player game, we let $\mathcal{S}_{\mathcal{D}}$ denote a minimal sink in $\mathcal{S}$, and $\Pi(\mathcal{S}_{\mathcal{D}}, \mathcal{S}_{\mathcal{D}})$ be the restriction of the game matrix over this set of strategies. Furthermore, given a graph G, let P_{in} and P_{out} be the in and out-degree distributions of G, i.e., $P_{in}(i)$ is the number of nodes with in-degree equal to i. We can now state the following theorem.

Theorem 2. *Consider a graph $G' = (V, E)$ with in and out-degree distributions, P_{in}, and P_{out}. There exists a class of games $((G \in \mathcal{F}$, LIS $=$ LTM, TBC1 $= R^{\prec}$, TBC2 $= R^{\prec}), \mathcal{M}, k)$, where $\mathcal{F}$ is a family of graphs with the same set of nodes as G', such that:*

i. each $G \in \mathcal{F}$ has degree distributions P_{in}^G, P_{out}^G such that for all i, $|P_{in}(i) - P_{in}^G(i)|/|V| \to 0$, as $|V| \to \infty$, and the same holds for P_{out} and P_{out}^G.
ii. For any $r \geq 3$, all sets of consistent improvement paths (including cycles) formed on a $r \times r$ matrix are realizable over the games played on $\mathcal{F}$ in $\Pi(\mathcal{S}_{\mathcal{D}}, \mathcal{S}_{\mathcal{D}})$, where $\mathcal{S}_{\mathcal{D}}$ (for each $G \in \mathcal{F}$) is a minimal sink with $|\mathcal{S}_{\mathcal{D}}| = r$.

This result discloses the richness of our games, but above all it severely mitigates the role of the widely studied degree distribution of networks to the stability of the involved games. Hence, we advocate, in the following, that one needs to take into account the effects of other properties as well.

4.2 Conditions for the Existence of a PNE

In this subsection, we use the notion of ordinal potentials to argue about existence of PNE. A function $P : \mathcal{S}^2 \mapsto \mathbb{R}$ is a **generalized ordinal potential** [16] (**GOP**) for a game Γ if $\forall i \in \mathcal{M}$, $\forall \mathbf{s_{-i}} \in \mathcal{S}^{m-1}$, and $\forall x, z \in \mathcal{S}$,

$$u_i(x, \mathbf{s_{-i}}) > u_i(z, \mathbf{s_{-i}}) \Rightarrow P(x, \mathbf{s_{-i}}) > P(z, \mathbf{s_{-i}}).$$

If Γ admits a GOP and is also finite (as our games are), the FIP property holds (see Section 4.1), and all improvement paths terminate at a PNE [16]. On the other hand, in our games the existence of PNE is not equivalent with the FIP property; we can construct games that possess PNE, but do not admit a GOP. We omit these due to lack of space. Instead, we continue with a set of necessary conditions for the existence of a GOP. To this end, we shall say that a set X of nodes is **reachable** from a strategy S if and only if $X \subseteq H_S$, where H_S is the ideal spread of S.

Lemma 1. *The game* $((G, \texttt{LIS}, \texttt{TBC1} = R^{\prec}, \texttt{TBC2}), \mathcal{M}, k)$ *cannot admit a generalized ordinal potential if*

i. $\exists (S_1, S_2) \in \mathcal{S}^2$, $S_1 \neq S_2$, *such that* S_1 *is reachable from* S_2, *and* S_2 *is reachable from* S_1.
ii. $\exists (S_1, S_2) \in \mathcal{S}^2$, $S_1 \neq S_2$, *such that* $|H_{S_1}| = |H_{S_2}|$, *and* S_1 *is reachable from* S_2, *or* S_2 *is reachable from* S_1.

The Lemma suggests that many classes of our games may not admit a GOP — in the next subsection we shall approximate how close to admitting a GOP these games are. The fact is elucidated further through the following corollary, where we assume $k = 1$, i.e., as in [1], each player has to pick a single node, therefore, the only reasonable strategies are the nodes u for which there is at least one edge (u, v) such that $w_{uv} \geq \theta_v$.

Corollary 1. *If the game* $((G, \texttt{LIS} = \texttt{LTM}, \texttt{TBC1} = R^{\prec}, \texttt{TBC2} = R^{\prec}), \mathcal{M}, k = 1)$ *admits a generalized ordinal potential, then*

i. G *contains a DAG that includes the set* $\{u | \exists v \in V, \text{such that } w_{uv} \geq \theta_v\}$.
ii. If $w_{uv} \geq \theta_v$ *for every edge* $(u, v) \in E$, *then* G *has to be a DAG.*

Corollary 1 shows that for the case of $k = 1$, the conditions on the ideal spreads implied by Lemma 1, enforce the graph to have the special structure of a DAG. But clearly not all DAGs admit a GOP as has been demonstrated in Example 1.

We now move on to derive a sufficient condition for the existence of a GOP.

Theorem 3. *Consider a game* $((G, \texttt{LIS}, \texttt{TBC1} = R^{\prec}, \texttt{TBC2}), \mathcal{M}, k)$*, and suppose that we order the set of the available strategies so that* $|H_{S_1}| \geq \ldots \geq |H_{S_{|\mathcal{S}|}}|$*. If for all* $i \in \{1, \ldots, |\mathcal{S}| - 1\}$ *it holds that*

$$|H_{S_{i+1}}| \leq \left\lfloor \frac{|H_{S_i}| + max\left\{\gamma_1(S_i, S_{i+1}), \gamma_2(S_i, S_{i+1})\right\}}{2} \right\rfloor \quad (3)$$

then the game admits a generalized ordinal potential. Moreover, all its PNE have the form (S_{max}, S_2)*, where* $S_{max} \equiv argmax_{S \in \mathcal{S}} \{|H_S|\}$.

For an interpretation of Theorem 3, consider a game where the max term is zero in (3). Then, a GOP exists if all ideal spreads are well separated, and Player 2 can never hope to take more than half of the nodes that Player 1 would get ideally (we refer the reader to the introductory example for this theorem in our full version).

The condition of Theorem 3 can be relaxed so that not all ideal spreads need to be well separated, e.g., in certain cases where there is no overlap between the ideal spreads of some strategies. For example, when G is a full and complete d-ary tree, $d \geq 2$, then (3) does not hold but using similar arguments as in the proof of Theorem 3 we have:

Corollary 2. *The games of the form* $((G, \texttt{LIS} = \texttt{LTM}(w_{uv} \geq \theta_v, \forall (u,v) \in E), \texttt{TBC1} = R^{\prec}, \texttt{TBC2} = R^{\prec}), \mathcal{M}, k = 1)$*, where G is a full and complete* d*-ary tree, admit a GOP.*

4.3 Quantifying Instability

The previous sections on existence and complexity motivate our next discussion on approximate PNE. Overall, the main conclusion of this subsection is that even though PNE do not always exist, we do have in certain cases approximate equilibria with a good quality of approximation, and we can also compute them in polynomial time.

A strategy profile s is an ϵ**-PNE**, if no agent can benefit more than ϵ by unilaterally deviating to a different strategy, i.e., for every $i \in \mathcal{M}$, and $S'_i \in \mathcal{S}$ it holds that $u_i(S'_i, \mathbf{s_{-i}}) \leq u_i(\mathbf{s}) + \epsilon$. Recall that in our case, utilities are integers in $\{0, ..., |V|\}$, and ϵ also takes integer values[4]. Additionally, a function $P : \mathcal{S}^2 \mapsto \mathbb{R}$ is an ϵ**-generalized ordinal potential** (ϵ**-GOP**) for a game $\varGamma$ (see [7]) if $\forall i \in \mathcal{M}$, $\forall \mathbf{s_{-i}} \in \mathcal{S}^2$, $\forall x, z \in \mathcal{S}$, $u_i(x, \mathbf{s_{-i}}) > u_i(z, \mathbf{s_{-i}}) + \epsilon \Rightarrow P(x, \mathbf{s_{-i}}) > P(z, \mathbf{s_{-i}})$. Such a function P yields directly the existence of ϵ-PNE. We first obtain such a potential function for games that have diffusion depth $D = 1$, based on the ideal spread of the players' strategies and on the quantification of the utility functions in the beginning of Section 4 (Definition 4).

Theorem 4. *Any game* $\varGamma = ((G, \texttt{LIS}, \texttt{TBC1} = R^{\prec}, \texttt{TBC2} = R^{\prec}), \mathcal{M}, k)$*, where* $D(\varGamma) = 1$*, admits the function* $P(\mathbf{s}) = (1 + \beta_{max} + \gamma_{max})|H_{S_1}| + |H_{S_2}| - \beta_2(\mathbf{s}) - \gamma_2(\mathbf{s})$*, as a* k*-GOP. Moreover, a* k*-PNE can be computed in polynomial time.*

The last assertion of Theorem 4 is easy to see since the value of the function $P(\mathbf{s})$ is at most $O(|V|^2)$. Therefore, by following an improvement path (with improvements of more than k), we can find an approximate PNE quite efficiently.

[4] We could normalize the utilities by dividing by $|V|$, and then ϵ would take values in the set $\{1/|V|, 2/|V|, ..., 1\}$. We present the theorems without the normalization so as to be consistent with all other sections.

Note that this holds for any local interaction scheme, and not just the linear threshold model. Theorem 4 implies that when $D(\Gamma) = 1$ and k is small, we can have a good quality of approximation. E.g., for $k = O(1)$, or $k = o(|V|)$, and as $|V| \to \infty$, we can have approximate equilibria where any node can additionally gain only a negligible fraction of the graph by deviating.

For games with higher diffusion depth, we define below an important parameter that captures the quality of approximation we can achieve in worst case via ϵ-GOP.

Definition 6. *i. Given a 2-player game, and two strategy profiles* $\mathbf{s} = (S_1, S_2), \mathbf{s}' = (S_1', S_2)$*, the* ***diffusion collision factor*** *of player* 1 *for strategy* S_1' *compared to* S_1*, given* S_2*, is defined as* $DC_1(S_1', S_1|S_2) \equiv (\alpha_1(\mathbf{s}') + \gamma_1(\mathbf{s}')) - (\alpha_1(\mathbf{s}) + \gamma_1(\mathbf{s}))$*.*

ii. Similarly, for $\mathbf{s} = (S_1, S_2), \mathbf{s}' = (S_1, S_2')$*, the diffusion collision factor of Player 2 for* S_2' *compared to* S_2*, given* S_1*, is defined as* $DC_2(S_2', S_2|S_1) \equiv (\alpha_2(\mathbf{s}') + \gamma_2(\mathbf{s}')) - (\alpha_2(\mathbf{s}) + \gamma_2(\mathbf{s}))$*.*

In order to understand this new notion, recall from Equation (1) that, given a profile $\mathbf{s}$, $\alpha_1(\mathbf{s}) + \gamma_1(\mathbf{s})$ denotes the number of nodes that Player 1 does not infect due to the presence of Player 2 in the market; this fact directly yields some intuition for the definition of DC_1. This is not exactly the case for DC_2, as the β_2-term is missing (see Eq. (2)); nonetheless, it turns out that it suffices to define DC_2 in a uniform manner as DC_1, when using $R^{\prec}$ for ties. Finally, we set DC_{max} to be the maximum possible diffusion collision factor.

Theorem 5. *Any game* $\Gamma = ((G, \texttt{LIS}, \texttt{TBC1} = R^{\prec}, \texttt{TBC2} = R^{\prec}), \mathcal{M}, k)$*, where* $D(\Gamma) \geq 2$*, admits the function* $P(\mathbf{s}) = x_1|H_{S_1}| + |H_{S_2}| - \beta_2(\mathbf{s})$*, as a* DC_{max}*-GOP, where* x_1 *is any number satisfying* $x_1 > \beta_{max}$*. Moreover, a* DC_{max}*-PNE can be computed in polynomial time.*

The approximations of k and DC_{max} are *tight* for `LIS=LTM`, and we provide the corresponding examples in our full version.

5 Quantifying Inefficiency

5.1 Price of Anarchy and Stability

Given an m-player game, and a strategy profile $\mathbf{s}$, the sum $SW(\mathbf{s}) = \sum_{j=1}^{m} u_j(\mathbf{s})$ is the **social welfare** of $\mathbf{s}$. The *Price of Anarchy* (PoA), for a family of games, is the worst possible ratio of $SW(\mathbf{s})/SW(\mathbf{s}')$, where $\mathbf{s}$ is a social optimum, and $\mathbf{s}'$ is a Nash equilibrium. Similarly, the *Price of Stability* (PoS) is defined as the best such ratio.

Suppose now that $|V|$ is sufficiently large, so that players will never play overlapping strategies at a PNE, e.g., this is ensured if $|V| \geq mk$. In that case we would have $1 \leq PoA \leq |V|/(mk)$. The question of interest then is whether PoA can be much lower than this upper bound.

The following theorem exhibits that for diffusion depths greater than one, competition can severely hurt social welfare. This can be detrimental both to the firms, and the network users, since it implies that in worst case the firms will have a very low utility, and the service offered by these competing products will reach only a small fraction of

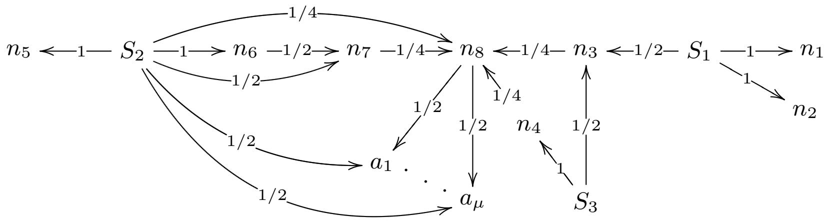

Fig. 2: The network for the proof of Theorem 7(ii): All nodes have threshold 1, except of node n_3 that has $\theta_{n_3} = 1/2$, and node n_8 that has $\theta_{n_8} = 1/2$.

the nodes. This is in agreement with the worst case scenario in the model of [10][5]. On the contrary, this is not always the case when the diffusion depth is one.

Theorem 6. *i. For the family of games* $((G, \texttt{LIS} = \texttt{LTM}, \texttt{TBC1} = R^{\prec}, \texttt{TBC2} = R^{\prec}), \mathcal{M}, k)$*,* $PoA = |V|/(mk)$*, and* $PoS \geq \frac{k}{k+1}\frac{|V|}{mk}$*, even for* $D = 2$*.*

ii. For the family of games $((G, \texttt{LIS} = \texttt{LTM}, \texttt{TBC1} = R^{\prec}, \texttt{TBC2} = R^{\prec}), \mathcal{M} = \{1, 2\}, k = 1)$*, with* $D = 1$*, we have* $PoS = 1$*, and* $PoA \leq SW(\mathbf{s})/(SW(\mathbf{s}) - 1)$*, where* $\mathbf{s}$ *is a social optimum. Moreover, if there exist at least two nodes with nonzero out-degree, then* $PoA = 1$*.*

The negative effect of competition on the players' utilities is further illustrated in the next subsection from the perspective of the best quality player.

5.2 Worst-case scenarios for the best quality player

We end our presentation with identifying a different form of inefficiency for PNE, which arises from the following question: Consider games with the reputation ordering $R^{\prec}$ as the tie-breaker. Does the firm with the best quality product ensure the maximum spread among all the players at any PNE? Theorem 7 illustrates that this may not always be the case for games with at least three players (but it is so for 2-player games). In fact, the payoff of the best quality player may be arbitrarily lower than the player with the highest market share at a PNE. We consider this as a form of inefficiency since in a socially desirable outcome, one would expect that the product with the best quality/reputation should have the largest market share. This surprising result dictates the necessity for quantifying such effects in PNE.

Theorem 7. *Consider the class of games* $((G, \texttt{LIS}, \texttt{TBC1} = R^{\prec}, \texttt{TBC2}), \mathcal{M}, k)$*.*

i. If $m = 2$*, then for all PNE* $\mathbf{s}$*, it is* $u_1(\mathbf{s}) \geq u_2(\mathbf{s})$*.*

ii. If $m \geq 3$*,* $\texttt{LIS} = \texttt{LTM}$*, and* $\texttt{TBC2} = R^{\prec}$*, then a game exists with a PNE* $\mathbf{s}$ *such that* $u_i(\mathbf{s}) < u_j(\mathbf{s})$*, although* $i \succ j$ *with regard to* $R^{\prec}$*.*

[5] In the stochastic process of [10], PoA can be very high when their so-called switching function is not concave.

Proof. i. Assume that a PNE $\mathbf{s} = (S_1, S_2)$ exists such that $u_1(\mathbf{s}) < u_2(\mathbf{s})$. Then, Player 1 can deviate to $S'_1 = S_2$, and obtain utility $u_1(S_2, S_2) \geq u_2(\mathbf{s}) > u_1(\mathbf{s})$. Thus, $\mathbf{s}$ cannot be a PNE.

ii. Note that $R^{\prec} = 1 \succ 2 \succ 3$, and consider the social network in Figure 2: As n_i, $\forall i \in \{1, \ldots, 8\}$, and as a_i, $\forall i \in \{1, \ldots, \mu\}$, where $\mu > k$, we denote single nodes. We assume that all of them have threshold 1, except of nodes n_3, and n_8 that have $\theta_{n_3} = 1/2$, and $\theta_{n_8} = 1/2$. As S_i, $\forall i \in \{1, 2, 3\}$, we denote sets of k nodes. Finally, the edges between single nodes are annotated with their corresponding weight. On the other hand, the edges that emanate from a set S_i are annotated with the *accumulated* corresponding weight of the underlying edges between each of the nodes in S_i and the involved end-node (e.g., $\forall v \in S_1$, it is $w_{vn_3} = \theta_{n_3}/k$).
One can now verify that the profile $\mathbf{s} \equiv (S_1, S_2, S_3)$ constitutes a PNE, even though it is $u_2(\mathbf{s}) = k + \mu + 4$, $u_1(\mathbf{s}) = k + 3$, and $u_3(\mathbf{s}) = k + 1$ — i.e., $u_2(\mathbf{s}) > u_1(\mathbf{s}) > u_3(\mathbf{s})$. We omit the details.

In the network of Figure 2, observe that at the PNE $\mathbf{s} = (S_1, S_2, S_3)$, $u_2(\mathbf{s}) = k + \mu + 4 > u_1(\mathbf{s}) + u_3(\mathbf{s}) = 2k + 4$, since $\mu > k$ — note that μ can be arbitrarily large. Thereby, if firm 1 is affiliated with firm 3, while their products are marketed as competing and incompatible (e.g. airline merges), firm 1 is incentivized to withdraw firm 3 from the game: the resulting 2-player game between firm 1 and firm 2, has a unique PNE, namely (S_2, S_1), in which firm 1 achieves the maximum possible utility — $u_1(S_2, S_1) = k + \mu + 4$. Moreover, notice that in this 2-player game, firm 1 initiates only k nodes to achieve this utility. On the other hand, in the original 3-player game, firms 1 and 3 initiate k nodes *each*, and still they achieve a lower sum of utilities at $\mathbf{s}$.

Our discussion indicates the necessity to capture the motivation of a player to either merge with other players, or to divide itself to several new ones that, although affiliated, they are still non-cooperative within the induced game. For example, given the network of Figure 2 Player 1 faces the question: Should I play alone against the others, since I am the best firm, or should I merge even with the weakest? We believe this aspect of PNE is worth further investigation and we leave it as an open direction for future work.

6 Conclusions and Future Work

We have studied a competitive diffusion process from a non-cooperative game-theoretic viewpoint. We have investigated several aspects related to the stability of such games and we have unveiled some important parameters that have met no previous investigation. We believe that our work motivates primarily further empirical research on social networks with regard to the following questions: Can we identify a range of typical values for decisive structural features such as the diffusion depth, the ideal spread, and the maximum diffusion collision factor? This could quantify the instability of the induced games, in light of Theorems 4 and 5, as well as the results in Section 4.2.

Other interesting questions have to do with resolving some of the remaining open problems from our work. It is still open if the complexity of determining that a PNE exists is Σ_2^p-complete, or not. The Price of Anarchy is also not yet completely determined when $D = 1$ and k is arbitrary. Finally, additional compelling questions may concern the robustness to network changes.

References

1. Alon, N., Feldman, M., Procaccia, A.D., Tennenholtz, M.: A note on competitive diffusion through social networks. Information Processing Letters (IPL) 110(6), 221–225 (2010)
2. Apt, K.R., Markakis, E.: Diffusion in social networks with competing products. In: 4th International Symposium on Algorithmic Game Theory (SAGT). pp. 212–223 (2011)
3. Bharathi, S., Kempe, D., Salek, M.: Competitive influence maximization in social networks. In: International Workshop on Internet and Network Economics (WINE). pp. 306–311 (2007)
4. Borodin, A., Filmus, Y., Oren, J.: Threshold models for competitive influence in social networks. In: 6th International Workshop on Internet and Network Economics (WINE). pp. 539–550 (2010)
5. Carnes, T., Nagarajan, C., Wild, S.M., van Zuylen, A.: Maximizing influence in a competitive social network: A follower's perspective. In: 9th International Conference on Electronic Commerce (ICEC). pp. 351–360 (2007)
6. Chen, N.: On the approximability of influence in social networks. SIAM Journal on Discrete Mathematics (SIDMA) 23(3), 1400–1415 (2009)
7. Chien, S., Sinclair, A.: Convergence to approximate Nash equilibria in congestion games. Games and Economic Behavior 71(2), 315–327 (2011)
8. Domingos, P., Richardson, M.: Mining the network value of customers. In: 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD). pp. 57–66 (2001)
9. Goldenberg, J., Libai, B., Muller, E.: Talk of the network: A complex systems look at the underlying process of word-of-mouth. Marketing Letters 12(3), 211–223 (2001)
10. Goyal, S., Kearns, M.: Competitive contagion in networks. In: 44th ACM Symposium on Theory of Computing (STOC). pp. 759–774 (2012)
11. Granovetter, M.S.: Threshold models of collective behavior. The American Journal of Sociology 83(6), 1420–1443 (1978)
12. Immorlica, N., Kleinberg, J., Mahdian, M., Wexler, T.: The role of compatibility in the diffusion of technologies through social networks. In: Proceedings of the 8th ACM Conference on Electronic Commerce (EC). pp. 75–83 (2007)
13. Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence in a social network. In: 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD). pp. 137–146 (2003)
14. Kostka, J., A., O.Y., Wattenhofer, R.: Word of mouth: Rumor dissemination in social networks. In: 15th International Colloquium on Structural Information and Communication Complexity (SIROCCO). pp. 185–196 (2008)
15. Leskovec, J., Adamic, L., Huberman, B.: The dynamics of viral marketing. ACM Transactions on the Web (TWEB) 1(1), Article No. 5 (2007)
16. Monderer, D., Shapley, L.S.: Potential games. Games and Economic Behavior 14, 124–143 (1996)
17. Morris, S.: Contagion. The Review of Economic Studies 67(1), 57–78 (2000)
18. Mossel, E., Roch, S.: On the submodularity of influence in social networks. In: Symposium on Theory of Computing (STOC). pp. 128–134 (2007)
19. Schelling, T.C.: Micromotives and Macrobehavior. Norton, New York, NY (1978)
20. Simon, S., Apt, K.R.: Choosing products in social networks. Manuscript, available at `http://arxiv.org/abs/1202.2209` (2012)
21. Takehara, R., Hachimori, M., Shigeno, M.: A comment on pure-strategy Nash equilibria in competitive diffusion games. Information Processing Letters (IPL) 112(3), 59–60 (2010)