

Optimal Iterative Pricing over Social Networks

(Extended Abstract)

Hessameddin Akhlaghpour¹, Mohammad Ghodsi^{1,2} *, Nima Haghpanah³, Vahab S. Mirrokni⁴, Hamid Mahini¹, and Afshin Nikzad⁵

¹ Computer Engineering Department, Sharif University of Technology,
akhlaghpour@ce.sharif.edu, ghodsi@sharif.edu, mahini@ce.sharif.edu

² Institute for Research in Fundamental Sciences (IPM), P.O. Box: 19395-5746, Tehran, Iran

³ Northwestern University, EECS Department, nima.haghpanah@eecs.northwestern.edu

⁴ Google Research NYC, 76 9th Ave, New York, NY 10011, mirrokni@gmail.com

⁵ Carnegie Mellon University, Tepper School of Business, anikzad@andrew.cmu.edu

Abstract. We study the optimal pricing for revenue maximization over social networks in the presence of positive network externalities. In our model, the value of a digital good for a buyer is a function of the set of buyers who have already bought the item. In this setting, a decision to buy an item depends on its price and also on the set of other buyers that have already owned that item. The revenue maximization problem in the context of social networks has been studied by Hartline, Mirrokni, and Sundararajan [4], following the previous line of research on optimal viral marketing over social networks [5–7].

We consider the Bayesian setting in which there are some prior knowledge of the probability distribution on the valuations of buyers. In particular, we study two iterative pricing models in which a seller iteratively posts a new price for a digital good (visible to all buyers). In one model, re-pricing of the items are only allowed at a limited rate. For this case, we give a FPTAS for the optimal pricing strategy in the general case. In the second model, we allow very frequent re-pricing of the items. We show that the revenue maximization problem in this case is inapproximable even for simple deterministic valuation functions. In the light of this hardness result, we present constant and logarithmic approximation algorithms when the individual distributions are identical.

1 Introduction

Despite the rapid growth, online social networks have not yet generated significant revenue. Most efforts to design a comprehensive business model for monetizing such social networks [9, 10], are based on contextual display advertising [12]. An alternative way to monetize social networks is viral marketing, or advertising through word-of-mouth. This can be done by understanding the *externalities* among buyers in a social network. The increasing popularity of these networks has allowed companies to collect and use information about inter-relationships among users of social networks. In particular, by designing certain experiments, these companies can determine how users influence each others' activities.

* This author's work has been partially supported by IPM School of CS (contract: CS1388-2-01)

Consider an item or a service for which one buyer's valuation is influenced by other buyers. In many settings, such influence among users are positive. That is, the purchase value of a buyer for a service increases as more people use this service. In this case, we say that buyers have *positive externalities* on each other. Such phenomena arise in various settings. For example, the value of a cell-phone service that offers extra discounts for calls among people using the same service, increases as more friends buy the same service. Such positive externality also appears for any high-quality service through positive reviews or the word-of-mouth advertising.

By taking into account the positive externalities, sellers can employ forward-looking pricing strategies that maximize their long-term expected revenue. For this purpose, there is a clear trade-off between the revenue extracted from a buyer at the beginning, and the revenue from future sales. For example, the seller can give large discounts at the beginning to convince buyers to adopt the service. These buyers will, in turn, influence other buyers and the seller can extract more revenue from the rest of the population, later on. Other than being explored in research papers [4], this idea has been employed in various marketing strategies in practice, e.g., in selling TiVo digital video recorders [11].

Preliminaries. Consider a case of selling multiple copies of a digital good (with no cost for producing a copy) to a set V of n buyers. In the presence of network externality, the valuation of buyer i for the good is a function of buyers who already own that item, $v_i : 2^V \rightarrow R$, i.e., $v_i(S)$ is the value of the digital good for buyer i , if set S of buyers already own that item. We say that users have *positive externality* on each other, if and only if $v_i(S) \leq v_i(T)$ for each two subsets $S \subseteq T \subseteq V$. In general, we assume that the seller is not aware of the exact value of the valuation functions, but she knows the distribution $f_{i,S}$ with an accumulative distribution $F_{i,S}$ for each random variable $v_i(S)$, for all $S \in V$ and any buyer i . Also, we assume that each buyer is interested only in a single copy of the item. The seller is allowed to post different prices at different time steps and buyer i buys the item in a step t if $v_i(S_t) - p_t \geq 0$, where S_t is the set of buyers who own the item in step t , and p_t is the price of the item in that step. Note that $v_i(\emptyset)$ does not need to be zero; in fact $v_i(\emptyset)$ is the value of the item for a user before any other buyer owns the item and influence him.

We study optimal iterative pricing strategies without price discrimination during k time steps. In particular, we assume an *iterative posted price* setting in which we post a *public price* p_i at each step i for $1 \leq i \leq k$. The price p_i at each step i is visible to all buyers, and each buyer might decide to buy the item based on her valuation for the item and the price of the item in that time step. We consider myopic or impatient buyers who buy an item at the first time in which the offered price is less than their valuations. In order to formally define the problem, we should also define each time step. A time step can be long enough in which the influence among users can propagate completely, and we can not modify the price when there is a buyer who is interested to buy the item at the current price. On the other extreme, we can consider settings in which the price of the item changes fast enough that we do not allow the influence amongst buyers to propagate in the same time step. In this setting, as we change the price per time step, we assume the influence among buyers will be effective on the next time step (and not on the same time step). In the following, we define these two problems formally.

Definition 1. The Basic(k) Problem: *In the Basic(k) problem, our goal is to find a sequence p_1, \dots, p_k of k prices in k consecutive time steps or days. A buyer decides to buy the item during a time step as soon as her valuation is more than or equal to the price offered in that time step. In contrast to the Rapid(k) problem, the buyer's decision in a time step immediately affects the valuations of other buyers in the same time step. More precisely, a time step is assumed to end when no more buyers are willing to buy the item at the price at this time step.*

Definition 2. The Rapid(k) Problem: *Given a number k , the Rapid(k) problem is to design a pricing policy for k consecutive days or time steps. In this problem, a pricing policy is to set a public price p_i at the start of time step (or day) i for each $1 \leq i \leq k$. At the start of each time step, after the public price p_i is announced, each buyer decides whether to buy the item or not, based on the price offered on that time step and her valuation. In the Rapid(k) problem, the decision of a buyer during a time step is not affected by the action of other buyers in the same time step* ⁶.

One insight about the Rapid(k) model is that buyers react slowly to the new price and the seller can change the price before the news spreads through the network. On the other hand, in the Basic(k) model, buyers immediately become aware of the new state of the network (the information spreads fast), and therefore respond to the new state of the world before the seller is capable of changing prices. Note that in the Basic(k) problem, the price sequence will be decreasing. If the price posted at any time step is greater than the previous price, no buyer would purchase the product at that time step.

A common assumption studied in the context of network externalities is the assumption of *submodular influence functions*. This assumption has been explored and justified by several previous work in this framework [3–5, 7]. In the context of revenue maximization over social networks, Hartline et. al. [4] state this assumption as follows: suppose that at some time step, S is the set of buyers who have bought the item. We use the notion of *optimal (myopic) revenue* of a buyer for S , which is $R_i(S) = \max_p p \cdot (1 - F_{i,S}(p))$. Following Hartline et.al [4], we consider the optimal revenue function as the *influence function*, and assume that the optimal revenue functions (or influence functions) are submodular, which means that for any two subsets $S \subset T$, and any element $j \notin S$, $R_i(S \cup \{j\}) - R_i(S) \geq R_i(T \cup \{j\}) - R_i(T)$. In other words, submodularity corresponds to a diminishing return property of the optimal revenue function which has been observed in the social network context [3, 5, 7].

Definition 3. *We say that all buyers have identical initial distributions if there exists a distribution F_0 so that the valuation of a player is the sum of two independent random variables, one from F_0 , and another one from $F_{i,S}$, with $F_{i,\emptyset} = 0$.*

Definition 4. *A probability distribution f with accumulative distribution F satisfies the monotone hazard rate condition if the function $h(p) = f(p)/(1 - F(p))$ is monotone non-decreasing.*

Our Contributions. We first show that the deterministic Basic(k) problem is polynomial-time solvable. Moreover, for the Bayesian Basic(k) problem, we present a fully polynomial-

⁶ We use the terms time step and day interchangeably.

time approximation scheme. We study the structure of the optimal solution by performing experiments on randomly generated preferential attachment networks. In particular, we observe that using a small number of price changes, the seller can achieve almost the maximum achievable revenue by many price changes. In addition, this property seems to be closely related to the role of externalities. In particular, the density of the random graph, and therefore the role of network externalities increases, fewer number of price changes are required to achieve almost optimal revenue. We show our experiments in the full version.

Next we show that in contrast to the Basic(k) problem, the Rapid(k) problem is intractable. For the Rapid(k) problem, we show a strong hardness result: we show that the Rapid(k) problem is not approximable within any reasonable approximation factor even in the deterministic case unless $P=NP$. This hardness result holds even if the influence functions are submodular and the probability distributions satisfy the monotone hazard rate condition. In the light of this hardness result, we give an approximation algorithm using a minor and natural assumption. We show that the Rapid(k) problem for buyers with submodular influence functions and probability distributions with the monotone hazard rate condition, and *identical initial distributions* admits logarithmic approximation if k is a constant and a constant-factor approximation if $k \geq n^{\frac{1}{c}}$ for any constant c .

Related work. Optimal viral marketing over social networks have been studied extensively in the computer science literature [6]. For example, Kempe, Kleinberg and Tardos [5] study the following algorithmic question (posed by Domingos and Richardson [3]): How can we identify a set of k influential nodes in a social network to influence such that after convincing this set to use this service, the subsequent adoption of the service is maximized? Most of these models are inspired by the dynamics of adoption of ideas or technologies in social networks and only explore influence maximization in the spread of a *free* good or service over a social network [3, 5, 7]. As a result, they do not consider the effect of pricing in adopting such services. On the other hand, the pricing (as studied in this paper) could be an important factor on the probability of adopting a service, and as a result in the optimal strategies for revenue maximization.

In an earlier work, Hartline, Mirrokni, and Sundararajan [4] study the optimal marketing strategies in the presence of such positive externalities. They study optimal adaptive ordering and pricing by which the seller can maximize its expected revenue. However, in their study, they consider the marketing settings in which the seller can go to buyers one by one (or in groups) and offer a price to those specific buyers. Allowing such price discrimination makes the implementation of such strategies hard. Moreover, price discrimination, although useful for revenue maximization in some settings, may result in a negative reaction from buyers [8].

2 The Basic(k) Problem

We define $B^1(S, p) := \{i | v_i(S) \geq p\} \cup S$. Assume a time step where at the beginning, we set the global price p , and the set S of players already own the item. So $B^1(S, p)$ specifies the set of buyers who immediately want to buy (or already own) the item. As $B^1(S, p)$ will own the item before the time step ends, we can recursively define

$B^k(S, p) = B^1(B^{k-1}(S, p), p)$ and use induction to reason that $B^k(S, p)$ will own the item in this time step. Let $B(S, p) = B^{\hat{k}}(S, p)$, where $\hat{k} = \max\{k | B^k(S, p) - B^{k-1}(S, p) \neq \emptyset\}$, knowing that all buyers in $B(S, p)$ will own the item before the time step ends. One can easily argue that the set $B(S, p)$ does not depend on the order of users who choose to buy the item.

Solving Deterministic Basic(1): In the Basic(1) problem, the goal is to find a price p_1 such that $p_1 \cdot |B(\emptyset, p_1)|$ is maximized. Let $\beta_i := \sup\{p | i \in B(\emptyset, p)\}$ and $\beta := \{\beta_i | 1 \leq i \leq n\}$. WLOG we assume that $\beta_1 > \beta_2 \dots > \beta_n$. Player i will buy the item if and only if the price is set to be less than or equal to β_i .

Lemma 1. *The optimal price p_1 is in the set β .*

Now we provide an algorithm to find p_1 by finding all elements of the set β and considering the profit $\beta_i \cdot |B(\emptyset, \beta_i)|$ of each of them, to find the best result. Throughout the algorithm, we will store a set S of buyers who have bought the item and a global price g . In the beginning $S = \emptyset$ and $g = \infty$. The algorithm consists of $|\beta|$ steps. At the i -th step, we set the price equal to the maximum valuation of remaining players, considering the influence set to be S . We then update the state of the network until it stabilizes, and moves to the next step. Our main claim is as follows. At the end of the i -th step, the set who own the item is $B(\emptyset, \beta_i)$, and the maximum valuation of any remaining player is equal to β_{i+1} .

Generalization to Deterministic Basic(k): We attempt to solve the Basic(k) problem by executing the Basic(1) algorithm. We are looking for an optimal sequence (p_1, p_2, \dots, p_k) in order to maximize $\sum_{i=1}^k |B(\emptyset, p_i) - B(\emptyset, p_{i-1})| \cdot p_i$. We claim that an optimal sequence exists such that for every i , $p_i = \beta_j$ for some $1 \leq j \leq |\beta|$. This can be shown by a proof similar to that of lemma 1. Thus the problem Basic(k) can be solved by considering the subproblem $A[k', m]$ where we must choose a non-increasing sequence π of k' prices from the set $\{\beta_1, \beta_2, \dots, \beta_m\}$, to maximize the profit, and setting the price at the last day to β_m . This subproblem can be solved using the following dynamic program: $A[k', m] = \max_{1 \leq t < m} A[k' - 1, t] + |B(\emptyset, \beta_m) - B(\emptyset, \beta_t)| \cdot \beta_m$.

FPTAS for the Bayesian setting: For the Bayesian (or probabilistic) Basic(k) problem, we run a similar dynamic program, but the main difficulty for this problem is that the space of prices is continuous, and we do not have the same set of candidate prices as we have for the deterministic case. To overcome this issue, we employ a natural idea of discretizing the space of prices. Then we estimate the expected revenue by a sampling technique.

3 The Rapid(k) Problem

As we will see in theorem 2, the Rapid(k) problem is hard to approximate even with submodular influence functions and probability distributions satisfying the monotone hazard rate condition. So we consider the Rapid(k) problem with submodular influence functions and probability distributions satisfying the monotone hazard rate condition, and buyers have identical initial distributions. For this problem, we present an approximation algorithm whose approximation factor is logarithmic for a constant k and its approximation factor is constant for $k \geq n^{\frac{1}{c}}$ for any constant $c > 0$ (See Algorithm 1).

Algorithm 1 Approximation algorithm for Rapid(k) problem

- 1: Compute a price p_0 which maximizes $p(1 - F_0(p))$ and let R_0 be this maximum value.
 - 2: Compute a price $p_{1/2}$ such that $F_0(p_{1/2}) = 0.5$.
 - 3: With probability $\frac{1}{2}$, let $c = 1$, otherwise $c = 2$.
 - 4: **if** $c = 1$ **then**
 - 5: Set the price to the optimal myopic price of F_0 (i.e. p_0) on the first time step and terminate the algorithm after the first time step.
 - 6: **else** $\{c = 2\}$
 - 7: Post the price $p_{1/2}$ on the first time step.
 - 8: Let \bar{S} be the set of buyers that do not buy in the first day, and let their optimal revenues be $R_1(V - \bar{S}) \geq R_2(V - \bar{S}) \geq \dots \geq R_{|\bar{S}|}(V - \bar{S})$.
 - 9: Let p_j be the price which achieves $R_j(V - \bar{S})$, and Pr_j be the probability with which j accepts p_j for any $1 \leq j \leq |\bar{S}|$. Thus we have $R_j(V - \bar{S}) = p_j Pr_j$.
 - 10: Let $d_1 < d_2 < \dots < d_{k-1}$ be the indices returned by lemma 6 as an approximation of the area under the curve $R(V - \bar{S})$ maximizing $\sum_{j=1}^{k-1} (d_j - d_{j-1}) \cdot R_{d_j}(V - S)$.
 - 11: Sort prices $\frac{p_{d_j}}{e}$ for $1 \leq j \leq k - 1$, and offer them in non-increasing order in days 2 to k .
 - 12: **end if**
-

To analyze the expected revenue of the algorithm, we need the following lemmas:

Lemma 2. *Let S be the set formed by sampling each element from a set V independently with probability at least p . Also let f be a submodular set function defined over V , i.e., $f : 2^V \rightarrow R$. Then we have $E[f(S)] \geq pf(V)$ [4].*

Lemma 3. *If the valuation of a buyer is derived from a distribution satisfying the monotone hazard rate condition, she will accept the optimal myopic price with probability at least $1/e$ [4].*

Lemma 4. *Suppose that f is a probability distribution satisfying the monotone hazard rate condition, with expected value μ and myopic revenue $R = \max_p p(1 - F(p))$. Then we have $R(1 + e) \geq \mu$.*

Lemma 5. *Let i be the index maximizing ia_i in the set $\{a_1, a_2, \dots, a_m\}$. Then we have $ia_i \geq \sum_{j=1}^m a_j / (\lceil \log(m + 1) \rceil)$.*

Lemma 6. *For a set $\{a_1 \geq a_2 \geq \dots \geq a_n\}$, let $D = \{d_1 \leq d_2 \leq \dots \leq d_k\}$ be the set of indices maximizing $S(D) = \sum_{j=1}^k (d_j - d_{j-1})a_{d_j}$ (assuming $d_0 = 0$), over all sequences of size k . Then we have $S(D) \in \Theta(\frac{\sum_i a_i}{\log_k n})$.*

Proof idea: We present an algorithm that iteratively selects rectangles, such that after the m -th step the total area covered by the rectangles is at least $m / \log n$ using $4^m - 1$ rectangles. At the start of the m -th step, the uncovered area is partitioned into 4^{m-1} independent parts. In addition, the length of the lower edge of each of these parts is e_p which is at most $n / (2^{m-1})$. The algorithm solves each of these parts independently as follows. We use 3 rectangles for each part in each step. First, using lemma 5 we know that we can use a single rectangle to cover at least $1 / \log e_p$ of the total area of part. Then, we cover the two resulting uncovered parts by two rectangles, which each equally divide the lower edge of the corresponding part.

Theorem 1. *The expected revenue of the algorithm 1 is at least $\frac{1}{8e^2(e+1)\log_k n}$ of the optimal revenue.*

Proof. For simplicity assume that we are allowed to set $k + 1$ prices. In case $c = 1$, we set the optimal myopic price of all players and therefore achieve the expected revenue of nR_0 . If $c = 2$, consider the second day of the algorithm. By lemma 3, we know that each remaining buyer accepts her optimal myopic price with probability at least $1/e$, so for every j we have $Pr_j \geq 1/e \geq Pr_i/e$. In addition, we know that for each $j \leq i$, $R_j(V - \bar{S}) \geq R_i(V - \bar{S}) \geq p_i/e$. We also know that $R_j(V - \bar{S}) \leq p_j$. As a result, $p_j \geq p_i/e$, for each $j \leq i$. Therefore, if we offer the player $j \leq i$ the price p_i/e , she will accept it with probability at least Pr_i/e (she would have accepted p_j with probability at least $Pr_j \geq Pr_i/e$; offering a lower price of p_i/e will only increase the probability of acceptance).

For now suppose that we are able to partition players to k different groups, and offer each group a distinct price. Ignore the additional influence that players can have on each other. In that case, we can find a set $d_1 < d_2 < \dots < d_k$ maximizing $\sum_{j=1}^k (d_j - d_{j-1}) \cdot R_{d_j}(V - \bar{S})$. Assume that D_i is the set of players y with $d_{i-1} < y \leq d_i$. As we argued above, if we offer each of these players the price p_{d_i}/e , she will accept it with probability at least Pr_{d_i}/e . So the expected value of each of the players in D_i when offered p_{d_i}/e is at least $Pr_{d_i}/e \cdot p_{d_i}/e = R_{d_i}(V - \bar{S})/e^2$. The total expected revenue in this case will be $\sum_{j=1}^k (d_j - d_{j-1}) \cdot R_{d_j}(V - \bar{S})/e^2$, which, using lemma 6 is at least $\sum_i R_i(V - \bar{S})/(e^2 \log_k n)$. An important observation is that, if the expected revenue of a player when she is offered a price p is R , her expected revenue will not decrease when she is offered a non-increasing price sequence P which contains p . As a result, we can sort the prices that are offered to different groups, and offer them to *all players* in non-increasing order.

Finally, using Lemma 2, and since every player buys at the first day independently with probability $1/2$, we conclude that any buyer i that remains at the second day observe an expected influence of $R_i(V)/2$ from all other buyers.

As a result, the expected revenue of our algorithm is $nR_0/2$ (from setting p_0 with probability $1/2$ in the first day) plus $\sum_i R_i(V) \cdot (1/8) \cdot (1/(e^2 \log_k n))$. Since we set $p_{1/2}$ with probability $1/2$, a player does not buy at first day with probability $1/2$, and we achieve $1/(e^2 \log_k n)$ of the value of remaining players in the second day. We also know that the expected revenue that can be extracted from any player i is at most $E(F_0) + E(F_{i,V})$. Thus, using lemma 4, we conclude that the approximation factor of the algorithm is $8e^2(e + 1) \log_k n$.

At last, we prove the hardness of the Rapid(k) problem even in the deterministic case with additive (modular) valuation functions. Specifically, we consider the following special case of the problem: (i) $k = n$; (ii) The valuations of the buyers are deterministic, i.e., $f_{i,S}$ is an impulse function, and its value is nonzero only at $v_i(S)$; and finally (iii) The influence functions are additive; $\forall i, j, S$ such that $i \neq j$ and $i, j \notin S$ we have $v_i(S \cup \{j\}) = v_i(S) + v_i(\{j\})$, also each two buyers $i \neq j$, $v_i(\{j\}) \in \{0, 1\}$, and each buyer has a non-negative initial value, i.e, $v_i(\emptyset) \geq 0$.

We use a reduction from the *independent set* problem; We show that using any $\frac{1}{n^{1-\epsilon}}$ -approximation algorithm for the specified subproblem of Rapid(k), any instance of the

independent set problem can be solved in polynomial time. We discard details here and show how to construct an instance of Rapid(k) from an instance of the independent set problem in the full version.

Theorem 2. *The Rapid(k) problem with additive influence functions can not be approximated within any multiplicative factor unless $P=NP$.*

4 Concluding Remarks

In this paper, we introduce new models for studying the optimal pricing and marketing problems over social networks. We study two specific models and show a major difference between the complexity of the optimal pricing in these settings. This paper leaves many problems for future studies.

- We presented results for myopic buyers, but many problems remain open for strategic buyers. Studying optimal pricing strategies for strategic or patient buyers is an interesting problem. In fact, one can model the pricing problem for the seller and the optimal strategy for buyers as a game among buyers and the seller, and study equilibria of such a game. Two possible models have been proposed in [1, 2].
- We studied a monopolistic setting in which a seller does not compete with other sellers. It would be nice to study this problem in the non-monopolistic settings in which other sellers may provide similar items over time, and the seller should compete with other sellers to attract parts of the market.

References

1. N. AhmadiPourAnari, S. Ehsani, M. Ghodsi, N. Haghpanah, N. Immorlica, H. Mahini, and V. S. Mirrokni and. Equilibrium pricing with positive externalities. In *WINE*, 2010.
2. O. Candogan, K. Bimpikis, and A. Ozdaglar. Optimal pricing in the presence of local network effects. In *WINE*, 2010.
3. Pedro Domingos and Matt Richardson. Mining the network value of customers. In *KDD '01*, pages 57–66, New York, NY, USA, 2001. ACM.
4. J. Hartline, V. S. Mirrokni, and M. Sundararajan. Optimal marketing strategies over social networks. In *WWW*, pages 189–198, 2008.
5. David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *KDD '03*, pages 137–146, New York, NY, USA, 2003. ACM.
6. J. Kleinberg. *Cascading behavior in networks: algorithmic and economic issues*. Cambridge University Press, 2007.
7. Elchanan Mossel and Sebastien Roch. On the submodularity of influence in social networks. In *STOC '07*, pages 128–134, New York, NY, USA, 2007. ACM.
8. Richard L. Oliver and Mikhael Shor. Digital redemption of coupons: Satisfying and dissatisfying effects of promotion codes. *Journal of Product and Brand Management*, 12:121–134, 2003.
9. Ed Oswald. http://www.betanews.com/article/Google_Buy_MySpace_Ads_for_900m/1155050350.
10. Katharine Q. Seelye. www.nytimes.com/2006/08/23/technology/23soft.html.
11. Rob Walker. <http://www.slate.com/id/1006264/>.
12. Tim Weber. <http://news.bbc.co.uk/1/hi/business/6305957.stm?lsf>.