On the Efficiency of Influence-and-Exploit Strategies for Revenue Maximization under Positive Externalities[☆,☆☆]

Dimitris Fotakis[a,*], Paris Siminelakis[b,1]

[a]National Technical University of Athens, School of Electrical and Computer Engineering, Greece
[b]Stanford University, Electrical Engineering Department, United States

Abstract

The mitigated effectiveness of traditional forms of advertising along with winner-take-all phenomenam caused by globalization and the Internet, necessitate a new approach in marketing. Hartline et.al[10] introduced a marketing model for social networks, where a seller is trying to exploit positive externalities between buyers to maximize his revenue, by designing an intelligent series of individualized offers. Under this setting, we study the problem of revenue maximization and mostly focus on Influence-and-Exploit (IE) marketing strategies. We show that in undirected social networks, revenue maximization is NP-hard not only when we search for a general optimal marketing strategy, but also when we search for the best IE strategy. Rather surprisingly, we observe that allowing IE strategies to offer prices smaller than the myopic price in the exploit step leads to a significant improvement on their performance. Thus, we show that the best IE strategy approximates the maximum revenue within a factor of 0:911 for undirected and of roughly 0:553 for directed networks. Our main contribution lies in designing an IE strategy, based on semi-definite programming, that exploits the global network structure and significantly improves the approximation ratio from 0.6667(resp. 0.3333) to 0.8229(resp. 0.5011) for undirected(resp. directed) networks. We also generalize IE strategies for multiple classes to deal with issues of implementing IE strategies in practice.

Keywords: Pricing, Externalities, Social Network Monetization, Influence and Exploit Marketing, Approximation Algorithms

1. Introduction

Understanding the flow of information, influence and epidemics through the social fabric has become increasingly important due to the high interconnectedness brought about by technological advances. The digitization of communications(cell phones, emails, text messages)

☆A preliminary version of this paper appeared in the Proceedings of the 8th Workshop on Internet & Network Economics(WINE), Liverpool, UK, December 10-12, 2012.
☆☆This work was supported by the project AlgoNow, co-financed by the European Union (European Social Fund - ESF) and Greek national funds, through the Operational Program Education and Lifelong Learning, under the research funding program Thales, and by an NTUA Basic Research Grant (PEBE 2009).
*Corresponding author.
Email addresses: fotakis@cs.ntua.gr (Dimitris Fotakis), psimin@stanford.edu (Paris Siminelakis)
[1]This author is partially supported by a scholarship from the Onassis Foundation.

and social interaction(Facebook, Twitter, Second Life) has provided researchers not only with a strong empirical footing upon which to base their theories and test their predictions, but also has opened the frontier of direct algorithmic applications. Particularly, there has been a shift from aggregate descriptive theories in the spirit of *Diffusion of Innovations* to models incorporating the micro-structure of social networks, culminating with the algorithmic paradigm of *Influence Maximization*.

Firms operating in such a reticular environment, where information about products and services diffuses rapidly between individuals, have acknowledged the importance of revisiting their approach. The availability of information about users and the mitigated effectiveness of traditional forms of marketing occasion the need for more intelligent marketing strategies. Towards realizing that goal there are three challenges: mining individual preferences, quantifying the influence that buyers exert upon each other and finally fusing these information along a promotion process. The ideal is an algorithm that intelligently adjusts itself(prices, individuals to approach) based on the current state of the network aiming to maximize the sellers revenue.

In this work, we are interested in the latter goal of designing a marketing strategy that exploits the influence between buyers to maximize the seller's revenue. We focus on the setting where the utility of the product depends inherently on the scale of the product's adoption(e.g. operating systems, VCR format) and particularly on the case where buyers' valuations depend on the specific set of their friends owning/using the product(e.g., cell phones, on-line gaming).

Modelling these externalities is inarguably a non-trivial matter. The networked nature of the process along with the corresponding combinatorial explosion necessitate a trade-off between realism and tractability. The purpose of any mathematical model is to capture the few details that really matter and abstract the rest. Towards that goal, we believe that the crucial features are: *(i)Individuation:* buyers should be allowed to have different valuations for the same product and be influenced in different amounts by their social contacts. *(ii)Uncertainty:*social interaction is itself stochastic in nature and the seller in practice can only estimate the model parameters(buyer's valuations). *(iii)Succinct representation:* buyers do not have a hard time deciding how much they are willing to pay for the product, which means that their internal representation of the product's value should have a simple form.

The Marketing Model. We adopt the model of Hartline, Mirrokni, and Sundararajan [10], where a digital product is sold to a set of potential buyers under positive externalities. We assume an unlimited supply of the product and that there is no production cost for it. A (possibly directed) weighted social network $G(V, E, w)$ on the set V of potential buyers models how buyers' value of the product is affected by other buyers who already own the product. Specifically, an edge $(j, i) \in E$ denotes that the event that j owns the product has a positive influence on i's value of the product. The strength of this influence is quantified by a non-negative weight w_{ji} associated with edge (j, i) (we assume that $w_{ji} = 0$ if $(j, i) \notin E$). Also, buyer i may have an intrinsic value of the product, quantified by a non-negative weight w_{ii}. The model incorporates the aforementioned features as follows:

(i) *Individuation:* The product's value to each buyer i is given by a non-decreasing function $v_i : 2^{N_i} \mapsto \mathbb{R}_+$, which depends on w_{ii} and on the set $S \subseteq N_i$ of i's neighbors who already own the product, where $N_i = \{j \in V \setminus \{i\} : (j, i) \in E\}$.

(ii) *Uncertainty:* The exact values $v_i(S)$ are unknown and are treated as random variables of which only the distributions $F_{i,S}$ are known to the seller. In particular, we assume that for each buyer i and each set $S \subseteq N_i$, the seller only knows the probability distribution $F_{i,S}(x) = \mathbb{P}\mathrm{r}[v_i(S) < x]$ that buyer i rejects an offer of price x for the product.

(iii) *Succinct representation:* Regarding the distribution of $v_i(S)$'s, the most interesting cases outlined in [10] are the *Concave Graph Model*(CGM), where the weights w_{ji} are random variables, and each $v_i(S)$ is a concave function of the total influence $M_{i,S} = \sum_{j \in S \cup \{i\}} w_{ji}$ perceived by buyer i from the set S of her neighbors owning the product, and the *Uniform Additive Model*(UAM), where weights w_{ji} are deterministic, and each $v_i(S)$ is uniformly distributed in $[0, M_{i,S}]$.

Marketing Strategies. We assume that the seller approaches each potential buyer once and makes an individualized offer to him. Thus, a *marketing strategy* $(\vec{\pi}, \vec{x})$ consists of a permutation $\vec{\pi}$ of the buyers and a pricing vector $\vec{x} = (x_1, \ldots, x_n)$, where $\vec{\pi}$ determines the order in which the buyers are approached and $\vec{x}$ the prices offered to them. Given the set S of i's neighbors who own the product when the seller approaches her, buyer i accepts the offer with probability $1 - F_{i,S}(x_i)$, in which case she pays the price x_i, or rejects it, with probability $F_{i,S}(x_i)$, in which case she pays nothing and never receives an offer again. The seller's goal is to compute a marketing strategy $(\vec{\pi}, \vec{x})$ that maximizes her expected revenue, namely the total amount paid by the buyers who accept the offer.

Previous Work. This setting was introduced by Hartline, Mirrokni and Sundararajan in [10]. They showed that for a complete graph and identical buyers, the problem can be solved optimally by an polynomial-time dynamic programming algorithm. However, using a transformation from Maximum Acyclic Subgraph, the authors showed that computing a revenue maximizing ordering is NP-Hard for directed networks under the Uniform Additive Model, even with complete knowledge of buyers' valuations. Combined with the result of [11], this suggests an upper bound of 0.5 on the approximability of revenue maximization for directed networks and deterministic additive valuations. Motivated by hardness the authors then turn to approximation strategies.

Their work revolves around a class of elegant marketing strategies called Influence-and-Exploit(IE). An IE strategy first offers the product for free to a selected subset of buyers, aiming to increase the value of the product for the remaining buyers(influence step). Then, in the exploit step, it approaches the remaining buyers, in a random order, and offers the product at the so called myopic price, which just aims to extract the maximum revenue from each buyer ignoring the influence on future buyers. They first discuss IE strategies that construct the Influence set by including ever vertex independently with a uniform probability. They showed that under the Uniform Additive Model for undirected networks, this strategy results in a $2/3$ approximation ratio. Whereas, for the Concave Graph Model and under some mild assumptions on buyers' distributions, they proved that it results in an $\frac{e}{4e-2} \simeq 0.306$-approximation strategy. Furthermore, they showed that if the revenue functions(resulting from the distributions F_i) are non-negative, monotone and submodular, then the total revenue function of an IE strategy is also non-negative, submodular and that the Optimal IE strategy generates revenue at least $1/4$ of the optimal. Based on these results they showed that if we have an α-approximation algorithm for maximizing non-negative submodular functions then we can use it to construct an IE strategy, i.e. find an Influence set, with approximation ratio $\alpha/4$. This suggest that for the submodular case the approximation ratio is 0.125, using the recent result[FOCS'12].

1.1. Contribution and Techniques.

Although IE strategies are simple, elegant, and promising in terms of efficiency, their properties are not well understood. Furthermore, the absence of tight lower bounds(best is $1/4$) on the performance of the Optimal IE strategy for submodular valuation functions, suggests looking into simpler models. The example of Influence Maximization has shown that studying simple models can provide many insights into the structure of the problem, which then can inform our understanding of the general setting as well. In this work, we focus on the Uniform Additive Model(see Section 2), which can be seen as an extension of the widely accepted Linear Threshold Model(LTM) of Social Influence in order to incorporate pricing dynamics. We first show that under the UAM both computing the optimal marketing strategy and the optimal IE strategy are NP-Hard for undirected networks. We then embark on a systematic investigation of the approximability and efficiency of IE strategies and manage to significantly improve the best known approximation ratio for revenue maximization.

One of the messages conveyed by [10] is that promotional offers coupled with value-based pricing can go a long way. We supplement this message by showing(Section 3) that there is a significant gain in the expected revenue if we appropriately reduce the prices offered in the exploit step. Particularly, by offering prices lower than the myopic(see Section 2), we slightly improved the approximation ratio to 0.686(resp. 0.343) for undirected(resp. directed networks) from 0.666(resp. 0.333). The reason for offering slightly lower prices is to utilize the influence between buyers belonging in the exploit set.

Discriminative pricing in Social Networks can be seen as a process of partitioning buyers in different classes(according to prices offered[2]). Exploiting this intuition, we propose a natural generalization of IE strategies that use more than two pricing classes. We show that a simple random partitioning of the buyers in six pricing classes further improves the approximation ratio for the maximum revenue to 0.7032 for undirected networks and to 0.3516 for directed social networks. In the final section we discuss the implications of the success of such strategies.

Next, we explore the limits of IE strategies in approximating the optimal revenue. We provide lower bounds for the revenue of an Optimal IE strategy, by proving the existence of an IE strategy which approximates the maximum revenue within a factor of 0.911 for undirected(resp. 0.55289 for directed) networks. The proof assumes an optimal pricing vector $\mathbf{p}$, and constructs such a strategy by applying randomized rounding to $\mathbf{p}$. The analysis exploits the fact that in the UAM, the expected revenue is a linear functions of edge weights. An interesting consequence is that the upper bound of 0.5 on the approximation ratio of the maximum revenue for directed networks does not apply to the Uniform Additive Model. We discuss the technical reasons behind this and show a pair of upper bounds on the approximation ratio for directed networks. Specifically, assuming the Unique Games conjecture, we show that it is $\mathbf{NP}$-hard to approximate the maximum revenue within a factor greater than $27/32$, if we can use any algorithm, and greater than $3/4$, if we are restricted to IE strategies with pricing probability $2/3$.

The strategies proposed so far were either oblivious(random sampling) or with only augmented local awareness(submodular maximization). Our main contribution lies in designing an IE strategy that exploits the global network structure and significantly improves the approximation ratio. The main hurdle in obtaining better approximation guarantess for the maximum revenue problem is the lack of any strong upper bounds on it. We circumvent this issue by approximating instead the revenue of an an Optimal IE strategy. In Section 6, we introduce a strong Semidefinite Programming (SDP) relaxation for the problem of computing the best IE strategy with any given pricing probability. Our approach exploits the resemblance between computing the best IE strategy and the problems of MAX-CUT and MAX-DICUT, and builds on the elegant approach of Goemans and Williamson [8] and Feige and Goemans [7]. Solving the SDP relaxation and using randomized rounding, we obtain a 0.9032 (resp. 0.9064) approximation for the

[2]This is a coarse statement, in reality the correct parameter is the probability of accepting(see Section 2).

best IE strategy with pricing probability 0.586 for undirected networks (resp. $2/3$ for directed networks). Combining these results with the bounds on the fraction of the maximum revenue extracted by IE strategies, we significantly improve on the best known approximation ratio for revenue maximization to 0.8229(from 0.6666) for undirected networks and 0.5011 for directed networks (from 0.3333). To the best of our knowledge, this is the first time an (approximate) SDP relaxation for a pricing model under positive externalities is suggested and exploited to improve the approximation ratio for the corresponding revenue (or welfare) maximization problem. We believe that our approach may find applications to other pricing models under externalities.

1.2. Other Related Work

Hartline et al. [10] were the first to consider social influence in the framework of revenue maximization. Since then, relevant research has focused either on posted price strategies, where there is no price discrimination, or on game theoretic considerations, where the buyers act strategically according to their perceived value of the product. To the best of our knowledge, our work is the first that considers the approximability of the revenue extracted by an optimal strategy and by the best IE strategy, the central problems in [10]. Arthur et al. [4] considered a model that is similar in nature to the one discussed here, generalizing the LTM and the Independent Cascade Model. However, they assume that the seller cannot approach individual buyers, but recommendations about the product cascade through the network from early adopters. The crucial feature is that, since each buyer offered the product to a non-zero price has a constant probability of declining, the cascades have a finite length and therefore there are only finite range correlations. Exploiting this fact they presented an IE $O(1)$-approximation algorithm based on the construction of a max-leaf spanning tree.

In the posted price setting, Akhlaghpour et al. [1] study a model for monotone submodular buyers valuations. For the problem of designing a price trajectory for k days, they provided an IE algorithm, where the first price is selected such that a large portion of buyers will participate and then the remaining prices are selected so as to optimize the expected revenue. For a special case of buyers valuations, equivalent to the LTM, and for the case, where the rate which prices change is limited, they present an FPTAS based on dynamic programming. Anari et al. [3] considered a posted price setting where the product exhibits historical externalities. They assume there is a continuum of buyers each belonging to some type(e.g. beta testers, professionals, average users) and their valuation depends on the fraction of buyers that already own the product. Given a fixed price trajectory, the buyers decide when to buy the product. In this setting, they studied existence and uniqueness of equilibria, and presented an FPTAS for special cases of revenue maximization, based on a rectangular covering problem.

In a complementary direction, Chen et al. [6] investigated the Bayesian-Nash equilibria in the presence of positive externalities. They consider a model of linear additive externalities and focused on two classes of equilibria, pessimistic and optimistic ones. They showed how to compute these equilibria and computed revenue-maximizing prices, based on a line-sweep algorithm. Candogan et al. [5] investigated a scenario where a monopolist sells a divisible good to buyers under positive externalities. They considered a two-stage game where the seller first sets a price for each buyer, and then the buyers decide on their consumption level. They proved that the optimal price for each buyer is proportional to her Bonacich centrality, and that if the buyers are partitioned into two pricing classes (which is conceptually similar to IE), the problem is reducible to MAX-CUT.

2. Model and Preliminaries

The Influence Model. The social network is a (possibly directed) weighted network $G(V, E, w)$ on the set V of potential buyers. There is a positive weight w_{ij} associated with each edge $(i, j) \in E$ (we assume that $w_{ij} = 0$ if $(i, j) \notin E$). A social network is undirected (or symmetric) if $w_{ij} = w_{ji}$ for all $i, j \in V$, and directed (or asymmetric) otherwise. There may exist a non-negative weight w_{ii} associated with each buyer i[3]. Every buyer has a value $v_i : 2^{N_i} \mapsto \mathbb{R}_+$ of the product, which depends on w_{ii} and on the set $S \subseteq N_i$ of his neighbors who already own the product, where $N_i = \{j \in V \setminus \{i\} : (j, i) \in E\}$. In the Uniform Additive Model, the values $v_i(S)$ are random variables drawn from the uniform distribution in $[0, M_{i,S}]$, where $M_{i,S} = \sum_{j \in S \cup \{i\}} w_{ji}$ is the total influence perceived by i, given the set S of his neighbors who own the product at the time of the offer. Then, the probability that i rejects an offer of price x is $F_{i,S}(x) = x/M_{i,S}$.

Myopic Pricing. The *myopic price* disregards any externalities imposed by i on his neighbors, and simply maximizes the expected revenue extracted from buyer i, given that S is the current set of i's neighbors who own the product. For the Uniform Additive Model, the myopic price is $M_{i,S}/2$, the probability that buyer i accepts it is $1/2$, and the expected revenue extracted from him with the myopic price is $M_{i,S}/4$, which is the maximum revenue one can extract from buyer i alone.

Pricing with Probabilities and Revenue Maximization. We can usually extract more revenue from G by employing a marketing strategy that exploits the positive influence between the buyers. We observe that for any buyer i and any probability p that i accepts an offer, there is an (essentially unique) price x_p such that i accepts an offer of x_p with probability p. For the Uniform Additive Model, $x_p = (1 - p)M_{i,S}$ and the expected revenue extracted from buyer i with such an offer is $p(1-p)M_{i,S}$. Throughout this paper, we equivalently regard marketing strategies as consisting of a permutation π of the buyers and a vector $\mathbf{p} = (p_1, \dots, p_n)$ of pricing probabilities. We note that if $p_i = 1$, i gets the product for free, while if $p_i = 1/2$, the price offered to i is (the myopic price of) $M_{i,S}/2$. We assume that $p_i \in [1/2, 1]$, since any expected revenue in $[0, M_{i,S}/4]$ can be achieved with such pricing probabilities. Then, the expected revenue of a marketing strategy $(\pi, \mathbf{p})$ is:

$$R(\pi, \mathbf{p}) = \sum_{i \in V} p_i(1 - p_i) \left(w_{ii} + \sum_{j : \pi_j < \pi_i} p_j w_{ji} \right) \tag{1}$$

The problem of *revenue maximization* under the Uniform Additive Model is to find a marketing strategy $(\pi^*, \mathbf{p}^*)$ that extracts a maximum revenue of $R(\pi^*, \mathbf{p}^*)$ from a given social network $G(V, E, w)$.

[3] For simplicity, we ignore w_{ii}'s for directed social networks. This is without loss of generality, since we can replace each w_{ii} by an edge (i', i) of weight w_{ii} from a new node i' with a single outgoing edge (i', i) and no incoming edges.

Bounds on the Maximum Revenue. Let $N = \sum_{i \in V} w_{ii}$ and $W = \sum_{i<j} w_{ij}$, if the social network G is undirected, and $W = \sum_{(i,j) \in E} w_{ij}$, if G is directed. Then an upper bound on the maximum revenue of G is $R^* = (W + N)/4$, and follows by summing up the myopic revenue over all edges of G [10, Fact 1]. For a lower bound on the maximum revenue, if G is undirected (resp. directed), approaching the buyers in any order (resp. in a random order) and offering them the myopic price yields a revenue of $(W + 2N)/8$ (resp. $(W + 4N)/16$). Thus, myopic pricing achieves an approximation ratio of 0.5 for undirected networks and of 0.25 for directed networks.

Ordering and $\mathbf{NP}$*-Hardness..* Revenue maximization exhibits a dual nature involving optimizing both the pricing probabilities and the sequence of offers. For directed networks, finding a good ordering π of the buyers bears a resemblance to the Maximum Acyclic Subgraph problem, where given a directed network $G(V, E, w)$, we seek for an acyclic subgraph of maximum total edge weight. In fact, any permutation π of V corresponds to an acyclic subgraph of G that includes all edges going forward in π, i.e, all edges (i, j) with $\pi_i < \pi_j$. [10, Lemma 3.2] shows that given a directed network G and a pricing probability vector $\mathbf{p}$, computing an optimal ordering of the buyers (for the particular $\mathbf{p}$) is equivalent to computing a Maximum Acyclic Subgraph of G, with each edge (i, j) having a weight of $p_i p_j (1 - p_j) w_{ij}$. Consequently, computing an ordering π that maximizes $R(\pi, \mathbf{p})$ is NP-hard and Unique-Games-hard to approximate within a factor greater than 0.5 [9].

On the other hand, we show that in the undirected case, if the pricing probabilities are given, we can easily compute the best ordering of the buyers (see also Section A.1, in the Appendix of the full paper, for a simple example about the importance of a good ordering in the undirected case).

Lemma 1. *Let $G(V, E, w)$ be an undirected social network, and let $\mathbf{p}$ be any pricing probability vector. Then, approaching the buyers in non-increasing order of their pricing probabilities maximizes the revenue extracted from G under $\mathbf{p}$.*

PROOF. We consider an optimal ordering π (wrt. $\mathbf{p}$) that minimizes the number of buyers' pairs appearing in increasing order of their pricing probabilities, namely, the number of pairs i_1, i_2 with $p_{i_1} < p_{i_2}$ and $\pi_{i_1} < \pi_{i_2}$. If there is such a pair in π, we can find a pair of buyers i and j with $p_i < p_j$ such that i appears just before j in π. Then, switching the positions of i and j in π changes the expected revenue extracted from G under $\mathbf{p}$ by $p_i p_j w_{ij} (p_j - p_i) \geq 0$, a contradiction. $\square$

A consequence of Lemma 1 is that [10, Lemma 3.2] does not imply the NP-hardness of revenue maximization for undirected social networks. The following lemma employs a reduction from monotone One-in-Three 3-SAT, and shows that revenue maximization is NP-hard for undirected networks.

Lemma 2. *The problem of computing a marketing strategy that extracts the maximum revenue from an undirected social network is NP-hard.*

PROOF. In monotone One-in-Three 3-SAT, we are given a set V of n items and m subsets $T_1, \ldots, T_m$ of V, with $2 \leq |T_j| \leq 3$ for each $j \in \{1, \ldots, m\}$. We ask for a subset $S \subset V$ such that $|S \cap T_j| = 1$ for all $j \in \{1, \ldots, m\}$. Monotone One-in-Three 3-SAT is shown $\mathbf{NP}$-complete in [11]. In the following, we let m_2 (resp. m_3) denote the number of 2-item (resp. 3-item) sets T_j in an instance $(V, T_1, \ldots, T_m)$ of monotone One-in-Three 3-SAT.

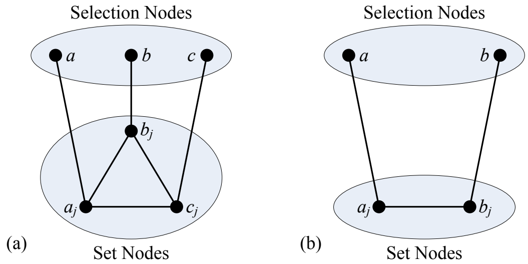

Figure 1: Examples of (a) an extended triangle and (b) a 3-path, used in the proof of Lemma 2. We create an extended triangle for each 3-item set T_j and a 3-path for each 2-item set T_j. The set nodes are different for each set T_j, while the selection nodes are common for all sets.

Given $(V, T_1, \ldots, T_m)$, we construct an undirected social network G. The network G contains a *selection-node* corresponding to each item in V. There are no edges between the selection nodes of G. For each 3-item set $T_j = \{a, b, c\}$, we create an *extended triangle* consisting of a triangle on three *set nodes* a_j, b_j, and c_j, and three additional edges that connect a_j, b_j, c_j to the corresponding selection nodes a, b, and c (see also Fig. 2.a). For each 2-item set $T_j = \{a, b\}$, we create a 3-*path* consisting of an edge connecting two set nodes a_j and b_j, and two additional edges connecting a_j and b_j to the corresponding selection nodes a and b (see also Fig. 2.b). Therefore, G contains $n + 2m_2 + 3m_3$ nodes and $3m_2 + 6m_3$ edges. The weight of all edges of G is 1. We next show that $(V, T_1, \ldots, T_m)$ is a YES-instance of monotone One-in-Three 3-SAT iff the maximum revenue of G is at least $\frac{177}{128}\, m_3 + \frac{3}{4}\, m_2$.

By Lemma 1, the revenue extracted from G is maximized if the nodes are approached in non-increasing order of their pricing probabilities. Therefore, we can ignore the ordering of the nodes, and focus on their pricing probabilities. The important property is that if each extended triangle (Fig. 2.a) is considered alone, its maximum revenue is $177/128$, and is obtained when exactly one of the selection nodes a, b, c has a pricing probability of $1/2$ and the other two have a pricing probability of 1. More specifically, since the selection nodes a, b, c have degree 1, the revenue of the extended triangle is maximized when they have a pricing probability of either 1 or $1/2$. If all a, b, c have a pricing probability of 1, the best revenue of the extended triangle is ≈ 1.196435, and is obtained when one of a_j, b_j, and c_j has a pricing probability of ≈ 0.7474, the other has a pricing probability of ≈ 0.5715, and the third has a pricing probability of $1/2$. If all a, b, c have a pricing probability of $1/2$, the best revenue of the extended triangle is again ≈ 1.196435, and is obtained with the same pricing probabilities of a_j, b_j, and c_j. If two of a, b, c (say a and b) have a pricing probability of $1/2$ and c has a pricing probability of 1, the best revenue of the extended triangle is $\frac{21}{16} = 1.3125$, and is obtained when one of a_j and b_j has a pricing probability of 1, the other has a pricing probability of $3/4$, and c_j has a pricing probability of $1/2$. Finally, if two of a, b, c (say b and c) have a pricing probability of 1 and a has a pricing probability of $1/2$, we extract a maximum revenue from the extended triangle, which is $\frac{177}{128} = 1.3828125$ and is obtained when a_j has a pricing probability of 1, one of b_j and c_j has a pricing probability of $9/16$, and the other has a pricing probability of $1/2$.

Similarly, if each 3-path (Fig. 2.b) is considered alone, its maximum revenue is $3/4$, and is obtained when exactly one of the selection nodes a, b has a pricing probability of $1/2$ and the other has a pricing probability of 1. In fact, since the 3-path is a bipartite graph, Proposition **??**

implies that the maximum revenue, which is $3/4$, is extracted when a_j and b have a pricing probability of 1 and b_j and a have a pricing probability of $1/2$ (or the other way around). If both a and b have a pricing probability of 1, the best revenue of the 3-path is $41/64$ and is obtained when one of a_j and b_j has a pricing probability of $5/8$, and the other has a pricing probability of $1/2$. If both a and b have a pricing probability of $1/2$, the best revenue of 3-path is again $41/64$ and is obtained when one of a_j and b_j has a pricing probability of 1, and the other has a pricing probability of $5/8$.

If $(V, T_1, \dots, T_m)$ is a YES-instance of monotone One-in-Three 3-SAT, we assign a pricing probability of $1/2$ to the selection nodes in S and a pricing probability of 1 to the selection nodes in $V \setminus S$, where S is a set with exactly one element of each T_j. Thus, we have exactly one selection node with pricing probability $1/2$ in each extended triangle and in each 3-path. Then, we can set the pricing probabilities of the set nodes as above, so that the revenue of each extended triangle is $177/128$ and the revenue of each 3-path is $3/4$. Thus, the maximum revenue of G is at least $\frac{177}{128}\, m_3 + \frac{3}{4}\, m_2$.

For the converse, we recall that the edges of G can be partitioned into m_3 extended triangles and m_2 3-paths. Consequently, if the maximum revenue of G is at least $\frac{177}{128}\, m_3 + \frac{3}{4}\, m_2$, each extended triangle contributes exactly $177/128$ and each 3-path contributes exactly $3/4$ to the revenue of G. Thus, by the analysis on their revenue above, each extended triangle and each 3-path includes exactly one selection node with a pricing probability of $1/2$. Therefore, if we let S consist of the selection nodes with pricing probability $1/2$, we have that $|S \cap T_j| = 1$ for all $j \in \{1, \dots, m\}$. □

3. Influence & Exploit Strategies

An *Influence-and-Exploit* (IE) strategy $IE(A, p)$ consists of a set of buyers A receiving the product for free and a pricing probability p offered to the remaining buyers in $V \setminus A$, who are approached in a random order. We slightly abuse the notation and let $IE(q, p)$ denote an IE strategy where each buyer is selected in A independently with probability q. $IE(A, p)$ extracts an expected (wrt the random ordering of the exploit set) revenue of:

$$RIE(A, p) = p(1-p) \sum_{i \in V \setminus A} \left(w_{ii} + \sum_{j \in A} w_{ji} + \sum_{j \in V \setminus A,\ j \neq i} \frac{p\, w_{ji}}{2} \right) \tag{2}$$

The problem of finding the best IE strategy is to compute a subset of buyers A^* and a pricing probability p^* that extract a maximum revenue of $RIE(A^*, p^*)$ from a given social network $G(V, E, w)$. The following lemma employs a reduction from monotone One-in-Three 3-SAT, and shows that computing the best IE strategy is NP-hard.

Lemma 3. *Let $p \in [1/2, 1)$ be any fixed pricing probability. The problem of finding the best IE strategy with pricing probability p is NP-hard, even for undirected social networks.*

PROOF. We recall that in monotone One-in-Three 3-SAT, we are given a set V of n items and m subsets $T_1, \dots, T_m$ of V, with $2 \leq |T_j| \leq 3$ for each $j \in \{1, \dots, m\}$. We ask for a subset $S \subset V$ such that $|S \cap T_j| = 1$ for all $j \in \{1, \dots, m\}$.

Given $(V, T_1, \dots, T_m)$, we construct an undirected social network G on V. For each 3-item set $T_j = \{a, b, c\}$, we create a *set-triangle* on nodes a, b, and c with 3 edges of weight 1. For each 2-item set $T_j = \{a, b\}$, we add a *set-edge* $\{a, b\}$ of weight $2 + p$, where p is the pricing probability. To avoid multiple appearances of the same edge, we let the weight of each edge be the total weight of its appearances. Namely, if an edge e appears in k_3 set-triangles and in k_2 set-edges, e's weight is $k_3 + (2+p)k_2$. We observe that for any $p \in [1/2, 1)$, the maximum revenue extracted from any set-triangle and any set-edge is $p(1-p)(2+p)$, by giving the product for free to exactly one of the nodes of the set-triangle (resp. the set-edge).

We next show that $(V, T_1, \ldots, T_m)$ is a YES-instance of monotone One-in-Three 3-SAT iff there is an influence set A in G such that $R_{\mathrm{IE}}(A, p) \geq mp(1-p)(2+p)$. If $(V, T_1, \ldots, T_m)$ is a YES-instance of monotone One-in-Three 3-SAT, we let the influence set $A = S$, where S is a set with exactly one element of each T_j. Then, we extract an expected revenue of $p(1-p)(2+p)$ from each set-triangle and each set-edge in G, which yields an expected revenue of $mp(1-p)(2+p)$ in total. For the converse, if there is an influence set A in G such that $R_{\mathrm{IE}}(A, p) \geq mp(1-p)(2+p)$, we let $S = A$. Since $R_{\mathrm{IE}}(A, p) \geq mp(1-p)(2+p)$, and since the edges of G can be partitioned into m set-triangles and set-edges, each with a maximum revenue of at most $p(1-p)(2+p)$, each set-triangle and each set-edge contributes exactly $p(1-p)(2+p)$ to $R_{\mathrm{IE}}(A, p)$. Therefore, for all set-triangles and all set-edges, there is exactly one node in A. Thus, we have that $|S \cap T_j| = 1$ for all $j \in \{1, \ldots, m\}$. □

Exploiting Large Cuts. A natural idea towards obtaining approximation algorithms is to exploit the apparent connection between a large cut in the social network and a good IE strategy. For example, in the undirected case, an IE strategy $IE(q, p)$ is conceptually similar to the randomized 0.5-approximation algorithm for MAX-CUT, which puts each node in set A with probability $1/2$. However, in addition to a revenue of $p(1-p)w_{ij}$ from each edge $\{i, j\}$ in the cut $(A, V \setminus A)$, $IE(q, p)$ extracts a revenue of $p^2(1-p)w_{ij}$ from each edge $\{i, j\}$ between nodes in the exploit set $V \setminus A$. Thus, to optimize the performance of $IE(q, p)$, we carefully adjust the probabilities q and p so that $IE(q, p)$ balances between the two sources of revenue. Hence, we obtain the following:

Proposition 4. *Let $G(V, E, w)$ be an undirected social network, and let $q = \max\{1 - \frac{\sqrt{2}(2+\lambda)}{4}, 0\}$, where $\lambda = N/W$. Then, $IE(q, 2-\sqrt{2})$ approximates the maximum revenue extracted from G within a factor of at least $2\sqrt{2}(2-\sqrt{2})(\sqrt{2}-1) \approx 0.686$.*

PROOF. The proof extends the proof of [10, Theorem 3.1]. We start with calculating the expected (wrt to the random choice of the influence set) revenue of $\mathrm{IE}(q, p)$. The expected revenue of $\mathrm{IE}(q, p)$ from each loop $\{i, i\}$ is $(1-q)p(1-p)w_{ii}$. In particular, a revenue of $p(1-p)w_{ii}$ is extracted from $\{i, i\}$ if buyer i is included in the exploit set, which happens with probability $1-q$. The expected revenue of $\mathrm{IE}(q, p)$ from each edge $\{i, j\}$, $i < j$, is $(2q(1-q)p(1-p) + (1-q)^2p^2(1-p))w_{ij}$. More specifically, if one of i, j is included in the influence set and the other is included in the exploit set, which happens with probability $2q(1-q)$, a revenue of $p(1-p)w_{ij}$ is extracted from edge $\{i, j\}$. Otherwise, if both i and j are included in the exploit set, which happens with probability $(1-q)^2$, a revenue of $p^2(1-p)w_{ij}$ is extracted from edge $\{i, j\}$ (note that since $\{i, j\}$ is an undirected edge, the order in which i and j are considered in the exploit set is insignificant). By linearity of expectation, the expected revenue of $\mathrm{IE}(q, p)$ is:

$$R_{\mathrm{IE}}(q, p) = (1-q)p(1-p)\sum_{i \in V} w_{ii} + (1-q)p(1-p)\sum_{i<j}(2q + p(1-q))w_{ij}$$

Using that $N = \sum_{i \in V} w_{ii}$ and $W = \sum_{i<j} w_{ij}$, and setting $N = \lambda W$, we obtain that:

$$R_{\mathrm{IE}}(q, p) = (1-q)p(1-p)(\lambda + 2q + p(1-q))W$$

Differentiating with respect to q, we obtain that the optimal value of q is

$$q^* = \max\left\{\frac{1-p-\lambda/2}{2-p}, 0\right\}$$

We recall that $R^* = (1+\lambda)W/4$ is an upper bound on the maximum revenue of G. Therefore, the approximation ratio of IE(q,p) is:

$$\frac{4(1-q)p(1-p)(\lambda+2q+p(1-q))}{1+\lambda} \tag{3}$$

Using $p = 1/2$ and $q = \max\left\{\frac{1-\lambda}{3}, 0\right\}$ in (3), we obtain the IE strategy of [10, Theorem 3.1], whose approximation ratio is at least $2/3$, attained at $\lambda = 0$. Assuming small values of λ, so that $q^* > 0$, and differentiating with respect to p, we obtain that the best value of p for IE(q^*, p) is $p^* = 2-\sqrt{2}$. Using $p = 2-\sqrt{2}$ and $q = \max\left\{1 - \frac{\sqrt{2}(2+\lambda)}{4}, 0\right\}$, we obtain an IE strategy with an approximation ratio of at least $2\sqrt{2}(2-\sqrt{2})(\sqrt{2}-1) \approx 0.686$, attained at $\lambda = 0$. □

Proposition 5. *Let $G(V,E,w)$ be a directed social network. Then, $IE\left(1-\frac{\sqrt{2}}{2}, 2-\sqrt{2}\right)$ approximates the maximum revenue of G within a factor of $\sqrt{2}(2-\sqrt{2})(\sqrt{2}-1) \approx 0.343$.*

PROOF. The proof is similar to the proof of Proposition 4. We recall that for the directed case, we can ignore loops (i,i). Since the social network G is directed, the expected (wrt to the random choice of the influence set and the random order of the exploit set) revenue of IE(q,p) is:

$$\begin{aligned} R_{\mathrm{IE}}(q,p) &= (1-q)p(1-p)\sum_{(i,j)\in E}(q+p(1-q)/2)w_{ij} \\ &= (1-q)p(1-p)(q+p(1-q)/2)W \end{aligned}$$

More specifically, if i is included in the influence set and j is included in the exploit set, which happens with probability $q(1-q)$, a revenue of $p(1-p)w_{ij}$ is extracted from each edge (i,j). Furthermore, if both i and j are included in the exploit set $V \setminus A$ and i appears before j in the random order of $V \setminus A$, which happens with probability $(1-q)^2/2$, a revenue of $p^2(1-p)w_{ij}$ is extracted from edge (i,j).

Using the upper bound of $W/4$ on the maximum revenue of G, we have that the approximation ratio of IE(q,p) is at least $4(1-q)p(1-p)(q+p(1-q)/2)$. Setting $q = 1/3$ and $p = 1/2$, we obtain the IE strategy of [10, Theorem 3.1], whose approximation ratio for directed networks is $1/3$. Using $q = 1-\frac{\sqrt{2}}{2}$ and $p = 2-\sqrt{2}$, we obtain an IE strategy with an approximation ratio of $\sqrt{2}(2-\sqrt{2})(\sqrt{2}-1) \approx 0.343$. □

4. Generalized Influence & Exploit

Building on the idea of generating revenue from large cuts between different pricing classes, we obtain a class of generalized IE strategies, which employ a refined partition of buyers in more than two pricing classes. We first analyze the efficiency of generalized IE strategies for undirected networks, and then translate our results to the directed case. The analysis generalizes the proof of Proposition 4.

A generalized IE strategy consists of K pricing classes, for some appropriately large integer $K \geq$

2. Each class k, $k = 1, \ldots, K$, is associated with a pricing probability of $p_k = 1 - \frac{k-1}{2(K-1)}$. Each buyer is assigned to the pricing class k independently with probability q_k, where $\sum_{k=1}^{K} q_k = 1$, and is offered a pricing probability of p_k. The buyers are considered in non-increasing order of their pricing probabilities, i.e., the buyers in class k are considered before the buyers in class $k+1$, $k = 1, \ldots, K-1$. The buyers in the same class are considered in random order. In the following, we let $IE(\mathbf{q}, \mathbf{p})$ denote such a generalized IE strategy, where $\mathbf{q} = (q_1, \ldots, q_K)$ is the assignment probability vector and $\mathbf{p} = (p_1, \ldots, p_K)$ is the pricing probability vector.

Using linearity of expectation and grouping similar terms together we get that the expected revenue extracted by the generalized IE strategy $IE(\mathbf{q}, \mathbf{p})$ from an undirected social network $G(V, E, w)$ is:

$$RIE(\mathbf{q}, \mathbf{p}) = N \sum_{k=1}^{K} q_k p_k (1 - p_k) + W \sum_{k=1}^{K} q_k p_k (1 - p_k) \left(q_k p_k + 2 \sum_{\ell=1}^{k-1} q_\ell p_\ell \right) \tag{4}$$

where $N = \sum_{i \in V} w_{ii}$ and $W = \sum_{i<j} w_{ij}$. Since $R^* = (N + W)/4$ is an upper bound on the maximum revenue of G, the approximation ratio of $IE(\mathbf{q}, \mathbf{p})$ is at least:

$$\min \left\{ 4 \sum_{k=1}^{K} q_k p_k (1 - p_k),\ \ 4 \sum_{k=1}^{K} q_k p_k (1 - p_k) \left(q_k p_k + 2 \sum_{\ell=1}^{k-1} q_\ell p_\ell \right) \right\} \tag{5}$$

We can now select the assignment probability vector $\mathbf{q}$ so that (5) is maximized. We note that with the pricing probability vector $\mathbf{p}$ fixed, this involves maximizing a quadratic function of $\mathbf{q}$ over linear constraints. Thus, we obtain the following:

Theorem 6. *For any undirected social network G, the generalized IE strategy with* $K = 6$ *pricing classes and assignment probabilities* $\mathbf{q} = (0.183, 0.075, 0.075, 0.175, 0.261, 0.231)$ *approximates the maximum revenue of G within a factor of* 0.7032.

Remark 1. We note that the approximation ratio can be improved to 0.706 by considering more pricing classes. Furthermore, if we disregard terms of the form $q_k p_k (1 - p_k) q_k p_k$, i.e. intra-class influence, there are assignment probabilities that achieve 0.652(resp. 0.7057) approximation ratio for $K = 10$(resp. sufficiently large $K > 1000$) pricing classes.

By the same approach, we show that for directed social networks, the approximation ratio of $IE(\mathbf{q}, \mathbf{p})$ is at least half the quantity in (5). Therefore:

Corollary 7. *For any directed social network G, the generalized IE strategy with* $K = 6$ *pricing classes and assignment probabilities* $\mathbf{q} = (0.183, 0.075, 0.075, 0.175, 0.261, 0.231)$ *approximates the maximum revenue of G within a factor of* 0.3516.

PROOF. Similarly to the proof of Theorem 6, we calculate the expected (wrt the random partition of buyers into pricing classes and the random order of buyers in the pricing classes) revenue extracted by the generalized IE strategy IE($\vec{p}, \vec{q}$) from a directed social network $G(V, E, w)$. We recall that for directed social networks, we can ignore loops (i, i). The expected revenue of IE($\vec{p}, \vec{q}$) from each edge (i, j) is:

$$w_{ij} \sum_{k=1}^{K} q_k p_k (1 - p_k) \left(\frac{q_k p_k}{2} + \sum_{\ell=1}^{k-1} q_\ell p_\ell \right)$$

More specifically, for each class k, if both i, j are included in the pricing class k and i appears before j in the random order of the buyers in k, which happens with probability $q_k^2/2$, the revenue extracted from each edge (i,j) is $p_k^2(1-p_k)w_{ij}$. Furthermore, for each pair ℓ, k of pricing classes, $1 \le \ell < k \le K$, if i is included in ℓ and j is included in k, which happens with probability $q_\ell q_k$, the revenue extracted from (i,j) is $p_\ell p_k(1-p_k)w_{ij}$.

Using linearity of expectation and setting $W = \sum_{(i,j)\in E} w_{ij}$, we obtain that the expected revenue of $\mathrm{IE}(\vec{q},\vec{p})$ is:

$$R_{\mathrm{IE}}(\vec{q},\vec{p}) = W\sum_{k=1}^{K} q_k p_k(1-p_k)\left(\frac{q_k p_k}{2} + \sum_{\ell=1}^{k-1} q_\ell p_\ell\right)$$

Since $W/4$ is an upper bound on the maximum revenue of G, the approximation ratio of $\mathrm{IE}(\vec{q},\vec{p})$ is at least:

$$4\sum_{k=1}^{K} q_k p_k(1-p_k)\left(\frac{q_k p_k}{2} + \sum_{\ell=1}^{k-1} q_\ell p_\ell\right), \qquad (6)$$

namely at least half of the approximation ratio in the undirected case.

Using $\vec{q} = (0.183, 0.075, 0.075, 0.175, 0.261, 0.231)$ in (7), we obtain an approximation ratio of at least 0.3516. ☐

5. On the Efficiency of Influence-and-Exploit Strategies

In the previous sections we showed that the inherent randomness of the model has brought about a continuous search space in which we are seeking an optimal solution. Moreover, this freedom of choices has rendered the problem NP-Hard and the process of obtaining any non-trivial upper bound intangible. The Influence-and-exploit idea simply makes a very rough discretization of this space, allowing only two prices. This is a natural choice since the combinatorial properties of partitioning vertices into two sets have been extensively studied. Nevertheless, we are left with the daunting task of correlating the performance of an optimal solution using only two prices with that when any price is allowed. In that direction, the lack of upper(lower) bounds for the optimal(optimal IE) strategy combined with NP-Hardness of both problems is discouraging. We resolve this shortcoming by exploiting linearity of the revenue in both cases with respect to the edges by utilizing Local Analysis. We show that the best IE strategy manages to extract a significant fraction of the maximum revenue.

Theorem 8. *For any undirected social network, there is an IE strategy with pricing probability* 0.586 *whose revenue is at least* 0.9111 *times the maximum revenue.*

PROOF. PROOF. Similarly to the proof of Theorem 6, we calculate the expected (wrt the random partition of buyers into pricing classes and the random order of buyers in the pricing classes) revenue extracted by the generalized IE strategy $\mathrm{IE}(\vec{p},\vec{q})$ from a directed social network $G(V,E,w)$. We recall that for directed social networks, we can ignore loops (i,i). The expected revenue of $\mathrm{IE}(\vec{p},\vec{q})$ from each edge (i,j) is:

$$w_{ij}\sum_{k=1}^{K} q_k p_k(1-p_k)\left(\frac{q_k p_k}{2} + \sum_{\ell=1}^{k-1} q_\ell p_\ell\right)$$

More specifically, for each class k, if both i, j are included in the pricing class k and i appears before j in the random order of the buyers in k, which happens with probability $q_k^2/2$, the revenue extracted from each edge (i, j) is $p_k^2(1-p_k)w_{ij}$. Furthermore, for each pair ℓ, k of pricing classes, $1 \le \ell < k \le K$, if i is included in ℓ and j is included in k, which happens with probability $q_\ell q_k$, the revenue extracted from (i, j) is $p_\ell p_k(1-p_k)w_{ij}$.

Using linearity of expectation and setting $W = \sum_{(i,j)\in E} w_{ij}$, we obtain that the expected revenue of IE$(\vec{q}, \vec{p})$ is:

$$R_{\text{IE}}(\vec{q}, \vec{p}) = W \sum_{k=1}^{K} q_k p_k (1-p_k) \left(\frac{q_k p_k}{2} + \sum_{\ell=1}^{k-1} q_\ell p_\ell \right)$$

Since $W/4$ is an upper bound on the maximum revenue of G, the approximation ratio of IE$(\vec{q}, \vec{p})$ is at least:

$$4 \sum_{k=1}^{K} q_k p_k (1-p_k) \left(\frac{q_k p_k}{2} + \sum_{\ell=1}^{k-1} q_\ell p_\ell \right) , \tag{7}$$

namely at least half of the approximation ratio in the undirected case.

Using $\vec{q} = (0.183, 0.075, 0.075, 0.175, 0.261, 0.231)$ in (7), we obtain an approximation ratio of at least 0.3516. □

This result is to be expected as in case that the optimal revenue was attainable for any instance by using only two prices would imply a surprising degeneracy of the search space. Turning to directed networks, we show that the optimal IE strategy performs surprisingly well, given that in the full information setting the best approximation ratio would be 0.5.

Theorem 9. *For any directed social network, there is an IE strategy with pricing probability $2/3$ whose expected revenue is at least 0.55289 times the maximum revenue.*

As before, we consider an arbitrary directed social network $G(V, E, w)$, start from an arbitrary pricing probability vector $\vec{p}$, and obtain an IE strategy IE$(A, \hat{p})$ by applying randomized rounding to $\vec{p}$. We show that for $\hat{p} = 2/3$, the expected (wrt the randomized rounding choices) revenue of IE$(A, \hat{p})$ is at least 0.55289 times the revenue extracted from G under the best ordering for $\vec{p}$ (which ordering is Unique-Games-hard to approximate within a factor less than 0.5!).

We recall that in the directed case, we can, without loss of generality, ignore loops (i, i). Let $\vec{\pi}$ be the best ordering $\vec{\pi}$ for $\vec{p}$. Then, the maximum revenue extracted from G with pricing probabilities $\vec{p}$ is $R(\vec{\pi}, \vec{p}) \le \sum_{(i,j)\in E} p_i p_j (1-p_j) w_{ij}$.

As in the proof of Theorem 8, we assign each buyer i to the influence set A independently with probability $I(p_i) = \alpha(p_i - 0.5)$, for some $\alpha \in [0, 2]$, and to the exploit set with probability $E(p_i) = 1 - I(p_i)$. By linearity of expectation, the expected (wrt the randomized rounding choices) revenue extracted by IE$(A, \hat{p})$ is:

$$R_{\text{IE}}(A, \hat{p}) = \sum_{(i,j)\in E} \hat{p}(1-\hat{p})(I(p_i)E(p_j) + 0.5\, \hat{p}\, E(p_i)E(p_j))w_{ij}$$

Specifically, IE$(A, \hat{p})$ extracts a revenue of $\hat{p}(1-\hat{p})w_{ij}$ from each edge (i, j), if i is included in the influence set and j is included in the exploit set, and a revenue of $\hat{p}^2(1-\hat{p})w_{ij}$ if both i and j are included in the exploit set $V \setminus A$ and i appears before j in the random order of $V \setminus A$.

The approximation ratio is derived as the minimum ratio between any pair of terms in $R(\vec{\pi}, \vec{p})$ and $R_{\text{IE}}(A, \hat{p})$ corresponding to the same edge (i, j). Thus, we select $\hat{p}$ and α so that the following quantity is maximized:

$$\min_{0.5 \leq x, y \leq 1} \frac{\hat{p}\,(1-\hat{p})(I(x)\,E(y) + 0.5\,\hat{p}\,E(x)\,E(y))}{x\,y\,(1-y)}$$

We observe that for $\hat{p} = 2/3$ and $\alpha = 1.0$, this quantity is simplified to $\min_{y \in [0.5,1]} \frac{2(3-2y)}{27y(1-y)}$. The minimum value is ≈ 0.55289 at $y = \frac{3-\sqrt{3}}{2}$. $\square$

5.1. On the Approximability of the Maximum Revenue for Directed Networks

The results of [10, Lemma 3.2] and [9] suggest that given a pricing probability vector $\mathbf{p}$, it is Unique-Games-hard to compute a vertex ordering π of a directed network G for which the revenue of $(\pi, \mathbf{p})$ is at least 0.5 times the maximum revenue of G under $\mathbf{p}$. An interesting consequence of Theorem 9 is that this inapproximability bound of 0.5 does not apply to revenue maximization in the Uniform Additive Model. In particular, given a pricing probability vector $\mathbf{p}$, Theorem 9 constructs, in linear time, an IE strategy with an expected revenue of at least 0.55289 times the maximum revenue of G under $\mathbf{p}$. This does not contradict the results of [10, 9], because the pricing probabilities of the IE strategy are different from $\mathbf{p}$. Moreover, in the Uniform Additive Model, different acyclic (sub)graphs (equivalently, different vertex orderings) allow for a different fraction of their edge weight to be translated into revenue , while in the reduction of [10, Lemma 3.2], the weight of each edge in an acyclic subgraph is equal to its revenue. The following propositions establish a pair of inapproximabity results for revenue maximization in the Uniform Additive Model. We defer the proofs to the full version of the paper.

Proposition 10. *Assuming the Unique Games conjecture, it is NP-hard to compute an IE strategy with pricing probability $2/3$ that approximates within a factor greater than $3/4$ the maximum revenue of a directed social network in the Uniform Additive Model.*

PROOF. Let $G(V, E, w)$ be a directed social network, and let π^* be a vertex ordering corresponding to an acyclic subgraph of G with a maximum edge weight of W^*. Then, approaching the buyers according to π^* and offering a pricing probability of $2/3$ to each of them, we extract a revenue of $4W^*/27$. Therefore, the maximum revenue of G is at least $4W^*/27$.

Now, we assume an influence set A so that $IE(A, 2/3)$ approximates the maximum revenue of G within a factor of r. Thus, $RIE(A, 2/3) \geq 4rW^*/27$. Let π be the order in which $IE(A, 2/3)$ approaches the buyers, and let (i, j) be any edge with $\pi_i < \pi_j$, namely, any edge from which $IE(A, 2/3)$ extracts some revenue. Since the revenue extracted from each such edge (i, j) is at most $2w_{ij}/9$, the edge weight of the acyclic subgraph defined by π is at least $\frac{9}{2} RIE(A, 2/3) \geq \frac{2r}{3} W^*$.

Hence, given an r-approximate $IE(A, 2/3)$, we can approximate W^* within a ratio of $2r/3$. The proposition follows from [9, Theorem 1.1], which assumes the Unique Games conjecture and shows that it is NP-hard to approximate W^* within a ratio greater than $1/2$. $\square$

Proposition 11. *Assuming the Unique Games conjecture, it is NP-hard to approximate within a factor greater than $27/32$ the maximum revenue of a directed social network in the Uniform Additive Model.*

PROOF. The proof is similar to the proof of Proposition 10. Let $G(V, E, w)$ be a directed social network, and let π^* be a vertex ordering corresponding to an acyclic subgraph of G with a maximum edge weight of W^*. Using π^* and a pricing probability of $2/3$ for all buyers, we obtain that the maximum revenue of G is at least $4W^*/27$.

We assume a marketing strategy $(\pi, \mathbf{p})$ that approximates the maximum revenue of G within a factor of r. Thus, $R(\pi, \mathbf{p}) \geq 4rW^*/27$. Let (i, j) be any edge with $\pi_i < \pi_j$, namely, any edge from which $(\pi, \mathbf{p})$ extracts some revenue. Since the revenue extracted from each such edge (i, j) is at most $w_{ij}/4$, the edge weight of the acyclic subgraph defined by π is at least $4R(\pi, \mathbf{p}) \geq \frac{16r}{27} W^*$

Thus, given an r-approximate marketing strategy $(\pi, \mathbf{p})$, we can approximate W^* within a ratio of $16r/27$. Now, the proposition follows from [9, Theorem 1.1]. $\square$

6. Influence & Exploit via Semidefinite Programming

The main hurdle in obtaining better approximation guarantees for the maximum revenue problem is the loose upper bound of $(N + W)/4$ on the optimal revenue. We do not know how to obtain a stronger upper bound on the maximum revenue. However, in this section, we obtain a strong Semidefinite Programming (SDP) relaxation for the problem of computing the best IE strategy with any given pricing probability $p \in [1/2, 1)$. Our approach exploits the resemblance between computing the best IE strategy and the problems of MAX-CUT (for undirected networks) and MAX-DICUT (for directed networks), and builds on the elegant approach of Goemans and Williamson [8] and Feige and Goemans [7]. Solving the SDP relaxation and using randomized rounding, we obtain, in polynomial time, a good approximation to the best influence set for the given pricing probability p. Then, employing the bounds of Theorem 8 and Theorem 9, we obtain strong approximation guarantees for the maximum revenue problem for both directed and undirected networks.

We start with the case of a directed social network $G(V, E, w)$, which is a bit simpler, because we can ignore loops (i, i) without loss of generality. We observe that for any given pricing probability $p \in [1/2, 1)$, the problem of computing the best IE strategy $IE(A, p)$ is equivalent to solving the following Quadratic Integer Program:

$$\begin{array}{lll} \max \ \frac{p(1-p)}{4} \sum\limits_{(i,j) \in E} w_{ij} [1 + \frac{p}{2} + (1 - \frac{p}{2}) y_0 y_i - (1 + \frac{p}{2}) y_0 y_j - (1 - \frac{p}{2}) y_i y_j] & & \\ \text{s.t.} \qquad\qquad y_i \in \{-1, 1\} & & \forall i \in V \cup \{0\} \end{array}$$

In (Q1), there is a variable y_i for each buyer i and an additional variable y_0 denoting the influence set. A buyer i is assigned to the influence set A, if $y_i = y_0$, and to the exploit set, otherwise. For each edge (i, j), $1 + y_0 y_i - y_0 y_j - y_i y_j$ is 4, if $y_i = y_0 = -y_j$ (i.e., if i is assigned to the influence set and j is assigned to the exploit set), and 0, otherwise. Moreover, $\frac{p}{2}(1 - y_0 y_i - y_0 y_j + y_i y_j)$ is $2p$, if $y_i = y_j = -y_0$ (i.e., if both i and j are assigned to the exploit set), and 0, otherwise. Therefore, the contribution of each edge (i, j) to the objective function of (Q1) is equal to the revenue extracted from (i, j) by $IE(A, p)$.

Following the approach of [8, 7], we relax (Q1) to a Semidefinite Program, where the variables y_i are replaced by vectors $v_i \in S^{n+1}$ and products are replaced by inner products, which can be solved within any precision ε in time polynomial in n and in $\ln \frac{1}{\varepsilon}$ (see e.g. [2]). Thus, given a directed social network $G(V, E, w)$, a pricing probability p, and a parameter $\gamma \in [0, 1]$, the algorithm SDP-IE(p, γ) first computes an optimal solution $v_0, v_1, \ldots, v_n$ to the semidefinite relaxation. Then, following [7], the algorithm maps each vector v_i to a rotated vector v'_i which is coplanar with v_0 and v_i. Finally, the algorithm computes a random vector r uniformly distributed on the unit $(n+1)$-sphere, and assigns each buyer i to the influence set A, if $\mathrm{sgn}(v'_i \cdot r) = \mathrm{sgn}(v_0 \cdot r)$, and to the exploit set $V \setminus A$, otherwise.Extending the analysis of [7], we show that:

Theorem 12. *For any directed social network G, SDP-IE$(2/3, 0.722)$ approximates the maximum revenue extracted from G by the best IE strategy with pricing probability $2/3$ within a factor of 0.9064.*

PROOF. In the following, we let $v_0, v_1, \ldots, v_n$ be an optimal solution to (S1), let $\theta_{ij} = \arccos(v_i \cdot v_j)$ be the angle of any two vectors v_i and v_j, and let $\theta_i = \arccos(v_0 \cdot v_i)$ be the angle of v_0 and any vector v_i. Similarly, we let $\theta'_{ij} = \arccos(v'_i \cdot v'_j)$ be the angle of any two rotated vectors v'_i and v'_j, and let $\theta'_i = \arccos(v_0 \cdot v'_i)$ be the angle of v_0 and any rotated vector v'_i. We first calculate the expected revenue extracted from each edge $(i, j) \in E$ by the IE strategy of SDP-IE(p, γ).

Lemma 13. *The IE strategy of SDP-IE(p, γ) extracts from each edge (i, j) an expected revenue of:*

$$w_{ij}\, p(1-p)\, \frac{(1-\frac{p}{2})\,\theta'_{ij} - (1-\frac{p}{2})\,\theta'_i + (1+\frac{p}{2})\,\theta'_j}{2\pi} \tag{8}$$

PROOF. We first define the following mutually disjoint events:

$$\begin{aligned}
B^{ij} &: \ \mathrm{sgn}(v'_i \cdot r) = \mathrm{sgn}(v'_j \cdot r) = \mathrm{sgn}(v_0 \cdot r) \\
B^{i}_{j} &: \ \mathrm{sgn}(v'_i \cdot r) = \mathrm{sgn}(v_0 \cdot r) \neq \mathrm{sgn}(v'_j \cdot r) \\
B^{j}_{i} &: \ \mathrm{sgn}(v'_j \cdot r) = \mathrm{sgn}(v_0 \cdot r) \neq \mathrm{sgn}(v'_i \cdot r) \\
B_{ij} &: \ \mathrm{sgn}(v'_i \cdot r) = \mathrm{sgn}(v'_j \cdot r) \neq \mathrm{sgn}(v_0 \cdot r)
\end{aligned}$$

Namely, B^{ij} (resp. B_{ij}) is the event that both i and j are assigned to the influence set A (resp. to the exploit set $V \setminus A$), and B^i_j (resp. B^j_i) is the event that i (resp. j) is assigned to the influence set A and j (resp. i) is assigned to the exploit set $V \setminus A$. Also, we let $\mathbb{P}\mathrm{r}[B]$ denote the probability of any event B. Then, the expected revenue extracted from each edge (i, j) is:

$$w_{ij}\, p(1-p)\, \left(\mathbb{P}\mathrm{r}[B^i_j] + \tfrac{p}{2}\, \mathbb{P}\mathrm{r}[B_{ij}]\right) \tag{9}$$

To calculate $\mathbb{P}\mathrm{r}[B^i_j]$ and $\mathbb{P}\mathrm{r}[B_{ij}]$, we use that if r is a vector uniformly distributed on the unit sphere, for any vectors v_i, v_j on the unit sphere, $\mathbb{P}\mathrm{r}[\mathrm{sgn}(v_i \cdot r) \neq \mathrm{sgn}(v_j \cdot r)] = \theta_{ij}/\pi$ [8, Lemma 3.2]. For $\mathbb{P}\mathrm{r}[B^i_j]$, we calculate the probability of the event $B^i_j \cup B^j_i$ that i and j are in different sets, of the event $B^i_j \cup B^{ij}$ that i is in the influence set, and of the event $B^j_i \cup B^{ij}$ that j is in the influence set.

$$\begin{aligned}
\mathbb{P}\mathrm{r}[B^i_j] + \mathbb{P}\mathrm{r}[B^j_i] &= \mathbb{P}\mathrm{r}[B^i_j \cup B^j_i] = \mathbb{P}\mathrm{r}[\mathrm{sgn}(v'_i \cdot r) \neq \mathrm{sgn}(v'_j \cdot r)] = \theta'_{ij}/\pi && (10)\\
\mathbb{P}\mathrm{r}[B^i_j] + \mathbb{P}\mathrm{r}[B^{ij}] &= \mathbb{P}\mathrm{r}[B^i_j \cup B^{ij}] = \mathbb{P}\mathrm{r}[\mathrm{sgn}(v'_i \cdot r) = \mathrm{sgn}(v_0 \cdot r)] = 1 - \theta'_i/\pi && (11)\\
\mathbb{P}\mathrm{r}[B^j_i] + \mathbb{P}\mathrm{r}[B^{ij}] &= \mathbb{P}\mathrm{r}[B^j_i \cup B^{ij}] = \mathbb{P}\mathrm{r}[\mathrm{sgn}(v'_j \cdot r) = \mathrm{sgn}(v_0 \cdot r)] = 1 - \theta'_j/\pi && (12)
\end{aligned}$$

Subtracting (12) from (10) plus (11), we obtain that:

$$\mathbb{P}\mathrm{r}[B^i_j] = \tfrac{1}{2\pi}(\theta'_{ij} - \theta'_i + \theta'_j) \tag{13}$$

For $\mathbb{P}\mathrm{r}[B_{ij}]$, we also need the probability of the event $B_i^j \cup B_{ij}$ that i is in the exploit set, and of the event $B_j^i \cup B_{ij}$ that j is in the exploit set.

$$\mathbb{P}\mathrm{r}[B_i^j] + \mathbb{P}\mathrm{r}[B_{ij}] = \mathbb{P}\mathrm{r}[B_i^j \cup B_{ij}] = \mathbb{P}\mathrm{r}[\mathrm{sgn}(v_i' \cdot r) \neq \mathrm{sgn}(v_0 \cdot r)] = \theta_i'/\pi \tag{14}$$

$$\mathbb{P}\mathrm{r}[B_j^i] + \mathbb{P}\mathrm{r}[B_{ij}] = \mathbb{P}\mathrm{r}[B_j^i \cup B_{ij}] = \mathbb{P}\mathrm{r}[\mathrm{sgn}(v_j' \cdot r) \neq \mathrm{sgn}(v_0 \cdot r)] = \theta_j'/\pi \tag{15}$$

Subtracting (10) from (14) plus (15), we obtain that:

$$\mathbb{P}\mathrm{r}[B_{ij}] = \tfrac{1}{2\pi}(-\theta_{ij}' + \theta_i' + \theta_j') \tag{16}$$

Substituting (13) and (16) in (9), we obtain (8), and conclude the proof of the lemma. $\square$

Since (S1) is a relaxation of the problem of computing the best IE strategy with pricing probability p, the revenue of an optimal $\mathrm{IE}(A, p)$ strategy is at most:

$$\tfrac{p(1-p)}{4} \sum_{(i,j)\in E} w_{ij} \left(1 + \tfrac{p}{2} + (1 - \tfrac{p}{2}) \cos\theta_i - (1 + \tfrac{p}{2}) \cos\theta_j - (1 - \tfrac{p}{2}) \cos\theta_{ij}\right) \tag{17}$$

On the other hand, by Lemma 13 and linearity of expectation, the IE strategy of SDP-IE(p, γ) generates an expected revenue of:

$$\tfrac{p(1-p)}{2\pi} \sum_{(i,j)\in E} w_{ij} \left((1 - \tfrac{p}{2})\, \theta_{ij}' - (1 - \tfrac{p}{2})\, \theta_i' + (1 + \tfrac{p}{2})\, \theta_j'\right) \tag{18}$$

We recall that for each i, $\theta_i' = f_\gamma(\theta_i)$. Moreover, in [7, Section 4], it is shown that for each i, j,

$$\theta_{ij}' = g_\gamma(\theta_{ij}, \theta_i, \theta_j) = \arccos\left(\cos f_\gamma(\theta_i)\, \cos f_\gamma(\theta_j) + \frac{\cos\theta_{ij} - \cos\theta_i\, \cos\theta_j}{\sin\theta_i\, \sin\theta_j}\, \sin f_\gamma(\theta_i)\, \sin f_\gamma(\theta_j)\right)$$

The approximation ratio of SDP-IE(p, γ) is derived as the minimum ratio of any pair of terms in (18) and (17) corresponding to the same edge (i, j). Thus, the approximation ratio of SDP-IE(p, γ) is:

$$\begin{aligned}
\rho(p, \gamma) = \frac{2}{\pi} \min_{0\le x,y,z\le\pi} & \frac{(1 - \frac{p}{2})\, g_\gamma(x, y, z) - (1 - \frac{p}{2}) f_\gamma(y) + (1 + \frac{p}{2}) f_\gamma(z)}{1 + \frac{p}{2} + (1 - \frac{p}{2}) \cos y - (1 + \frac{p}{2}) \cos z - (1 - \frac{p}{2}) \cos x} \\
\text{s.t.} \quad & \cos x + \cos y + \cos z \ge -1 \\
& \cos x - \cos y - \cos z \ge -1 \\
& -\cos x - \cos y + \cos z \ge -1 \\
& -\cos x + \cos y - \cos z \ge -1
\end{aligned}$$

It can be shown numerically, that $\rho(2/3, 0.722) \ge 0.9064$. $\square$

Combining Theorem 12 and Theorem 9, we conclude that:

Theorem 14. *For any directed social network G, the IE strategy computed by SDP-IE*$(2/3, 0.722)$ *approximates the maximum revenue of G within a factor of* 0.5011.

Undirected Social Networks. We apply the same approach to an undirected network $G(V, E, w)$. For any given pricing probability $p \in [1/2, 1)$, the problem of computing the best IE strategy $\mathrm{IE}(A, p)$ for G is equivalent to solving the following Quadratic Integer Program:

$$\max \tfrac{p(1-p)}{2} \sum_{i \in V} w_{ii}\,(1 - y_0 y_i) + \tfrac{p(1-p)}{4} \sum_{i<j} w_{ij}\,(2 + p - p y_0 y_i - p y_0 y_j - (2-p) y_i y_j) \tag{Q2}$$

$$\text{s.t.} \qquad y_i \in \{-1, 1\} \qquad \forall i \in V \cup \{0\}$$

In (Q2), there is a variable y_i for each buyer i and an additional variable y_0 denoting the influence set. A buyer i is assigned to the influence set A, if $y_i = y_0$, and to the exploit set, otherwise. For each loop $\{i, i\}$, $1 - y_0 y_i$ is 2, if i is assigned to the exploit set, and 0, otherwise. For each edge $\{i, j\}$, $i < j$, $2 - 2y_i y_j$ is 4, if i and j are assigned to different sets, and 0, otherwise. Also, $p(1 - y_0 y_i - y_0 y_j + y_i y_j)$ is $4p$, if both i and j are assigned to the exploit set, and 0, otherwise. Therefore, the contribution of each loop $\{i, i\}$ and each edge $\{i, j\}$, $i < j$, to the objective function of (Q2) is equal to the revenue extracted from them by $\mathrm{IE}(A, p)$. The next step is to relax (Q1) to the following Semidefinite Program:

$$\begin{aligned} \max \tfrac{p(1-p)}{2} \sum_{i \in V} w_{ii}\,(1 - v_0 \cdot v_i) + \tfrac{p(1-p)}{4} \sum_{i<j} w_{ij}\,(2 + p - p\, v_0 \cdot v_i - p\, v_0 \cdot v_j - (2-p)\, v_i \cdot v_j) \\ \text{s.t.} \qquad v_i \cdot v_j + v_0 \cdot v_i + v_0 \cdot v_j \geq -1 \\ v_i \cdot v_j - v_0 \cdot v_i - v_0 \cdot v_j \geq -1 \\ -v_i \cdot v_j - v_0 \cdot v_i + v_0 \cdot v_j \geq -1 \\ -v_i \cdot v_j + v_0 \cdot v_i - v_0 \cdot v_j \geq -1 \\ v_i \cdot v_i = 1, \quad v_i \in \mathbb{R}^{n+1} \qquad \forall i \in V \cup \{0\} \end{aligned} \tag{S2}$$

The algorithm is the same as the algorithm for directed networks. Specifically, given an undirected social network $G(V, E, w)$, a pricing probability p, and a parameter $\gamma \in [0, 1]$, the algorithm SDP-IE(p, γ) first computes an optimal solution $v_0, v_1, \ldots, v_n$ to (S2). Then, it maps each vector v_i to a rotated vector v'_i which is coplanar with v_0 and v_i, lies on the same side of v_0 as v_i, and forms an angle $f_\gamma(\theta_i)$ with v_0, where $\theta_i = \arccos(v_0 \cdot v_i)$. Finally, the algorithm computes a random vector r uniformly distributed on the unit $(n+1)$-sphere, and assigns each buyer i to the influence set A, if $\mathrm{sgn}(v'_i \cdot r) = \mathrm{sgn}(v_0 \cdot r)$, and to the exploit set $V \setminus A$, otherwise. We prove that:

Theorem 15. *For any undirected network G, SDP-IE*$(0.586, 0.209)$ *approximates the maximum revenue extracted from G by the best IE strategy with pricing probability* 0.586 *within a factor of* 0.9032*.*

PROOF. We employ the same approach, techniques, and notation as in the proof of Theorem 12. The expected revenue extracted from each loop $\{i, i\}$ is $w_{ii}\, p(1-p)$ times the probability that i is in the exploit set, which is equal to $\mathbb{P}\mathrm{r}[\mathrm{sgn}(v'_i \cdot r) \neq \mathrm{sgn}(v_0 \cdot r)] = \theta'_i/\pi$. Therefore, the algorithm extracts an expected revenue of $w_{ii}\, p(1-p)\, \theta'_i/\pi$ from each loop $\{i, i\}$. Next, we calculate the expected revenue extracted from each (undirected) edge $\{i, j\}$, $i < j$, by the IE strategy of SDP-IE(p, γ).

Lemma 16. *SDP-IE*(p,γ) *extracts from each edge* $\{i,j\}$, $i<j$, *an expected revenue of:*

$$w_{ij}\,p(1-p)\,\frac{(2-p)\,\theta'_{ij}+p\,\theta'_i+p\,\theta'_j}{2\pi}$$

PROOF. Let the events B^i_j, B^j_i, and B_{ij} be defined as in the proof of Lemma 13. In particular, $B^i_j\cup B^j_i$ is the event that i and j are in different sets, and B_{ij} is the event that both i and j are in the exploit set. Thus, the expected revenue extracted from edge $\{i,j\}$ is:

$$w_{ij}\,p(1-p)\left(\mathbb{P}\mathrm{r}[B^i_j\cup B^j_i]+p\,\mathbb{P}\mathrm{r}[B_{ij}]\right) \tag{19}$$

In the proof of Lemma 13, in (10) and (16) respectively, we show that $\mathbb{P}\mathrm{r}[B^i_j\cup B^j_i]=\theta'_{ij}/\pi$, and that $\mathbb{P}\mathrm{r}[B_{ij}]=(-\theta'_{ij}+\theta'_i+\theta'_j)/(2\pi)$. Substituting these in (19), we obtain the lemma. □

Therefore, by linearity of expectation, the expected revenue of SDP-IE(p,γ) is:

$$\tfrac{p(1-p)}{\pi}\sum_{i\in V}w_{ii}\,\theta'_i+\tfrac{p(1-p)}{2\pi}\sum_{i<j}w_{ij}\left((2-p)\,\theta'_{ij}+p\,\theta'_i+p\,\theta'_j\right)\,, \tag{20}$$

where $\theta'_i=f_\gamma(\theta_i)$, for each $i\in V$, and $\theta'_{ij}=g_\gamma(\theta_{ij},\theta_i,\theta_j)$, for each $i,j\in V$.

On the other hand, since (S2) relaxes the problem of computing the best IE strategy with pricing probability p, the revenue of the best IE(A,p) strategy is at most:

$$\tfrac{p(1-p)}{2}\sum_{i\in V}w_{ii}(1-\cos\theta_i)+\tfrac{p(1-p)}{4}\sum_{i<j}w_{ij}\left(2+p-p\cos\theta_i-p\cos\theta_j-(2-p)\cos\theta_{ij}\right) \tag{21}$$

The approximation ratio of SDP-IE(p,γ) is derived as the minimum ratio of any pair of terms in (20) and (21) corresponding either to the same loop $\{i,i\}$ or to the same edge $\{i,j\}$, $i<j$. Therefore, the approximation ratio of SDP-IE(p,γ) for undirected social networks is the minimum of $\rho_1(\gamma)$ and $\rho_2(p,\gamma)$, where:

$$\begin{aligned}
\rho_1(\gamma) &= \frac{2}{\pi}\min_{0\le x\le\pi}\frac{f_\gamma(x)}{1-\cos x}\qquad\text{and}\\
\rho_2(p,\gamma) &= \frac{2}{\pi}\min_{0\le x,y,z\le\pi}\frac{(2-p)\,g_\gamma(x,y,z)+pf_\gamma(y)+pf_\gamma(z)}{2+p-p\cos y-p\cos z-(2-p)\cos x}\\
&\qquad\text{s.t.}\qquad \cos x+\cos y+\cos z\ge -1\\
&\qquad\qquad\qquad \cos x-\cos y-\cos z\ge -1\\
&\qquad\qquad\quad -\cos x-\cos y+\cos z\ge -1\\
&\qquad\qquad\quad -\cos x+\cos y-\cos z\ge -1
\end{aligned}$$

It can be shown numerically, that $\rho_1(0.209)\ge 0.9035$ and that $\rho_2(0.586,0.209)\ge 0.9032$. □

Combining Theorem 15 and Theorem 8, we conclude that:

Theorem 17. *For any undirected social network G, the IE strategy computed by SDP-IE*$(0.586,0.209)$ *approximates the maximum revenue of G within a factor of* 0.8229.

Remark. We can use $\rho(p,\gamma)$ and $\min\{\rho_1(\gamma), \rho_2(p,\gamma)\}$, and compute the approximation ratio of SDP-IE(p,γ) for the best IE strategy with any given pricing probability $p \in [1/2, 1)$. We note that $\rho_1(\gamma)$ is ≈ 0.87856, for $\gamma = 0$ (see e.g. [8, Lemma 3.5]), and increases slowly with γ. Viewed as a function of p, the value of γ maximizing $\rho(p,\gamma)$ and $\rho_2(p,\gamma)$ and the corresponding approximation ratio for the revenue of the best IE strategy increase slowly with p (see also Fig **??** about the dependence of γ and the approximation ratio as a function of p). For example, for directed social networks, the approximation ratio of SDP-IE$(0.5, 0.653)$ (resp. SDP-IE$(0.52, 0.685)$ and SDP-IE$(0.52, 0.704)$) is 0.8942 (resp. 0.8955 and 0.9005). For undirected networks, the ratio of SDP-IE$(0.5, 0.176)$ (resp. SDP-IE$(0.52, 0.183)$ and SDP-IE$(2/3, 0.425)$) is 0.899 (resp. 0.9005 and 0.907). □

7. Discussion

In our paper, we have provided a series of theoretical and algorithmic results for the problem of Revenue Maximization. This is a problem with strong practical and commercial appeal. Companies willing to adopt such strategies, besides estimating model parameters, face the challenge of both computing them over huge social networks and implementing them in practice. The assumption of sequential pricing seriously encumbers the latter endeavour, as it requires time linear in the number of buyers. The way to alleviate this constraint is to make offers to buyers in parallel.

We propose a promotion process proceeding in stages that captures all the essential features while being implementable. The strategy consists of two phases: the Influence phase and the Exploit phase. The Influence phase consists of ℓ stages, where at each stage i buyers in a set S_i receive an offer in parallel. The offer for each buyer is computed(according to the UAM) by taking into account only the influence of buyers that have accepted at previous stages and the pricing probability of the stage. Buyers then have a finite time to respond, after which the offer ceases to apply and we proceed to the next stage. The Exploit phase is implemented, after the end of the Influence phase, via a posted price mechanism for k time intervals. We call this strategy *H*ybridIE$(\mathbf{S}, \mathbf{p})$, where $\mathbf{S}$ is the vector of the ℓ influence sets and $\mathbf{p}$ is the vector of the k prices. The running time of the process is therefore $\Theta(\ell + k)$, where the constant depends on how much time buyers are given to respond to individual offers and how often we can repost a new price.

We argue that *H*ybridIE is also efficient in terms of revenue. The Influence phase aims in seeding the network in more a refined way, while still extracting significant revenue. Making offers in parallel essentially disallows us to exploit the influence between buyers considered in the same stage. Nevertheless, using the generalized IE strategy to construct the sets S_i and if we also implement the Exploit phase via individualized offers(computed via the myopic pricing probability), the remark at Section 4 provides a lower bound on the revenue. Making the transition to the posted price setting at the Exploit phase is important to account for buyers that we do not have information about, and to keep in tradition with retail practices. The price trajectory can be computed by the FPTAS-algorithm of Akhlaghpour et.al for the symmetric variant of the DAILY(k) PROBLEM[1, Section 2.1]. *H*ybridIE fuses three approaches IE, Generalized-IE and posted price mechanisms into an realistic marketing practice. We conclude with some interesting directions for future research.

Clustering. Consider a strategy where the social network is partitioned into a large number of clusters and inside each cluster we assign roles to buyers depending on their potential as influencers or revenue providers. Generalized IE, is an analyzable, though oblivious, way to implement such a strategy, by assuming that any vertex is equally likely to play any role independently of its neighbours. The success of Gen-IE paves the way for a multitude of reasonable heuristics based on clustering.

Eigenvector Centrality. The reason that our SDP-IE strategy performed so well is that, in deciding on the Infleunce set, it algorithmically accounts for the whole network structure. However, SDP based strategies are complex to implement and have high running time for networks with millions of users. This suggests looking into scalable alternatives that provide global awareness. Eigenvector methods are a promising direction to pursue, since their inherent recursive structure enables global reasoning. They could be used as part of marketing practices where an individual is offered a product at a time and price according to his social index(independent of the product under promotion and updated at intervals) and the products base value, similar to the way web search is conducted. Preliminary results show that eigenvector methods combined with local search strategies are extremely successful.

Hybrid Influence Models. In the context of product adoption, network externalities come into play by two complementary ways. One facet is the *information propagation* aspect, exemplified by the *Influence Maximization* setting, and the other is the *network value* aspect of the product, exploited in the *Revenue Maximization* setting. Research, up to now, has focused on the two aspects disjointly and a number of results have been provided. Nevertheless a more pragmatic approach would be to construct a model where both aspects are jointly taken into account. Such an approach might yield interesting new insights unattainable in settings where the two aspects are disunited.

[1] H. Akhlaghpour, M. Ghodsi, N. Haghpanah, V. Mirrokni, H. Mahini, and A. Nikzad. Optimal iterative pricing over social networks. In *Proc. of the 6th Workshop on Internet and Network Economics (WINE '10)*, volume 6484 of *LNCS*, pages 415–423. Springer-Verlag, 2010.

[2] F. Alizadeh. Interior point methods in Semidefinite Programming with applications to combinatorial optimization. *SIAM J. on Optimization*, 5:13–51, 1995.

[3] N. Anari, S. Ehsani, M. Ghodsi, N. Haghpanah, N. Immorlica, H. Mahini, and V. Mirrokni. Equilibrium pricing with positive externalities. In *Proc. of the 6th Workshop on Internet and Network Economics (WINE '10)*, volume 6484 of *LNCS*, pages 424–431. Springer-Verlag, 2010.

[4] D. Arthur, R. Motwani, A. Sharma, and Y. Xu. Pricing strategies for viral marketing on social networks. In *Proc. of the 5th Workshop on Internet and Network Economics (WINE '09)*, volume 5929 of *LNCS*, pages 101–112. Springer-Verlag, 2009.

[5] O. Candogan, K. Bimpikis, and A. Ozdaglar. Optimal pricing in the presence of local network effects. In *Proc. of the 6th Workshop on Internet and Network Economics (WINE '10)*, volume 6484 of *LNCS*, pages 118–132. Springer-Verlag, 2010.

[6] W. Chen, P. Lu, X. Sun, Y. Wang, and Z. Zhu. Pricing in social networks: Equilibrium and revenue maximization. *CoRR*, abs/1007.1501, 2010.

[7] U. Feige and M. Goemans. Aproximating the value of two prover proof systems, with applications to MAX 2SAT and MAX DICUT. In *Proc. of the 3rd Israel Symposium on Theory of Computing and Systems*, pages 182–189. IEEE Computer Society, 1995.

[8] M. Goemans and D. Williamson. Improved approximation algorithms for Maximum Cut and Satisfiability problems using Semidefinite Programming. *J. Assoc. Comput. Mach.*, 42:1115–1145, 1995.

[9] V. Guruswami, R. Manokaran, and P. Raghavendra. Beating the random ordering is hard: Inapproximability of Maximum Acyclic Subgraph. In *Proc. of the 49th IEEE Symposium on Foundations of Computer Science (FOCS '08)*, pages 573–582. IEEE Computer Society, 2008.

[10] J. Hartline, V. Mirrokni, and M. Sundararajan. Optimal marketing strategies over social networks. In *Proc. of the 17th International Conference on World Wide Web (WWW '08)*, pages 189–198. ACM, 2008.

[11] T. Schaefer. The complexity of Satisfiability problems. In *Proc. of the 10th ACM Symposium on Theory of Computing (STOC '78)*, pages 216–226. ACM, 1978.