

Threshold Models for Competitive Influence in Social Networks

Allan Borodin*, Yuval Filmus†, Joel Oren‡
Department of Computer Science
University of Toronto

October 3, 2010

Abstract

The problem of influence maximization deals with choosing the optimal set of nodes in a social network so as to maximize the resulting spread of a technology (opinion, product-ownership, etc.), given a model of diffusion of influence in a network. A natural extension of this would be to introduce a competitive setting, in which the goal is to maximize the spread of our technology in the presence of one or more competitors. We suggest several natural extensions to the well-studied linear-threshold model that was used in the single-technology case, and show that the original greedy approach cannot be used. Furthermore, we show that for a broad family of competitive influence models, it is NP-hard to achieve an approximation that is better than a square root of the optimal solution. Also, we show that the same proof of hardness of approximation can also be applied to give a negative result for a conjecture in [2] about a general cascade model for competitive diffusion. Finally, we suggest a natural model that is amenable to the greedy approach.

1 Introduction

The problem of influence maximization has long been the focus of study in social science (e.g. [5]). It has been formally defined and addressed in [6] and [3] as follows: given a social-network as a directed graph, and a prescribed number k , pick the k most “influential” nodes that will function as early adopters of a particular influence, so as to maximize the final number number of *infected*, or *activated* nodes (the two terms are used in this paper interchangeably), subject to a specified model of influence diffusion. This problem begs the natural extension of a competitive version: given the competitor’s choice of early adopters of technology B , maximize the spread of technology A by choosing a set of early adopters such that the expected spread of technology A will be maximal. Indeed, this problem has been the subject of interest in subsequent papers ([2],[7],[1]) which present competitive extensions for the independent cascade model presented in [6]. In this paper we suggest a few natural extensions to a well-studied model that was also given in [6] for the diffusion of social influence in a social network — the linear threshold model. Formally, an instance of the problem would be composed of a directed, edge-weighted graph $G = (V, E)$, a set of technology B ’s initial adopters $I_B \subseteq V$, and an integer k . The computational problem is how to choose a set $I_A \subseteq V - I_B$

*bor@cs.toronto.edu

†yuvalf@cs.toronto.edu

‡oren@cs.toronto.edu

of k nodes such that the *expected* number of A -active nodes at the end of the process, $\sigma(I_A, I_B)$, is maximized, given the specific model for competitive diffusion of technologies (when I_B is known from context we omit it and simply write $\sigma(I_A)$). It is important to note that all of the presented models can be motivated by natural processes. Our models reduce to the original linear threshold model formalized in [6] whenever I_B is the empty set. For simplicity of notation, the models presented are stated in terms of only two competing technologies. However, all of the models and results can be easily extended for when there are several competing technologies.

A well-known greedy $(1 - e^{-1})$ -approximation given in [4] is used extensively for problems of maximizing set-functions, and in particular has been applied to both the original problem and competitive extensions ([2], [1]). The approximation algorithm requires that the set function $\sigma(\cdot)$ at hand, which assigns a real-value to subsets of a ground set U , uphold two basic properties¹.

- **Monotonicity:** the value of the function increases as more items are added to the set: $\sigma(S) \leq \sigma(T)$ for any two sets $S \subseteq T$;
- **Submodularity:** the impact of adding an element to a set decreases as the set is extended (diminishing returns): $\sigma(S \cup \{x\}) - \sigma(S) \geq \sigma(T \cup \{x\}) - \sigma(T)$, for any $S \subseteq T \subseteq U$ and $x \in U - T$;

Except for the last model, described in section 6 — the *OR* model, all of the models do not satisfy submodularity. In fact, one of them is not even monotone.

Outline The remainder of this document is organized as follows. Sections 2 and 3 describe two competitive threshold models. Section 4 shows that even when applying a final step that A -activates more nodes, the process remains non-submodular. Section 5 shows that the last two models are in general hard to approximate, regardless of the final forcing step or any tie-breaking rules that are used when a node is reached by more than one technology. On a more positive note, in section 6 we suggest a fairly natural and simple model for which the approximation algorithm given in [4] is applicable. Finally, section 7 summarizes our main results along with a few open problems and possible directions for future research.

2 The Weight-Proportional Competitive Linear Threshold Model

As in the non-competitive case, the process unfolds in discrete steps, during which new nodes become “activated” for a single technology² The infection of a node by a technology represents an individual in the social network that has assumed the influence of that technology. The process is *progressive*: a node that is infected by a technology remains infected by it. As in the non-competitive case, every edge (u, v) is assigned a weight $w_{u,v} \in [0, 1]$ which roughly characterizes the weight of influence that u has over v (i.e the impact that u ’s infection will have over v ’s likelihood to be infected with the same technology as u). Also, the total weight of incoming edges to every node is bounded: for every $v \in V$ we have: $\sum_u w_{u,v} \in [0, 1]$. Each node u initially chooses a threshold

¹There are constant-factor approximations algorithms for cases when submodularity holds but monotonicity does not hold

²The term “technology” stands for any concept or influence that spreads in the social-network (car ownership, club membership, voting preference, etc.).

θ_u which represents the minimum fraction of active neighbours necessary for u 's activation. As in [6], in order to make up for our lack of knowledge about each node we assume that $\theta_u \in_R [0, 1]$ (uniformly at random), or $\theta_u \in_R [a, a']$ for $0 \leq a \leq a' \leq 1$ to reflect partial knowledge about a node.

In order to describe the process itself, we will use the following notation:

Definition 2.1. For a given step t in the process, let Φ^t denote the set of active nodes at the beginning of step t . Furthermore, let Φ_A^t and Φ_B^t be the sets of A -active and B -active nodes in step t , respectively.

Given the sets I_A, I_B of early technology adopters, the process unfolds as follows. First, each node chooses its threshold value at step 0. Then, in each step t , every inactive node v checks the set of edges incoming from its active neighbours. If their collective weight exceeds the threshold values by having $\sum_{u \in \Phi^t} w_{u,v} \geq \theta_v$, the node becomes active. In that case, the node will adopt technology A with probability equal to the ratio between the collective weight of edges outgoing from A -active neighbours and the total collective weight of edges outgoing from all active neighbours; that is,

$$Pr[v \in \Phi_A^t | v \in \Phi^t \setminus \Phi^{t-1}] = \frac{\sum_{u \in \Phi_A^t} w_{u,v}}{\sum_{u \in \Phi^t} w_{u,v}} \quad (1)$$

Otherwise, it will adopt technology B . Since this problem can be reduced to the single-technology linear threshold model whenever I_B is the empty set, we notice that this problem is NP-hard — as proved in [6].

Intuitively, it appears that by adding a new node to the set of initial A -adopters, the spread of technology A in the social network can only increase (or remain unchanged). However, this is in fact not always the case, even for some binary rooted trees. We will formalize this somewhat counter-intuitive behaviour.

Theorem 2.1. *There exists an instance of the weight-proportional competitive linear threshold problem for which monotonicity does not hold.*

Also, it can be shown that submodularity fails to hold in some cases, as the following theorem shows:

Theorem 2.2. *There exists a graph G , for which the expected influence of technology A is not submodular.*

The proof of the above two theorems is given in appendix A.

3 The Separated-Threshold Model for Competing Technologies

In the previous model, the node activation step regarded active nodes as equal, so that a given node is activated by its active neighbours regardless of their technologies. That is, the sum of generally active nodes was used for activating a node. However, one could model the following notion of a spread process. Each individual has separate criteria for becoming active for each technology. A node can be activated by either its A -active or B -active neighbours whenever the sums of their respective edge-weights exceed the required thresholds specified for their technologies.

Formally, consider the following model. For a given network $G = (V, E)$, every edge $(u, v) \in E$ is assigned a real-valued weight corresponding to each technology $w_{u,v}^A, w_{u,v}^B \in [0, 1]$ such that $\sum_u w_{u,v}^A, \sum_u w_{u,v}^B \in [0, 1]$, which reflects node u 's impact on v . Two *disjoint* sets $I_A^0, I_B^0 \subseteq V$ denote the sets of initially A -active and B -active nodes, respectively. At step 0, each node $v \in V$ picks two threshold values $\theta_v^A, \theta_v^B \in_R [0, 1]$. For step t , denote I_A^{t-1}, I_B^{t-1} as the sets of A -active and B -active nodes. During every step t , an inactive node v will become A -active if $\sum_{u \in I_A^{t-1}} w_{u,v}^A \geq \theta_v^A$, and will become B -active if $\sum_{u \in I_B^{t-1}} w_{u,v}^B \geq \theta_v^B$. If for the node v both thresholds are exceeded during the same step t , then a technology would be picked uniformly at random (we can either use a simple coin-flip or employ a more involved tie-breaking function).

In contrast to the previous model, this model *is* monotone. Its key property, which distinguishes it from the previous model, is that the probability that technology B will propagate cannot increase as a result of A -activating additional nodes. This stems from the definition of the model, in which each set of technology specific neighbours relate to a separate threshold value.

Let us use the following notation:

Definition 3.1. Given the sets I_A and I_B , and a node $x \notin I_B$, define $\alpha_v^t, \hat{\alpha}_v^t$ as the probabilities that a given node v will adopt technology A by step t for the initial sets of early adopters (I_A, I_B) and $(I_A \cup \{x\}, I_B)$, respectively. Similarly, define similar probabilities $\beta_v^t, \hat{\beta}_v^t$ for technology B .

Theorem 3.1. *For a given instance of the problem and a choice of early adopters: I_A, I_B and node x , $\hat{\alpha}_v^t \geq \alpha_v^t$ for any node v and for any step $t \geq 0$.*

The proof of theorem 3.1 is fairly straightforward, and is given in appendix B for completeness. The process is not submodular in general.

Theorem 3.2. *There exist instances of the competitive influence problem where the separated-threshold competitive model is not submodular.*

A corresponding counter-example along with the intuition behind it are given in appendix C.

4 Competitive Threshold Model with Forcing

We now introduce a modification which changes the concept of influence in a network: *forcing*. Specifically, at the end of the previous model, each inactive node v will choose a technology randomly (say, it will choose technology A with probability δ). This step is natural for cases where individuals have to eventually decide which influence to adopt (e.g. voting when abstentions are not allowed). For convenience we will assume that the “forcing” step occurs at step n (the spread can take up to $n - 1$ steps), whether or not the spread took $n - 1$ steps. Clearly this does not have any effect on the outcome of the process. We show that regardless of the forcing step, this variant does not help us achieve submodularity. However, the process remains monotone as the following theorem can be proven by extending lemma 3 (in appendix C) with a case analysis for the forcing step.

Theorem 4.1. *For a given instance of the competitive influence with forcing problem, a choice of early adopters I_A, I_B and node x , $\hat{\alpha}_v^t \geq \alpha_v^t, \hat{\beta}_v^t \leq \beta_v^t$ for any node v and for any t .*

4.1 Submodularity

The following theorem shows that not only is the given model non-submodular, but also regardless of the tie-breaking rule and the forcing rule (if any is used), the model remains non-submodular.

Theorem 4.2. *For any tie-breaking rule, and any forcing rule, the separated-threshold competitive model is non-submodular.*

A corresponding counter-example is given in appendix D.

5 Hardness of Approximation

We show that in any model with separate edge-weights and separate threshold values for each technology the problem is hard to approximate.

Theorem 5.1. *It is NP-hard to give an approximation with a ratio better than $\Omega(N^{\frac{1}{2}-\epsilon})$, for all $\epsilon > 0$, for the Separated-Threshold Competitive Influence problem, where N is the number of nodes in the graph.*

The proof is supplied in appendix E.

It is important to note the proof of theorem can be applied to similar competitive cascade models as well. Namely, in [2] it was conjectured that when allowing 2 sets of edge weights for each edge — one for each technology, the process will remain monotone and submodular. The above hardness of approximation result can be modified in order to apply for the separate edge-weights case of the Wave Propagation model suggested by Carnes et al., thereby giving a negative answer to their conjecture.

Theorem 5.2. *It is NP-hard to give an approximation with a ratio better than $\Omega(N^{\frac{1}{2}-\epsilon})$, for all $\epsilon > 0$, for the Wave Propagation Competitive Influence problem given by Carnes when edges are allowed to have technology-specific probabilities.*

6 The OR Model

We now introduce a different way of extending the original threshold model, in which each technology diffuses unhindered by the competing technology. Here, the tie-breaking stage will take place after all technologies finish spreading. This model can be considered natural for cases in which individuals have the liberty of being influenced separately and independently by different technologies, but have to choose a single technology eventually.

We will define the *OR* model as follows. As before, an instance of the model is a graph $G = (V, E)$, a set of edge weights for each technology: $W_A = \{w_{u,v}^A\}_{(u,v) \in E}$, $W_B = \{w_{u,v}^B\}_{(u,v) \in E}$ (with the same constraints as before), and initial adopters: $I_A, I_B \subseteq V$. Additionally, for each node $v \in V$ two “decision” functions $f_v^A : 2^V \times 2^V \rightarrow [0, 1]$, $f_v^B : 2^V \times 2^V \rightarrow [0, 1]$ are assigned. Let each technology propagate separately throughout the graph w.r.t the original non-competitive linear threshold propagation process, and let $R_A, R_B \subseteq V$ be the sets of nodes reached by the

technologies (independently). As a final step, a node $v \notin I_A \cup I_B$ will assume technology A with probability $f_v^A(R_A, R_B)$, technology B with probability $f_v^B(R_A, R_B)$, and no technology with probability $1 - f_v^A(R_A, R_B) - f_v^B(R_A, R_B)$, respectively ($f_v^A(R_A, R_B) + f_v^B(R_A, R_B) \leq 1$). We only require the functions $f_v^A(\cdot, \cdot)$, for every $v \in V$, to be monotone and submodular with respect to the set of initial A nodes.

The following theorem shows that one can efficiently find an approximation for the set that maximizes the spread of one's own technology, given a competitor[s] choice of initial adopters:

Theorem 6.1. *Given technology B 's early adopters I_B , one can find an $(1 - e^{-1} - \epsilon)$ -approximation for the competitive OR process in a polynomial number of steps, for any $\epsilon > 0$.*

The proof follows immediately from the following two lemmas which prove the properties required in [4]. We will show that this process is monotone and submodular whenever the function $f_v(\cdot, \cdot)$ is monotone and submodular w.r.t. technology A , for all $v \in V$.

Lemma 1. The OR model is monotone with respect to the number of nodes influenced by technology A .

Proof. Let $r_A(I_A), r_B(I_B)$ denote an outcome for a run of the independent propagation processes of the two technologies. Monotonicity w.r.t technology A is satisfied if for any two sets $S \subseteq T \subseteq V - I_B$:

$$\mathbb{E}[f_v^A(r_A(S), r_B(I_B))] \leq \mathbb{E}[f_v^A(r_A(T), r_B(I_B))] \quad (2)$$

Since until the decision step the two technologies' propagations are independent, we can fix the outcome of technology B , and show that the expected propagation of technology A is monotone. This is immediate since first, the propagation of technology A until the decision step is clearly monotone (follows from the non-competitive threshold model in [6]). Second, the decision functions $f_v^A(\cdot, \cdot)$ and $f_v^B(\cdot, \cdot)$ are monotone with respect to technologies A and B , which along with the previous argument yields monotonicity. \square

Lemma 2. The OR model is submodular with respect to the number of nodes influenced by technology A .

Proof. In order to prove this, we will use a technique given in [6] that suggests an alternative and equivalent model for the propagation of a single technology. For each node $v \in V$, instead of choosing a threshold in $[0, 1]$, choose an incoming edge (u, v) with respective probability $w_{u,v}$, and no incoming edge with probability $\sum_u w_{u,v}$. A node will become infected if and only if there is a path from the initially infected nodes that consists strictly of such chosen edges. Originally, this equivalent process was used for showing submodularity for the non-competitive case by simply fixing an instantiation of chosen edges. We will do the same in our case. Fix an instantiation R_B of the outcome of the propagation of technology B (independent of the propagation of technology A) and a set of chosen edges for the propagation process of technology A . For a set I_A of initial A nodes, as before, let $R_A(I_A)$ denote the set of nodes reachable from I_A in the sub-graph induced by the set of chosen edges. In order to show that the process is submodular, we need to show that for all $S \subseteq T \subseteq V - I_B$:

$$f_v^A(R_A(S \cup \{x\}), R_B) - f_v^A(R_A(S), R_B) \geq f_v^A(R_A(T \cup \{x\}), R_B) - f_v^A(R_A(T), R_B), \quad (3)$$

for all $v \in V$. We will simply use the monotonicity property of the independent propagation process and the submodularity of $f_v^A(\cdot, \cdot)$. Let $R_A(S \cup \{x\}) = R_A(S) \cup \Delta_1$, and similarly, $R_A(T \cup \{x\}) = R_A(T) \cup \Delta_2$. From the monotonicity and submodularity we get that $R_A(S) \subseteq R_A(T)$ and $\Delta_2 \subseteq \Delta_1$. Therefore:

$$\begin{aligned} & f_v^A(R_A(S) \cup \Delta_1, R_B) - f_v^A(R_A(S), R_B) \\ & \geq f_v^A(R_A(T) \cup \Delta_1, R_B) - f_v^A(R_A(T), R_B) \\ & \geq f_v^A(R_A(T) \cup \Delta_2, R_B) - f_v^A(R_A(T), R_B) \end{aligned} \quad (4)$$

The first inequality and second inequality follow from the submodularity and the monotonicity of $f_v^A(\cdot, \cdot)$, respectively. Taking all possible instantiations gives submodularity since a positive linear combination of submodular functions is submodular. \square

Mossel et al. [8] show that if we generalize the propagation process by replacing the linear sum (used to decide whether an uninfected node exceeds its threshold) with an arbitrary monotone submodular function, then the resulting process (under any monotone submodular objective function) is again monotone and submodular. This result generalizes to the corresponding competitive process, which we call the *generalized OR process*.

Theorem 6.2. *Given technology B 's early adopters I_B , one can find an $(1 - e^{-1} - \epsilon)$ -approximation for the generalized competitive OR process in a polynomial number of steps, for any $\epsilon > 0$.*

Proof. Use the objective function $\varphi(R_A) = \mathbb{E}_{R_B} \sum_v f_v^A(R_A, R_B)$ in the main result of [8]. The function φ counts the expected number of A -adopters at the end of the process. It is monotone and submodular because the f_v^A are. \square

6.1 Repeating OR processes

Finally, we give a natural extension of the *OR* model. There are cases in which the independent propagation process will repeat several times (e.g. every day, for ℓ days). The process can be thought of as being run iteratively, where during each iteration i the previous iteration's turnouts R_A and R_B are used as the initial adopters for each technology. At the end of the ℓ 'th iteration, *and only then*, the decision step is invoked by using the functions $f_v^A(\cdot, \cdot)$ and $f_v^B(\cdot, \cdot)$, for all $v \in V$. One may notice that this formulation simply defines a composition of ℓ *OR* processes (with a single execution of the decision step at the end).

We can give a natural motivation for such a process: during the course of an election race, voters will spread the word each day. However, once in while, an unaffected voter may change her mind (her threshold value) and thus the process of "rumor spread" and social-based recommendation will run again, infecting additional voters as a result.

With this in mind, the following general theorem follows from a simple generalization of the proof in [8].

Theorem 6.3. *A process based on the repetitive execution of the generalized OR process with a single decision step at the end is monotone and submodular.*

Note that theorem 6.3 holds even if the edge weights are modified between each iteration.

7 Conclusions

We have presented a number of fairly natural and general approaches for modelling competitive diffusion of influence in a social network, extending the known threshold model for the spread of a single technology. However, most of our suggested approaches have been shown to be unfit for the Nemhauser et al. [4] approximation technique. For some models, we can show NP-hardness of approximation, while for others we only show that they are not submodular (and not even monotone in one case), leaving open the question to whether an efficient approximation algorithm can be found.

We emphasize that all of the suggested models in this paper have reasonable, natural motivations, which implies that there is no single “true” model. Also, as suggested in [2], we believe that these models can be used in a more game theoretic setting, where players are the competing companies, who place bids on strategic nodes in hope for maximizing their outcome. We suggest the following directions for future research:

- Can the hardness-of-approximation result be extended to other models?
- Are there any other natural competitive models which can be approximated in polynomial time?
- Study some natural game-theoretic setting for the competitive models.
- Suggest models for cases where nodes may adopt more than one technology.

8 Acknowledgments

We would like to thank Justin Ward for reviewing an early draft of this paper. We also thank Akitoshi Kawamura for helping with the counter example in the proof of Theorem 2.1.

References

- [1] S. Bharathi, D. Kempe, and M. Salek. Competitive influence maximization in social networks. In *WINE*, pages 306–311, 2007.
- [2] T. Carnes, C. Nagarajan, S. M. Wild, and A. van Zuylen. Maximizing influence in a competitive social network: a follower’s perspective. In *ICEC '07: Proceedings of the ninth international conference on Electronic commerce*, pages 351–360, New York, NY, USA, 2007. ACM.
- [3] P. Domingos and M. Richardson. Mining the network value of customers. In *KDD '01: Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 57–66, New York, NY, USA, 2001. ACM.
- [4] M. F. G. Nemhauser, L. Wolsey. An analysis of the approximations for maximizing submodular set functions. In *Mathematical Programming*, 14, pages 265–294, 1978.
- [5] P. R. J. Brown. Social ties and word-of-mouth referral behavior. In *Journal of Consumer Research* 14:3, pages 350–362, 1987.




Figure 1: A counter-example for monotonicity

- [6] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *KDD '03: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, New York, NY, USA, 2003. ACM.
- [7] J. Kostka, Y. A. Oswald, and R. Wattenhofer. Word of mouth: Rumor dissemination in social networks. In *SIROCCO*, pages 185–196, 2008.
- [8] E. Mossel and S. Roch. Submodularity of influence in social networks: From local to global. *SIAM J. Comput.*, 39(6):2176–2188, 2010.

A Counter Examples for The Weight-Proportional Competitive Linear Threshold Model

In section 2 we gave two theorems concerning the monotonicity and submodularity of the model described. These theorems will be proven in this appendix.

Theorem A.1. *There exists an instance of the weight-proportional competitive linear threshold problem for which monotonicity does not hold.*

Proof. Consider the tree in figure A. We will explicitly calculate the values of α_v^t and $\hat{\alpha}_v^t$. First, consider the case in which $\theta_r \leq 0.5$. Then r will be A -active if one of the following occurs:

1. v becomes A -active during step 2.
2. v does not become active (hence r will be A -activated in step 4).

Before adding x , the first case cannot happen, and the second case happens only when $\theta_v \in (0.9, 1]$. Hence, $\alpha_r^4|_{\theta_r \leq 0.5} = \frac{1}{10}$. After A -activating x , the first case may occur only when $\theta_v \in (0.9, 1]$, and the second case cannot happen. Thus, $\hat{\alpha}_r^4|_{\theta_r \leq 0.5} = \frac{1}{10} \cdot \frac{1}{10} = \frac{1}{100}$.

Consider the case in which $\theta_r > 0.5$. The r will be A -active if one of the following occurs:

1. v becomes B -active in step 1, but r becomes A -active in step 4.

2. v becomes B -active in step 2, but r becomes A -active in step 4.
3. v becomes A -active in step 2.

Before A -activating node x , only the first case may occur, and only in cases where $\theta_v \in [0, 0.9]$. So $\alpha_r^4|_{\theta_r > 0.5} = \frac{9}{10} \cdot \frac{1}{2} = \frac{9}{20}$. After A -activating x , the first case may occur only when $\theta_v \in [0, 0.9]$, and cases (2) and (3) may occur only when $\theta_v \in (0.9, 1]$:

$$\hat{\alpha}_r^4|_{\theta_r > 0.5} = \frac{1}{2} \cdot \frac{9}{10} + \frac{1}{2} \cdot \frac{1}{10} \cdot \frac{9}{10} + \frac{1}{10} \cdot \frac{1}{10} = \frac{101}{200} \quad (5)$$

Putting the two cases together gives

$$\alpha_r^4 = \frac{1}{2} \cdot \alpha_v^r|_{\theta_r \leq 0.5} + \frac{1}{2} \cdot \alpha_r^4|_{\theta_r > 0.5} = \frac{1}{2} \cdot \frac{1}{10} + \frac{1}{2} \cdot \frac{9}{20} = \frac{11}{40} \quad (6)$$

$$\hat{\alpha}_r^4 = \frac{1}{2} \cdot \hat{\alpha}_r^4|_{\theta_r \leq 0.5} + \frac{1}{2} \cdot \hat{\alpha}_r^4|_{\theta_r > 0.5} = \frac{1}{2} \cdot \frac{1}{100} + \frac{1}{2} \cdot \frac{101}{200} = \frac{103}{400} \quad (7)$$

□

Essentially, this counter-example illustrates an interesting behaviour of the process. First, notice the high ratio of the weight of the edge connecting the B -subtree to the weight of the subtree containing x . Also, notice that the initially B -active node is located higher relatively to initially A -active nodes (only x in the given simple example). These two properties increase the probability that the parent node (v in our case) will adopt technology B , much more, and much more quickly than the probability that it will adopt technology A . This in effect, contributes to the decrease in the probability that the higher node (r in our case) will adopt A . Another point, is that this example shows that in less pathological cases, a node's probability of becoming A -active may increase, but this increase may come at the expense of a much higher increase in the probability of becoming B -active. In other words, by A -activating certain nodes, actually the competitor will benefit the most.

Theorem A.2. *There exists a graph G , for which the expected influence of technology A is not submodular.*

Proof. Consider the tree depicted in Figure A.

It can be easily shown that for $S = \{w\}, T = \{w, y\}$ (the set of early adopters of technology B is denoted in the diagram) submodularity does not hold as $\alpha_r^4(S) = \frac{3}{10}, \hat{\alpha}_r^4(S) = \frac{17}{60}, \alpha_r^4(T) = \frac{7}{10}, \hat{\alpha}_r^4(T) = \frac{17}{50}$. □

B Proof of The Monotonicity of The Separated Threshold Model

Theorem B.1. *For a given instance of the problem and a choice of early adopters: I_A, I_B and node x , $\hat{\alpha}_v^t \geq \alpha_v^t$ for any node v and for any step $t \geq 0$.*

In order to prove the theorem, we will fix the set I_B of early technology B adopters and consider a set of early technology A adopters I_A and a node x not in I_B .

We prove the monotonicity by fixing an arbitrary instantiation of the thresholds vectors, $\Theta_A \in [a_1^1, a_1^2] \times \dots \times [a_n^1, a_n^2], \Theta_B \in [b_1^1, b_1^2] \times \dots \times [b_n^1, b_n^2]$, and by choosing for every node technology A with




Figure 2: A counter-example for submodularity

probability 0.5 and choosing B otherwise. The outcomes of the tie-breaking fair coin flips will be revealed in cases where the two thresholds chosen for a particular node are exceeded simultaneously. Notice that this defines a deterministic instantiation of the process.

Denote by π_1, π_2 to be the deterministic processes using the same instantiations of the threshold values and coin-flips, and using (I_A, I_B) and $(I_A \cup \{x\}, I_B)$, respectively. Furthermore, let $N_A^t(\pi), N_B^t(\pi)$ denote the set of A and B active nodes at step t in process π , respectively. The following lemma implies theorem 3.1

Lemma 3. For a fixed instantiation of the thresholds values and the tie-breaking coin flips for each node in advance, the following holds for each node $v \in V$ and every step $t \geq 0$

1. If v is not B -active at step t in π_1 , then v is not B -active at any step $t' \leq t$ in π_2 .
2. If v is A -active at step t of π_1 , then v is activated in some step $t' \leq t$ in π_2 .

Proof. We will prove the lemma by induction on t . For step $t = 0$, the lemma trivially holds since $I_A \subseteq I_A \cup \{x\}$. Assume that the lemma holds up to step t , we will prove the lemma for step $t + 1$. For proving the first part of the lemma, consider any node v that is not B -active at step $t + 1$ in π_1 . By the induction hypothesis, node v was not B -active at step t in π_2 . Moreover, by the induction hypothesis, none of the neighbours of v that were not B -active at step t in π_1 are B -active at step t in π_2 . Therefore, $\sum_{u \in N_B^t(\pi_1)} w_{u,v}^B \geq \sum_{u \in N_B^t(\pi_2)} w_{u,v}^B$, which implies the first part of the lemma. For proving the second part of the lemma, let v be a node that is A -activated at step $t + 1$ in π_1 . Consider the following cases:

1. $\sum_{u \in N_B^t(\pi_1)} w_{u,v}^B < \theta_v^B$: then from the first part of the lemma $\sum_{u \in N_B^t(\pi_2)} w_{u,v}^B < \theta_v^B$ as well. On the other hand, since $\sum_{u \in N_A^t(\pi_1)} w_{u,v}^A \geq \theta_v^A$, the induction hypothesis implies that $\sum_{u \in N_A^t(\pi_2)} w_{u,v}^A \geq \theta_v^A$, which implies that v is A -active by step $t + 1$.
2. $\sum_{u \in N_B^t(\pi_1)} w_{u,v}^B \geq \theta_v^B$: This implies that both thresholds were exceeded at step $t + 1$. From part 1 of the lemma, θ_v^B cannot be exceeded prior to step $t + 1$ in π_2 . By the induction hypothesis $N_A^t(\pi_1) \subseteq N_A^t(\pi_2)$, which implies that $\sum_{u \in N_A^t(\pi_2)} w_{u,v}^A \geq \theta_v^A$. The above observations imply




Figure 3: A counter-example for submodularity for the case of separated thresholds.

that whereas θ_v^B may not be exceeded at, or before step $t + 1$ in π_2 , θ_v^A is exceeded by step $t + 1$ of π_2 . Since the same instance of coin-flip is used for π_1 and π_2 , v will be A -active even in the case that θ_v^A and θ_v^B are exceeded in step $t + 1$ of π_2 simultaneously. □

C Counter-Examples for The Separated Threshold Model

Theorem C.1. *There exist instances of the competitive influence problem where the separated-threshold competitive model is not submodular.*

Proof. Adding a node x to a set of initial A adopters may directly increase the probability that a node v will be A -activated. This can be seen in the alternative (equivalent) model case: it may lie on the A -edge path ending in v , thereby shortening the existing shortest path or create a path which didn't exist prior to the addition of x . However, it may indirectly increase the probability of A -activation by simply “eliminating” the competition as a result of blocking the propagation of technology B . Consider the tree in Figure 3 (the first number in the edge annotation is the A -weight, and the second one is the B weight).

Let the two sets of initial A -adopters be: $S = \{\}, T = \{v\}$. For this example we will consider the original threshold model.

- Before adding x to S and T : r had zero probability of becoming A -active when using S , and probability 0.5 when using T .
- After adding x to S and T : The probability that r will become A -active remains 0 when x is added to S (though its probability of becoming B active decreased to 0). When using T on the other hand, the probability is now 1, since by adding x technology B is blocked.

Note: The above counter-example also holds for a model in which the tie-breaking step is a little different: for example if for a node v the two thresholds are exceeded simultaneously, the technology with the highest sum of respective edge-weights is chosen. If all sums of edge-weights are equal, a technology is picked uniformly at random. □

D Counter-Examples for The Competitive Threshold Model with Forcing

Theorem D.1. *For any tie-breaking rule, and any forcing rule, the separated-threshold competitive model is non-submodular.*




Figure 4: A Counter-example for submodularity when applying a forcing step

Proof. We will give a counter-example in which there are no ties, and the node in question does not remain inactive. Consider the rooted tree in Figure 4. Let $S = \{u\}, T = \{u, y\}$. The initially B -activated nodes are given in the diagram.

Notice that no ties are possible — any uninfected node may be infected by at most one node at any given step. Also notice, that the forcing step at the end would not apply to the root since if its left child remains uninfected or is B -infected, its right child would eventually infect it with technology B . We get a “race condition” in which we have to block B ’s propagation in the left subtree of the root in order to get r A -infected.

When using S , $\alpha_r^6(S) = \hat{\alpha}_r^6(S) = 0$ since r ’s left child is infected with B first. On the other hand, $\alpha_r^6(T) = 0$ and $\hat{\alpha}_r^6(T) = 1$ since technology B ’s propagation in r ’s left child is completely blocked, allowing technology A to propagate to r and infect it before B reaches it through r ’s right child. Therefore, the probability jumps from 0 to 1, contradicting submodularity. \square

E Proof of The Hardness of Approximation Result

Theorem E.1. *It is NP-hard to give an approximation with a ratio better than $\Omega(N^{\frac{1}{2}-\epsilon})$, for all $\epsilon > 0$, for the Separated-Threshold Competitive Influence problem, where N is the number of nodes in the graph.*

Proof. We are motivated by the counter-example in theorem 4.2. We construct a reduction from vertex cover problem. Recall that in the vertex cover problem the problem is to find a minimal set S of vertices in the graph G such that every edge in G has at least one endpoint in S . In the separated-threshold competitive influence problem the task is to find a set S' of k' nodes that maximizes $\sigma(S', I_B)$ (expected number of A -infected nodes), given a prescribed number k' and the competitor’s set of nodes I_B .




Figure 5: The reduction. The diagram should be repeated once for each $t \in \{1, \dots, n^\alpha\}$ (except for the A_x vertices).

The reduction We are given an instance of the Vertex Cover problem, a graph $G = (V, E)$ and a number k . Let α, β be constants defined later. Let A_0 be a special source node. For every node $v \in V$ create a corresponding node A_v . For every edge $e = (i, j)$ and $t \in [1, n^\alpha]$ create nodes $B_0^{e,t}, X_0^{e,t}, X_1^{e,t}, M^{e,t}$. Additionally, for every $t \in [1, n^\alpha]$ we add nodes B_1^t, P_0^t, P_1^t . Extra nodes are needed for creating paths as depicted in Figure E. We create several edge sets $E_1 = \{(A_u, X_1^{e,t}), (A_v, X_2^{e,t}) | e = (u, v) \in E, t \in [1, n^\alpha]\}$, $E_2 = \{(X_1^{e,t}, X_2^{e,t}) | e \in E, t \in [1, n^\alpha]\}$. In addition, for every t we create a set of edges as depicted in Figure E. The edge annotation is as follows:

- Dotted edges have A -weight 1 and B -weight 0.
- Dashed edges have A -weight 0 and B -weight 1.
- Plain edges have both weights set to 1.
- Edges with a length annotation are paths of the annotated length; all edges are of the same type.

Finally, let I_B be composed of the set of nodes $B_0^{e,t}$ and B_1^t for every $e \in E$ and $t \in [1, n^\alpha]$.

Claim 1. *If there exists a k -cover for the original graph G there exists a set $I_A \subseteq V - I_B$ of size $k + 1$ that yields $\sigma(I_A) \geq n^{\alpha+\beta}$. Otherwise, for every $I_A \subseteq V - I_B$, $\sigma(I_A) = O(\max\{n^{\alpha+3}, n^{\beta+1}\})$.*

Consider the first part of the claim. Assume that there is a k -cover S for G . Then simply let $I_A = \{A_v | v \in S\} \cup \{A_0\}$. Since S is a vertex-cover, it implies that the spread of technology B emanating from the vertices B_0^t is completely blocked. Thus, every node on the path from A_0 to P_0^t , for all $t \in [1, n^\alpha]$ will be A -infected. This in turn implies that for all $t \in [1, n^\alpha]$, every node on the path from P_0^t to P_1^t will be A -infected. Thus, we have at least $n^{\alpha+\beta}$ A -active nodes, as required.

Now consider the second part of the claim. Let I_A be the optimal set of $k + 1$ initial A -adopters that either does not contain A_0 or $I_A \cap \{A_v\}_{v \in V}$ is not a vertex cover for G . Each vertex in I_A can contribute the most by either one of two optimal choices:

1. Choosing A_0 may contribute at most $(|E| + 5)n^\alpha$.
2. Choosing P_0^t will contribute n^β A -nodes.

Therefore, the contribution of vertices of the first type is at most $O(s \cdot n^{\alpha+2}) = O(n^{\alpha+3})$, and the vertices of the second type contribute at most $O(s \cdot n^\beta) = O(n^{\beta+1})$, which concludes the claim.

Set $\beta = \alpha + 2$. The total number of vertices in the reduced graph is $N = O(n^{\alpha+\beta} + |E| \cdot n^\alpha) = O(n^{2\alpha+2})$. Thus we get that if there is a k -cover for G then the optimal I_A yields $\sigma(I_A) = \Omega(N)$, whereas any I_A that does not correspond to a k -cover yields $\sigma(I_A) = O(N^{(\alpha+3)/(2\alpha+2)})$. Hence, any algorithm that gives an approximation ratio of $o(N^{1-(\alpha+3)/(2\alpha+2)})$ can solve the NP-complete vertex cover problem. Therefore we get that the approximation ratio of any poly-time algorithm is $\Omega(N^{1/2-\epsilon})$, for all $\epsilon > 0$, unless $P = NP$. \square