

Optimal Pricing in Social Networks with Incomplete Information

Wei Chen¹, Pinyan Lu¹, Xiaorui Sun^{3,*}, Bo Tang^{2,*},
Yajun Wang¹, and Zeyuan Allen Zhu^{4,*,**}

¹ Microsoft Research Asia
{weic,pinyanl,yajunw}@microsoft.com

² Shanghai Jiaotong University
tangbo1@sjtu.edu.cn

³ Columbia University
xiaoruisun@cs.columbia.edu

⁴ MIT CSAIL
zeyuan@csail.mit.edu

Abstract. In revenue maximization of selling a digital product in a social network, the utility of an agent is often considered to have two parts: a private valuation, and linearly additive influences from other agents. We study the incomplete information case where agents know a common distribution about others' private valuations, and make decisions simultaneously. The “rational behavior” of agents in this case is captured by the well-known Bayesian Nash equilibrium.

Two challenging questions arise: how to *compute* an equilibrium and how to *optimize* a pricing strategy accordingly to maximize the revenue assuming agents follow the equilibrium? In this paper, we mainly focus on the natural model where the private valuation of each agent is sampled from a uniform distribution, which turns out to be already challenging.

Our main result is a polynomial-time algorithm that can *exactly* compute the equilibrium and the optimal price, when pairwise influences are non-negative. If negative influences are allowed, computing any equilibrium even approximately is PPAD-hard. Our algorithm can also be used to design an FPTAS for optimizing discriminative price profile.

1 Introduction

In this paper, we study the problem of selling a digital product to agents in a social network. To incorporate social influence, we assume each agent's utility of having the product is the summation of two parts: the private intrinsic valuation and the overall influence from her friends who also have the product. In this paper, we study the linear influence case, i.e., the overall influence is simply the summation of influence values from her friends who have the product.

* Part of this work was done while the authors were visiting Microsoft Research Asia.

** A preliminary version of this work has appeared as a part of the B.Sci thesis of this author [18].

Given such assumption, the purchasing decision of one agent is not solely made based on her own valuation, but also on information about her friends' purchasing decisions. However, a typical agent does not have complete information about others' private valuations, and thus might make the decision based on her belief of other agents' valuations.

We study the case when this belief forms a public distribution, and rely on the solution concept of Bayesian Nash equilibrium [8]. Specifically, each agent knows her own private valuation (also referred to as her *type*); in addition, there is a distribution of this private valuation, publicly known by everyone in the network as well as the seller. We assume that the joint distribution is a product of uniform distributions, and the valuations for all agents are sampled from it.

Computing the Equilibria. Usually, there exist multiple equilibria in this game. We first study the case when all influences are *non-negative*. We show that there exist two special ones: the *pessimistic equilibrium* and the *optimistic equilibrium*, and all other equilibria are between these two. We then design a polynomial time algorithm to compute the pessimistic (resp. optimistic) equilibrium *exactly*.

The overall idea is to utilize the fact that the pessimistic (resp. optimistic) equilibrium is “monotonically increasing” when the price increases. However, the iterative method requires exponential number of steps to converge, just like many potential games which may well be PLS-hard. Our algorithm is based on the line sweep paradigms, by increasing the price p and computing the equilibrium on the way. There are several challenges we have to address to implement the line sweep algorithm. See Section 3.1 for more discussions on the difficulties.

On the negative side, when there exist negative influences among agents, the monotone property of the equilibria does not hold. In fact, we show that computing an approximate equilibrium is PPAD-hard for a given price, by a reduction from the two player Nash equilibrium problem.

Optimal Pricing Strategy. When the seller considers offering a uniform price, our proposed line sweep algorithm calculates the equilibrium as a function of the price. This closed form allows us to find the price for the optimal revenue.

We also discuss the extensions to discriminative pricing setting: agents are partitioned into k groups and the seller can offer different prices to different groups. Depending on whether the algorithm can choose the partition or not, we discuss the hardness and approximation algorithms of these extensions.

1.1 Related Work

Pricing with equilibrium models. When there is social influence, a large stream of literature is focusing on simultaneous games. This is also known as the “two-stage” game where the seller sets the price in the first stage, and agents play a one-shot game in their purchasing decisions. Agents' rational behavior in this case is captured by Nash equilibrium (or Bayesian Nash equilibrium).

The concept and existence of pessimistic and optimistic equilibria is not new. For instance, in analogous problems with externalities, Milgrom and Roberts [12] and Vives [17] have witnessed the existence of such equilibria in the *complete*

information setting. Notice that our pricing problem, when restricted to complete information, can be trivially solved by an iterative method.

In incomplete information setting, Vives and Van Zandt [16] prove a similar existential result using iterative methods. However, they do not provide any convergence guarantee. In our setting, such type of iterative methods may take exponential time to converge. (See the full version of this paper for an example.) Our proposed algorithm instead *exactly* computes the equilibrium, through a much more involved (but constructive) method. In parallel to this work, Sundararajan [15] also discovers the monotonicity of the equilibria, but for symmetry and limited knowledge of the structure (only the degree distribution is known).

It is worth noting that those works above have considered non-linear influences. Though our paper focuses on linear influences, our monotonicity results for equilibria do easily extend to non-linear ones. See Section 2.

When the influence is linear, Candogan, Bimpikis and Ozdaglar [4] study the problem with (uniform) pricing model for a divisible good on sale. It differs from our paper in the model: they are in complete information and divisible good setting; more over, they have relied on a diagonal dominant assumption, which simplifies the problem and ensures the uniqueness of the equilibrium.

Another paper for linear influence is by Bloch and Querou [3], which also studies the uniform pricing model. When the influence is small, they approximate the influence matrix by taking the first 3 layers of influence, and then an equilibrium can be easily computed. They also provide experiments to show that the approximation is numerically good for random inputs.

Pricing with cascading models. In contrast to the simultaneous-move game considered by us (and many others), another stream of work focuses on the cascading models with social influence.

Hartline, Mirrokni and Sundararajan [9] study the *explore and exploit* framework. In their model the seller offers the product to the agents in a sequential manner, and assumes all agents are *myopic*, i.e., each agent is making the decision based on the known results of the previous agents in the sequence. As they have pointed out, if the pricing strategy of the seller and the private value distributions of the subsequent agents are publicly known, the agents can make more “informed” decisions than the myopic ones. In contrast to them, we consider “perfect rational” agents in the simultaneous-move game, where agents make decisions *in anticipation* of what others may do given their beliefs to other agents’ valuations.

Arthur et al. [2] also use the explore and exploit framework, and study a similar problem; potential buyers do not arrive sequentially as in [9], but can choose to buy the product with some probability only if being recommended by friends.

Recently, Akhlaghpour et al. [1] consider the multi-stage model that the seller sets different prices for each stage. In contrast to [9], within each stage, agents are “perfectly rational”, which is characterized by the pessimistic equilibrium in our setting with *complete information*. As mentioned in [1], they did not consider

the case where a rational agent may defer her decision to later stages in order to improve the utility.

Other works. Another notable body of work in computer science is the *optimal seeding* problem (e.g. Kempe et al. [11] and Chen et al. [5]), in which a set of k seeds are selected to maximize the total influence according to some stochastic propagation model. If the value of the product does not exhibit social influence, the seller can maximize the revenue following the optimal auction process by the seminal work of Myerson [13]. Truthful auction mechanisms have also been studied for digital goods, where one can achieve constant ratio of the profit with optimal fixed price [7,10]. On computing equilibria for problems that guarantees to find an equilibrium through iterative methods, most of them, for instance the famous congestion game, is proved to be PLS-hard [6].

2 Model and Solution Concept

We consider the sale of one digital product by a seller with zero cost, to the set of agents $V = [n] = \{1, 2, \dots, n\}$ in a social network. The network is modeled as a simple *directed* graphs $G = (V, E)$ with no self-loops.

- **Valuation:** Agent i has a private value $v_i \geq 0$ for the product. We assume v_i is sampled from a uniform distribution with interval $[a_i, b_i]$ for $0 \leq a_i < b_i$, which we denote as $U(a_i, b_i)$. The values a_i and b_i are common knowledge.
- **Price:** We consider the seller offering the product at a uniform price p .
- **Revenue:** Let $\mathbf{d} = \{d_1, \dots, d_n\} \in \{0, 1\}^n$ be the decision vector the agents make, i.e., $d_i = 1$ if agent i buys the product and 0 otherwise. The revenue of the seller is defined as $\sum_i p \cdot d_i$. When the decisions are random variables, the revenue is defined as the expected payments received from the users.
- **Influence:** Let matrix $T = (T_{j,i})$ with $T_{j,i} \in \mathbb{R}$ and $i, j \in V$ represent the influences among agents, with $T_{j,i} = 0$ for all $(j, i) \notin E$. In particular, $T_{j,i}$ is the utility that agent i receives from agent j , if both of them buy the product. Except for the hardness result, we consider $T_{j,i}$ to be non-negative.
- **Utility:** Let \mathbf{d}_{-i} be the decision vector of the agents other than agent i . For convenience, we denote $\langle d'_i, \mathbf{d}_{-i} \rangle$ the vector by replacing the i -th entry of \mathbf{d} by d'_i . In particular, given the influence matrix T , the utility is defined as:

$$u_i(\langle d_i, \mathbf{d}_{-i} \rangle, v_i, p) = \begin{cases} v_i - p + \sum_{j \in [n]} d_j \cdot T_{j,i}, & \text{if } d_i = 1 \\ 0, & \text{if } d_i = 0 \end{cases} \quad (1)$$

Remark 2.1. In our algorithm later, the requirement $a_i < b_i$ is only for ease of presentation. It can be relaxed to $a_i \leq b_i$ to handle fixed value case as well.

We study the agents' *rational behavior* using the concept Bayesian Nash equilibrium (BNE).¹

¹ Given equilibrium \mathbf{q} in our definition, the strategy profile that each agent i "buys the product iff her valuation $v_i \geq p - \sum_{j \neq i} T_{j,i} q_j$ " is a BNE. See the full version of this paper for details.

Definition 2.2. *The probability vector $\mathbf{q} = (q_1, q_2, \dots, q_n) \in [0, 1]^n$ is an equilibrium at price p , if for all $i \in [n]$: (where med is the median function)*

$$q_i = \Pr_{v_i \sim U(a_i, b_i)} \left[v_i - p + \sum_{j \in [n]} T_{j,i} \cdot q_j \geq 0 \right] = \text{med} \left\{ 0, 1, \frac{b_i - p + \sum_{j \in [n]} T_{j,i} q_j}{b_i - a_i} \right\}. \quad (2)$$

Eq.(2) can be also defined in the language of a transfer function, which we will extensively reply on in the rest of the paper.

Definition 2.3 (Transfer function). *Given price p , we define the transfer function $f_p : [0, 1]^n \rightarrow [0, 1]^n$ as*

$$[f_p(\mathbf{q})]_i = \text{med}\{0, 1, [g_p(\mathbf{q})]_i\} \quad (3)$$

in which

$$[g_p(\mathbf{q})]_i = \frac{b_i - p + \sum_{j \in [n]} T_{j,i} q_j}{b_i - a_i}.$$

Notice that \mathbf{q} is an equilibrium at price p if and only $f_p(\mathbf{q}) = \mathbf{q}$.

Using Brouwer fixed point theorem, the existence of BNE is not surprising, even when influences are negative. However, we will show in Section 4 that computing BNE will be PPAD-hard with negative influences. We now define the pessimistic and optimistic equilibria based on the transfer function.

Definition 2.4. *Let $f_p^{(1)} = f_p$, and $f_p^{(m)}(\mathbf{q}) = f_p(f_p^{(m-1)}(\mathbf{q}))$ for $m \geq 2$. When all influences are non-negative, we define*

- **Pessimistic equilibrium:** $\underline{\mathbf{q}}(p) = \lim_{m \rightarrow \infty} f_p^{(m)}(\mathbf{0})$;
- **Optimistic equilibrium:** $\overline{\mathbf{q}}(p) = \lim_{m \rightarrow \infty} f_p^{(m)}(\mathbf{1})$.

We remark that both limits exist by monotonicity of f (see Fact 2.5 below), when all influences are non-negative. In addition, $\underline{\mathbf{q}}(p)$ and $\overline{\mathbf{q}}(p)$ are both equilibria themselves, because $f_p(\underline{\mathbf{q}}(p)) = \underline{\mathbf{q}}(p)$ and $f_p(\overline{\mathbf{q}}(p)) = \overline{\mathbf{q}}(p)$. We later show that $\underline{\mathbf{q}}(p)$ and $\overline{\mathbf{q}}(p)$ are the lower bound and upper bound for any equilibrium at price p respectively. Now we state some properties of equilibria, which we will use extensively later. See the full version of this paper for proofs.

For two vectors $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^n$, we write $\mathbf{v}_1 \geq \mathbf{v}_2$ if $\forall i \in [n], [\mathbf{v}_1]_i \geq [\mathbf{v}_2]_i$ and we write $\mathbf{v}_1 > \mathbf{v}_2$ if $\mathbf{v}_1 \geq \mathbf{v}_2$ and $\mathbf{v}_1 \neq \mathbf{v}_2$.

Fact 2.5. *When all influences are non-negative, given $p_1 \leq p_2, \mathbf{q}^1 \leq \mathbf{q}^2$, the transfer function satisfies $f_{p_2}(\mathbf{q}^1) \leq f_{p_1}(\mathbf{q}^1) \leq f_{p_1}(\mathbf{q}^2)$.*

Lemma 2.6. *When all influences are non-negative, equilibria satisfy the following properties:*

- a) *For any equilibrium \mathbf{q} at price p , we have $\underline{\mathbf{q}}(p) \leq \mathbf{q} \leq \overline{\mathbf{q}}(p)$.*
- b) *Given price p , for any vector $\mathbf{q} \leq \underline{\mathbf{q}}(p)$, we have $f_p^{(\infty)}(\mathbf{0}) = \underline{\mathbf{q}}(p) = f_p^{(\infty)}(\mathbf{q})$.*
- c) *Given price $p_1 \leq p_2$, we have $\underline{\mathbf{q}}(p_1) \geq \underline{\mathbf{q}}(p_2)$ and $\overline{\mathbf{q}}(p_1) \geq \overline{\mathbf{q}}(p_2)$.*
- d) *$\underline{\mathbf{q}}(p) = \lim_{\varepsilon \rightarrow 0^+} \underline{\mathbf{q}}(p + \varepsilon)$ and $\overline{\mathbf{q}}(p) = \lim_{\varepsilon \rightarrow 0^-} \overline{\mathbf{q}}(p + \varepsilon)$.*

In this paper, we consider the problem that whether we can exactly calculate the pessimistic (resp. optimistic) equilibrium, and whether we can maximize the revenue. The latter is formally defined as follows:

Definition 2.7 (Revenue maximization problem)

Assume the value of agent i is sampled from $U(a_i, b_i)$ and the influence matrix T is given. The revenue maximization problem is to compute an optimal price with respect to the pessimistic equilibrium (resp. optimistic equilibrium):

$$\arg \max_{p>0} \sum_{i \in [n]} p \cdot [\underline{\mathbf{q}}(p)]_i \quad (\text{resp. } \arg \max_{p>0} \sum_{i \in [n]} p \cdot [\overline{\mathbf{q}}(p)]_i).$$

Notice that the optimal revenue with respect to the pessimistic equilibrium is robust against equilibrium selection. By Lemma 2.6(a), no matter which equilibrium the agents choose, this revenue is a minimal guarantee from the seller’s perspective. The revenue guarantees for pessimistic and optimistic equilibria is an important objective to study; see for instance the *price of anarchy* and the *price of stability* in [14] for details.

3 The Main Algorithm

When all influences are non-negative, can we calculate $\underline{\mathbf{q}}(p)$ and $\overline{\mathbf{q}}(p)$ in polynomial time? We answer this question positively in this section by providing an efficient algorithm. Notice that it is possible to iteratively apply the transfer function Eq.(3) to reach the equilibria, but this may take exponential time. See the full version of this paper for a counter example.

3.1 Outline of Our Line Sweep Algorithm

We start to introduce our algorithm with the easy case where valuations of agents are fixed. Consider the pessimistic decision vector as a function of p . By monotonicity, there are at most $O(n)$ different such vectors when p varies from $+\infty$ to 0. In particular, at each price p , if we decrease p gradually to some threshold value, one more agent would change his decision to buy the product. Such kind of process can be casted in the “line sweep algorithm” paradigm.

When the private valuations of the agents are sampled from uniform distributions, the line sweep algorithm is much more complicated. We now introduce the algorithm to obtain the pessimistic equilibrium $\underline{\mathbf{q}}(p)$, while the method to obtain $\overline{\mathbf{q}}(p)$ is similar.² The essence of the line sweep algorithm is processing the events corresponding to some structural changes. We define the possible structures of a probability vector as follows.

Definition 3.1. Given $\mathbf{q} \in [0, 1]^n$, we define the structure function $S : [0, 1]^n \rightarrow \{0, \star, 1\}^n$ satisfying:

$$[S(\mathbf{q})]_i = \begin{cases} 0, & q_i = 0 \\ \star, & q_i \in (0, 1) \\ 1, & q_i = 1. \end{cases} \quad (4)$$

² We sweep the price from $+\infty$ to 0 to compute the pessimistic equilibrium, but we need to sweep from 0 to $+\infty$ for the optimistic one.

Our line sweep algorithm is based on the following fact: when p is sufficiently large, obviously $\underline{\mathbf{q}}(p) = \mathbf{0}$; with the decreasing of p , at some point $p = p_1$ the pessimistic equilibrium $\underline{\mathbf{q}}(p)$ becomes non-zero, and there exists some *structural change* at this moment. Due to the monotonicity of $\underline{\mathbf{q}}(p)$ in Lemma 2.6, such structural changes can happen at most $2n$ times. (Each agent i can contribute to at most two changes: $0 \rightarrow \star$ and $\star \rightarrow 1$.) Therefore, there exist threshold prices $p_1 > p_2 > \dots > p_m$ for $m \leq 2n$ such that within two consecutive prices, the structure of the pessimistic equilibrium remains unchanged and $\underline{\mathbf{q}}(p)$ is a linear function of p . This indicates that the total revenue, i.e., $p \cdot \sum_i [\underline{\mathbf{q}}(p)]_i$, and its maximum value is easy to obtain. If we can compute the threshold prices and the corresponding pessimistic equilibrium $\underline{\mathbf{q}}(p)$ as a function of p , it will be straightforward to determine the optimal price p .

There are several difficulties to address in this line sweep algorithm.

- First, degeneracies, i.e., more than one structural changes in one event, are intrinsic in our problem. Unlike geometric problems where degeneracies can often be eliminated by perturbations, the degeneracies in our problem are persistent to small perturbations.
- Second, to deal with degeneracies, we need to identify the next structural change, which is related to the eigenvector corresponding to the largest eigenvalue of a linear operator. By a careful inspection, we avoid solving eigen systems so that our algorithm can be implemented by pure algebraic computations.
- Third, after the next change is identified, the usual method of pushing the sweeping line further does not work directly in our case. Instead, we recursively solve a subproblem and combine the solution of the subproblem with the current one to a global solution. The polynomial complexity of our algorithm is guaranteed by the monotonicity of the structures.

We first design a line sweep algorithm for the problem with a diagonal dominant condition, which will not contain degenerate cases, in Section 3.2. Then we describe techniques to deal with the unrestricted case in Section 3.3.

3.2 Diagonal Dominant Case

Definition 3.2 (Diagonal dominant condition)

Let $L_{i,j} = T_{j,i}/(b_i - a_i)$ and $L_{i,i} = T_{i,i} = 0$. The matrix $I - L$ is strictly diagonal dominant, if $\sum_j L_{i,j} = \sum_j T_{j,i}/(b_i - a_i) < 1$.

This condition has some natural interpretation on the buying behavior of the agents. It means that the decision of any agent cannot be solely determined by the decisions of her friends. In particular, the following two situations cannot occur *simultaneously* for any agent i and price p : a) agent i will not buy the product regardless of her own valuation when none of her friends bought the product ($p \geq b_i$), and b) agent i will always buy the product regardless of her own valuation when all her friends bought the product ($\sum_j T_{j,i} + a_i \geq p$).

In our line sweep algorithm, we maintain a *partition* $Z \cup W \cup O = V = [n]$, and name Z the *zero set*, W the *working set* and O the *one set*. This corresponds to the structure $\mathbf{s} \in \{0, \star, 1\}^n$ as follows:

$$s_i = 0 (\forall i \in Z), \quad s_i = \star (\forall i \in W), \quad s_i = 1 (\forall i \in O).$$

We use \mathbf{x}_W to denote the restriction of vector \mathbf{x} on set W , and for simplicity we write $\langle \mathbf{x}_Z, \mathbf{x}_W, \mathbf{x}_O \rangle = \mathbf{x}$. Let $L_{W \times W}$ be the projection of matrix L to $W \times W$.

We start from the price $p = +\infty$ where the structure of the pessimistic equilibrium $\underline{\mathbf{q}}(p)$ is $\mathbf{s}^0 = \mathbf{0}$, i.e., $Z = [n]$ and $W = O = \emptyset$. The first event happens when p drops to $p_1 = \max_i b_i$ and $\underline{\mathbf{q}}(p)$ starts to become non-zero.

Assume now we have reached threshold price p_t , the current pessimistic equilibrium is $\mathbf{q}^t = \underline{\mathbf{q}}(p_t)$, and the structure in interval (p_t, p_{t-1}) (or $(p_t, +\infty)$ if $t = 1$) is \mathbf{s}^{t-1} . We define

$$\mathbf{x} = \left(\frac{b_1 - p_t}{b_1 - a_1}, \frac{b_2 - p_t}{b_2 - a_2}, \dots, \frac{b_n - p_t}{b_n - a_n} \right)^T, \quad \text{and } \mathbf{y} = \left(\frac{1}{b_1 - a_1}, \frac{1}{b_2 - a_2}, \dots, \frac{1}{b_n - a_n} \right)^T.$$

To analyze the pessimistic equilibrium in the next price interval, for price $p = p_t - \varepsilon$ where $\varepsilon > 0$, we write function $g_p(\cdot)$ (recall Eq.(3)) as:

$$g_{p_t - \varepsilon}(\mathbf{q}) = \mathbf{x} + \varepsilon \mathbf{y} + L\mathbf{q}.$$

For $p \in (p_t, p_{t-1})$, let partition $Z \cup W \cup O = [n]$ be consistent with the structure \mathbf{s}^{t-1} . According to Def. 3.1 and the right continuity $\mathbf{q}^t = \lim_{p \rightarrow p_t^+} \underline{\mathbf{q}}(p)$ (see Lemma 2.6d), we have

$$\begin{aligned} \forall i \in Z, [g_{p_t}(\mathbf{q}^t)]_i &= [\mathbf{x} + L\mathbf{q}^t]_i \leq 0 \\ \forall i \in W, [g_{p_t}(\mathbf{q}^t)]_i &= [\mathbf{x} + L\mathbf{q}^t]_i \in (0, 1] \\ \forall i \in O, [g_{p_t}(\mathbf{q}^t)]_i &= [\mathbf{x} + L\mathbf{q}^t]_i \geq 1 \end{aligned} \quad (5)$$

Step 1: For any $i \in Z$, if $[\mathbf{x} + L\mathbf{q}^t]_i = 0$, move i from zero set Z to working set W ; for any $i \in W$, if $[\mathbf{x} + L\mathbf{q}^t]_i = 1$, move i from working set W to one set O .

Notice that the structural changes we apply in Step 1 are exactly the changes defining the threshold price p_t . We will see in a moment that after the process in Step 1, the new partition will be the next structure \mathbf{s}^t for $p \in (p_{t+1}, p_t)$. In other words, there is no more structural change at price p_t .

In the next two steps, we calculate the next threshold price p_{t+1} . For notation simplicity, we assume Z, W and O remain unchanged in these two steps. When p decreases by ε , we show that the probability vector of agents in W , $[\underline{\mathbf{q}}(p)]_W$, increases linearly with respect to ε . (See $\mathbf{r}_W(\varepsilon)$ below.) However, this linearity holds until we reach some point, where the next structural change takes place.

Step 2: Define the vector $\mathbf{r}(\varepsilon) \in \mathbb{R}^n$, and let:

$$\begin{aligned} \mathbf{r}_W(\varepsilon) &= \varepsilon(I - L_{W \times W})^{-1} \mathbf{y}_W + \mathbf{q}_W^t \\ &= \varepsilon(I - L_{W \times W})^{-1} \mathbf{y}_W + [\mathbf{x} + L\mathbf{q}^t]_W \\ \mathbf{r}_Z(\varepsilon) &= \mathbf{x}_Z + \varepsilon \mathbf{y}_Z + L_{Z \times W} \mathbf{r}_W(\varepsilon) + L_{Z \times O} \mathbf{1}_O \\ &= \varepsilon(\mathbf{y}_Z + L_{Z \times W}(I - L_{W \times W})^{-1} \mathbf{y}_W) + [\mathbf{x} + L\mathbf{q}^t]_Z \\ \mathbf{r}_O(\varepsilon) &= \mathbf{x}_O + \varepsilon \mathbf{y}_O + L_{O \times W} \mathbf{r}_W(\varepsilon) + L_{O \times O} \mathbf{1}_O \\ &= \varepsilon(\mathbf{y}_O + L_{O \times W}(I - L_{W \times W})^{-1} \mathbf{y}_W) + [\mathbf{x} + L\mathbf{q}^t]_O \end{aligned} \quad (6)$$

Clearly, $\mathbf{r}(\varepsilon)$ is linear to ε and we write $\mathbf{r}(\varepsilon) = \varepsilon \boldsymbol{\ell} + (\mathbf{x} + L\mathbf{q}^t)$ where $\boldsymbol{\ell} = \langle \ell_1, \ell_2, \dots, \ell_n \rangle \in \mathbb{R}^n$ is the linear coefficient derived from Eq.(6). When $I - L$ is strictly diagonal dominant, the largest eigenvalue of $L_{W \times W}$ is smaller than 1. Using this property one can verify (see full version) that $\boldsymbol{\ell}$ is strictly positive.

$$\text{Step 3} \quad \varepsilon_{\min} = \min \left\{ \min_{i \in Z} \left\{ \frac{0 - [\mathbf{x} + L\mathbf{q}^t]_i}{\ell_i} \right\}, \min_{i \in W} \left\{ \frac{1 - [\mathbf{x} + L\mathbf{q}^t]_i}{\ell_i} \right\} \right\} \quad (7)$$

Using the positiveness of vector $\boldsymbol{\ell}$ one can verify that $\varepsilon_{\min} > 0$. Also, the next threshold price $p_{t+1} = p_t - \varepsilon_{\min}$. (See full version for proofs.)

Lemma 3.3. $\forall 0 < \varepsilon \leq \varepsilon_{\min}$, $\underline{\mathbf{q}}(p_t - \varepsilon) = \langle \mathbf{0}_Z, \mathbf{r}_W(\varepsilon), \mathbf{1}_O \rangle$.

We remark here that the above lemma has confirmed that our structural adjustments in Step 1 are correct and complete. Now we let $p_{t+1} = p_t - \varepsilon_{\min}$, $\mathbf{q}^{t+1} = \langle \mathbf{0}_Z, \mathbf{r}_W(\varepsilon_{\min}), \mathbf{1}_O \rangle$. The next structural change will take place at $p = p_{t+1}$. This is because according to the definition of ε_{\min} (Eq.(7)), there must be some

$$i \in W \wedge [\mathbf{x} + \varepsilon_{\min} \mathbf{y} + L\mathbf{q}^{t+1}]_i = 1, \quad \text{or} \quad i \in Z \wedge [\mathbf{x} + \varepsilon_{\min} \mathbf{y} + L\mathbf{q}^{t+1}]_i = 0.$$

One can see that in the next iteration, this i will move to one set O or working set W accordingly. Therefore, we can iteratively execute the above three steps by sweeping the price further down.

The return value of our constrained line sweep method is a function $\underline{\mathbf{q}}$ which gives the pessimistic equilibrium for any price $p \in \mathbb{R}$, and $\underline{\mathbf{q}}(p)$ is a piecewise linear function of p with no more than $2n + 1$ pieces. All three steps in our algorithm can be done in polynomial time. Since there are only $O(n)$ threshold prices, we have the following result.

Theorem 3.4. *When the matrix $I - L$ is strictly diagonal dominant, we can calculate the pessimistic equilibrium $\underline{\mathbf{q}}(p)$ (resp. $\overline{\mathbf{q}}(p)$) for any given price p in polynomial time, together with the optimal revenue.*

3.3 General Case

After relaxing the diagonal dominance condition, the algorithm becomes more complicated. This can be seen from this simple scenario. There are 2 agents, with $[a_1, b_1] = [a_2, b_2] = [0, 1]$, and $T_{1,2} = T_{2,1} = 2$. One can verify that $\underline{\mathbf{q}}(p) = (0, 0)^T$ when $p \geq 1$; $\underline{\mathbf{q}}(p) = (1, 1)^T$ when $p < 1$.

In this example, there is an *equilibrium jump* at price $p = 1$, i.e., $\underline{\mathbf{q}}(1) \neq \lim_{p \rightarrow 1^-} \underline{\mathbf{q}}(p)$. Our previous algorithm essentially requires that both the left and the right continuity of $\underline{\mathbf{q}}(p)$. However, only the right continuity is unconditional by Lemma 2.6d. More importantly, degeneracies may occur: the new structure \mathbf{s}^t when $p = p_t$ cannot be determined all in once in Step 1. When p goes from $p_t + \varepsilon$ to $p_t - \varepsilon$, there might take place even two-stage jumps: some index i might leave Z for O , without being in the intermediate state.

Let $\rho(L)$ be the largest norm of the eigenvalues in matrix L . The ultimate reason for such degeneracies, is $\rho(L_{W \times W}) \geq 1$ and $(I - L_{W \times W})^{-1} \neq \lim_{m \rightarrow \infty} (I + L_{W \times W} + \dots + L_{W \times W}^{m-1})$. We will prove shortly in such cases, those structural changes in Step 1 are *incomplete*, that is, as p sweeps across p_t , at least one more structural change will take place. We derive a method to identify one *pivot*, i.e. an additional structural change, in polynomial time. Afterwards, we recursively solve a subproblem with set O taken out, and combine the solution from the subproblem with the current one. The follow lemma shows that whether $\rho(L) < 1$ can be determined efficiently.

Lemma 3.5. *Given non-negative matrix M , if $I - M$ is reversible and $(I - M)^{-1}$ is also non-negative, then $\rho(M) < 1$; on the contrary, if $I - M$ is degenerate or if $(I - M)^{-1}$ contains negative entries, $\rho(M) \geq 1$.*

Finding the pivot. When $\rho(L_{W \times W}) < 1$ for the new working set W , one can find the next threshold price p_{t+1} following Step 2 and 3 in the previous subsection. Now, we deal with the case that $\rho(L_{W \times W}) \geq 1$ by showing that there must exists some additional agent $i \in W$ such that $[\underline{\mathbf{q}}(p)]_i = 1$ for any p smaller than the current price. We call such agent a *pivot*.

Since $\rho(L_{W \times W}) \geq 1$, we can always find a non-empty set $W_1 \subset W$ and $W_2 = W_1 \cup \{w\} \subset W$, satisfying $\rho(L_{W_1 \times W_1}) < 1$ but $\rho(L_{W_2 \times W_2}) \geq 1$. The pair (W_1, W_2) can be found by ordering the elements in W and add them to W_1 one by one. We now show that there is a pivot in W_2 .

As $L_{W_2 \times W_2}$ is a non-negative matrix, based on knowledge from spectral theory, exists a non-zero eigenvector $\mathbf{u}_{W_2} \geq \mathbf{0}_{W_2}$ such that $L_{W_2 \times W_2} \mathbf{u}_{W_2} = \lambda \mathbf{u}_{W_2}$ and $\lambda = \rho(L_{W_2 \times W_2}) \geq 1$. \mathbf{u}_{W_2} can be extended to $[n]$ by defining $\mathbf{u}_{[n] \setminus W_2} = \mathbf{0}_{[n] \setminus W_2}$. Let

$$k = \arg \min_{k \in W_2, u_k \neq 0} \frac{1 - q_k^t}{u_k} = \arg \min_{k \in [n], u_k \neq 0} \frac{1 - q_k^t}{u_k} \quad (8)$$

Now we prove that k is a pivot. Intuitively, if we slightly increase the probability vector $\mathbf{q}_{W_2}^t$ by $\delta \mathbf{u}_{W_2}$, where δ is a small constant, by performing the transfer function only on agents in W m times, their probability will increase by $\delta(1 + \lambda + \dots + \lambda^m) \mathbf{u}_{W_2}$, while $\lambda \geq 1$. Therefore, after performing the transfer function sufficiently many times, agent $k \in W_2$'s probability will hit 1 first.

Lemma 3.6. *$\forall W_2 \subset W$ s.t. $\rho(L_{W_2 \times W_2}) \geq 1$, we have $\forall \varepsilon > 0$, $[\underline{\mathbf{q}}(p_t - \varepsilon)]_k = 1$.*

We remark that if we can exactly estimate the eigenvector (which may be irrational), then the above lemma has already determined that the k defined in Eq.(8) is a pivot. To avoid the eigenvalue computation, we find a quasi-eigenvector \mathbf{u} in the following manner.

$$\mathbf{u} = \begin{cases} \mathbf{u}_{W_1} = (I - L_{W_1 \times W_1})^{-1} L_{W_1 \times \{w\}}; \\ u_w = 1; \\ \mathbf{u}_{Z \cup O \cup W \setminus W_2} = \mathbf{0}_{Z \cup O \cup W \setminus W_2}. \end{cases} \quad (9)$$

The meaning of the above vector is as follows. If we raise agent w 's probability by δ , those probabilities of agents in W_1 increase proportionally to $L_{W_1 \times \{w\}} \delta$.

Assuming that we ignore the probability changes outside W_2 (which will even increase the probabilities in W_2), the probability of agents in W_1 will eventually converge to $(I + L_{W_1 \times W_1} + L_{W_1 \times W_1}^2 + \dots)L_{W_1 \times \{w\}}\delta = (I - L_{W_1 \times W_1})^{-1}L_{W_1 \times \{w\}}\delta$.

We will see that the real probability vector increases at least “as much as if we increase in the direction of \mathbf{u} ”. In other words, we pick a pivot in the same way as Eq.(8). The following is the critical lemma to support our result.

Lemma 3.7. *For \mathbf{u} in Eq.(9) and k in Eq.(8), we have $\forall \varepsilon > 0, [\underline{\mathbf{q}}(p_t - \varepsilon)]_k = 1$.*

Recursion on the subproblem. Let $W' = W \setminus \{k\}$, $O' = O \cup \{k\}$, and we consider a subproblem with $n' = n - |O'| < n$ agents, where k is the pivot identified in the previous section. This subproblem is a projection of the original one, assuming that the agents in O' always tend to buy the product.

$$\forall i \in Z \cup W', \quad [a'_i, b'_i] = [a_i + \sum_{j \in O'} T_{j,i}, b_i + \sum_{j \in O'} T_{j,i}]. \quad (10)$$

By recursively solving this new instance, we can solve the pessimistic equilibrium of the subproblem for any given price p . This recursive procedure will eventually terminate because every invocation reduces the number of agents by at least 1. The following lemma tells us that for any $p < p_t$, the pessimistic equilibrium of the original problem and the subproblem are one-to-one.

Lemma 3.8. *Let $\underline{\mathbf{q}}'(p)$ be the pessimistic equilibrium function in the subproblem. We have:*

$$\forall p < p_t, \underline{\mathbf{q}}(p) = \langle \underline{\mathbf{q}}'(p), \mathbf{1}_{O'} \rangle.$$

At this moment we have solved the pessimistic equilibrium $\underline{\mathbf{q}}(p)$ for $p < p_t$, and thus solved the original problem. Again $\underline{\mathbf{q}}(p)$ is a piecewise linear function of p with no more than $2n + 1$ pieces.

Theorem 3.9. *For matrix T satisfying $T_{i,i} = 0$ and $T_{i,j} \geq 0$, in polynomial time we can calculate the pessimistic equilibrium $\underline{\mathbf{q}}(p)$ (resp. $\overline{\mathbf{q}}(p)$) at any price p , together with the optimal revenue.*

4 Extensions

In the full version of this paper, we also prove the following theorems. When the influence values can be negative, it is actually PPAD-hard to compute an *approximate* equilibrium. We define a probability vector \mathbf{q} to be an ε -approximate equilibrium for price p if:

$$q_i \in (q'_i - \varepsilon, q'_i + \varepsilon),$$

where $q'_i = \text{med} \left\{ 0, 1, \frac{b_i - p + \sum_{j \in [n]} T_{j,i} q_j}{b_i - a_i} \right\}$. We have the following theorem:

Theorem 4.1. *It is PPAD-hard to compute an n^{-c} -approximate equilibrium of our pricing system for any $c > 1$ when influences can be negative.*

In discriminative pricing setting, we study the revenue maximization problem in two natural models. We assume the agents are partitioned into k groups.

The seller can offer different prices to different groups. The first model we consider is the fixed partition model, i.e., the partition is predefined. In the second model, we allow the seller to partition the agents into k groups and offer prices to the groups respectively. We have the following two theorems:

Theorem 4.2. *There is an FPTAS for the discriminative pricing problem in the fixed partition case with constant k .*

Theorem 4.3. *It is NP-hard to compute the optimal pessimistic discriminative pricing equilibrium in the choosing partition case.*

References

1. Akhlaghpour, H., Ghodsi, M., Haghpanah, N., Mirrokni, V.S., Mahini, H., Nikzad, A.: Optimal Iterative Pricing over Social Networks. In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 415–423. Springer, Heidelberg (2010)
2. Arthur, D., Motwani, R., Sharma, A., Xu, Y.: Pricing Strategies for Viral Marketing on Social Networks. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 101–112. Springer, Heidelberg (2009)
3. Bloch, F., Quérou, N.: Pricing with local network externalities. Technical report (July 2009)
4. Candogan, O., Bimpikis, K., Ozdaglar, A.E.: Optimal Pricing in the Presence of Local Network Effects. In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 118–132. Springer, Heidelberg (2010)
5. Chen, W., Wang, Y., Yang, S.: Efficient influence maximization in social networks. In: SIGKDD 2009, pp. 199–208 (2009)
6. Fabrikant, A., Papadimitriou, C., Talwar, K.: The complexity of pure nash equilibria. In: STOC 2004, pp. 604–612 (2004)
7. Goldberg, A.V., Hartline, J.D., Karlin, A.R., Saks, M., Wright, A.: Competitive auctions. *Games and Economic Behavior* 55(2), 242–269 (2006)
8. Harsanyi, J.C.: Games with incomplete information played by “bayesian” players, i-iii. part i. the basic model. *Management Science* 14(3), 159–182 (1967)
9. Hartline, J., Mirrokni, V., Sundararajan, M.: Optimal marketing strategies over social networks. In: WWW 2008, pp. 189–198 (2008)
10. Hartline, J.D., McGrew, R.: From optimal limited to unlimited supply auctions. In: ACM-EC 2005, pp. 175–182 (2005)
11. Kempe, D., Kleinberg, J., Tardos, É.: Maximizing the spread of influence through a social network. In: SIGKDD 2003, pp. 137–146 (2003)
12. Milgrom, P., Roberts, J.: Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* 58(6), 1255–1277 (1990)
13. Myerson, R.B.: Optimal auction design. *Math. Oper. Res.* 6(1), 58 (1981)
14. Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V.V.: *Algorithmic game theory*. Cambridge University Press (2007)
15. Sundararajan, A.: Local network effects and complex network structure. *The B.E. Journal of Theoretical Economics* 7(1) (2008)
16. Van Zandt, T., Vives, X.: Monotone equilibria in bayesian games of strategic complementarities. *Journal of Economic Theory* 134(1), 339–360 (2007)
17. Vives, X.: Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics* 19(3), 305–321 (1990)
18. Zhu, Z.A.: Two Topics on Nash Equilibrium in Algorithmic Game Theory. Bachelor’s thesis. Tsinghua University (June 2010)