

Eeniko Metiobio Пonฯtexneio<br>TMHMA H^EKTPO $\cap \Gamma \Omega$ N MHXANIK $\Omega$ N KAI MHXANIK $\Omega \mathrm{N} \Upsilon \Pi O \Lambda О Г I \Sigma T \Omega \mathrm{~N}$<br><br>

#  $\pi \alpha \rho O \cup \sigma i \alpha \varepsilon \xi \omega \gamma \varepsilon \nu \omega \dot{\nu} \varepsilon \pi \iota \delta \rho \alpha ́ \sigma \varepsilon \omega \nu$ 

## $\Delta$ IП $\Lambda \Omega$ MATIKH ЕРГАГIA

ఆoठшри́ $\Lambda$ uxoúp $\eta$

##  <br> ^е́ктораs



# Eeniko Metsobio ПОлฯtexneio 

 TMHMA HAEKTPOAOГתN MHXANIK $\Omega$ N KAI MHXANIK $\Omega$ N $\Upsilon \Pi O \Lambda O Г I \Sigma T \Omega N$


#  $\pi \alpha \rho O \nu \sigma i \alpha \varepsilon \xi \omega \gamma \varepsilon \nu \omega \dot{\nu} \varepsilon \pi \iota \delta \rho \alpha ́ \sigma \varepsilon \omega \nu$ 

## $\Delta$ IП $\Lambda \Omega$ МАТІКН ЕРГАГIA

ఆoठшри́ $\Lambda$ uxoúp


^е́кторац

$\Delta$ nuйтons Ф $\omega \tau \alpha \dot{\alpha} \times n s$ ^є́хторац

K $\alpha \vartheta \eta \gamma \eta \tau \eta \dot{\rho}$
'Apns Паүoupt弓ńs
Entx. K $\alpha \vartheta \eta \gamma \eta \tau \dot{\eta} s$

## ©oठwpйs $\Lambda$ uxoúpクs





# E@NIKO METsOBIO ПO^ฯTEXNEIO TMHMA H $\Lambda$ EKTPOAOГתN MHXANIK $\Omega \mathrm{N}$ KAI 

 MHXANIK $\Omega$ N $\Upsilon П O \Lambda О Г I \Sigma T \Omega N$


Copyright © Cooíwpŕs Auxoúprs, 2012





 тov ourypacéa.




## Euxapıoties








 тท $\begin{gathered}\text { ouaía. }\end{gathered}$










 Eтıбтท́uns $\Upsilon \pi о \lambda о \gamma เ \sigma \tau \omega \dot{ }$.










#### Abstract

Nowadays, monopolies rarely show up as many companies compete for the crucial mass of most potential services. Moreover, the advent of the Internet and the rapid growth of social networking facilitates the spread of information among the customers. Hence, when determining their selling strategy, companies should take into account these interactions to evade getting overtaken by their competitors.

Such strategies might include offering the product to some customers for free so that they can influence their friends or even paying to attract their interest, either by bribing them or through advertisement. Another possible decision is selling the product at a low price at first and later increasing the price, as new customers are rather more eager to pay for it.

In this thesis we survey the main directions in this area. We first elaborate on the diffusion of a product, should an initial influence set be targeted. Subsequently, we extend our viewpoint to the framework where competing products spread their influence. In addition to the existing approaches, our work in computing and characterizing equilibria on such a setting is concisely described. Afterwards, we move to pricing models and the vital role that externalities among people play in the firms' revenue maximization. We conclude this document with a brief review of the algorithmic work in the most fundamental competition pricing models.


## Contents

1 Introduction ..... 11
I Competitive Diffusion in Networks ..... 14
2 Influence Maximization Problem ..... 15
2.1 Modeling the world ..... 16
2.2 Hardness results ..... 17
2.2.1 Exact utility computation ..... 18
2.2.2 Approximating the best response in the general case ..... 18
2.3 Local submodularity of influence is extended globally ..... 19
3 Competitive Diffusion from an algorithmic point of view ..... 21
3.1 The first attempts ..... 21
3.2 Threshold models ..... 23
3.3 Leader-Follower models ..... 25
4 Competitive Diffusion from a game-theoretic perspective ..... 29
4.1 The switching-selection model ..... 30
4.1.1 Price of Anarchy ..... 31
4.1.2 Budget Multiplier ..... 33
4.2 Attempts of computing Nash equilibria ..... 33
4.3 Compatibility in competitive diffusion ..... 34
5 Our results ..... 36
5.1 Pure Nash Equilibria ..... 36
5.2 Equivalent Threshold Model ..... 37
5.3 Exact computation of the utility ..... 38
5.4 Approximate computation of the utility ..... 39
5.5 Hardness of best response computation ..... 40
II Competitive Pricing under externalities ..... 42
6 Pricing without externalities ..... 43
6.1 Auctions vs Pricing ..... 43
6.2 Unknown distributions ..... 45
7 Influence and exploit ..... 46
7.1 Marketing Strategies ..... 46
7.2 Submodular externalities ..... 47
7.3 Extensions ..... 48
8 Price trajectory under externalities ..... 49
8.1 Modeling and exploiting externalities ..... 49
8.2 Competitive Pricing under externalities ..... 50
9 Competition Pricing Models ..... 51
9.1 Bertrand competition ..... 52
9.2 Cournot competition ..... 53

## 1 Introduction

Everybody has, at some point of their life, faced the dilemma of whether an investment is worthwhile. There is a tug of war between the happiness that the investment offers and the money that should be spent in it. The winner of this game determines the eventual decision.

In Algoritmic Game Theory, these notions are formalized. The happiness is called valuation and is quantified by a virtual price $v$. On the other hand, the money is usually depicted by the price $p$ of the product. Utility $u=v-p$ shows the eventual satisfaction of the buyer, considering the pleasure that the product brings to them but also the disappointment due to the money spent. It is typical that players in Game Theory try to maximize their utility.

Nevertheless, there are many products that are far more useful when friends have also adopted them. Consider, for instance, the mobile phone industry. When you and your friends have the same package, you may have unlimited calls or other provisions. Another example comes from the Internet. Many web games or social networking sites would be less amusing if you were surrounded by strangers.

These positive externalities are taken into account and increase the valuation for a product. As a result, the valuation is no longer a pure number but is instead a function of the set $S$ of friends that already own it. Hence, the utility is $u=v(S)-p$.

In order to model such social network interactions explicitly, graphs are usually used. The vertices are the players of the game, i.e. the potential buyers in our case. The edges point the friendship between the vertices that are adjacent to them. The graph is not necessarily undirected as the influence may be one way. For example, Brad Pitt's aroma choices will probably influence John Mimic but the opposite direction is rather unlikey. Moreover the edges can be weighted, which quantifies the influence they depict.

Of course, when deciding to buy a product, the influence that a friend's suggestion will exert is usually not independent to the previous recommendations. For example, a friend who will propose you the new product first might tempt you more marginally than if he was the tenth to speak about it to you. This feature that makes the mere influence of a node being greater when you have already been approached by some set than when a
superset of it has tried to convince you is called submodularity.
The firms try to use the properties of these externalities among the potential buyers in order to maximize their revenue. This direction has become more urgent with the expansion of social networks and the rapidity with which information moves. Indeed, most of the marketing power still remains unexploited. Hence, modeling the way that influence spreads and designing efficient algorithms that approximate the best solutions is of crucial importance for them.

However, this radical shift in the technologies has brought trouble to some firms as the competition became larger. Now advertising a product is easy and often at low cost and the dispersion of digital goods, i.e. goods that are produced in zero cost such as music songs, is facilitated. As a result, in most of the areas, monopolies tend to become extinct and are replaced by competitive firms that seek for the same customers.

This competitive marketing setting was the motivation behind this thesis. We try to regard how much it determines the firms' profit and, consequently, the prices that the customers are about to pay. In addition, we survey the neat techniques that are suggested in modeling influence and checking its expansion. This gives an exciting viewpoint towards the deeper understanding of how social networks work.

The structure of this thesis contains two parts. In the first part, we examine competitive diffusion models, where the firms' strategy is to select some initial set of buyers to attract and later exploit its expansion. In this setting, the goal of the firms is to gain the most ultimate buyers possible and do not have any control on the price of the product. On the other hand, the second part is more interested in the pricing models. There firms select not only the way they will approach the potential buyers but also the prices they will offer them.

In section 2, we regard the monopoly Influence Maximization Problem. There, the firm needs to select the $k$ initial nodes that will ensure the maximum number of eventual adopters. Many models have been suggested that differ on how the influence spreads and whether it propagates independently (cascade models) or in an accumulative way (threshold models). The Influence Maximization Problem is proved to be inapproximable with any constant factor in each general case, however if there are local submodular conditions then there is a constant approximation for it.

In section 3, we move to competitive diffusion models and ask algorith-
mic questions towards them. Firstly, we are interested in the extension of monopoly models to the competitive case. Then we survey a broad class of threshold models with high proximity to reality. We finish the section posing best response questions in the Stackeberg problem, where there is a leader who establishes her dominance and followers that take decisions after this.

In section 4, we emphasise on a competitive contagion model, Switchingselection model and set game-theoretic questions on it. Firstly, the inefficiency of equilibria (Price of Anarchy) and the superiority of the firm with the highest budget (Budget Multiplier) is examined. Afterwards, we move on attempts towards computing and characterizing equilibria in such a setting, which is the focal point of our own ongoing work. Last but not least, we examine the role of compatibility in the diffusion of products.

In section 5, we present our own results on computing and characterizing equilibria in the Switching-selection model. Firstly, we prove that pure Nash equilibria do not always exist. Then we try to approach these equlibria, when they exist. We provide an equivalent threshold model that is more easily tractable. Moreover, we give exact and approximation algorithms for computing a firm's utility in an instance. Finally, we move to best response questions trying to understand the nature of these equilibria.

In section 6 , we begin the second part, by adding another weapon to the firm, price. We focus on models with no externalities, where the already adoptions do not shift users' valuation. This is a very active aspect of recent research in Algorithmic Mechanism Design.

In section 7, we initiate externalities and study Influence-and-exploit strategies. There some people are given the product for free and later firms try to extract the maximum profit possible by choosing the sequence in which they will approach the customers and the prices they will offer them. We interest more in some simple cases of models where the externalities are submodular. Closing this section, we review the case where the price addressed to the customers is necessarily undifferentiated.

In section 8, we care about price trajectories under the existence of externalities. There we examine both posted pricing strategies that capture this notion and a competitive scenario.

Section 9 , which is the concluding section, examines competition pricing models. After a concise overview in the most prevalent models, we review some algorithmic extensions on them.

Part I
Competitive Diffusion in
Networks

## 2 Influence Maximization Problem

As previously mentioned, one strategy that the companies implement in order to attract more customers is to offer the product for free to some initial influential set of them, so that they can urge their friends to buy it. The problem of finding the initial set, posit $S$, that maximizes the eventual number of adoptions of the product in the monopoly case, posit $\sigma(S)$, was first addressed by Domingos and Richardson in [23, 45]. Following this work, Kempe, Kleinberg and Tardos $[33,34]$ have shown that this problem is NP-hard to approximate within a factor better than $1-\frac{1}{e}$ even for simple models. However the outcome can be approximately maximized for these models using a hill-climbing algorithm when $\sigma$ is a monotone and submodular function of $S$ (definition given below). Inspired by this work, Mossel and Roch [40] have extended the latter result for all the models where the influence functions are monotone and submodular, showing that such local conditions imply the global monotonicity and submodularity of $\sigma$.

Definition 1. (Monotonicity) The function $f: 2^{V} \rightarrow \mathbb{R}$ is monotone if

$$
f(S) \leq f(T) \text { for all } S \subseteq T \subseteq V
$$

Definition 2. (Submodularity) The function $f: 2^{V} \rightarrow \mathbb{R}$ is submodular if

$$
f(S \cup v)-f(S) \geq f(T \cup v)-f(T) \text { for all } S \subseteq T \subseteq V, v \in V
$$

Most of the positive results referenced above are based on the following theorem [20, 44, 33]:

Theorem 1. Let $f$ be a non-negative, monotone, submodular function on sets.

1. The greedy algorithm, which always picks the element $v$ with largest marginal gain $f(S \cup\{v\})-f(S)$, is a $\left(1-\frac{1}{e}\right)$-approximation algorithm for maximizing $f$ on $k$-element sets $S$.
2. A greedy algorithm which always picks on element $v$ within $1-\epsilon$ of the largest marginal gain results in a $1-\frac{1}{e}-\epsilon^{\prime}$ approximation, for some $\epsilon^{\prime}$ depending polynomially on $\epsilon$.

Intuitively, the monotonicity rule implies that the addition of an element in set $S$ can't decrease the value of $f(S)$, whereas the submodularity
rule implies that this addition offers more marginal gain to $S$ than to any of its supersets.

Applying the theorem above, the existence of an approximation algorithm of $1-\frac{1}{e}$ (or $1-\frac{1}{e}-\epsilon$ when the exact greedy computation of the function's maximizer is not efficient) is proved by just showing that $\sigma(S)$ is monotone and submodular.

### 2.1 Modeling the world

The problem space is modeled as a graph whose nodes are the potential buyers and where the existence of an edge implies the friendship of two people. A node $v$ 's decision of whether to buy the product or not (becoming active or staying inactive) depends on an influence function for each vertex/edge that is based on the set of customers that have already adopted the product(in most cases, the node is just influenced by its relatives). Hence an initial set of active nodes helps the product propagate through the network. The goal of the company is to choose the initial $k$-set that maximizes the eventual number of active nodes. In most models, this diffusion is progressive, which means that a node can't get deactivated, should it become active once. This is rational as we might imagine the activation of a node as buying the product. The models proposed differ on either the condition that leads nodes to adopt the product or the properties of the influence functions. The most sensible models of this form are the following:

One of the first models proposed is the Linear Threshold Model. In this model, each edge $e=(v, w)$ has a weight $b_{v, w}$ that quantifies the influence node $v$ wields to node $w$, supposing $v$ is already active. Each node $v$ chooses a threshold $\theta_{v}$ uniformly at random from $[0,1]$, which declares the minimum total amount of influence that should be exerted to this node to activate it. Choosing an initial influence $k$-set $S$ to maximize the expected $\sigma(S)$ is NP-hard, which was proved in [33] by a reduction from Vertex Cover. What is more interesting is the other result about this model in the same paper, which shows that $\sigma(S)$ is monotone and submodular for this model. To prove that, the authors use a technique that proved to be very useful for future work in the same area. As $\sum_{w} b_{v, w}<1$ for all $v \in V$, they throw a biased coin to decide which of the internal edges becomes live, while all the others become blocked (with probability $1-\sum_{w} b_{v, w}$,
no internal edge becomes live). By this technique, they transfer the u.a.r. selection of the threshold in a much more easy-handling form. Hence, the problem is reduced to finding live-edge paths from the initial set. As the cardinality of these nodes, for each of the live-edge paths' decision, is a submodular function of the initial set, $\sigma(S)$ is also submodular (as a linear combination of them).

Another widely studied model is the Independent Cascade Model. In this model, each node $v$ has a probability $p_{v, w}$ to influence another node $w$. The influence is exerted at most once, when node $v$ gets activated. In [33], the authors prove that the influnce maximization problem is NP-hard, even for this simple model using a reduction from the set cover problem. Moreover, by a similar technique as in the previous model (throwing a biased coin for each edge at the beginning of the process), $\sigma(S)$ is proved to be a monotone and submoduar function of $S$, which shows the existence of an efficient approximation algorithm (as already mentioned).

Extending this in [34], the Decreasing Cascade Model preserves the order-independence of the influence functions but introduces a notion of submodularity to the probabilities of the model. The influence probability of each edge now depends on the set of nodes already adopted the product, however it decreases as the set of adopters increases. Supposing $S$ is the initial influence set and $B, B^{\prime}$ are the sets of adopters at some time, this means that:

$$
p_{v}(u, S \cup B) \geq p_{v}\left(u, S \cup B^{\prime}\right) \text { for all } v, u \in V, B \subseteq B^{\prime}
$$

The equation clearly implies that, if $u$ is the only contagious node at a time step, its contribution will be submodular. By getting a threshold identical viewpoint of the problem and by delaying the influence exertion of newly activated nodes, the authors isolate each node's contagion and, as a result, the desired submodularity follows.

### 2.2 Hardness results

Although the models already presented hold some nice properties that lead them to allow efficient approximation algorithms, the general case is not as easily tractable. In fact, the exact utility computation is $\# P$-complete and the best response is NP-hard to approximate within any constant factor.

### 2.2.1 Exact utility computation

Theorem 2. The exact computation of the utility in the Linear Threshold Model is \#P-complete.

This was proved by Chen et al in [18]. The reduction comes from the the problem of counting the number of simple paths in a directed graph, which is known to be $\# P$-complete. On the positive side, Kempe et al [34] hinted the existence of a FPRAS for the problem. We will present an extension of it in section 5 .

### 2.2.2 Approximating the best response in the general case

On the other hand, the best response is not easy to handle, as implies a reduction from the set cover problem in [33]:

Theorem 3. It is NP-hard to approximate the influence maximization problem within a factor of $n^{1-\epsilon}$ for any $\epsilon>0$

As this proof is highly descriptive for such inapproximable results, we will present it here.

Proof. Firstly, we should remind the set cover problem: Given a set of $n$ elements $E=\left\{e_{1} \ldots e_{n}\right\}$ and $m$ sets $S_{1} \ldots S_{m}$, each of which contains a subset of $E$. Find whether there exist $k$ of the sets $S_{i}$ such that their union equals $E$.

Given an instance of this problem, we create the following instance of Influence Maximization Problem. For every element of $E$, we introduce a vertex $u_{i}$ and, for every set $S_{j}$, a vertex $s_{j}$. If $e_{i} \in S_{j}$ then a directed edge $\left(s_{j}, u_{i}\right)$ is put. In addition, $N=n^{1-\epsilon}$ vertices are created and we create directed edges from all the elements of $E$ to all of those $N$ vertices. Last but not least, we require for each of the $N$ nodes to adopt the product that all their neighbours have already adopted it, whilst for the nodes $e_{i}$, just one neighbour is needed.

As is pretty obvious, the only way to make all the $N$ nodes eventually adopt is through the adoption of all the elements of $E$. There is no reason to select one element of $E$ in the initial $k$-set, as we can attract this node through one of $s_{i}$. As a result, we need to find a k-set of $\left\{s_{1}, \ldots, s_{m}\right\}$ such that all the elements of $E$ have a neighbour in it. However, if we can efficiently solve this problem, we can solve the initial general case of the
set cover problem. As this problem is NP-hard to solve, any other solution can't assure us to have the $N=n^{1-\epsilon}$ nodes in our eventual adopters. As a result, the problem is NP-hard to approximate with ratio better than $n^{1-\epsilon}$.

The above theorem is based on very extreme deterministic thresholds. In section 5, we will show that even without so extreme thresholds, the same hardness is preserved.

### 2.3 Local submodularity of influence is extended globally

It is easy to figure out that the model used in the previous reduction doesn't have the property of submodularity, which is natural as the thresholds are deterministically chosen. For example, the addition of $u_{n}$ offers more marginal gain in the eventual adopters in set $T=\left\{u_{1}, \ldots, u_{m-1}\right\}$ than in its subset $S=\left\{u_{1}, \ldots, u_{i}\right\}$ where $i<m-1$ as in the first case it leads all the $N$ nodes to adopt. This is expected as we have already pointed the existence of an approximation algorithm with better approximation ratio than the one of the inapproximable result. What is more interesting is that, in this model, the influence functions are not submodular as well. This hinted the authors of [33] to conjecture a correlation between the submodularity of the threshold functions of the vertices and the submodularity of eventual adopters. In [40], the authors answered the aforementioned hypothesis and reached to the following result:

The diffusion, starting with initial set $S$, is modeled as a process $S_{t=0}^{n-1}$ where $S_{0}=S$, each vertex $v$ has a threshold function $f_{v}: 2^{V} \rightarrow[0,1]$ in order to decide when to adopt and selects u.a.r. one threshold $\theta_{v}$ in $[0,1]$ and, at any time $t>0$, a vertex $v$ is added to $S_{t}$ if and only if $f_{v}\left(S_{t-1}\right) \geq \theta_{v}$.

Theorem 4. If $F=\left(f_{v}\right)_{v \in V}$ is monotone and submodular with respect to the set of already adopters then the eventual number of adopters is monotone and submodular with respect to the initial set of adopters.

The proof shares some similar ideas with the proof presented for the Linear Threshold Model as it unravels information for the threshold step by step, combining the cascade and the threshold models effectively (need-toknow representation). Another interesting technique used is the antisense
coupling, which uses both $\theta_{v}$ and $1-\theta_{v}$ in the coupling process as the latter is equivalent as $\theta_{v}$ is selected u.a.r.

## 3 Competitive Diffusion from an algorithmic point of view

The work on the monopoly influence maximization broadened the horizons in the understanding of social networks and created questions on how such an environment would behave when the setting contains more than one firms that compete for the same target group. Many models have been suggested since then to capture the competitive nature of the aforementioned problem. In most of the cases, the principal direction of work is based on the firms' strategy for maximizing their eventual adopted nodes, supposing that they know the strategy of their competitor, which is a normal generalization of the problem discussed in the previous subsection. In this section, we will refer to some of the most interesting models and how they handle this target. In the next section, we will continue with more game-theoretic questions.

### 3.1 The first attempts

The first attempt to deal with this issue was made in [9]. The model proposed was a simple extension of the Independent Cascade Model. Similarly to this model, there is a directed graph (V,E), each vertex is either active or inactive and each edge $e=(v, u)$ has a probability $p_{e}$ to activate $u$, once $v$ is infected (the activation takes place after $T_{e}$ steps, where $T_{e}$ is an independent and exponentially distributed random variable). The innovation of this model is that now there are $b$ firms, instead of one. Each firm $i$ selects an initial set $S_{i}$ and the diffusion process evolves as previously mentioned. The goal of each firm $i$ is to select the initial set $S_{i}$ that maximizes the eventual number of adopters of product $i$. Adapting the idea of live paths presented in [33], the authors prove that the last firm can approximate his optimal solution to a factor of $1-\frac{1}{e}$, using the hill-climbing brute-force algorithm, as finding the payoff of a firm $i$ when all the initial sets of the other firms are fixed is a monotone and submodular function of $S_{i}$.

In the meantime, Carnes et al have focused on the competition among two firms proposing two models. Given a graph $G(V, E)$ their models create a subgraph of it, where each edge $e$ remains with probability $p_{e}$ measuring the potential influence exerting on it (a common used to technique to avoid
the stochastic part of the model during the diffusion process). As in the previous case, the main question examined is the selection of $k$-influence set by a firm A, supposing that their competitor B has fixed their influence set and those two sets must be disjoint. The goal is again the maximization of the expected number of eventual adopters $\left(\rho\left(I_{A} \mid I_{B}\right)\right.$ for the first model and $\pi\left(I_{A} \mid I_{B}\right)$ for the second).

Their Distance-based model focused more on the proximity between the node that is about to adopt and the nodes that influence it. Thus, for each edge, there was a distance $d_{e}$ which measure how far the adjacent nodes are. The game is played in steps and at each step $i$ each node adopts a product if it is at a distance $i$ from a node in some initial set (if this happens for both products, there is a tie-breaking rule, for instance the number of its neighbours that have adopted product A over the number of its neighbours that have adopted either of the products). For this model, selecting $I_{A}$ that maximizes $\rho\left(I_{A} \mid I_{B}\right)$ is NP-hard, which was proved by a reduction from the set cover problem, however this function is monotone and submodular and thus the optimal result can be approximated to a factor of $\left(1-\frac{1}{e}\right)$. This model takes into account the global initial coloring of the graph and not only the local neighbourhood of a node. This is, in fact, far fetched from reality as one usually gets influenced just by people that are close to them.

On the other hand, their Wave propagation model envisages this problem by interesting only in the node's neighbours. In fact, the first nodes to get influenced are the neighbours of the initially selected influential nodes, with probability proportional to their influence in contrast to the influence of their competitor. The diffusion process, afterwards, continues with the newly colored nodes influencing their neighbours. For this model, as well, the authors prove that the decision problem of finding the initial sets is NP-hard, using the same idea as above but $\pi\left(I_{A} \mid I_{B}\right)$ is a monotone and submodular function of $I_{A}$ and, as a result, there can be found a good approximation of this problem. The authors supposed that their result could be extended in the case of weighted edges, however some years laters, Borodin et al showed in [11] that this model is NP-hard to approximate within a factor of $o\left(n^{\frac{1}{2}-\epsilon}\right)$.

### 3.2 Threshold models

In the same paper, Borodin et al suggested some models that focus on the total influence exerted in one node, combining the General Threshold Model with the competitive feature. Each node $v$ selects a threshold $\theta_{v}$ and its choice of whether to adopt the product depends on the total influence on this node and whether it surpasses $\theta_{v}$. The firms' strategy is again to select an initial set of influential nodes. Their Weighted-Proportional Competitive Linear Threshold Model takes into account the total influence exerted in a node, when it decides whether to buy, no matter which firm the influencers have adopted. Afterwards, the node decides which firm to follow proportionally to the influence of each firm or using another tie-breaking rule. On the other hand, their Seperated Threshold Model picks different threshold values for each firm (node $v$ picks $\theta_{v}^{A}$ and $\theta_{v}^{B}$ respectively for firms $A$ and $B$ ). Each edge has different influence weights $w_{A}$ and $w_{B}$ corresponding to each firm and in the decision of adopting one firm's product, the nodes examine just the influence exerted by the same firm.

Unlike the previous models, these threshold models handle the monotonicity and submodularity properties in a different way. The first model is not monotone as one firm's additional influence can help their competitors win some key nodes with high threshold and then dominate in the graph, ruining their own firm's fate. On the contrary, the second model cannot show such side effects as the addition of new influence by a firm cannot help the other firm surpass its threshold. Nevertheless, both models are not submodular, as adding a single node in a firm's initial set can devastate their competitor's strategy and hugely increase its own payoff. As a result, these rather logical models don't have the nice approximation that the property of submodularity provides.

To strengthen the previous argument, the authors have proved that:

Theorem 5. It is NP-hard to give an approximation with a ratio better than $\Omega\left(N^{\frac{1}{2}-\epsilon}\right)$ for any $\epsilon>0$ for the Seperated-Threshold Competitive Influence problem, where $N$ is the number of nodes in the graph.

Proof. The idea of the proof is similar to the one of the same result for the Wave Propagation Model and comes from a reduction from the Vertex Cover Problem.


Figure 1: Gadget of the reduction

Firstly, we should remind the description of the Vertex Cover Problem. Given a graph $G$ and an integer $k$, we are asked to decide whether a selection of $k$ nodes so that each edge has at least one of its endpoints in the selection can be found.

From this instance, we create an instance of the Seperated Threshold Competitive Problem using the gadget above. Suppose that the number of edges of $G$ is $n^{\alpha}$. The dotted lines demonstrate influence just for player A, whereas the dashed lines demonstrate influence just for player B. The gadget is repeated for $t=1, \ldots, n^{\alpha}$. The influential set of player $B$ contains all the nodes of the form $B_{0}^{*}$ for all the edges and $t \in\left[1, n^{\alpha}\right]$ as well as the nodes of the form $B_{1}^{*}$ and the goal is to maximize the payoff of $A$ when he is about to select $k+1$ nodes.

If a $k$ Vertex Cover exists in $G$ then A can block the influence from the $B_{0}$ nodes and, by selecting $A_{0}$, it can influence $P_{0}^{t}$ for all $t \in\left[1, n^{\alpha}\right]$ and thus acquire a total payoff of at least $n^{\alpha} \times n^{\beta}$. Else, it can not influence $P_{0}^{t}$ through $A_{0}$ and at some $j, M_{j}$ will be taken by B. Thus, A can get influence at most $(|E|+5) n^{\alpha}=O\left(n^{\alpha+2}\right)$ from selecting $A_{0}$ or an $M^{*}$ that is upper to $M^{j}$ and at most $O\left(n^{\beta}\right)$ from selecting $P_{0}$ or an $M$ node
that is downer than $M_{j}$. As a result, its eventual cascade will be at most $\max \left(n^{\alpha+3}, n^{\beta+1}\right)$. The total numbers of vertices in $G^{\prime}$ is $N=O\left(n^{2 \alpha+2}\right)$. Supposing $\beta=\alpha+2$, this means that, if a Vertex Cover does not exist then we can get at most $O\left(N^{\frac{\alpha+3}{2 \alpha+2}}\right)$.

As a result, a polynomial time approximation algorithm that gives a better approximation ratio that the one in the theorem implicitly solves the Vertex Cover problem which is not solvable in polynomial time unless $P=N P$.

On a positive note, the authors have also suggested a model that is both monotone and submodular, the OR model, combining features from previous work on cascade models. There, the two firms act independently and, given an istance ( $R_{A}, R_{B}$ ) denoting the set of nodes already infected by either firm, there exist probabilities $f_{v}^{A}\left(R_{A}, R_{B}\right), f_{v}^{A}\left(R_{A}, R_{B}\right), f_{v}^{B}\left(R_{A}, R_{B}\right)$ that a node $v$ gets infected. If those functions are monotone and submodular for all vertices then the total expected number of eventual adopters is monotone and submodular as well, meaning that the approximation algorithm described in the previous section can apply to the model as well.

Another attempt to model the competition among firms using thresholds is presented in [4]. Unlike the previous models, the thresholds are not chosen uniformly at random but are selected deterministically and quantify the reluctance of the nodes to adopt a product under multiple influence. If a node $v$ is targeted by just one firm then it adopts its product in a straightforward manner. Else it is urgent that a total influence of at least $\theta_{v}$ is addressed to it so that the product is eventually adopted. In this model, the authors study whether there are instances where one firm can dominate in the whole network and whether such an outcome is unavoidable. Furthermore, they study the possible outcomes and whether they are uniquely determined. Should this not occur, they give upper and lower bounds for the eventual number of adoptions a firm can ensure under different possible outcomes.

### 3.3 Leader-Follower models

One broad class of models are based on the Stackelberg competitition model which constitutes of the following setting. There is a leader who decides its own strategy. Afterwards, based on the strategy chosen by the leader, the followers decide their own strategy. This model makes sense
when there is a great firm which has pioneering ideas and can initially sustain their influnce, whereas others will follow, strategizing to those not already influenced by the leader. More specifically for the Competitive Diffusion Model, the leader will choose an initial influence set, whereas the follower will select their influence set from the vertices that are not selected by the leader. This problem was examined in both [35] and [19] for the case of one leader and one follower.

The first paper to deal with this issue was [35] which suggested that the decision of the firms' strategies was the answer in two problems. The leader shall answer the $(r \mid p)$-centroid problem, which consists of choosing $p$ nodes to influence at first knowing that the follower will select afterwards $r$ nodes in order to maximize its eventual influence. On the other hand, the follower needs to face the $(r \mid p)$-medianoid problem of maximizing its eventual adopters by deciding $r$ nodes to infect, already knowing which $p$ nodes the leader has selected. The propagation model used is a deterministic one as any firm getting influenced by the first neighbour to influence it and, whenever there is a collision among two different firms trying both to influence it, the tie breaking rule declares that the vertex refuses both influences and never gets influenced again. Both problems are NP-hard to get solved but whilst the first problem is NP-hard even to get approximated by a constant factor even for $r=1$, the second problem has a ( $1-\frac{1}{e}$ ) when the hill climbing's algorithm is impelemented. Last but not least, the authors present a counter example showing that the leader might be condemned to get outplayed by the follower although this seems rather obscure.


Figure 2: Gadget for the Centroid problem's reduction

In order to prove that the $(1 \mid p)$ centroid problem is NP-hard, the
authors use a reduction from the Vertex Cover problem. More precisely, each edge $(u, v)$ of the instance of the vertex cover problem is replaced by an instance of the gadget shown above (the figure shows two edges $(u, v)$ and $(v, w)$ ). If there is a vertex cover of size $p$ in the initial problem then the follower will eventually have at most 2 vertices of the new graph influenced (supposing that the leader picks the vertices of the vertex cover). On the other hand, unless there is a vertex cover of size $p$, there will be at least one whole gudget with no leader vertices in it. Hence the follower will get 3 or more eventually influenced nodes. As a result, if the existence of a $p$-node selection minimizing the follower's eventual influence could be selected efficiently, then it could be used to efficiently solve the Vertex Cover problem, which leads to contradiction unless $P=N P$. The proof that this problem cannot give efficiently neither an $\alpha$-approximation uses the same idea but a more complex gadget.


Figure 3: Graph used for the Medianoid problem's reduction

The reduction used to prove that the $(r \mid p)$ medianoid problem is NPhard comes from the Dominating set problem which asks whether there is a subset $V^{\prime}$ of the vertices of the initial graph $G(V, E)$ such that all the vertices of the graph are in the set or are straight neighbours of a member of $V^{\prime}$. Posit that the graph is the one shown above, where, for $1 \leq i \leq n, s_{i}$ has edges both to $v_{i}$ and each neighbours, as well as to $s_{n+1}$. Furhtermore, $X_{1}=\left\{s_{n+1}\right\}$. It is obvious that, should there exists a dominating set, the selection of the respective $s_{i}$ will give an eventual adoption for the follower of $|V|_{r}$. Unless this occurs, this size cannot be acquired as the diffusion of the leader will influence both the rest of $s$-nodes and the vertices of $G$
that are at distance at least 2 from the initially selected. As a result, it is NP-hard to solve the follower's problem.

Sharing a similar perspective, Clark et al suggest in [19] a Monte Carlo model to examine the leader-follower strategies. Their Dynamic Influence in Competitive Environment model differentiates from the previous models in that the nodes, once colored by a firm, can shift to the other firm or to no firm. This Markov process, under some sensible conditions, converges to a unique stationary distribution. Furthermore, the insinuated utility function both for the leader and for the follower has the property of submodularity which leads to good approximation algorithms for the firms' strategies.

## 4 Competitive Diffusion from a game-theoretic perspective

Until now, the firms' strategies have been regarded independently as if their decisions don't alter the strategy of their competitors. However, in many occasions, the firms change their strategy regarding their opponents' behavior. This procedure continues till no firm has an incentive to deviate from its predetermined strategy unilaterally. This situation is called Nash equilibrium, was introduced in [43] and exists in every finite game.

Although in many cases more than one such situations can exist, they might not be equally prefered by the players. The social welfare of an instance captures the total happiness of the game's players, given that each player chooses a particular strategy. It might be the case that sometimes the Nash equilibrium differs from the situation where the social welfare is maximized, as players tend to play selfishly, trying to maximize their own welfare. Hence, an interesting unit for describing a game is how bad an equilibrium can become, should we let the game unravel. This notion is captured by the Price of Anarchy, introduced in [36], which measures the inefficiency of the worst-case equilibrium comparing to the set of strategies that lead to the maximum social welfare. Essentially, it describes how much the system's efficiency deteriorates due to the selfishness of the players.

From a game-theoretic point of view, we are interested in two directions which we will investigate in our problem as well. The first direction concerns the difficulty of efficiently discovering the equilibria of a game in order to predict the system's behavior, given the rules of a game, as well as studying the characteristics of those equilibria. The second focal point is how we can alter the rules of the game in order to lead players to a strategies' trajectory that maximizes the social welfare or another goal. In this case, we are interested in designing mechanisms that urge the system to behave in a desired way.

In this section, we will deal with such problems in the case of competitive diffusion in networks. We will first focus in the switching-selection model, recently introduced in [27]. Afterwards, other attempts to stress competitive diffusion from a game theoretic point of view will be regarded.

### 4.1 The switching-selection model

In this model, there are two firms (Red and Black) that tend to maximize their total eventual adopters of a social network, which is modeled as a graph $G(V, E)$. In order to achieve that, they have some initial budgets $K_{R}$ and $K_{B}$ respectively. Their strategy is to spare this money in some of the nodes of the graph, suppose $\left(a_{p 1}, a_{p 2}, \ldots, a_{p n}\right)$ is this distribution of the budget, where $a_{p i}$ denotes the money firm $p$ provides in order to make node $i$ adopt its product $\left(\sum_{i} a_{p i}=K_{p}\right)$. If node $i$ is targeted by just one firm then it adopts their product. Else, it adopts product $p$ w.p. $\frac{a_{p i}}{a_{R i}+a_{b i}}$. This is the first step of the procedure that leads to the initial influential sets of each firm.

Afterwards, the diffusion process begins. This can happen either sequentially (each node is asked in a random order) or in steps (each node gets influenced just by nodes that have already adopted the product in the previous steps). We will focus on the latter way. The game is thus played in steps. The initial part that leads in the creation of the influential sets is step 0 . In any step $k$, any noninfected node can adopt a product if enough of his friends urge him to do so.

Hence, there is a switching function $f$, which gives the probability that a node will choose to adopt some product (with probability $f\left(\alpha_{R i}+\alpha_{B i}\right)$ node $i$ adopts some product where $\alpha_{i p}$ is the ratio of node i's neighbours that have adopted product $p$ until this step). The two products are judged as complementary (for example choosing an operating system with Windows and Mac as possible options). Thus the switching function expresses the pressure induced by a node's friend to buy the product, independently of its special features.

Once a node decides to buy the product, she tends to decide which of the different firms she should choose. This feature is captured by the selection function $g$, which gives the probability that a node chooses to adopt product $p$, given that he has already decided to buy one of the products (with probability $g\left(\frac{\alpha_{p i}}{\alpha_{R i}+\alpha_{B i}}\right)$ node $i$ adopts product $p$ ).

The update process can occur in many different ways, according to the authors (such as after every step for a particular period). We choose to treat function $f$ as the total probability that node $i$ will have bought the product by time step $k$ and consider that each neighbor influences just once. Hence, given that $\alpha^{(k-1)}$ neighbors of $i$ have adopted till step $k-2$
and $\alpha^{k-1}$ have adopted it at step $k-1$, the probability that $i$ adopts the product at step $k$ is $\frac{f\left(\alpha^{k-1}\right)-f\left(\alpha^{(k-1)}\right)}{\left(1-f\left(\alpha^{(k-1)}\right)\right)}$ as only the new adopters influence $i$ to adopt. It is obvious that $f$ is a monotone increasing function and it is sensible to suppose $f(0)=0$ and $f(1)=1$ (no infection leads to lack of information and thus non adoption and full peer pressure leads to certain adoption).

For the selection between the firms, no such modification is urgent as this biased coin will fall just once. Again, if all the influence of a node comes from the same firm then he will get this product $(g(0)=0$ and $g(1)=1)$ and, as the node will get some product in any case, once she has decided to buy, $g(y)+g(1-y)=1$.

Recall that the strategy of a firm $p, \sigma_{p}$, is to select a distribution of each initial budget. We suppose that the firms play just pure strategies, i.e. this distribution is deterministically given. The payoff of this firm $\Pi_{p}$, given both the strategies $\sigma_{R}$ and $\sigma_{B}$, is as previously the expected number of eventual adopters of its product, that is $\Pi_{p}\left(\sigma_{R}, \sigma_{B}\right)=\mathbf{E}\left[\chi_{p} \mid\left(\sigma_{R}, \sigma_{B}\right)\right]$ where $\chi_{p}$ denotes the random variable of the number of eventually infected nodes by $p$. If no firm has incentive to deviate from their selected strategy and win more, that is, for any $s_{i}$ and any $p, \Pi_{p}\left(s_{i}, \sigma_{-i}\right) \leq \Pi_{p}\left(\sigma_{i}, \sigma_{-i}\right)$ then $\left(\sigma_{R}, \sigma_{B}\right)$ is a Nash equilibrium of the system ( $\sigma_{-i}$ denotes that all other players apart from $i$ keep playing the same strategy).

### 4.1.1 Price of Anarchy

The Price of Anarchy in this game quantifies the effect of the non-collaboration of the firms in their total eventual adopters. More precisely, suppose $\left(a_{R}^{*}, a_{B}^{*}\right)$ is the allocation tha maximizes their total payoff, that is $\mathbf{E}\left[\chi_{R}+\right.$ $\left.\chi_{B} \mid\left(\sigma_{R}, \sigma_{B}\right)\right]$ and $\left(\sigma_{R}, \sigma_{B}\right)$ is the equilibrium minimizing the total payoff. Then the Price of Anarchy is defined as $\frac{\mathbf{E}\left[\chi_{R}+\chi_{B} \mid\left(a_{R}^{*}, a_{B}^{*}\right]\right.}{\mathbf{E}\left[\chi_{R}+\chi_{B} \mid\left(\sigma_{R}, \sigma_{B}\right)\right]}$. The denominator interests just in Nash equilibria as there are the only stable outcomes of the system, whereas the nominator cares for any possible outcome as, should an alliance between the firms was decided, this would be a possible outcome.

The authors go on by giving upper and lower bounds for the Price of Anarchy. They focus in linear selection functions. As for the switching function, they use functions of the form $f(x)=x^{r}$. The value of $r$ can make this function either concave $(r<1)$, linear $(r=1)$ or convex $(r>1)$.

Succinctly, their results are given in the following theorem:
Theorem 6. Let the switching function be $f(x)=x^{r}$ and let the selection function be linear $g(y)=y$. Then:

- For any $r \leq 1$, the Price of Anarchy is at most 4 for any graph $G$
- For any $r>1$ and any $V$, there exists a graph $G$ for which the Price of Anarchy is greater than $V$

We will now give a sketch of the proof for the upper bound. First observe that the concavity of the switching function and the linearity of the selection function imply that a firm $p$ is more likely to infect a node in the absense of a competitor, if we suppose that the number of neighbors that have adopted its product remains unaltered. This occurs as in such a case $\frac{f(x+y)}{x+y} \leq \frac{f(x)}{x}$ as the slope of a concave function reduces as the argument increases. Using this observation and, running a coupling simulation of the process where the competitor is absent (solo Red process) and where the competitor exists (joint process), we get that $\mathbf{E}\left[\chi_{R} \mid\left(A_{R}, \emptyset\right)\right] \geq \mathbf{E}\left[\chi_{R} \mid\right.$ $\left.\left(A_{R}, A_{B}\right)\right]$ and $\mathbf{E}\left[\chi_{B} \mid\left(\emptyset, A_{B}\right)\right] \geq \mathbf{E}\left[\chi_{B} \mid\left(A_{R}, A_{B}\right)\right]$. The coupling ensures that the same nodes are asked and that at each time step no node of the joint process will be painted red if this doesn't happen to the solo process as well (as previously mentioned, the latter probability is greater). Hence, we have the following Lemma, which states that the total payoff of the firms when competing each other is less than what both could get if their competitor didn't exist:

Lemma 1. $\mathbf{E}\left[\chi_{R} \mid\left(A_{R}, \emptyset\right)\right]+\mathbf{E}\left[\chi_{B} \mid\left(\emptyset, A_{B}\right)\right] \geq \mathbf{E}\left[\chi_{R}+\chi_{B} \mid\left(A_{R}, A_{B}\right)\right]$
Using a similar coupling technique, the authors show that the total payoff of the firms when competing each other is greatet that what each of them would get if its competitor didn't exist.

Lemma 2. $\mathbf{E}\left[\chi_{R} \mid\left(A_{R}, \emptyset\right)\right] \leq \mathbf{E}\left[\chi_{R}+\chi_{B} \mid\left(A_{R}, A_{B}\right)\right]$
Combining these lemmas, we get the following corollary that proves the upper bound:

Corollary 1. For any Nash equilibrium $\left(S_{R}, S_{B}\right)$,

$$
\mathbf{E}\left[\chi_{R}+\chi_{B} \mid\left(S_{R}, S_{B}\right)\right] \geq \frac{\mathbf{E}\left[\chi_{R}+\chi_{B} \mid\left(S_{R}^{*}, S_{B}^{*}\right)\right]}{4}
$$

where $\left(S_{R}^{*}, S_{B}^{*}\right)$ is the strategy trajectory leading to the maximum total payoff.

### 4.1.2 Budget Multiplier

A significant contribution of the same paper is the notion of the Budget Multiplier, which quantifies the extent to which the initial difference in budgets determines the eventual dominance of a firm in the outcome of the game. More precisely, supposing that $K_{R}>K_{B}$ the Budget Mutiplier is the ratio $\frac{\Pi_{R}\left(\sigma_{R} \cdot \sigma_{B}\right)}{\Pi_{B}\left(\sigma_{R} \cdot \sigma_{B}\right)} \times \frac{K_{B}}{K_{R}}$ for the Nash equilibrium $\left(\sigma_{R}, \sigma_{B}\right)$ that maximizes this ratio. Intuitively, this quantity declares whether the initial suppremacy in budgets is reflected in the outcome (bigger Budget Multiplier implies more suppremacy of the initially dominant firm).

Similarly to the Price of Anarchy, the authors continue with offering upper and lower bounds for the Budgett Multiplier. For the selection function, they focus on Tullock context functions of the form $\frac{y^{s}}{y^{s+(1-y)^{s}}}$. For $s=1$ we have the linear selection function discussed above. For $s>1$, there is a polarizing selection function that favors a winner-take-it-all conception, whereas for $s<1$, there is an equalizing selection function that normalizes the initial inequalities. When either the selection function has a polarizing behavior or the switching function is strongly convex $\left(f\left(\frac{1}{2}\right)=0\right)$, the authors show instances where the budget multiplier, for any $V$ is greater than $V$.

On the other hand, for linear switching and selection functions, it is proved that the budget multiplier is at most 2 as the minority party (posit B) can copy the $K_{B}$ most profitable nodes of $R$. From those nodes, the red node gets $\sum_{i=1}^{K_{B}} \mathbf{E}\left[\chi_{R_{i}} \mid\left(S_{R}, \emptyset\right)\right] \geq \frac{K_{B}}{K_{R}} \mathbf{E}\left[\chi_{R} \mid\left(S_{R}, \emptyset\right)\right]$, supposing it plays alone, where $\chi_{R_{i}}$ is a random variable declaring the expected contribution in the payoff of this node. As a result, as all the functions are linear, playing this strategy, the Blue firm has at least $\frac{1}{2} \frac{K_{B}}{K_{R}} \mathbf{E}\left[\chi_{R} \mid\left(S_{R}, S_{B}\right)\right]$. Thus, in Nash equilibrium it will have at least this payoff (or else it would have incentive to deviate to the decribed strategy: $\mathbf{E}\left[\chi_{B} \mid\left(S_{R}, S_{B}\right)\right] \geq$ $\frac{1}{2} \frac{K_{B}}{K_{R}} \mathbf{E}\left[\chi_{R} \mid\left(S_{R}, S_{B}\right)\right]$. As a result, the budget multiplier is at most 2.

### 4.2 Attempts of computing Nash equilibria

Probably the first attempt to address equilibria questions arose in [2]. Their model included $n$ firms that target (for simplicity reasons) one node each. In the later steps, the nodes influence their neighbors. If a node gets influence by just one firm, it gets this product but if it gets multiple competitive influence, it abstains from buying it for the rest of the
game (as the authors claim there is great negative advertisement from the competing firms). For this simple model, they prove that Pure Nash Equilibrium always exists when the maximum distance of two nodes is 2 , but there are cases where it does not exist when it is greater than 2. This result, although seen as far fetched, seems to approach the form of social networks. Recently, Goel et al in [26] have presented a number of experiments on real social networks (Yahoo!, Zync, Twitter News) showing that the vast majority of diffusion occurs in the the first few steps, challenging the cascading models.

### 4.3 Compatibility in competitive diffusion

Until now, we have limited our approaches to cases where the potential buyers choose only one firm's product and can communicate just with those having picked the same firm. In [31], Immorlica et al waive these constraints and study the effect of multi-firm adoption and compatibility among those firms. Their model is an extension of Morris' Contagion threshold model presented in [39]

In his model, the firms are located in the ends of a line and there is a point, whose distance to a firm shows its superiority in quality comparing to its opponent. More precisely, a node's utility is $q\left|S_{A}\right|$ where $q \in[0,1]$ is a quality factor and $S_{A}$ the set of its neighbors that have selected product A, if it selects product A and $(1-q)\left|S_{B}\right|$, (where $S_{B}$ is the set of its neighbors that have selected product B ), if it selects product B . Morris interests in the significance of the quality in the dominance of a firm and searches the behavior of the system for different network structures and different quality factors, trying to figure out how much quality superiority is needed so that a newly introduced firm can dominate fully in the network, by endowing a small fraction of the total nodes in order to attract them. If it can dominate in the whole network, then the firm is called epidemic

Immorlica et al extend his model, permitting the buyers to take both products. In that case, the buyers enjoy full communication with all their neighbors but have a cost $c$ due to preserving both products (their utility function is $|S|-c$ where $S$ is the set of its neighbors). Like previously, they study the case that a new firm bribes some small fraction of the total nodes to join it. Afterwards, each node may change either by adopting just one product or both until no node has reason to deviate and increase
its payoff, where a Nash Equilibrium is reached. They prove, through a contradiction technique, that when a node decides to adopt the new firm or discard the old firm, this decision can never get recalled, which means that a Nash Equilibrium can always appear (it works like a potential function [38]) and, if there are infinite time steps, the order with which the update process occurs does not influence the eventual result. Continuing their results, they prove that, even in this generalized setting, a firm can never become epidemic if the other firm is of higher quality and, even when it is of higher quality, there are cases where it will stay non-epidemic

The second extension they make is letting firms have some limited compatibility among them (meaning that users of different firms have some limited utility from their communication). Although the results remain unchanged for the case of 2 firms under this new addition, this option offers the existent firms a powerful weapon to resist, should a new firm tries to enter the market, being superior in quality than them. In fact, the coalition of the existent firms is proved often capable to prevent such a firm from becoming epidemic and eliminating them.

## 5 Our results

One interesting extension of [27] contains studying how the firms will behave in the Switching selection model, presented in subsection 4.1, and the strategies they will implement. We first show a counterexample proving that Pure Nash Equilibria do not exist in all the instances of the problem. Afterwards we suggest a stochastically equivalent threshold model, where we give a Dynamic Programming Algorithm for computing a firm's payoff given the strategies of the firms. Although this algorithm is polynomial when the graph is a DAG(Directed Acyclic Graph), it is exponential in general graphs. We conjecture that the exact payoff computation for general graphs is NP-Hard. However, using sampling, we can $\epsilon$-approximate it $\forall \epsilon>0$. This means that there is an efficient way for a firm to decide its best response, given the strategy of its opponent (if the budgets are relatively small). Trying to approach the best response, we extend the hardness result of [33] and prove an inapproximability result when the thresholds take any deterministic value. Open directions contain answering questions regarding the computation of best response with thresholds chosen uniformally at random. This is a joint work with Dimitris Fotakis and Vangelis Markakis.

### 5.1 Pure Nash Equilibria

A Pure Nash Equilibrium, i.e. a Nash Equilibrium where the players have pure strategies (deterministically defined), is not always existent in our game. For instance, suppose there are two components $C_{1}$ and $C_{2}$, each of which is a star with $N$ nodes and both the switching and selection functions are linear. Moreover posit $K_{R}=3$ and $K_{B}=1$. It is obvious that the Red player will put some budget to the centre of both components, as it will ensure all the $N$ nodes of this component if the Blue player puts no budget there and some of it if the Blue player's budget is put to that componennt. Hence 2 of the Red player's budget can be put to the centre of $C_{1}$ and 1 to the centre of $C_{2}$. The Blue player will search to avoid the component where the Red player has put more budget and will select to put its budget to the centre $C_{2}$ to ensure payoff $\frac{N}{2}$. However, the Red player will have incentive to move one unit of budget to $C_{2}$ in order to increase its payoff from $N+\frac{N}{2}$ to $N+\frac{2 N}{3}$. Sequentially, the Blue player will have incentive to alter its
distribution and this never stops. It is easy to see that any other selection of initial nodes than targeting the central nodes of the components is not best response for a player. Hence, these games do not always possess Pure Nash Equilibria.

### 5.2 Equivalent Threshold Model

In order to collect all the randomness in the beginning of the process and then treat everything in a deterministic way, we suggest an equivalent threshold model to the switching-selection model, under the switching function we have presented in subsection 4.1.

We first remind the details of this function. We choose to give a cummulative pattern in the switching function. More precisely, $f\left(\alpha^{(k)}\right)$ shows the probability that a node would have bought the product before time step $k$, providing that $\alpha^{(k)}$ percentage of her neighbours have already adopted till then. Hence, at each step, the probability that a node adopts, posit $x$, is influenced only by the marginal influence of neighbors that have not influenced her in the past. As a result, the probability that the node will have adopted after $k$ steps is equal to the probability that she has adopted before the $k$-th step plus the conditional probability to adopt at this step. Formally:

$$
f\left(\alpha^{k}\right)=f\left(\alpha^{(k-1)}\right)+\left(1-f\left(\alpha^{(k-1)}\right)\right) x \Rightarrow x=\frac{f\left(\alpha^{k}\right)-f\left(\alpha^{(k-1)}\right)}{\left(1-f\left(\alpha^{(k-1)}\right)\right)}
$$

We observe that for every node there is a threshold behavior during the diffusion process, as there is some percentage of neighbors that lead to her eventual adoption. This is decided during the process as it is a cascade model. In order to collect all the randomness in the initial step, we pick this threshold in the beginning, randomly from the distribution $f^{-1}$.

As for the selection part, we number the neighbors of each node and pick a number $\theta$ for each node uniformally at random. When the switching function decides that the node adopts, then each of her live neighbors (already adopters) takes a part of $[0,1]$ according to the distribution $g$ and the relevant position in the initial numbering. Obviously, $\theta$ is in one of these parts and this shows who will influnce her to decide which firm to follow.

### 5.3 Exact computation of the utility

In order to deal with the computation of the utility in general cases, we first present a dynamic programming algorithm for directed acyclic graphs (DAG). For every node and every step, we keep $R_{i, k}\left(B_{i, k}\right)$ denoting the probability that this node was infected strictly before step $k-1$ by red (blue) firm. Furthermore, we keep $R_{c, k}\left(B_{c, k}\right)$ denoting the probability that this node was infected in the step $k-1$ and is thus contagious. The game ends after at most $N$ steps as, in order to continue, a shift in at least one node must occur. This happens because the nodes are $N$, they are colored at most once and a shift in a step demands a shift in the previous step so that there are contagious nodes.

Initially, we set:

- $R_{c, 0}=0, B_{c, 0}=0$ if no firm has targeted node $i$ in its initial set
- $R_{c, 0}=1, B_{c, 0}=0$ if just the red firm has targeted node $i$
- $R_{c, 0}=0, B_{c, 0}=1$ if just the blue firm has targeted node $i$
- $R_{c, 0}=\frac{a_{R}}{a_{R}+a_{B}}, B_{c, 0}=\frac{a_{B}}{a_{R}+a_{B}}$ if the red firm has put a budget of $a_{R}$ to node $i$ and the blue firm has put a budget of $a_{B}$ to it.

In step $k$, node $j$ is influenced just by the contagious nodes. The probability of that influence is conditional to the fact that it was not previously colored by another node. Hence, this probability is:

$$
x=\frac{f\left(\sum_{i \in N(j)}\left(R_{i, k-1}+B_{i, k-1}+R_{c, k-1}+B_{c, k-1}\right)\right)-f\left(\sum_{i \in N(j)}\left(R_{i, k-1}+B_{i, k-1}\right)\right)}{1-f\left(\sum_{i \in N(j)}\left(R_{i, k-1}+B_{i, k-1}\right)\right)}
$$

If it is decided to get colored then it is colored red w.p.

$$
y=g\left(\frac{\sum_{i \in N(j)}\left(R_{i, k-1}+R_{c, k-1}\right)}{\sum_{i \in N(j)}\left(R_{i, k-1}+B_{i, k-1}+R_{c, k-1}+B_{c, k-1}\right)}\right)
$$

and blue w.p. $1-y$.
Following to that, we update $R_{i}, B_{i}, R_{c}, B_{c}$ :

- $R_{i, k}=R_{i, k-1}+\left(1-R_{i, k-1}\right) R_{c, k}$
- $B_{i, k}=B_{i, k-1}+\left(1-B_{i, k-1}\right) B_{c, k}$
- $R_{c, k}=x y$
- $B_{c, k}=x(1-y)$

The expected utility of a firm is the sum of the probabilities of each node being colored by firm's color at step $N: \Pi_{p}=\sum_{i} p_{i, N}$

The above algorithm is polynomial to the number of nodes and has complexity $O\left(N^{3}\right)$.

The reason that it works is because the influence has a certain direction and hence the increase in the probability that a node $i$ is colored due to $j$ cannot enhance $j$ 's probability of being colored. In order to extend the result to general graph we need to keep, for each probability, a history of those that participated in its augmentation and exclude phenomena of influencing the influencers. This demands an additional $2^{N}$ for every such probability and makes the algorithm exponential.

We conjecture that this cycle phenomenon cannot be faced in order to compute exactly the payoff in polynomial time. However, as we show in the next subsection, it can be approximately computed with as good approximation as we desire using sampling.

### 5.4 Approximate computation of the utility

In order to compute the utility in polynomial time, we use sampling. From Hoeffding bounds, it is known that when $X_{1}, X_{2}, \ldots, X_{n}$ are independent random variables that are bounded in $\left[a_{i}, b_{i}\right]$ then, by sampling, we can approximate the expected value. In our case, these independent random variables are the payoffs of the process. Thus, if $S=X_{1}+X_{2}+\cdots+X_{n}$ then:

$$
\operatorname{Pr}(|S-E[S]| \geq t) \leq 2 \exp \left(-\frac{2 t^{2}}{\sum_{i=1}^{n}\left(b_{i}-a_{i}\right)^{2}}\right)
$$

We need to approximate $E[S]$, which is the expected value of the sum of $n$ variables. As a result, in order to achieve $\epsilon$ additive approximation, we need to set $t=n \epsilon$. Then, we shall find the $n$ that makes this approximation almost certain (sets the right part of the comparison in an extremely low value).

As the possible values of the random variables are numbers of nodes, they will be between $a_{i}=0$ and $b_{i}=N$. As a result, we have:
$\operatorname{Pr}(|S-E[S]| \geq n \epsilon) \leq 2 \exp \left(-\frac{2 n^{2} \epsilon^{2}}{n N^{2}}\right) \Rightarrow \operatorname{Pr}(|S-E[S]| \geq n \epsilon) \leq 2 \exp \left(-\frac{2 n \epsilon^{2}}{N^{2}}\right)$
As a result, in order to achieve an $\epsilon$-approximation of the utility w.p. $1-\epsilon^{\prime}$, we need $n$ samples s.t. $\epsilon^{\prime}=2 \exp \left(-\frac{2 n \epsilon^{2}}{N^{2}}\right) \Rightarrow-\frac{2 n \epsilon^{2}}{N^{2}}=\ln \left(\frac{\epsilon^{\prime}}{2}\right) \Rightarrow \frac{2 n \epsilon^{2}}{N^{2}}=$ $\ln \left(\frac{2}{\epsilon^{\prime}}\right) \Rightarrow n=\ln \left(\frac{2}{\epsilon^{\prime}}\right) \frac{N^{2}}{2 \epsilon^{2}}$.

With appropriate selection of $\epsilon, \epsilon^{\prime}$ and thus of the number of samples, we have as good an approximation as we wish.

### 5.5 Hardness of best response computation

In [33], the authors give an inapproximability result (presented in subsection 2.2) which states that finding the initial influential set is NP-hard to approximate within $O\left(n^{1-\epsilon}\right)$ for all $\epsilon>0$. Their reduction is from Set Cover but demands extreme deterministic thresholds. In the one layer, a node buys the product if at least one neighbor has bought it, whilst in the other it buys if all her neighbors have bought it. In this subsection we extend their result for every possible deterministic thresholds and not only for extreme cases.

In fact, this shows that the best response in our problem is NP-hard to approximate within $O\left(n^{1-\epsilon}\right)$ when the thresholds are deterministically chosen, even in the absense of the opponent. We will restrict our reduction to the case that all the threshols are deterministically chosen at $\frac{1}{2}$, i.e. a node adopts the product if at least half of her neighbors have already adopted it.

Proof. Suppose an instance of the Set Cover problem. There exist $m$ sets $S_{i} \in S$ and $n$ elements $e_{j} \in E$. Each set is a subset of the elements $S \subset E$. We need to find $k$ sets such that all the elements exist in their union.

We will create an instance of our problem. As in [33], we first create a node $S_{i}$ for each set and a node $e_{j}$ for each element. We put a directed edge from $\left(S_{i}, e_{j}\right)$ if and only if $e_{j} \in S_{i}$. Moreover, we create $N$ extra nodes and put directed edges from all $e_{j} \in E$ to all those $N$ nodes.

Now we need to formulate the $\frac{1}{2}$ threshold for the layer of the elements. Hence, we create a universal node $u$. Furthermore, for every $e_{i} \in E$ we create in_degree $(i)-1$ nodes $v_{i, j}$ (one less than the number of her total appearances in the sets of $S$ ). We add the edges ( $u, v_{i, j}$ ) for all those nodes and the edges $\left(v_{i, j}, e_{i}\right)$ for the element node to which they refer.

Following to that, we fix the $\frac{1}{2}$ threshold for the layer of the extra nodes. We create $n+1$ dummy nodes and put directed edges from them to all the $N$ extra nodes. Last but not least, we add an edge between the universal node and each of the $N$ extra nodes.

We need to select $k^{\prime}=k+1$ initial influential nodes in order to maximize the number of eventual adoptions. We will show that, if a $k$-set cover exists then we can make all the $N$ extra nodes adopt. Unless this happens, we can take none of them if we do not target it in the initial set. Hence, with appropriate choice of $N$, the inapproximability result will follow.

If there exists a $k$-set cover then, by selecting those nodes and the universal nodes as our initial influential set, we can make all the element nodes adopt. The universal node makes all $v_{i, j}$ nodes buy (as their indegree is 1 , their adoption depends only on the adoption of $u$ ). In order for an element node to reach the $\frac{1}{2}$ threshold, one more adopting neighbor is needed. This comes from the set cover as each element is in at least one of the selected sets. Afterwards, all the $N$ extra nodes buy as half of their neighbors (the element nodes and the universal node have already adopted). Hence, $N+n+k+1$ nodes adopt the product.

If we pick the universal node and there is no $k$-set cover then, in the solution that maximizes the eventual number of nodes adopting the product, at least on element node does not buy (posit $e_{j}$ ). As there is no $k$-set cover, we cannot pick set nodes such that all the elements are in them. Hence, $e_{j}$ will have less than the half of her neighbors adopting the product and, as a result, will not adopt. This means that the $N$ extra nodes will not adopt the product, as they will have at most the influence of the universal node and $n-1$ from the element nodes, leading to $\frac{n}{2 n+2}<\frac{1}{2}$ of her neighbors. It is obvious that we have none of the dummy nodes adopts the product because, if one of them has adopted, we can pick a set containing $e_{j}$ instead and win at least as much.

If we do not pick the universal node, then even if we manage to take all the element nodes ( $k+1$-set cover), we cannot take any of the $N$ extra nodes as they will receive again influence less than their threshold's. As previously, we cannot have a dummy node selected as it would be more profitable to replace it with either the universal node or one set node containing an element node not selected.

Thus, unless a $k$-set cover exists, we can make at most $(n-1)+k+1=$ $n+k$ adopt whilst, if a $k$-set cover exists, we can make $N+n+k+1$ nodes adopt. The total number of nodes is $M=N+m+n+(n+1)+(m-n)+1=$ $N+2 m+n+2$. Fixing $N=\Omega\left(M^{1-\epsilon}\right)$, the inapproximability result follows.

## Part II

## Competitive Pricing under externalities

## 6 Pricing without externalities

In the previous part, the main objective of the firms was to spread their influence to the maximum possible number of buyers. Hence, the problem was to find the initial influential set that would enable the maximum eventual expansion.

However, seeking to maximize their revenue, firms take advantage of another degree of freedom, price. Sometimes it is more profitable to have less diffusion when each buyer pays more for the product.

In this section, we will thus neglect the importance of the externalities in convincing potential buyers to adopt a product and will focus on the power of price and how it can be used to make them pay the maximum amount of money they are willing to. The goal of the section is to introduce the reader to the area of pricing and discuss some correlations with gametheoretic and algorithmic areas. Hence we will not cover all the litterature of the area.

### 6.1 Auctions vs Pricing

This direction was widely introduced in [28]. Supposing that many different objects are offered, each potential buyer has a valuation for each subset of them. Guruswami et al are interested in designing pricing algorithms that are envy-free, i.e. allocations of bundles to buyers at some prices such that no buyer would prefer another bundle and its price to what he has received. As the general case is intractable due to the vast variety of different possible allocations and the limited supply of objects, the authors specify to two cases: when each bundle contains just one object (unitdemand bidders) and when each buyer has positive valuation just for one object (single-parameter bidders). Even these special cases are APX-hard and the authors give logarithmic approximation algorithms. Dealing with the unit-demand bidders' problem, Guruswami et al regard the solution as a matching between the objects and the customers. In order to ensure envy-free pricing algorithms, they cut low-priced edges of the matching, putting reserve prices to the objects (prices under which the object is given to nobody). This elucidates a correlation between optimal algorithmic pricing and algorithmic multi-parameter mechanism design.

Algorithmic mechanism design is a very active research area. The focal
point in such problems is to find a mechanism that decides the allocation of the products to buyers and the prices they will pay, ensuring the maximum amount of profit to the auctioneer. Buyers might misreport their true valuation for the objects, trying to increase their own utility and thus such mechanisms usually have to be truthful, giving no incentive to the buyers to deviate from reporting their true valuations. Single-parameter mechanism design, i.e. the case where each buyer reports a single value, is solved by Myerson. In his seminal paper [41], he considers the case that the values of each buyer are private and come from a distribution, known to the auctioneer. He implements an auction on his virtual valuations, that take into account both this distribution and the declared buyers' bid and using a technique called ironing in order to make the virtual valuation function non-decreasing. The item is afterwards given to the buyer with the biggest virtual valuation in a price that comes from his own virtual valuation function and the second bigger virtual valuation. This algorithm solves the single-parameter problem optimally. Much research is done to the direction of addressing the multi-dimensional problem but, apart from its relations with the Pricing problem, is out of the scope of this thesis.

In [14] and [15], the Algorithmic Pricing problem and the Algorithmic mechanism design problem are directly related. In [14], Chawla et al use Myerson's technique to approximate the optimal solution of the aforementioned Unit-demand Pricing Problem with a constant factor of 3, when the valuations of the players are independent random variables. They also suggest Polynomial Approximation Schemes for the case where the valuations are regular (the virtual valuations are non-decreasing). Extending their work in [15], the authors arrive at good approximations of the optimal single-dimensional pricing solution by using sequential posted-pricing mechanisms, where buyers in sequence are given a take-it-or-leave it price. This result is generalized in some multi-dimensional settings, showing the implicit connection among these two problems.

This framework is discussed and extended in some other papers with positive and negative results. In [13], Cai and Daskalakis stress that there exists a PTAS for computing a price vector whose revenue is a $(1+\epsilon)$ approximation of the optimal revenue both for the regular case and the Monotone Hazard Condition case (analyzed in subsection 7.1). In both cases, a constant number of distinct prices are needed. As the authors analyze some structural properties of the problem, they reduce the vec-
tor space of possible solutions and, as a result, do not use a loose lower bound for the optimal solution, unlike previous work on the same problem. On the opposite side of the coin, optimal pricing in the unit-demand bidder is hard when the bidders' prices are correlated [12] and in various other Bayesian settings when prices follow some specific distributions [22], making resorting to approximation schemes vital.

### 6.2 Unknown distributions

While the previous work considers that the distribution of the valuations is known, this is not always the case. In many cases, the seller should extract this information during the selling process through a sampling step in order to be able then to design an effective pricing policy. This problem has a direct correlation to the Secretary problem where the potential employees come online for an interview and the employer should inform them for the outcome just after the interview. In [6], Babaioff et al address the onedimensional problem of designing such a seller's strategy when the bidders' utility come from a distribution which is unknown and the bidders are unwilling to reveal. Focusing on sequential posted-prices, they show that the optimal revenue in such a case is $\Omega\left(\frac{\operatorname{logn}}{\log \log n}\right)$ to the optimal online revenue when the distribution is known. However, when we restrict the possible distributions to those that satisfy the Monotone hazard rate condition, this ratio diminishes radically to just a constant factor, when the number of players is significantly bigger than the support of the distributions.

## 7 Influence and exploit

In this section, we will focus on monopoly Influence-and-exploit strategies that, in fact, bind the aforementioned directions. At the first step, the firm offers for free the good to some selected nodes (the influencers) and subsequently exploits the externalities they induce to the rest in order to sell the product to them in a higher price.

The idea was introduced in the seminal paper [30]. There, Hartline, Mirrokni and Sundararajan suggest two main models to capture the influence exerted: the Uniform Additive Model where the influence of each neighbor that has already bought the product is independent and works in an additive way and the Concave Graph Model where this influence is a submodular function of the set influencing the potential buyer. In both models the authors have suggested some results that were later ameliorated. Another direction that this paper suggested was the idea of fixedprices where the firms have no power on defining different prices for each buyer.

### 7.1 Marketing Strategies

The marketing strategy of a firm depends on some decisions it must take. At first, it should decide the sequence with which it will approach the potential buyers and then offer them such prices that will maximize its expected revenue.

In [30], the authors simplify the second step by assuming that sellers are myopic. As a result, when selecting the price, they try to maximize the revenue from each buyer neglecting its potential influence on the following buyers. This price is called myopic price.

Challenging this assumption, Fotakis and Siminelakis [25], focusing on the Undirected Additive Model, proved that offering prices smaller that the myopic is often more profitable. Selling more inexpensively, sellers augment their influential set and this influence exerted on future buyers compensates for the myopic loss.

Approximating the optimal solution of these problems was a main issue in both papers. In [30], the authors prove the NP-hardness of choosing a myopic marketing strategy, i.e. choosing the sequense that maximizes the revenue. The approximation algorithm they suggest lacks a tight upper
bound. Following this work, the authors of [25], after proving the NPhardness of the non-myopic case, improve the approximation ratio using semidefinite programming and observing the relation between large cuts and good influence-and-exploit strategies.

### 7.2 Submodular externalities

In the Concave Graph Model, Hartline et al consider a useful property, the monotone hazard rate condition, in order to create their approximation algorithms.

Definition 3. A distribution, with density function $f$ and distribution function $F$, satisfies the monotone hazard rate condition if $h(t)=\frac{f(t)}{1-F(t)}$ is monotone non-decreasing.

Supposing that this condition holds and that the revenue functions for each player are submodular, they arrive at a $\frac{1}{4}$-approximation algorithm of the optimal revenue in the Concave Graph Model. In fact, the algorithm is very simple. Each potential buyer is chosen in the Influencers w.p. $\frac{1}{2}$. Then the rest pay at least half what they would pay if they were globally influenced, due to the submodularity of the revenue function. Hence, the approximation ratio comes straightforward.

In [29], Haghpanah et al consider the problem of auctions with positive network externalities, inspired by [30]. Using the submodularity of the expected revenue they can extract from one set and recent results in submodular maximization, they increase this result.

An interesting class of submodular externalities where the authors of [29] focused is the step-function externalities. In this case, a node buys the product if at least one of its neighbors has already bought it. This ressembles to a graph-theoretic problem, where we need to find a maximumweight subset of (possibly negative-weighted) vertices, whose induced subgraph contains no isolated vertices.

Facing the problem we discussed previously with the loose upper bound, they made a LP-based auction and rounded its solution. The result was a 0.73 -approximation for the initial problem, leading to a huge increase to the ratio.

### 7.3 Extensions

Based on the Influence-and-exploit model, many research works tried to pose questions towards a better understanding of the area.

In [5], Arthur et al enrichened the model, rewarding already adopters that influence their friends with a cashback offer for every customer they convince (to buy the product). This offers them incentive to personalize their influence and make their advertisement more successful. As a result, the offer of the seller to a prospective buyer includes both a price for adopting the product and a commitment in a cashback price for every subsequent adopter influenced by her. As was previously described, marketing in such a setting is the result of many biased coinflips (which decide whether a buyer will purchase the product or not). As a result, it is sensible to assume that the firms should better adapt their strategy depending on the state of the random process. Surprisingly, the authors suggest a seller strategy that, despite being non-adaptive, incurs revenue just a constant factor away from what could have been acquired by an adaptive strategy.

Another interesting extension of the original model is found in [37] where the authors permit just a fixed-price in the exploit-step for all the buyers, again under the existence of positive network externalities. The seller's strategy is to select the initial influential set and the fixed-price $p$. Bounding the possible values of $p$, they sample in order to generate oracle calls for each of them, influenced by [33], and receive a $\frac{1}{2}$ approximation for the maximum revenue. Then they prove that the expected revenue from each user in a price $p$ is a sumodular function of the influence set. Then, the approximation ratio comes from a recent result in submodular maximization.

## 8 Price trajectory under externalities

In this section, we allow firms to alter their price as time passes. The original Influence-and-exploit model includes a setting where a different price is destined to every customer and the game is one-shot. On the contrary, in posted pricing, the seller's strategy is a price trajectory that is addressed to all the customers and the game evolves in many stages.

The reasoning behind these decisions are based on the way advertising works and the significant role externalities play. Firstly, a new product can take some time to get best known and, as a result, its initial pricing should be low in order to urge people adopt it. What is more, some trial editions might include bugs that justify a significant discount in their early adoption. As the process goes by, the product becomes more known and better designed, which induces a higher price. Last but not least, the products are usually promoted through advertisement so the offers are usually universal and not customized to every possible buyer.

### 8.1 Modeling and exploiting externalities

The idea of avoiding price discrimination was introduced in [1]. Akhlaghpour et al supposed that the strategy of the firms was to choose a $k$-price trajectory for all the steps of the game. In both their models, they assumed that each customer has a valuation for the product that is increased in an additive way due to the externalities posed by their friends. Their Basic ( $k$ ) model suggests that the externality acts immediately and, as a result, the valuations alter during the same step. This obliges the firms to adopt a decreasing price trajectory that seems weird (as the only reason not to adopt a product is that she does not value it so highly, after all externalities are exerted, and this means that a higher price in the next step can offer nothing to the seller). On the other hand, their $\operatorname{Rapid}(k)$ model is more sensible as it assumes that the interactions take some time to occur and hence the valuation is updated in the next step. Finding an optimal strategy in such a setting is NP-hard to approximate within a constant factor. This follows by a neat reduction from the Independent Set.

A slightly different perspective is given by Anari et al in [3], where the buyers are not considered myopic. There, the price trajectory is explicitly given to them initially (for instance, trial editions' price, early adoption,
etc) and the customers adopt the product if and when they estimate that their utility will get larger. The essential difference is that the externalities in their model do not just increase the buyers' opinion towards the model but also enhance the quality of the model, as bugs may be fixed and the product is better tested in subsequent versions. In order to deal with this problem, the authors seperate the customers to types (each one describing people with similar valuation functions) and study the equilibria of this game and approximations on the best revenue strategy in some special cases (that differ in the nature of the valuation functions).

In some cases, the externalities do not act in a positive way. For instance, discriminating in a company by the products acquired is something that often provides joy to some people. In order to model this feature in [10], Bhattacharya et al examine a game where two products are offered (a cheaper with fewer assets and a more expensive one). The valuation of somebody towards the expensive product is increasing for every friend that has not adopted it, as this offers them the opportunity to boast for this acquisition. This effect of negative externalities is an interesting direction towards a better understanding of the interactions among people.

### 8.2 Competitive Pricing under externalities

In the previous part, we have discussed competition in their effort to disseminate their propagation and buyers' behavior was a result of the influence exerted to them. In [8], the problem is discussed from the buyer's point of view. Many competing products are offered in different prices. The valuation of the buyer towards each product is a combination of his initial (intrinsic valuation) and the increase made due to the externalities. The authors take into consideration simple cases of externalities (concave, convex and step-function), model them by linear programs through very elegant transformations and round their solutions in order to receive good approximation ratios for the welfare in such a game. Moreover, they prove that pure Nash equilibria always exist in such a game by showing the existence of a potential function. Last but not least, they study revenue maximizing mechanisms, using ideas already discussed in section 6. This is the first approach towards competitive pricing, taking into account the effect of externalities.

## 9 Competition Pricing Models

In this last section, we will focus on the two most prominent competition pricing models. In the Cournot model, the firms strategize on the quantities produced and the prices are implicitly given by them. On the other hand, in the Bertrand model, the prices are directly chosen by the firms and the competition is on them. Our main scope will be to elaborate on the algorithmic game-theoretic handling of those models but first it is useful to present some details about the models.

According to [46], in 1838, Cournot, who is considered as the pioneer of the mathematical formulation in the economics, tried to capture the competition among the mineral water production firms [21]. Each firm's strategy was to decide the quantity produced and the price was inversely related to the total quantity produced. Its goal was to cover the marginal cost of the production and ensure the maximum possible profit. It is obvious that firms may have incentive to alter their strategy, regarding their competitors' strategies in order to win more, which is very similar to the notion of equilibria, later formalized by Nash.

In 1883, Bertrand, reviewing this model [7], demonstrated its inefficiency when the marginal costs are zero and suggested that the main firms' strategy is the price and not the quantities. The buyers select the cheapest product and, again, the firms aim to maximize their profit. This model was later formalized by Edgeworth and catches the notion of perfect competition even in the duopoly case.

Nevertheless, the fact that even a slight price discrimination in the case of undifferentiated goods leads to a total dominance of one firm led many researchers to judge this model as unrealistic. In fact, in the duopoly setting, firms tend to lower their bids in order to surpass their competitor and the prices can end up to be equal to the initial marginal cost, offering zero profit to the firms. This is known as the Bertrand paradox and is the main critisism against the latter model.

To sum up, both models seem to provide some insight in the oligopoly situation but not a full understanding of how competition works. However, they both proved to be influential for the area as many interesting results have used them as their starting point.

### 9.1 Bertrand competition

In order to deal with the Bertrand paradox, Nadav and Piliouras [42] have used a regret minimization technique. In their model, the game is played in steps and, in each step, buyers prefer the cheapest product (if there are more than one, they are selected in equal proportion). At each step, firms select online, through an algorithm $A$, the distribution from which their bid comes from, that is their strategy. If an alteration of this strategy cannot, ultimately, improve their total payoff significantly $(\Omega(T)$, where $T$ is the number of steps), the algorithm $A$ is said to have no regret. The authors suggest a no regret algorithm that offers coarse correlated equilibrium which means that, supposing that all other firms follow a probability distribution, one cannot expect to gain more by following a single strategy instead of it. They produce such a distribution and prove that all single strategies offer less expected payoff. Each firm's payoff is significantly greater than the marginal cost, hence the Bertrand paradox is overcome. However, as the number of firms increases, the expected payoff rapidly returns to the marginal cost.

From another perspective, Chawla, Niu and Roughgarden [17],[16] study Bertrand competition in network setting, emphasizing on the inefficiency of the equilibria. In their model, each edge gets priced by its owner in the first step and afterwards the users decide the path that will lead them from their source to their sink in the minimum cost. The Prices of Anarchy and Stability differ much depending on the number of monopolies, that is edges that are cuts for this flow. In the absense of monopolies, the social welfare is approximately reached in every equilibrium, whereas the sellers' profit is devastated as the competition draws down the prices. On the other hand, should a monopoly exist, these quantities are inverse, as the monopoly dominates and imposes its own prices. These results were proven for the single-source single-sink model and were extended for the multiple-source single-sink model, where the Prices of Anarchy and Stability are bounded by the sparsity of the network. If at least two monopolies exist, the price of anarchy can become unbounded, however if the Monotone Hazard Rate Condition holds then it is bounded again by a quantity depending on the network's sparsity and the number of monopolies. As far as the multiple-source multiple-sink model is concerned, they prove that the price of anarchy can become unbounded.

### 9.2 Cournot competition

In the linear Cournot setting, each firm decides the quantity it is about to produce and the price of the oligopoly is inversely related to the total quantity produced. The utility of each firm is then the quantity it produced times the price in which it was sold. In fact, this procedure leads to a unique equilibrium that is usually refered as Cournot-Nash equilibrium.

Inspired by relevant economic work, Immorlica et al [32] pose questions on the features of this equilibrium in the symmetric case, should the firms are allowed to make coalitions. In their Coalition formation game, the utility of each member of a coalition is its equal proportional share to the coalition. They suggest all the possible deviations that can be made, including some members joining an existing coalition, dissolving it or inauguring another coalition and show the necessary conditions such that a partition of firms in coalitions is stable, that is no one has incentive to deviate and increase its utility. Using these conditions, they prove upper and lower bounds for the inefficiency of a possible equilibrium, showing that the Price of Anarchy is $\Theta\left(n^{2 / 5}\right)$.

Another direction appears in [24], where Fiat et al waive the myopic restriction of the Cournot model towards best response. In their model, firms can choose either to maximize profit (PM), which is equal to the myopic Cournot strategy, or revenue (RM), considering that the virtual cost of production equals to zero. Although seemingly absurd, seeking to maximize revenue is often more profitable. As firms' strategy is either PM or RM, some sequences of stretegies' alterations can never occur under best response behavior. Consequently, the authors prove that any succession of best response moves converges to a pure Nash equilibrium in at most linear number of steps, which shows an efficient algorithm for its computation. As for the quality of this equilibrium, the authors derive the conclusion that non-myopic action induces lower prices for the products, causing a significant reduction in the firms' profit.

## References

[1] Hessameddin Akhlaghpour, Mohammad Ghodsi, Nima Haghpanah, Vahab S. Mirrokni, Hamid Mahini, and Afshin Nikzad. Optimal iterative pricing over social networks. In Proceedings of the 6th international conference on Internet and network economics, WINE'10, pages 415-423, Berlin, Heidelberg, 2010. Springer-Verlag.
[2] Noga Alon, Michal Feldman, Ariel D. Procaccia, and Moshe Tennenholtz. A note on competitive diffusion through social networks. Inf. Process. Lett., 110(6):221-225, February 2010.
[3] Nima Anari, Shayan Ehsani, Mohammad Ghodsi, Nima Haghpanah, Nicole Immorlica, Hamid Mahini, and Vahab S. Mirrokni. Equilibrium pricing with positive externalities. In Proceedings of the 6th international conference on Internet and network economics, WINE'10, pages 424-431, Berlin, Heidelberg, 2010. Springer-Verlag.
[4] Krzysztof R. Apt and Evangelos Markakis. Diffusion in social networks with competing products. In Proceedings of the 4 th international conference on Algorithmic game theory, SAGT'11, pages 212-223, Berlin, Heidelberg, 2011. Springer-Verlag.
[5] David Arthur, Rajeev Motwani, Aneesh Sharma, and Ying Xu. Pricing strategies for viral marketing on social networks. In Proceedings of the 5th International Workshop on Internet and Network Economics, WINE '09, pages 101-112, Berlin, Heidelberg, 2009. Springer-Verlag.
[6] Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. Posting prices with unknown distributions. In In Innovations in Computer Science (ICS), 2011. [BH08] L. Blumrosen and.
[7] J. Bertrand. Book review of theorie mathematique de la richesse sociale and of recherches sur les principles mathematiques de la theorie des richesses. 1883.
[8] Anand Bhalgat, Sreenivas Gollapudi, and Kamesh Munagala. Mechanisms and allocations with positive network externalities. In Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12, pages 179-196, New York, NY, USA, 2012. ACM.
[9] Shishir Bharathi, David Kempe, and Mahyar Salek. Competitive influence maximization in social networks. In Proceedings of the 3rd in-
ternational conference on Internet and network economics, WINE'07, pages 306-311, Berlin, Heidelberg, 2007. Springer-Verlag.
[10] Sayan Bhattacharya, Janardhan Kulkarni, Kamesh Munagala, and Xiaoming Xu. On allocations with negative externalities. In Proceedings of the 7th international conference on Internet and Network Economics, WINE'11, pages 25-36, Berlin, Heidelberg, 2011. SpringerVerlag.
[11] Allan Borodin, Yuval Filmus, and Joel Oren. Threshold models for competitive influence in social networks. In Proceedings of the 6th international conference on Internet and network economics, WINE'10, pages 539-550, Berlin, Heidelberg, 2010. Springer-Verlag.
[12] Patrick Briest. Uniform budgets and the envy-free pricing problem. In Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I, ICALP '08, pages 808-819, Berlin, Heidelberg, 2008. Springer-Verlag.
[13] Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS '11, pages 522-531, Washington, DC, USA, 2011. IEEE Computer Society.
[14] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In Proceedings of the 8th ACM conference on Electronic commerce, EC '07, pages 243-251, New York, NY, USA, 2007. ACM.
[15] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In Proceedings of the 42nd ACM symposium on Theory of computing, STOC '10, pages 311-320, New York, NY, USA, 2010. ACM.
[16] Shuchi Chawla and Feng Niu. The price of anarchy in bertrand games. In Proceedings of the 10th ACM conference on Electronic commerce, EC '09, pages 305-314, New York, NY, USA, 2009. ACM.
[17] Shuchi Chawla and Tim Roughgarden. Bertrand competition in networks. In Proceedings of the 1st International Symposium on Algo-
rithmic Game Theory, SAGT '08, pages 70-82, Berlin, Heidelberg, 2008. Springer-Verlag.
[18] Wei Chen, Yifei Yuan, and Li Zhang. Scalable influence maximization in social networks under the linear threshold model. In Proceedings of the 2010 IEEE International Conference on Data Mining, ICDM '10, pages 88-97, Washington, DC, USA, 2010. IEEE Computer Society.
[19] Andrew Clark and Radha Poovendran. Maximizing influence in competitive environments: a game-theoretic approach. In Proceedings of the Second international conference on Decision and Game Theory for Security, GameSec'11, pages 151-162, Berlin, Heidelberg, 2011. Springer-Verlag.
[20] Gerard Cornuejols, Marshall L. Fisher, and George L. Nemhauser. Location of Bank Accounts to Optimize Float: An Analytic Study of Exact and Approximate Algorithms. Management Science, 23(8):789810, 1977.
[21] A. A. Cournot. Recherches sur les principes mathématiques de la théorie des richesses/par Augustin Cournot. L. Hachette, 1838.
[22] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Optimal pricing is hard. In WINE, pages 298-308, 2012.
[23] Pedro Domingos and Matt Richardson. Mining the network value of customers. In Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining, KDD '01, pages 57-66, New York, NY, USA, 2001. ACM.
[24] Amos Fiat, Elias Koutsoupias, Katrina Ligett, Yishay Mansour, and Svetlana Olonetsky. Beyond myopic best response (in cournot competition). In Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '12, pages 993-1005. SIAM, 2012.
[25] Dimitris Fotakis and Paris Siminelakis. On the efficiency of influence-and-exploit strategies for revenue maximization under positive externalities. CoRR, abs/1110.1894, 2011.
[26] Sharad Goel, Duncan J. Watts, and Daniel G. Goldstein. The structure of online diffusion networks. In Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12, pages 623-638, New York, NY, USA, 2012. ACM.
[27] Sanjeev Goyal and Michael Kearns. Competitive contagion in networks. In Proceedings of the 44 th symposium on Theory of Computing, STOC '12, pages 759-774, New York, NY, USA, 2012. ACM.
[28] Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In Proceedings of the sixteenth annual ACMSIAM symposium on Discrete algorithms, SODA '05, pages 11641173, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.
[29] Nima Haghpanah, Nicole Immorlica, Vahab Mirrokni, and Kamesh Munagala. Optimal auctions with positive network externalities. In Proceedings of the 12th ACM conference on Electronic commerce, EC '11, pages 11-20, New York, NY, USA, 2011. ACM.
[30] Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In Proceedings of the 17th international conference on World Wide Web, WWW '08, pages 189198, New York, NY, USA, 2008. ACM.
[31] Nicole Immorlica, Jon Kleinberg, Mohammad Mahdian, and Tom Wexler. The role of compatibility in the diffusion of technologies through social networks. In Proceedings of the 8th ACM conference on Electronic commerce, EC '07, pages 75-83, New York, NY, USA, 2007. ACM.
[32] Nicole Immorlica, Evangelos Markakis, and Georgios Piliouras. Coalition formation and price of anarchy in cournot oligopolies. In Proceedings of the 6th international conference on Internet and network economics, WINE'10, pages 270-281, Berlin, Heidelberg, 2010. SpringerVerlag.
[33] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, KDD '03, pages 137-146, New York, NY, USA, 2003. ACM.
[34] David Kempe, Jon Kleinberg, and Éva Tardos. Influential nodes in a diffusion model for social networks. In Proceedings of the 32 nd international conference on Automata, Languages and Program-
ming, ICALP'05, pages 1127-1138, Berlin, Heidelberg, 2005. SpringerVerlag.
[35] Jan Kostka, Yvonne Anne Oswald, and Roger Wattenhofer. Word of mouth: Rumor dissemination in social networks. In Proceedings of the 15 th international colloquium on Structural Information and Communication Complexity, SIROCCO '08, pages 185-196, Berlin, Heidelberg, 2008. Springer-Verlag.
[36] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In in Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, pages 404-413, 1999.
[37] Vahb S. Mirrokni, Sebastien Roch, and Mukund Sundararajan. On fixed-price marketing for goods with positive network externalities. WINE 2012, 2012.
[38] Dov Monderer and Lloyd S. Shapley. Potential games. Games and Economic Behavior, 14(1):124-143, 1996.
[39] Stephen Morris. Contagion, 1997.
[40] Elchanan Mossel and Sebastien Roch. On the submodularity of influence in social networks. In Proceedings of the thirty-ninth annual ACM symposium on Theory of computing, STOC '07, pages 128-134, New York, NY, USA, 2007. ACM.
[41] R. Myerson. Optimal auction design. Mathematics of Operations Research, 6(1):58-73, 1981.
[42] Uri Nadav and Georgios Piliouras. No regret learning in oligopolies: cournot vs. bertrand. In Proceedings of the Third international conference on Algorithmic game theory, SAGT'10, pages 300-311, Berlin, Heidelberg, 2010. Springer-Verlag.
[43] J.F. Nash. Non-cooperative games. Annals of Mathematics, 54(2):286-295, 1951.
[44] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions-i. Mathematical Programming, 14:265-294, 1978. 10.1007/BF01588971.
[45] Matthew Richardson and Pedro Domingos. Mining knowledge-sharing sites for viral marketing. In Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining, KDD '02, pages 61-70, New York, NY, USA, 2002. ACM.
[46] Xavier Vives. Cournot and the oligopoly problem. European Economic Review, 33(2-3):503-514, 1989.

