Influence Maximization in Switching-Selection Threshold Models

Dimitris Fotakis[1], Thodoris Lykouris[2],
Evangelos Markakis[3], and Svetlana Obraztsova[1]

[1] National Technical University of Athens, Athens, Greece
[2] Cornell University, Ithaca, NY, USA
[3] Athens University of Economics and Business, Athens, Greece

Abstract. We study influence maximization problems over social networks, in the presence of competition. Our focus is on diffusion processes within the family of threshold models. Motivated by the general lack of positive results establishing monotonicity and submodularity of the influence function for threshold models, we introduce a general class of switching-selection threshold models where the switching and selection functions may also depend on the node activation history. This extension allows us to establish monotonicity and submodularity when (i) the switching function is linear and depends on the influence by all active neighbors, and (ii) the selection function is linear and depends on the influence by the nodes activated only in the last step. This implies a $(1 - 1/e - \varepsilon)$-approximation for the influence maximization problem in the competitive setting. On the negative side, we present a collection of counterexamples establishing that the restrictions above are essentially necessary. Moreover, we show that switching-selection threshold games with properties (i) and (ii) are valid utility games, and thus their Price of Anarchy is at most 2.

1 Introduction

A large part of recent research on social networks concerns the design of marketing strategies for advertising new products over a network. The focus of these efforts is on exploiting viral effects for the spread of new ideas and technologies among networks of friends, colleagues, relatives or other circles. The algorithmic question that naturally arises under such diffusion processes is then the following: find a subset of "most influential" nodes to target (i.e., advertise the new product to or even give it for free), so as to maximize the expected number of product adoptions, subject to a budget constraint.

This problem was initially formalized and studied by Domingos and Richardson [3] and by Kempe et al. [9], who focused on two of the most popular families of stochastic diffusion processes, namely the so-called *threshold* models [6, 12] and *cascade* models [4]. Finding the optimal set of influential nodes under this framework is an NP-hard problem, and the work of [9] proposed an approximation algorithm, achieving a guarantee of $1-1/e$. The algorithm is based on the observation that the function quantifying the total influence of a set of early adopters is a monotone and submodular function, and thus, the classical greedy approach for maximizing such set functions applies [11].

The models above however, do not take into account the presence of multiple competing products in a market. In real networks, customers (i.e., nodes) end up choosing a product among various alternatives. To take the simplest possible scenario, suppose there are two firms, R and B (standing for the red and blue product respectively), trying to promote their product over a social network. A convenient way to model the process now is by viewing this setting as a 2-player game, with the strategy space being the subsets of nodes that can be targeted subject to each firm's budget constraint.

Within this game-theoretic framework, interesting research questions arise. First, one can have a natural extension of the problem studied in [9] for a single product, as follows: Given a strategy of firm B, find the best subset of nodes for firm R, so as to maximize the expected number of product adoptions in her favor. In other words, find an algorithm to compute the best response of a player to a strategy of her competitor. At first sight, it may appear that the problem under competition may not differ significantly from that without competition. For certain cascade models, this is indeed the case, see e.g. [1]. Interestingly enough however, this does not hold for threshold models. In [2], several extensions of the threshold model were presented where the best response function is nonmonotone and/or nonsubmodular and the techniques used in [9] cannot be employed to obtain a good approximation. It is still a major open problem in the area to understand for which diffusion models, one can compute (near) optimal strategies efficiently. Moreover, apart from best responses, another direction is to study further the properties of Nash equilibria of the game and quantify their performance, as was done recently in [7, 5]. For example, one can study the Price of Anarchy of such games, or other criteria, such as the Budget Multiplier, introduced in [5].

Our Contribution: Motivated by the lack of positive results establishing monotonicity and submodularity of the influence function in competitive threshold models, we embark on a more systematic study of this question. On the conceptual side, we introduce a fairly general class of threshold models that belong to the family of *switching-selection* models. Under these models, a node first makes a decision on whether to adopt some product (i.e., whether to switch to being activated) and then makes a separate decision on which product to adopt (selection process). These two steps are determined by a *switching* and a *selection* function. Our class is essentially a threshold version of the models studied recently in [5] and [7], generalizing at the same time some of their aspects. In particular, we do not restrict the switching and selection functions to depend only on the set of currently active neighbors. Instead, we let them depend on the whole activation history, i.e., on the sets of active nodes at every time step. This extension allows for a careful investigation of properties that lead to a monotone and submodular influence function, and we obtain both positive and negative results under this class.

On the positive side, our main technical contribution is a set of conditions on the switching and selection functions that lead to monotonicity and submodularity and thus, enable us to obtain an $(1-1/e-\varepsilon)$-approximation for the influence maximization problem in the competitive setting, for any $\epsilon > 0$. Specifically, our main result (Section 3) is that the best response of a switching-selection threshold model is monotone and submodular if (i) the switching function is linear and depends on the weight of all active neighbors of a node, and (ii) the selection function is linear and depends on the weight of the nodes activated in the last step (i.e., the most recent buyers are the ones to influence the actual product selection). For the proof, we first establish the equivalence of such models with a generalization of the “live edges” approach [9], applicable to this particular setting, and then we develop quite delicate coupling arguments for establishing monotonicity and submodularity. Moreover, we conjecture that our positive results extend to the case where the switching function is any nondecreasing concave function of the weight of all active neighbors (see the discussion at the end of Section 3).

On the negative side, we present (Section 4) a comprehensive collection of counterexamples establishing that the restrictions above are essentially necessary. Regarding the switching process, we present examples showing that the influence function may not be monotone and submodular if the switching function is either not monotone or not concave, or it allows for the influence to decrease over time. For the selection process, we have analogous counterexamples when the selection function depends on the weight of neighbors activated in steps before the last one, or when it deviates from linearity.

Finally, we also study the performance of Nash equilibria of the underlying game, motivated by the properties established for the models in [5, 7]. We show (Section 5) that switching-selection threshold games with the properties identified above are valid utility games, and thus their Price of Anarchy is at most 2.

2 The Model

In this section, we define the class of *Switching-Selection Threshold Models*. This is essentially a ”threshold” version of the *Switching-Selection Model* introduced in [5], generalizing at the same time some of its aspects, as we clarify later within this section.

Social Networks. We model a social network by a directed graph $G(V, E)$, $|V| = n$. Each edge (u, v) has a weight $w_{uv} \in [0, 1]$, specifying the degree of influence of node u towards node v. For any node v, we denote by $N(v)$ the set of in-neighbors of v, and we require that the sum of the weights of the edges towards v is no more than 1: $\sum_{u \in N(v)} w_{uv} \leq 1$.

We consider a 2-player game between two competing firms that try to promote their product over the network (in fact our results generalize to games with more players, as we state later on, but for simplicity the presentation in Section 3 is for 2 players). We denote the two players by R and B standing for the red and blue product respectively. Each player $p \in \{R, B\}$ has a budget $K_p \in \mathbb{N}_+$, which they will use to target selected nodes in the network. The decision that the firms need to make is to choose how to disperse their budget to the n nodes, hence the strategy space for each firm p consists of all vectors (i.e., multisets) in the form $a_p = (a_{1p}, a_{2p}, \ldots, a_{np})$, where $a_{jp} \in \mathbb{N}$ and $\sum_{j=1}^{n} a_{jp} \leq K_p$.

Once the firms make a choice, the spread of the two products is modeled by a stochastic diffusion process that takes as input the strategies of the 2 firms, a_R, a_B. We describe next a family of such processes that we are interested in.

Switching-Selection Diffusion Processes. The process that determines the eventual adoptions, takes place in discrete steps. The state $s_{ut} \in \{$*R(ed)*, *B(lue)*, *U(ncolored)*$\}$, of node u, denotes whether node u has adopted a product at step t and, if yes, which product it adopted. As with the majority of the literature, we assume that the process is progressive, i.e., once a node is colored, it never changes its state afterwards. The process evolves as follows:

- At time step $t = 0$, the initialization takes place. For every node u:
  - A threshold θ_u is selected uniformly at random in $[0, 1]$.
  - Given the strategies a_R, a_B of the 2 firms, if $(a_{uR} = 0 \wedge a_{uB} = 0)$ then $s_{u0} = U$.
  - Otherwise $s_{u0} = R$ with probability $\frac{a_{uR}}{a_{uR}+a_{uB}}$ and $s_{u0} = B$ w.p. $\frac{a_{uB}}{a_{uR}+a_{uB}}$
- At any time step $t > 0$, each uncolored node u decides:
  1. whether to adopt some product based on the decisions of its neighbors up until step $t-1$, on its threshold, θ_u, and on a *switching function* f, described below.
  2. which product to adopt, in case that it decided to adopt some product. The choice of product is determined by a *selection function* g, also described below.

  Clearly, the process can last for at most $n - 1$ steps. We allow for fairly general functions f and g. In particular, let R_t (respectively B_t) denote the set of red (blue) nodes at step t and let $A_t = R_t \cup B_t$. Similarly, let $W_{ut}(R)$ (respectively $W_{ut}(B)$) denote the total weight of the edges (v, u) such that $v \in R_t$ (resp. B_t). Let also $W_{ut} = W_{ut}(R) + W_{ut}(B)$.
  1. The switching function applied at step t to node u, takes as argument the vector $\boldsymbol{C}_{ut} = (W_{u0},\ldots, W_{u,t-1})$, i.e., the whole history of how the cumulative weight of active neighbors has evolved in the previous steps. Node u switches from uncolored to colored at time t if
  $$f(\boldsymbol{C}_{ut}) \geq \theta_u$$
  2. The selection function takes as arguments the vectors $\boldsymbol{C}_{ut}(R) = (W_{u0}(R),\ldots, W_{u(t-1)}(R))$, and $\boldsymbol{C}_{ut}(B) = (W_{u0}(B), \ldots, W_{u,t-1}(B))$, i.e., the histories for the total weight of red and blue neighbors in the previous steps. Then with probability
  $$g(\boldsymbol{C}_{ut}(R), \boldsymbol{C}_{ut}(B)),$$
  node u selects the red product and $s_{ut} = R$. Else $s_{ut} = B$.

Note that the model can be easily extended to the case of $k > 2$ players.

Comparisons with related models: The model encompasses some families that have already been described before. For example, for linear f and g, and with $f(\boldsymbol{C}_{ut}) := f(W_{u,t-1})$, and $g(\boldsymbol{C}_{ut}(R), \boldsymbol{C}_{ut}(B)) := g(\frac{W_{u,t-1}(R)}{W_{u,t-1}})$, we have the *Weight-Proportional Competitive Linear Threshold Model* studied in [2].

Our model can be viewed as a threshold version of the models studied in [7, 5]. We allow more general switching and selection functions, in the sense that these functions can depend on how the total weight evolves over time. In [7, 5], these functions depend only on the active nodes at step $t-1$, when applied for step t. Finally, another technical difference is that we do not have any update schedule determining the order of updates. Instead, we consider that at each step any node that can switch to a colored state will do so by taking into account what has happened up until time $t - 1$.

Best Response Computation. As with other competitive diffusion models, such as [1, 2], our primary focus is on the problem of computing the best strategy for a firm,

given its opponent's strategy. Suppose we take the viewpoint of the Red firm. Given strategies a_R, a_B, we let $\sigma(a_R, a_B)$ denote the expected number of red nodes at the end of the diffusion process. The expectation here is over both the tie-breaking rule in the initialization phase and over the probabilistic choice of thresholds. We take this as the utility function of the red firm under this game.The problem we are interested then is:
The Influence Maximization Problem: Given a diffusion process, specifying the functions f and g, and given the strategy of the blue firm, a_B, find a strategy a_R for the red firm so as to maximize $\sigma(a_R, a_B)$.

3 Dependence of Selection Function only on New Influencers

In this section, we will focus on the case where

- The switching function f depends only on the aggregate weight of all the colored neighbors, up until the previous step. Hence, to check if a node u becomes colored at step t, we check if $f(W_{u,t-1}) \geq \theta_u$.
- The selection function depends on the set of nodes that became active exactly at the previous step of the process. In particular, at step t, the function g depends on the aggregate weights of colored nodes at the previous 2 steps, in the form $g := g(\frac{W_{u,t-1}(R) - W_{u,t-2}(R)}{W_{u,t-1} - W_{u,t-2}})$, and we also require that g is a linear function.

To see the motivation behind these types of switching and selection functions, one can think of the competition between two smartphones. The choice of the switching function is quite natural, and follows the recent works in the literature. E.g., the decision on whether to buy a smartphone or not, is affected by the set of all neighbors who have already bought one, regardless of which of the two products they have chosen. As for the selection function, the rationale is that a node may be more heavily influenced by the most recent buyers, i.e., the nodes that became active at the previous step in our model. If in the recent past more people made a choice towards one of the two products, then the node will have a higher probability to select the same product as well.

As we will see in Section 4, significant deviations from these assumptions make the algorithmic considerations that we are interested in more challenging.

Linear Switching Functions.

Our positive results concern the case where f is a linear function. In fact, we can assume WLOG that f and g are the identity function. We will refer to this as the LSM-STM model (*Linear Switching-Marginal Selection Threshold Model*). We conclude this section with a discussion regarding non-linear switching functions.

From now on, fix a strategy of the blue firm, say $a_B = (a_{1B}, a_{2B}, \ldots, a_{nB})$. We want to find a strategy $a_R = (a_{1R}, a_{2R}, \ldots, a_{nR})$ so as to maximize $\sigma(a_R, a_B)$. We will provide an approximation algorithm to this problem by using the standard tools of optimizing monotone and submodular functions.

Definition 1. *Consider a function $h : \mathbb{Z}^n \to \mathbb{R}$. Let $x, y \in \mathbb{Z}^n$ be two vectors with $x_j \leq y_j$ for every $j = 1, ..., n$. Let also $e_j \in \mathbb{R}^n$ be the unit vector with $e_j(j) = 1$ and $e_j(k) = 0$, for $k \neq j$. We will say that h is*

– *monotone, if* $h(x) \leq h(y)$*,*
– *submodular, if* $h(x + e_j) - h(x) \geq h(y + e_j) - h(y)$ *for* $j = 1, ..., n$*.*

Note that this is a generalization of the standard definition of submodularity, to the case of functions defined over multisets rather than sets, as defined also in [8].

We are interested in the expected number of red nodes at the end of the diffusion process as a function of the red firm's strategy, i.e., $\sigma(a_R, a_B)$ viewed as a function of a_R only. Our main result is the following:

Theorem 1. *Under the LSMSTM model, and for any given strategy* a_B *of the blue firm, the function* $\sigma(a_R, a_B)$ *is monotone and submodular.*

In order to use the machinery of [11] or [8] (for multisets), we also need to be able to compute the expectation $\sigma(a_R, a_B)$, for any strategies a_R, a_B. We can use sampling methods to approximate this value within any accuracy and as explained in [9], this suffices for the greedy algorithm of [11]. This implies the following corollary:

Corollary 1. *Under the LSMSTM model, and for any* $\epsilon > 0$*, there is a* $(1 - 1/e - \epsilon)$*-approximation algorithm for computing the best response of any player against her competitor.*

To prove Theorem 1, we will begin by showing that LSMSTM is equivalent to another model, which we will refer to as *Single Incoming Edge Analog* (SIEA). This is in a similar spirit as the approach via "live edges" in [9].

Definition 2. *(SIEA) Under this stochastic process, given* a_R, a_B*, the initialization phase is exactly the same as in LSMSTM. Then, for each node* u *we preserve at most one incoming edge. Node* u *selects the edge* $e = (v, u)$ *with probability* $w_{v,u}$ *and no edge w.p.* $1 - \sum_{v \in N(u)} w_{v,u}$*. We refer to the selected edges as live edges. Afterwards the contagion process works deterministically. At step* $t = 1$*, any node that has an incoming live edge from a colored neighbor, obtains the color of its neighbor. Continuing in this manner, at step* t*, any node that has an incoming edge from a colored node, becomes colored with the color of that node.*

A crucial observation is the following:

Lemma 1. *Given a pair of strategies* a_R, a_B*, the distributions over red-colored sets and blue-colored sets derived from running LSMSTM are the same as the distributions produced by SIEA.*

The proof of Lemma 1 is based on similar techniques as the proof of Claim 2.6 in [9]. From now on and till the end of the proof of Theorem 1, we will work only with the SIEA model. We first prove monotonicity[4].

Lemma 2. *Let* $a_R, a'_R \in \mathbb{Z}^n$ *such that* $a_R \leq a'_R$*. Under SIEA, and for any* a_B*,* $\sigma(a_R, a_B) \leq \sigma(a'_R, a_B)$*.*

[4] Note that in the case of a single product, monotonicity is trivial. This is not always the case in threshold models with at least two competing products. See e.g. [2] for some examples.

Proof. Consider 2 SIEA processes, π_1 and π_2 with $a_R^{\pi_1} = a_R$, $a_R^{\pi_2} = a'_R$, and $a_B^{\pi_1} = a_B^{\pi_2} = a_B$. We will prove that the expected number of red nodes at π_2 is at least as high as that in π_1.

We define a coupling between π_1 and π_2, and prove the lemma using induction on the number of steps. We consider the following coupled processes, which by slight abuse of notation, we will keep denoting by π_1 and π_2: We first pick randomly the set of live edges, as described in the SIEA model, which we take to be the same for both processes. At step $t = 0$, for every node u, where $a_{uR}^{\pi_1} + a_{uB} > 0$, we pick a number uniformly at random in $[0, 1]$ and we decide on the color of u at each process, based on the following 3 intervals of $[0, 1]$.

– with probability $\frac{a_{uB}}{a_{uR}^{\pi_2}+a_{uB}}$, we color u blue in both processes.
– with probability $\frac{a_{uB}}{a_{uR}^{\pi_1}+a_{uB}} - \frac{a_{uB}}{a_{uR}^{\pi_2}+a_{uB}}$, we color u blue in π_1 and red in π_2.
– with probability $1 - \frac{a_{uB}}{a_{uR}^{\pi_1}+a_{uB}}$, we color u red in both processes.

Any other node can be colored with no ambiguity in π_1 and π_2 or remain uncolored in one or both of the processes (e.g., if $a_{uR}^{\pi_2} = a_{uB} = 0$). The next steps in both processes continue as in the original SIEA processes (but note that both processes will use the same set of live edges).

It is quite straightforward to see that this is a valid coupling, since it produces the same distribution of blue and red nodes at each step t, as if we run the original processes. Indeed, at step $t = 0$, the probability that in π_1 a node u is colored blue is the probability that the result of the coin flip falls in one of the first two cases and hence equal to:

$$\frac{a_{uB}}{a_{uR}^{\pi_2} + a_{uB}} + \left(\frac{a_{uB}}{a_{uR}^{\pi_1} + a_{uB}} - \frac{a_{uB}}{a_{uR}^{\pi_2} + a_{uB}}\right) = \frac{a_{uB}}{a_{uR}^{\pi_1} + a_{uB}}$$

This is precisely the same for the original π_1 process without coupling. The same is true for the process π_2 and by induction we can then prove that the distributions of red and blue nodes is the same as in the uncoupled processes.

Coupling helps us in establishing the following claim, which trivially then implies monotonicity:

Claim. For the coupled processes π_1 and π_2, for every step t and for every node u, it holds that if $s_{ut}^{\pi_1} = R$, then $s_{ut}^{\pi_2} = R$.

Proof. We proceed by induction on the number of steps.
Induction basis: This is trivial by the construction of the coupling.
Induction step: Suppose that the claim holds until step $t - 1$. We will show that it holds for step t. For an arbitrary node u, suppose $s_{u,t}^{\pi_1} = R$. If it is the case that the node was colored in previous steps, then we would also have $s_{u,t-1}^{\pi_1} = R$. But by the induction hypothesis, then $s_{u,t-1}^{\pi_2} = R$, and hence, $s_{u,t}^{\pi_2} = R$. Now consider the case where node u becomes red in π_1 exactly at step t. This means that there is a live edge from a node v, and also $s_{v,t-1}^{\pi_1} = R$. But then by the induction hypothesis, $s_{v,t-1}^{\pi_2} = R$. Recall now that the coupled processes use the same set of live edges, and also that there can be at most one incoming live edge to a node u. Hence, node u cannot have possibly been colored in π_2 by some other live edge before step t. Thus u is uncolored in π_2 at step $t - 1$, and it will become red in π_2 as well, at step t. □

We established that for any random selection of live edges, the number of red nodes at the end of π_2 is at least as high as those in π_1. Hence the expected number of red nodes will also have the same property, i.e., the SIEA model satisfies monotonicity. $\square$

We now proceed to prove submodularity for our model.

Lemma 3. *Let $a_R, a'_R \in \mathbb{Z}^n$ such that $a_R \leq a'_R$. Under SIEA, and for any a_B, and any $j \in \{1, ..., n\}$, $\sigma(a_R + e_j, a_B) - \sigma(a_R, a_B) \geq \sigma(a'_R + e_j, a_B) - \sigma(a'_R, a_B)$.*

Proof. The proof is based on more involved coupling arguments than the case of monotonicity.

Consider 4 processes, π_1,π_2, π_3 and π_4 with the following features:

- $a_B^{\pi_1} = a_B^{\pi_2} = a_B^{\pi_3} = a_B^{\pi_4} = a_B$,
- $a_R^{\pi_1} = a_R$, and $a_R^{\pi_2} = a'_R$,
- $a_R^{\pi_3} = a_R^{\pi_1} + e_j$, and $a_R^{\pi_4} = a_R^{\pi_2} + e_j$.

Let $p_i = \frac{a_{uR}^{\pi_i}}{a_{uR}^{\pi_i} + a_{uB}}$ be the probability that node i is colored red at the initialization phase of process π_i, $i \in \{1, 2, 3, 4\}$. We consider now the following coupling between these processes: We pick at random a set of live edges as described under the SIEA model, which will be the same for all the processes. Then at step $t = 0$, for a node u with $a_{uR}^{\pi_1} + a_{uB} > 0$, we pick uniformly at random a number in [0, 1] and we decide on the color of u at each of the coupled processes, based on whether the number falls in one of 5 subintervals of $[0, 1]$, with lengths as defined below. In particular,

- With probability p_1, we paint node u red in all processes.
- With probability $(p_2 + p_3) - (p_1 + p_4) : s_{u0}^{\pi_2} = s_{u0}^{\pi_3} = s_{u0}^{\pi_4} = R \wedge s_{u0}^{\pi_1} = B$.
- With probability $p_4 - p_3 : s_{u0}^{\pi_2} = s_{u0}^{\pi_4} = R \wedge s_{u0}^{\pi_1} = s_{u0}^{\pi_2} = B$.
- With probability $p_4 - p_2 : s_{u0}^{\pi_3} = s_{u0}^{\pi_4} = R \wedge s_{u0}^{\pi_1} = s_{u0}^{\pi_3} = B$.
- With probability $1 - p_4$: we color u blue in all processes.

We can easily see that the probabilities above sum up to 1. It is also easy to check that this is indeed a valid coupling that produces the same distribution of blue and red nodes at each step t as if we run the original processes. For example, at step $t = 0$, the probability that in π_4 a node u is colored red is the probability that the result of the coin flip falls in one of the first four cases above and hence equal to:

$$p_1 + (p_2 + p_3) - (p_1 + p_4) + (p_4 - p_3) + (p_4 - p_2)$$

The above is equal to p_4, as desired. The same holds for the other processes as well. For nodes where, $a_{uR}^{\pi_1} + a_{uB} = 0$, we need to have an analogous (but simpler) construction, and the same holds for the case where $a_{uR}^{\pi_2} + a_{uB} = 0$. We omit the details for handling these simpler cases from this version.

The claim that we need in order to conclude our proof is the following:

Claim. For the coupled processes, for every step t and for every node u, it holds:

- $(s_{ut}^{\pi_4} = R) \Rightarrow (s_{ut}^{\pi_2} = R) \vee (s_{ut}^{\pi_3} = R)$.
- $(s_{ut}^{\pi_1} = R) \Rightarrow (s_{ut}^{\pi_2} = R) \wedge (s_{ut}^{\pi_3} = R)$.

Proof. Induction basis: The properties hold by the construction of the coupling.
Inductive step: Suppose the claim holds for step $t-1$. To see the first part of the claim, consider a node u with $s^{\pi_4}_{u,t} = R$. If it is the case that the node was colored in previous steps, then we would also have $s^{\pi_4}_{u,t-1} = R$. But by the induction hypothesis, then either $s^{\pi_2}_{u,t-1} = R$, and hence, $s^{\pi_2}_{u,t} = R$ or $s^{\pi_3}_{u,t-1} = R$, and hence, $s^{\pi_3}_{u,t} = R$. Now consider the case where node u becomes red in π_4 exactly at step t. This means that there is a live edge from a node v, and also $s^{\pi_4}_{v,t-1} = R$. But then by the induction hypothesis, $s^{\pi_2}_{v,t-1} = R$, or $s^{\pi_3}_{v,t-1} = R$. Recall now that the coupled processes use the same set of live edges, and also that there can be at most one incoming live edge to a node u. This means that node u cannot have possibly been colored in both π_2 and π_3 by some other live edge up until step $t-1$. Hence u will become red in π_2 or π_3 at t. This establishes the first part of the claim. The second part is established in a very similar way. □

It is easy to see that the claim implies submodularity of $\sigma(a_R, a_B)$. Hence this completes the proof. □

Remark 1. We can generalize the above results (the equivalence to SIEA as well as monotonicity and submodularity) for the case of $k > 2$ players. The selection function would still retain the same form, taking into account in the denominator the weight of all neighbors that were colored in the last step. To prove the same results say for player 1, we only need to consider that there is one Blue opponent with budget for node u equal to the sum of all other players' budgets for $i = 2, \ldots, k$. The intuition behind this is that for each player, the identity of her opponents does not make a difference. Hence, it is as if playing versus one Blue player that is the union of all other players.

Discussion about non-linear switching functions In the absence of competition, when the switching function is concave (and there is no selection function), monotonicity and submodularity hold [10]. This gives some indication that with such a switching function and with a linear selection function that is implemented just on the new adopters, the same properties may also hold. However, in the competitive setting, concave switching functions make the problem more challenging.

Firstly, the *live-edge* technique cannot be used in the case of a concave switching function. The reason for this is that the model ceases to be equivalent to *SIEA*. The activation time is more crucial now, and the unconditional probability of a node influencing a neighbor, depends on the order with which it will become active. The later it becomes active, the smaller the influence it will exert.

Secondly, the technique used by Mossel and Roch in [10] for the single product case cannot apply here. Their proof relies on the so called *antisense coupling* technique. A crucial point for the technique to apply is that the ordering with which the neighbors will get colored does not affect the outcome. This is not the case in the competitive setting as the nodes might get painted with different colors and the ordering affects the probability of a node getting colored with a particular color.

Despite the technical difficulty of dealing with this case, we conjecture that monotonicity and submodularity hold in the case of concave switching functions along with a linear selection function depending solely on new adopters. This would provide an interesting generalization of [10] in the concave setting with competition.

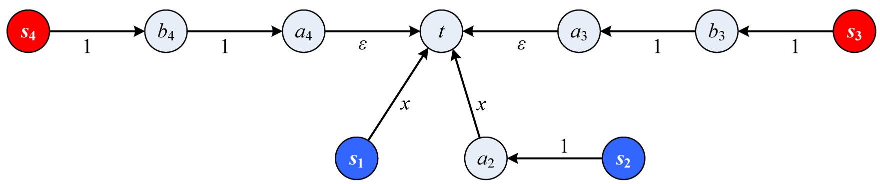

Fig. 1. Using this social network, we show that the utility function may not be submodular if the edge weights decrease in the diffusion process.

4 Necessity of Assumptions

Next, we justify the assumptions behind LSMSTM, by demonstrating that they are essentially necessary for the monotonicity and the submodularity of the influence functions. Specifically, we present examples showing that any significant deviation from LSMSTM yields a utility function that is nonmonotone or nonsubmodular (or both).

Monotonicity and Concavity of the Switching Function. Clearly, if the switching function is nonmonotone, the utility function need not be monotone. We also show here that the submodularity of the utility function requires that the switching function should be concave. For simplicity, we focus on the monopoly case with one product. Let the social network consist of 3 nodes s_1, s_2 and t and of two directed edges (s_1, t) and (s_2, t) with weights w_1 and w_2. Then, if the switching function f is strictly convex at some point, i.e., if there are w_1 and w_2 such that $f(w_1 + w_2) < f(w_1) + f(w_2)$, then the utility of the firm is not subadditive, and thus not submodular in such an instance.

Influence from the Neighbors in the Switching Function. Next, we show that if the edge weights decrease by an additive term of ε in the k-th step after their infection, the utility function is nonsubmodular. Thus, we demonstrate that submodularity requires that the edge weights, as taken into account by the switching function, should not decrease over time. Since we focus on models that do not depend on the node identities, we assume that this decrease takes place in any edge in the k-th step after its infection.

Let us consider the network in Fig. 1 where the blue firm selects nodes s_1 and s_2 and the red firm selects nodes s_3 and s_4. We assume that $k = 2$, i.e., the weight of each edge decreases by an additive term of ε in the second step after the edge's infection, that f satisfies $f(2x) < f(2x + \varepsilon)$, and that the selection function g is linear. Then, if t has not become blue by step 2 of the process, its threshold is larger than $f(2x)$. Then, in the third step, the weight of (s_1, t) decreases by ε and the total switching influence on t is $2x + \varepsilon$, if both s_3 and s_4 are selected by the red firm from the beginning, and at most $2x$, otherwise. Therefore, the probability that t becomes red is positive iff the red firm selects both s_3 and s_4 from the beginning. Thus, the utility function of the red firm is nonsubmodular in this case. Connecting s_2 to a_2 by a $(k - 1)$-chain of unit weight edges and connecting s_3 to a_3 and s_4 to a_4 by a k-chain of unit weight edges, we can generalize this example to the case where the edge weights decrease in the k-th step after their infection, for any $k \geq 2$. In fact, using similar in spirit (but more complicated) constructions, we can generalize this example to the case where the weight of each edge can decrease by a time dependent quantity in each step after the edge's infection.

Dependence of Selection Function on Previously Colored Nodes. Since we do not differentiate the nodes based on their identities, we can only differentiate them based on activation time. If the selection function considers not only the nodes colored in the last step, but also the nodes colored in previous steps, we can adjust the example in [2, Section 2] and show that the utility function may be nonmonotone and nonsubmodular.

(Almost) Linearity of the Selection Function. Finally, we observe that if the selection function g is highly convex at some point, i.e., if there exist some x_1, x_2, x_3 such that

$$g\left(\frac{x_1}{x_1+x_3}\right) f(x_1+x_3)+g\left(\frac{x_2}{x_2+x_3}\right) f(x_2+x_3) < g\left(\frac{x_1+x_2}{x_1+x_2+x_3}\right) f(x_1+x_2+x_3)\,, \quad (1)$$

then the utility function may not be submodular. This follows directly from (1) applied to a simple network with 4 nodes t, s_1, s_2, and s_3, and 3 directed edges (s_1, t), (s_2, t), and (s_3, t), with weights x_1, x_2, and x_3, respectively, where the blue firm selects s_3. The same argument shows that the selection function (of the red firm) $g(x)$ should not be highly concave, since otherwise, the selection function $1 - g(x)$ of the blue firm would be highly convex. Therefore, the selection function should be almost linear.

5 Performance of Equilibria

We conclude our work with a different and orthogonal question, namely studying the performance of Nash equilibria of the underlying game. We will present the analysis directly for an arbitrary number of competing firms, say k of them. For ease of presentation, we consider the case where the players choose a set rather than a multiset as their strategy to seed nodes.

Viewing the process as a game, we take as the utility of player i the expected number of nodes adopting product i at the end of the process. For a strategy profile $\boldsymbol{S} = (S_1, ..., S_n)$, we denote the payoff of i by $\sigma_i(\boldsymbol{S})$. Note that the nature of our switching function is such that the number of colored nodes at the end (independently of what color they chose), when starting from a strategy profile $\boldsymbol{S} = (S_1, ..., S_n)$ only depends on the set $S = \cup S_i$. Hence our social utility function can be defined simply over subsets of seeded nodes $S \subseteq V$, i.e., as $\gamma(S) = \gamma(\boldsymbol{S}) = \sum_i \sigma_i(\boldsymbol{S})$ where $\boldsymbol{S}$ can be any strategy profile that results in seeding S at step $t = 0$.

To quantify the Price of Anarchy of this game, we need to compare the values of $\gamma(\cdot)$ at the optimal seeding set against that at an equilibrium. For this we will use the approach of Vetta regarding utility games [13], also used by [1, 7]. We start with the definition of a utility game.

Definition 3. *Consider a game with k players, and a ground set V, so that the strategy space of each player are the subsets of V. Let $\gamma(S)$ be a social welfare function. A game is defined to be a utility game if it satisfies the following three properties:*

1. *The social utility function $\gamma(\cdot)$ is submodular.*
2. *Given a profile $\boldsymbol{S}$ resulting in a seeding set S, the total value for all the players is less than or equal to the total social value: $\sum \sigma_i(\boldsymbol{S}) \leq \gamma(S)$.*
3. *The value for a player i is at least her added value for the society: $\sigma_i(\boldsymbol{S}) \geq \gamma(\boldsymbol{S}) - \gamma(\boldsymbol{S}_{-i})$*

Theorem 2. *The LSMSTM model induces a utility game.*

The proof of this theorem is by establishing the three properties listed above. Note that Property 2 in Definition 3 is trivial. In fact, in our case it holds with equality. Hence, the main part of the proof is to ensure the first and the third property as well. For this we use the equivalence with the SIEA model, which facilitates the analysis (note that according to Remark 1, this equivalence holds for an arbitrary number of players). We omit further details from this version.

From the above theorem, using [13], we have:

Corollary 2. *The Price of Anarchy even for coarse correlated equilibria is at most* 2*.*

A modification of the tight example in [7] shows that our upper bound is tight as well.

Acknowledgements: We would like to thank Jon Kleinberg, Éva Tardos and Manolis Zampetakis for fruitful discussions and suggestions. This research has been cofinanced by the European Union (European Social Fund ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) Research Funding Program: **THALES**. Investing in knowledge society through the European Social Fund.

References

1. S. Bharathi, D. Kempe, and M. Salek. Competitive influence maximization in social networks. In *WINE 2007*, pages 306–311.
2. Allan Borodin, Yuval Filmus, and Joel Oren. Threshold models for competitive influence in social networks. In *WINE 2010*, pages 539–550.
3. Pedro Domingos and Matt Richardson. Mining the network value of customers. In *KDD 2001*, pages 57–66.
4. J. Goldenberg, B. Libai, and E. Muller. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing Letters*, page 211223, 2001.
5. Sanjeev Goyal and Michael Kearns. Competitive contagion in networks. In *STOC 2012*, pages 759–774.
6. M. S. Granovetter. Threshold models of collective behavior. *The American Journal of Sociology*, pages 1420–1443, 1978.
7. Xinran He and David Kempe. Price of anarchy for the n-player competitive cascade game with submodular activation functions. In *WINE 2013*, pages 232–248.
8. Michael Kapralov, Ian Post, and Jan Vondrák. Online submodular welfare maximization: Greedy is optimal. In *SODA 2013*, pages 1216–1225.
9. David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *KDD 2003*, pages 137–146.
10. Elchanan Mossel and Sbastien Roch. Submodularity of influence in social networks: From local to global. *SIAM J. Comput.*, pages 2176–2188, 2010.
11. G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, pages 265–294.
12. T. C. Schelling. *Micromotives and Macrobehavior*. 1978.
13. Adrian Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. pages 416–425.