Optimal Pricing in the Presence of Local Network Effects

Ozan Candogan[1], Kostas Bimpikis[2], and Asuman Ozdaglar[1]

[1] Laboratory for Information and Decision Systems, MIT
`candogan@mit.edu, asuman@mit.edu`
[2] Operations Research Center and Laboratory for Information and Decision Systems, MIT
`kostasb@mit.edu`

Abstract

We study the optimal pricing strategies of a monopolist selling a divisible good (service) to consumers that are embedded in a social network. A key feature of our model is that consumers experience a (positive) *local network effect*. In particular, each consumer's usage level depends directly on the usage of her *neighbors* in the social network structure. Thus, the monopolist's optimal pricing strategy may involve offering discounts to certain agents[3], who have a *central* position in the underlying network. Our results can be summarized as follows. First, we consider a setting where the monopolist can offer individualized prices and derive an explicit characterization of the optimal price for each consumer as a function of her network position. In particular, we show that it is optimal for the monopolist to charge each agent a price that is proportional to her *Bonacich centrality* in the social network. In the second part of the paper, we discuss the optimal strategy of a monopolist that can only choose a single uniform price for the good and derive an algorithm polynomial in the number of agents to compute such a price. Thirdly, we assume that the monopolist can offer the good in two prices, full and discounted, and study the problem of determining which set of consumers should be given the discount. We show that the problem is NP-hard, however we provide an explicit characterization of the set of agents that should be offered the discounted price. Finally, we describe an approximation algorithm for finding the optimal set of agents. We show that if the profit is nonnegative under any feasible price allocation, the algorithm guarantees at least 88 % of the optimal profit.

1 Introduction

Inarguably social networks, that describe the pattern and level of interaction of a set of agents, are instrumental in the propagation of information and act as conduits of influence among its members. Their importance is best exemplified by the overwhelming success of online social networking communities, such as Facebook and Twitter. The ubiquity of these internet based services, that are built around social networks, has made possible the collection of vast amounts of data on the structure and intensity of social interactions. The question that arises naturally is whether firms can intelligently use the available data to improve their business strategies.

In this paper, we focus on the question of using the potentially available data on network interactions to improve the pricing strategies of a seller, that offers a divisible good (service). A main feature of the products we consider is that they exhibit a local (positive) *network effect*: increasing the usage level of a consumer has a positive impact on the usage levels of her peers. As concrete examples of such goods, consider online games (e.g., World of Warcraft, Second Life) and social networking tools and communities (e.g., online dating services, employment websites etc.). More generally, the local network effect can capture *word of mouth* communication among agents:

---

[3] We use the terms "agent" and "consumer" interchangeably.

agents typically form their opinions about the quality of a product based on the information they obtain from their peers.

How can a monopolist exploit the above network effects and maximize her revenues? In particular, in such a setting it is plausible that an optimal pricing strategy may involve favoring certain agents by offering the good at a discounted price and subsequently exploiting the positive effect of their usage on the rest of the consumers. At its extreme, such a scheme would offer the product for free to a subset of consumers hoping that this would have a large positive impact on the purchasing decisions of the rest. Although such strategies have been used extensively in practice, mainly in the form of ad hoc or heuristic mechanisms, the available data enable companies to effectively target the agents to maximize that impact.

The goal of the present paper is to characterize optimal pricing strategies as a function of the underlying social interactions in a stylized model, which features consumers that are embedded in a given social network and influencing each other's decisions. In particular, a monopolist first chooses a pricing strategy and then consumers choose their usage levels, so as to maximize their own utility. We capture the local positive network effect by assuming that a consumer's utility is increasing in the usage level of her peers. We study three variations of the baseline model by imposing different assumptions on the set of available pricing strategies, that the monopolist can implement.

First, we allow the monopolist to set an individual price for each of the consumers. We show that the optimal price for each agent can be decomposed into three components: a fixed cost, that does not depend on the network structure, a markup and a discount. Both the markup and the discount are proportional to the *Bonacich centrality* of the agent's neighbors in the social network structure, which is a sociological measure of network influence. The Bonacich centrality measure, introduced by [2], can be computed as the stationary distribution of a random walk on the underlying network structure. Hence, the nodes with the highest centrality are the ones that are visited by the random walk most frequently. Intuitively, agents get a discount proportional to the amount they influence their peers to purchase the product, and they receive a markup if they are strongly influenced by other agents in the network. Our results provide an economic foundation for this sociological measure of influence.

Perfect price differentiation is typically hard to implement. Therefore, in the second part of the paper we study a setting, where the monopolist offers a single uniform price for the good. Intuitively, this price might make the product unattractive for a subset of consumers, who end up not purchasing, but the monopolist recovers the revenue losses from the rest of the consumers. We develop an algorithm that finds the optimal single price in time polynomial in the number of agents. The algorithm considers different subsets of the consumers and finds the optimal price provided that only the consumers in the given subset purchase a positive amount of the good. First, we show that given a subset S we can find the optimal price p_S under the above constraint in closed form. Then, we show that we only need to consider a small number of such subsets. In particular, we rank the agents with respect to a weighted centrality index and at each iteration of the algorithm we drop the consumer with the smallest such index and let S be the set of remaining consumers.

Finally, we consider an intermediate setting, where the monopolist can choose one of a small number of prices for each agent. For exposition purposes, we restrict the discussion to two prices, *full* and *discounted*. We show that the resulting problem, i.e., determining the optimal subset of consumers to offer the discounted price, is NP-hard[4]. We also provide an approximation algorithm that recovers (in polynomial time) at least 88 % of the optimal revenue.

As mentioned above, a main feature of our model is the positive impact of a consumer's purchasing decision to the purchasing behavior of other consumers. This effect, known as *network*

---

[4] The hardness result can be extended to the case of more than two prices.

externality, is extensively studied in the economics literature (e.g., [8], [15]). However, the network effects in those studies are of *global* nature, i.e., the utility of a consumer depends directly on the behavior of the whole set of consumers. In our model, consumers interact directly only with a subset of agents. Although interaction is local for each consumer, her utility may depend on the global structure of the network, since each consumer potentially interacts indirectly with a much larger set of agents than just her peers.

Given a set of prices, our model takes the form of a *network game* among agents that interact locally. A recent series of papers studies such games, e.g., [1], [3], [7], [9]. A key modeling assumption in [1], [3] and [7], that we also adopt in our setting, is that the payoff function of an agent takes the form of a linear-quadratic function. Ballester et al. in [1] were the first to note the linkage between Bonacich centrality and Nash equilibrium outcomes in a single stage game with local payoff complementarities. Our characterization of optimal prices when the monopolist can perfectly price differentiate is reminiscent of their results, since prices are inherently related to the Bonacich centrality of each consumer. However, both the motivation and the analysis are quite different, since ours is a two-stage game, where a monopolist chooses prices to maximize her revenue subject to equilibrium constraints. Also, [3] and [7] study a similar game to the one in [1] and interpret their results in terms of public good provision. A number of recent papers ([5], [10] and [18]) have a similar motivation to ours, but take a completely different approach: they make the assumption of limited knowledge of the social network structure, i.e., they assume that only the degree distribution is known, and thus derive optimal pricing strategies that depend on this first degree measure of influence of a consumer. In our model, we make the assumption that the monopolist has complete knowledge of the social network structure and, thus, obtain qualitatively different results: the degree is not the appropriate measure of influence but rather prices are proportional to the Bonacich centrality of the agents. On the technical side, note that assuming more global knowledge of the network structure increases the complexity of the problem in the following way: if only the degree of an agent is known, then essentially there are as many different *types* of agents as there are different degrees. This is no longer true when more is known: then, two agents of the same degree may be of different type because of the difference in the characteristics of their neighbors, and therefore, optimal prices charged to agents may be different.

Finally, there is a recent stream of literature in computer science, that studies a set of algorithmic questions related to marketing strategies over social networks. Kempe et al. in [16] discuss optimal *network seeding* strategies over social networks, when consumers act myopically according to a pre-specified rule of thumb. In particular, they distinguish between two basic models of diffusion: the *linear threshold model*, which assumes that an agent adopts a behavior as soon as adoption in her neighborhood of peers exceeds a given threshold and *independent cascade model*, which assumes that an adopter *infects* each of her neighbors with a given probability. The main question they ask is finding the optimal set of initial adopters, when their number is given, so as to maximize the eventual adoption of the behavior, when consumers behave according to one of the diffusion models described above. They show that the problem of *influence maximization* is NP-hard and provide a greedy heuristic, that achieves a solution, that is provably within 63 % of the optimal.

Closest in spirit with our work, is [13], which discusses the optimal marketing strategies of a monopolist. Specifically, they assume a general model of influence, where an agent's willingness to pay for the good is given by a function of the subset of agents that have already bought the product, i.e., $u_i : 2^V \to \mathbb{R}_+$, where u_i is the willingness to pay for agent i and V is the set of consumers. They restrict the monopolist to the following set of *marketing strategies*: the seller visits the consumers in some sequence and makes a take-it-or-leave-it offer to each one of them. Both the sequence of visits as well as the prices are chosen by the monopolist. They provide a dynamic programming algorithm that outputs the optimal pricing strategy for a symmetric setting, i.e., when the agents

are ex-ante identical (the sequence of visits is irrelevant in this setting). Not surprisingly the optimal strategy offers discounts to the consumers that are visited earlier in the sequence and then extracts revenue from the rest. The general problem, when agents are heterogeneous, is NP-hard, thus they consider approximation algorithms. They show, in particular, that *influence-and-exploit* strategies, that offer the product for free to a strategically chosen set A, and then offer the myopically optimal price to the remaining agents provably achieve a constant factor approximation of the optimal revenues under some assumptions on the influence model. However, this paper does not provide a qualitative insight on the relation between optimal strategies and the structure of the social network. In contrast, we are mainly interested in characterizing the optimal strategies as a function of the underlying network.

The rest of paper is organized as follows. Section 2 introduces the model. In Section 3 we begin our analysis by characterizing the usage level of the consumers at equilibrium given the vector of prices chosen by the monopolist. In Section 4 we turn attention to the pricing stage (first stage of the game) and characterize the optimal strategy for the monopolist under three different settings: when the monopolist can perfectly price discriminate (Subsection 4.1), when the monopolist chooses a single uniform price for all consumers (Subsection 4.2) and finally when the monopolist can choose between two exogenously given prices, the full and the discounted (Subsection 4.3). Finally, we conclude in Section 5. Due to space constraints all proofs are omitted and can be found in [6].

2   Model

The society consists of a set $\mathcal{I} = \{1, \ldots, n\}$ of agents embedded in a social network represented by the adjacency matrix $\mathbf{G}$. The ij-th entry of $\mathbf{G}$, denoted by g_{ij}, represents the *strength* of the influence of agent j on i. We assume that $g_{ij} \in [0, 1]$ for all i, j and we normalize $g_{ii} = 0$ for all i. A monopolist introduces a *divisible* good in the market and chooses a vector $\mathbf{p}$ of prices from the set of allowable *pricing strategies* $\mathbf{P}$. In its full generality, $\mathbf{p} \in \mathbf{P}$ is simply a mapping from the set of agents to $\mathbb{R}^n$, i.e., $\mathbf{p} : \mathcal{I} \to \mathbb{R}^n$. In particular, $\mathbf{p}(i)$ or equivalently p_i is the price that the monopolist offers to agent i for one unit of the divisible good. Then, the agents choose the amount of the divisible good they will purchase at the announced price. Their utility is given by an expression of the following form:

$$u_i(x_i, \mathbf{x_{-i}}, p_i) = f_i(x_i) + x_i h_i(\mathbf{G}, \mathbf{x_{-i}}) - p_i x_i,$$

where $x_i \in [0, \infty)$ is the amount of the divisible good that agent i chooses to purchase. Function $f_i : [0, \infty) \to \mathbb{R}$ represents the utility that the agent obtains from the good, assuming that there are no network externalities, and $p_i x_i$ is the amount agent i is charged for its consumption. The function $h_i : [0, 1]^{n \times n} \times [0, \infty)^{n-1} \to [0, \infty)$ is used to capture the utility the agent obtains due to the positive network effect (note the explicit dependence on the network structure).

We next describe the two-stage *pricing-consumption* game, which models the interaction between the agents and the monopolist:

Stage 1 (Pricing) : The monopolist chooses the pricing strategy $\mathbf{p}$, so as to maximize profits, i.e., $\max_{\mathbf{p} \in \mathbf{P}} \sum_i p_i x_i - c x_i$, where c denotes the marginal cost of producing a unit of the good and x_i denotes the amount of the good agent i purchases in the second stage of the game.

Stage 2 (Consumption) : Agent i chooses to purchase x_i units of the good, so as to maximize her utility given the prices chosen by the monopolist and $\mathbf{x}_{-i}$, i.e.,

$$x_i \in \arg \max_{y_i \in [0, \infty)} u_i(y_i, \mathbf{x_{-i}}, p_i).$$

We are interested in the *subgame perfect* equilibria of the two-stage pricing-consumption game.

For a fixed vector of prices $\mathbf{p} = [p_i]_i$ chosen by the monopolist, the equilibria of the second stage game, referred to as the consumption equilibria, are defined as follows:

Definition 1 (Consumption Equilibrium). *For a given vector of prices $\mathbf{p}$, a vector $\mathbf{x}$ is a consumption equilibrium if, for all $i \in \mathcal{I}$,*

$$x_i \in \arg \max_{y_i \in [0,\infty)} u_i(y_i, \mathbf{x_{-i}}, p_i).$$

We denote the set of consumption equilibria at a given price vector $\mathbf{p}$ by $C[\mathbf{p}]$.

We begin our analysis by the second stage (the consumption subgame) and then discuss the optimal pricing policies for the monopolist given that agents purchase according to the consumption equilibrium of the subgame defined by the monopolist's choice of prices.

3 Consumption Equilibria

For the remainder of the paper, we assume that the payoff function of agent i takes the following quadratic form:

$$u_i(x_i, \mathbf{x_{-i}}, p_i) = a_i x_i - b_i x_i^2 + x_i \cdot \sum_{j \in \{1,\cdots,n\}} g_{ij} \cdot x_j - p_i x_i, \tag{1}$$

where the first two terms represent the utility agent i derives from consuming x_i units of the good irrespective of the consumption of her peers, the third term represents the (positive) network effect of her social group and finally the last term is the cost of usage. The quadratic form of the utility function is essential for keeping the analysis tractable, but also serves as a second-order approximation of the broader class of concave utility functions.

For a given vector of prices $\mathbf{p}$, we denote by $\mathcal{G} = \{\mathcal{I}, \{u_i\}_{i \in \mathcal{I}}, [0,\infty)_{i \in \mathcal{I}}\}$ the second stage game where the set of players is $\mathcal{I}$, each player $i \in \mathcal{I}$ chooses her strategy (consumption level) from the set $[0,\infty)$, and her the utility function, u_i has the form in (1). The following assumption ensures that in this game the optimal consumption level of each agent is bounded.

Assumption 1 *For all $i \in \mathcal{I}$, $b_i > \sum_{j \in \mathcal{I}} g_{ij}$.*

The necessity of Assumption 1 is evident from the following example: assume that the adjacency matrix, which represents the level of influence among agents, takes the following simple form: $g_{ij} = 1$ for all i, j such that $i \neq j$, i.e., $\mathbf{G}$ represents a complete graph with unit weights. Also, assume that $0 < b_i = b < n - 1$ and $0 < a_i = a$ for all $i \in \mathcal{I}$. It is now straightforward to see that given any vector of prices $\mathbf{p}$ and assuming that $x_i = x$ for all $i \in \mathcal{I}$, the payoffs of all agents go to infinity as $x \to \infty$. Thus, if Assumption 1 does not hold, in the consumption game, consumers may choose to unboundedly increase their usage irrespective of the vector of prices.

Next, we study the second stage of the game defined in Section 2 under Assumption 1, and we characterize the equilibria of the consumption game among the agents for vector of prices $\mathbf{p}$. In particular, we show that the equilibrium is unique and we provide a closed form expression for it. To express the results in a compact form, we define the vectors $\mathbf{x}, \mathbf{a}, \mathbf{p} \in \mathbb{R}^n$ such that $\mathbf{x} = [x_i]_i$, $\mathbf{a} = [a_i]_i$, $\mathbf{p} = [p_i]_i$. We also define matrix $\Lambda \in \mathbb{R}^{n \times n}$ as:

$$\Lambda_{i,j} = \begin{cases} 2b_i & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

Let $\beta_i(\mathbf{x_{-i}})$ denote the best response of agent i, when the rest of the agents choose consumption levels represented by the vector $\mathbf{x_{-i}}$. From (1) it follows that:

$$\beta_i(\mathbf{x_{-i}}) = \max \left\{ \frac{a_i - p_i}{2b_i} + \frac{1}{2b_i} \sum_{j \in \mathcal{I}} g_{ij} x_j, 0 \right\}. \tag{2}$$

Our first result shows that the equilibrium of the consumption game is unique for any price vector.

Theorem 1. *Under Assumption 1, the game* $\mathcal{G} = \{\mathcal{I}, \{u_i\}_{i \in \mathcal{I}}, [0, \infty)_{i \in \mathcal{I}}\}$ *has a unique equilibrium.*

Intuitively, Theorem 1 follows from the fact that increasing one's consumption incurs a positive externality on her peers, which further implies that the game involves strategic complementarities and therefore the equilibria are ordered. The proof exploits this monotonic ordering to show that the equilibrium is actually unique.

We conclude this section, by characterizing the unique equilibrium of $\mathcal{G}$. Suppose that $\mathbf{x}$ is this equilibrium, and $x_i > 0$ only for $i \in S$. Then, it follows that

$$x_i = \beta_i(\mathbf{x_{-i}}) = \frac{a_i - p_i}{2b_i} + \frac{1}{2b_i} \sum_{j \in \mathcal{I}} g_{ij} x_j = \frac{a_i - p_i}{2b_i} + \frac{1}{2b_i} \sum_{j \in S} g_{ij} x_j \tag{3}$$

for all $i \in S$. Denoting by $\mathbf{x}_S$ the vector of all x_i such that $i \in S$, and defining the vectors $\mathbf{a}_S$, $\mathbf{b}_S$, $\mathbf{p}_S$ and the matrices G_S, Λ_S similarly, equation (3) can be rewritten as

$$\Lambda_S \mathbf{x}_S = \mathbf{a}_S - \mathbf{p}_S + G_S \mathbf{x}_S. \tag{4}$$

Note that Assumption 1 holds for the graph restricted to the nodes in S, hence $I - \Lambda_S^{-1} G_S$ is invertible (see [6]). Therefore, (4) implies that

$$\mathbf{x}_S = (\Lambda_S - G_S)^{-1}(\mathbf{a}_S - \mathbf{p}_S). \tag{5}$$

Therefore, the unique equilibrium of the consumption game takes the following form:

$$\begin{aligned} \mathbf{x}_S &= (\Lambda_S - G_S)^{-1}(\mathbf{a}_S - \mathbf{p}_S), \\ \mathbf{x}_{\mathcal{I} - \mathbf{S}} &= \mathbf{0}, \end{aligned} \tag{6}$$

for some subset S of the set of agents $\mathcal{I}$. This characterization suggests that consumptions of players (weakly) decrease with the prices. The following lemma, which is used in the subsequent analysis, formalizes this fact.

Lemma 1. *Let* $\mathbf{x}(\mathbf{p})$ *denote the unique consumption equilibrium in the game where each player* $i \in \mathcal{I}$ *is offered the price* p_i. *Then,* $x_i(\mathbf{p})$ *is weakly decreasing in* $\mathbf{p}$ *for all* $i \in \mathcal{I}$, *i.e, if* $\hat{\mathbf{p}}_j \geq \mathbf{p}_j$ *for all* $j \in \mathcal{I}$ *then* $x_i(\hat{\mathbf{p}}) \leq x_i(\mathbf{p})$.

4 Optimal Pricing

In this section, we turn attention to the first stage of the game, where a monopolist sets the vector of prices. We distinguish between three different scenarios. In the first subsection, we assume that the monopolist can *perfectly price discriminate* the agents, i.e., there is no restriction imposed on the prices. In the second subsection, we consider the problem of choosing a single uniform price, while in the third we allow the monopolist to choose between two exogenous prices, p_L and p_H, for each consumer. In our terminology, in the first case $\mathbf{P} = \mathbb{R}^{|\mathcal{I}|}$, in the second $\mathbf{P} = \{(p, \cdots, p)\}$, for $p \in [0, \infty)$ and finally in the third $\mathbf{P} = \{p_L, p_H\}^{|\mathcal{I}|}$.

4.1 Perfect Price Discrimination

For the remainder of the paper, we make the following assumption, which ensures that, even in the absence of any network effects, the monopolist would find it optimal to charge individual prices low enough, so that all consumers purchase a positive amount of the good.

Assumption 2 *For all $i \in \mathcal{I}$, $a_i > c$.*

Given Assumption 2, we are now ready to state Theorem 2, that provides a characterization of the optimal prices. We denote the vector of all 1's by $\mathbf{1}$.

Theorem 2. *Under Assumptions 1 and 2, the optimal prices are given by*

$$\mathbf{p} = \mathbf{a} - (\Lambda - G)\left(\Lambda - \frac{G + G^T}{2}\right)^{-1}\frac{\mathbf{a} - c\mathbf{1}}{2}. \tag{7}$$

The following corollary is an immediate consequence of Theorem 2.

Corollary 1. *Let Assumptions 1 and 2 hold. Moreover, assume that the interaction matrix G is symmetric. Then, the optimal prices satisfy*

$$\mathbf{p} = \frac{\mathbf{a} + c\mathbf{1}}{2},$$

i.e., the optimal prices do not depend on the network structure.

This result implies that when players affect each other in the same way, i.e., when the interaction matrix G is symmetric, then the graph topology has no effect on the optimal prices.

To better illustrate the effect of the network structure on prices we next consider a special setting, in which agents are symmetric in a sense defined precisely below and they differ only in terms of their network position.

Assumption 3 *Players are symmetric, i.e., $a_i = a_0$, $b_i = b_0$ for all $i \in \mathcal{I}$.*

We next provide the definition of Bonacich Centrality (see also [2]), and using it obtain an alternative characterization of the optimal prices.

Definition 2 (Bonacich Centrality). *For a network with (weighted) adjacency matrix G and scalar α, the Bonacich centrality vector of parameter α is given by $\mathcal{K}(G, \alpha) = (I - \alpha G)^{-1}\mathbf{1}$ provided that $(I - \alpha G)^{-1}$ is well defined and nonnegative.*

Theorem 3. *Under Assumptions 1, 2 and 3, the vector of optimal prices is given by*

$$\mathbf{p} = \frac{a_0 + c}{2}\mathbf{1} + \frac{a_0 - c}{8b_0}G\mathcal{K}\left(\frac{G + G^T}{2}, \frac{1}{2b_0}\right) - \frac{a_0 - c}{8b_0}G^T\mathcal{K}\left(\frac{G + G^T}{2}, \frac{1}{2b_0}\right).$$

The network $\frac{G+G^T}{2}$ is the average interaction network, and it represents the average interaction between pairs of agents in network G. Intuitively, the centrality $\mathcal{K}\left(\frac{G+G^T}{2}, \frac{1}{2b_0}\right)$ measures how "central" each node is with respect to the average interaction network.

The optimal prices in Theorem 3 have three components. The first component can be thought of as a nominal price, which is charged to all agents irrespective of the network structure. The second term is a markup that the monopolist can impose on the price of consumer i due to the utility the latter derives from her peers. Finally, the third component can be seen as a discount term, which is offered to a consumer, since increasing her consumption increases the consumption

level of her peers. Theorem 3 suggests that it is optimal to give each agent a markup proportional to the utility she derives from the central agents. In contrast, prices offered to the agents should be discounted proportionally to their influence on central agents. Therefore, it follows that the nodes which pay the most favorable prices are the ones, that *influence* highly central nodes.

Note that if Assumption 3 fails, then Theorem 3 can be modified to relate the optimal prices to centrality measures in the underlying graph. In particular, the price structure is still as given in Theorem 3, but when the parameters $\{a_i\}$ and $\{b_i\}$ are not identical, the discount and markup terms are proportional to a weighted version of the Bonacich centrality measure, defined below.

Definition 3 (Weighted Bonacich Centrality). *For a network with (weighted) adjacency matrix G, diagonal matrix D and weight vector $\mathbf{v}$, the weighted Bonacich centrality vector is given by $\tilde{\mathcal{K}}(G, D, \mathbf{v}) = (I - GD)^{-1}\mathbf{v}$ provided that $(I - GD)^{-1}$ is well defined and nonnegative.*

We next characterize the optimal prices in terms of the weighted Bonacich centrality measure.

Theorem 4. *Under Assumptions 1 and 2 the vector of optimal prices is given by*

$$\mathbf{p} = \frac{\mathbf{a} + c\mathbf{1}}{2} + G\Lambda^{-1}\tilde{\mathcal{K}}\left(\tilde{G}, \Lambda^{-1}, \tilde{\mathbf{v}}\right) - G^T\Lambda^{-1}\tilde{\mathcal{K}}\left(\tilde{G}, \Lambda^{-1}, \tilde{\mathbf{v}}\right),$$

where $\tilde{G} = \frac{G + G^T}{2}$ and $\tilde{\mathbf{v}} = \frac{\mathbf{a} - c\mathbf{1}}{2}$.

4.2 Choosing a Single Uniform Price

In this subsection we characterize the equilibria of the pricing-consumption game, when the monopolist can only set a single uniform price, i.e., $p_i = p_0$ for all i. Then, for any fixed p_0, the payoff function of agent i is given by

$$u_i(x_i, \mathbf{x_{-i}}, p_i) = a_i x_i - b_i x_i^2 + x_i \cdot \sum_{j \in \{1, \cdots, n\}} g_{ij} \cdot x_j - p_i x_i,$$

and the payoff function for the monopolist is given by

$$\max_{p_0 \in [0, \infty)} (p_0 - c) \sum_i x_i$$
$$s.t. \qquad \mathbf{x} \in C[\mathbf{p_0}],$$

where $\mathbf{p_0} = (p_0, \cdots, p_0)$. Note that Theorem 1 implies that even when the monopolist offers a single price, the consumption game has a unique equilibrium point. Next lemma states that the consumption of each agent decreases monotonically in the price.

Lemma 2. *Let $\mathbf{x}(\mathbf{p_0})$ denote the unique equilibrium in the game where $p_i = p_0$ for all i. Then, $x_i(\mathbf{p_0})$ is weakly decreasing in p_0 for all $i \in \mathcal{I}$ and strictly decreasing for all i such that $x_i(\mathbf{p_0}) > 0$.*

Next, we introduce the notion of the centrality gain.

Definition 4 (Centrality Gain). *In a network with (weighted) adjacency matrix G, for any diagonal matrix D and weight vector $\mathbf{v}$, the centrality gain of agent i is defined as*

$$H_i(G, D, \mathbf{v}) = \frac{\tilde{\mathcal{K}}_i(G, D, \mathbf{v})}{\tilde{\mathcal{K}}_i(G, D, \mathbf{1})}.$$

The following theorem provides a characterization of the consumption vector at equilibrium as a function of the single uniform price p.

Theorem 5. *Consider game $\bar{\mathcal{G}} = \{\mathcal{I}, \{u_i\}_{i \in \mathcal{I}}, [0, \infty)_{i \in \mathcal{I}}\}$, and define*

$$D_1 = \arg\min_{i \in \mathcal{I}} H_i\left(G, \Lambda^{-1}, \mathbf{a}\right) \quad and \quad p_1 = \min_{i \in \mathcal{I}} H_i\left(G, \Lambda^{-1}, \mathbf{a}\right).$$

Moreover, let $I_k = \mathcal{I} - \cup_{i=1}^{k} D_i$ and define

$$D_k = \arg\min_{i \in I_k} H_i\left(G_{I_k}, \Lambda_{I_k}^{-1}, \mathbf{a}_{I_k}\right) \quad and \quad p_k = \min_{i \in I_k} H_i\left(G_{I_k}, \Lambda_{I_k}^{-1}, \mathbf{a}_{I_k}\right),$$

for $k \in \{2, 3 \dots n\}$. Then,

(1) p_k strictly increases in k.
(2) Given a p such that $p < p_1$, all agents purchase a positive amount of the good, i.e., $x_i(p) > 0$ for all $i \in \mathcal{I}$, where $\mathbf{x}(\mathbf{p})$ denotes the unique consumption equilibrium at price p. If $k \geq 1$, and p is such that $p_k \leq p \leq p_{k+1}$, then $x_i(p) > 0$ if and only if $i \in I_k$. Moreover, the corresponding consumption levels are given as in (6), where $S = I_k$.

Theorem 5 also suggests a polynomial time algorithm for computing the optimal uniform price p_{opt}. Intuitively, the algorithm sequentially removes consumers with the lowest centrality gain and computes the optimal price for the remaining consumers under the assumption that the price is low enough so that only these agents purchase a positive amount of the good at the associated consumption equilibrium. In particular, using Theorem 5, it is possible to identify the set of agents who purchase a positive amount of the good for price ranges $[p_k, p_{k+1}]$, $k \in \{1, \dots\}$. Observe that given a set of players, who purchase a positive amount of the good, the equilibrium consumption levels can be obtained in closed form as a linear function of the offered price, and, thus, the profit function of the monopolist takes a quadratic form in the price. It follows that for each price range, the maximum profit can be found by solving a quadratic optimization problem. Thus, Theorem 5 suggests Algorithm 1 for finding the optimal single uniform price p_{opt}.

Algorithm 1 : Compute the optimal single uniform price p_{opt}

STEP 1. Preliminaries:

- Initialize the set of *active* agents: $S := \mathcal{I}$.
- Initialize $k = 1$ and $p_0 = 0$, $p_1 = \min_{i \in \mathcal{I}} H_i(G_{\mathcal{I}}, \Lambda_{\mathcal{I}}^{-1}, \mathbf{a}_{\mathcal{I}})$
- Initialize the monopolist's revenues with $Re_{opt} = 0$ and $p_{opt} = 0$.

STEP 2.

- Let $\hat{p} = \frac{\mathbf{1}^T (\Lambda_S - G_S)^{-1} \mathbf{a}_S - c\mathbf{1}^T (\Lambda_S - G_S)^{-1} \mathbf{1}}{\mathbf{1}^T \left((\Lambda_S - G_S)^{-1} + (\Lambda_S - G_S^T)^{-1}\right)\mathbf{1}}$
- **IF** $\hat{p} \geq p_k$, let $p = p_k$.
  ELSE IF $\hat{p} \leq p_{k-1}$, let $p = p_{k-1}$ **ELSE** $p = \hat{p}$.
- $Re = (p - c)\mathbf{1}^T \cdot (\Lambda_S - G_S)^{-1}(\mathbf{a}_S - p\mathbf{1})$.
- **IF** $Re > Re_{opt}$ **THEN** $Re_{opt} = Re$ and $p_{opt} = p$.
- $D = \arg\min_{i \in S} H_i(G_S, \Lambda_S^{-1}, \mathbf{a}_S)$ and $S := S - D$.
- Increase k by 1 and let $p_k = \min_{i \in S} H_i(G_S, \Lambda_S^{-1}, \mathbf{a}_S)$.
- Return to **STEP 2** if $S \neq \emptyset$ **ELSE Output** p_{opt}.

The algorithm solves a series of subproblems, where the monopolist is constrained to choose a price p in a given interval $[p_k, p_{k+1}]$ with appropriately chosen endpoints. In particular, from Theorem 5, we can choose those endpoints, so as to ensure that only a particular set S of agents purchase a positive amount of the good. In this case, the consumption at price p is given by

$(\Lambda_S - G_S)^{-1}(\mathbf{a}_S - p\mathbf{1})$ and the profit of the monopolist is equal to $(p-c)\mathbf{1}^T \cdot (\Lambda_S - G_S)^{-1}(\mathbf{a}_S - p\mathbf{1})$. The maximum of this profit function is achieved at $\hat{p} = \frac{\mathbf{1}^T (\Lambda_S - G_S)^{-1} \mathbf{a}_S - c\mathbf{1}^T (\Lambda_S - G_S)^{-1}\mathbf{1}}{\mathbf{1}^T \left((\Lambda_S - G_S)^{-1} + (\Lambda_S - G_S^T)^{-1}\right)\mathbf{1}}$, as can be seen from the first order optimality conditions. Then, the overall optimal price is found by comparing the monopolist's profits achieved at the optimal solutions of the constrained subproblems. The complexity of the algorithm is $O(n^4)$, since there are at most n such subproblems (again from Theorem 5) and each such subproblem simply involves a matrix inversion ($O(n^3)$) in computing the centrality gain and the maximum achievable profit.

4.3 The Case of Two Prices: Full and Discounted

In this subsection, we assume that the monopolist can choose to offer the good in one of two prices, p_L and p_H ($p_L < p_H$) that are exogenously defined. For clarity of exposition we call p_L and p_H the discounted and the full price respectively. The question that remains to be studied is to which agents should the monopolist offer the discounted price, so as to maximize her revenues. We state the following assumption that significantly simplifies the exposition.

Assumption 4 *The exogenous prices p_L, p_H are such that $p_L, p_H < \min_{i \in \mathcal{I}} a_i$.*

Note that under Assumption 4, Equation (2) implies that all agents purchase a positive amount of the good at equilibrium, regardless of the actions of their peers. As shown previously, the vector of consumption levels satisfies $\mathbf{x} = \Lambda^{-1}(\mathbf{a} - \mathbf{p} + G\mathbf{x})$, and hence $\mathbf{x} = (\Lambda - G)^{-1}(\mathbf{a} - \mathbf{p})$. An instance of the monopolist's problem can now be written as:

$$(OPT) \quad \max \quad (\mathbf{p} - c\mathbf{1})^T (\Lambda - G)^{-1}(\mathbf{a} - \mathbf{p})$$
$$st. \quad p_i \in \{p_L, p_H\} \quad \text{for all } i \in \mathcal{I},$$

where $\Lambda \succ 0$ is a diagonal matrix, G is such that $G \geq 0$, $diag(G) = 0$ and Assumption 1 holds.

Let $p_N \triangleq \frac{p_H + p_L}{2}$, $\delta \triangleq p_H - p_N$, $\hat{\mathbf{a}} \triangleq \mathbf{a} - p_N\mathbf{1}$ and $\hat{c} \triangleq p_N - c \geq \delta$. Using these variables, and noting that any feasible price allocation can be expressed as $\mathbf{p} = \delta\mathbf{y} + p_N\mathbf{1}$, where $y_i \in \{-1, 1\}$, OPT can alternatively be expressed as

$$\max \quad (\delta\mathbf{y} + \hat{c}\mathbf{1})^T (\Lambda - G)^{-1}(\hat{\mathbf{a}} - \delta\mathbf{y})$$
$$s.t. \quad y_i \in \{-1, 1\} \quad \text{for all } i \in \mathcal{I}. \tag{8}$$

We next show that OPT is NP-hard, and provide an algorithm that achieves an approximately optimal solution. To obtain our results, we relate the alternative formulation of OPT in (8) to the MAX-CUT problem (see [11, 12]).

Theorem 6. *Let Assumptions 1, 2 and 4 hold. Then, the monopolist's optimal pricing problem, i.e., problem OPT, is NP-hard.*

Theorem 7. *Let Assumptions 1 and 4 hold and W_{OPT} denote the optimal profits for the monopolist, i.e., W_{OPT} is the optimal value for problem OPT. Then, there exists a randomized polynomial time algorithm, that outputs a solution with objective value W_{ALG} such that $E[W_{ALG}] + m > 0.878(W_{OPT} + m)$, where*

$$m = \delta^2 \mathbf{1}^T A \mathbf{1} + \delta \mathbf{1}^T \left| A\hat{\mathbf{a}} - A^T \hat{c}\mathbf{1} \right| - \hat{c}\mathbf{1}^T A\hat{\mathbf{a}} - 2\delta^2 Trace(A),$$

and $A = (\Lambda - G)^{-1}$.

In the remainder of the section, we provide a characterization of the optimal prices in OPT. In particular, we argue that the pricing problem faced by the monopolist is equivalent to finding the cut with maximum weight in an appropriately defined weighted graph. For simplicity, assume that $b_i = b_0$ for all i and $\left((\Lambda - G)^{-1}\hat{\mathbf{a}} - \hat{c}(\Lambda - G)^{-T}\mathbf{1}\right) = 0$ (which holds, for instance when $\hat{\mathbf{a}} = \hat{c}\mathbf{1}$, or equivalently $\mathbf{a} - p_N\mathbf{1} = (p_N - c)\mathbf{1}$, and $G = G^T$). Observe that in this case, the alternative formulation of the profit maximization problem in (8), can equivalently be written as (after adding a constant to the objective function, and scaling):

$$\begin{aligned} \max \quad & \alpha - \mathbf{y}(\Lambda - G)^{-1}\mathbf{y} \\ s.t. \quad & y_i \in \{-1, 1\} \quad \text{for all } i \in \mathcal{I}, \end{aligned} \tag{9}$$

where $\alpha = \sum_{ij}(\Lambda - G)_{ij}^{-1}$. It can be seen that this optimization problem is equivalent to an instance of the MAX-CUT problem, where the cut weights are given by the off diagonal entries of $(\Lambda - G)^{-1}$ (see [11, 12]). On the other hand observe that $(\Lambda - G)^{-1}\mathbf{1} = \frac{1}{2b_0}(I - \frac{1}{2b_0}G)^{-1}\mathbf{1}$, hence, the ith row sum of the entries of the matrix $(\Lambda - G)^{-1}$ is proportional to the centrality of the ith agent in the network. Consequently, the (i,j)th entry of the matrix $(\Lambda - G)^{-1}$, gives a measure of how much the edge between i and j contributes to the centrality of node i. Since the MAX-CUT interpretation suggests that the optimal solution of the pricing problem is achieved by maximizing the cut weight, it follows that the optimal solution of this problem price differentiates the agents who affect the centrality of each other significantly.

5   Conclusions

The paper studies a stylized model of pricing of divisible goods (services) over social networks, when consumers' actions are influenced by the choices of their peers. We provide a concrete characterization of the optimal scheme for a monopolist under different restrictions on the set of allowable pricing policies when consumers behave according to the unique Nash equilibrium profile of the corresponding game. We consider a setting of static pricing: the monopolist first sets prices and then the consumers choose their usage levels. Moreover, the game we define is essentially of complete information, since we assume that both the monopolist, as well as the consumers, know the network structure and the utility functions of the population. Extending our analysis by introducing incomplete information is an interesting direction for future research. Concretely, consider a monopolist that introduces a new product of unknown quality to a market. Agents benefit the monopolist in two ways when purchasing the product; directly by increasing her revenues, and indirectly by generating information about the product's quality and making it more attractive to the rest of the consumer pool. What is the optimal (dynamic) pricing strategy for the monopolist?

Finally, note that in the current setup we consider a single seller (monopolist), so as to focus on explicitly characterizing the optimal prices as a function of the network structure. A natural departure from this model is studying a competitive environment. The simplest such setting would involve a small number of sellers offering a perfectly substitutable good to the market. Then, pricing may be even more aggressive than in the monopolistic environment: sellers may offer even larger discounts to "central" consumers, so as to subsequently exploit the effect of their decisions to the rest of the network. Potentially one could relate the *intensity of competition* with the network structure. In particular, one would expect the competition to be less fierce when the network consists of disjoint large subnetworks, since then sellers would segment the market at equilibrium and exercise monopoly power in their respective segments.

12

References

1. C Ballester, A Calvo-Armengol, and Y Zenou. Who's who in networks. wanted: the key player. *Econometrica*, 74(5):1403–1417, 2006.
2. P Bonacich. Power and centrality: A family of measures. *The American Journal of Sociology*, 92(5):1170–1182, 1987.
3. Y Bramoullé and R Kranton. Public goods in networks. *Journal of Economic Theory*, 135:478–494, 2007.
4. Y Bramoullé, R Kranton, and M D'Amours. Strategic interaction and networks. *Working paper*, 2009.
5. A Campbell. Tell your friends! word of mouth and percolation in social networks. *Working paper*, 2009.
6. O Candogan, K Bimpikis, and A Ozdaglar. Optimal pricing in the presence of local network effects. *Working paper*, 2010.
7. J Corbo, A Calvo-Armengol, and D C Parkes. The importance of network topology in local contribution games. *Proceedings of the 3rd international Workshop on Internet and Network Economics*, 2007.
8. J Farrell and G Saloner. Standardization, compatibility, and innovation. *RAND Journal of Economics*, 16(1):70–83, 1985.
9. A Galeotti and S Goyal. Influencing the influencers: a theory of strategic diffusion. *RAND Journal of Economics*, 40(3):509–532, 2009.
10. A Galeotti, S Goyal, M Jackson, F Vega-Redondo, and L Yariv. Network games. *Review of Economic Studies*, 77(1):218–244, 2010.
11. M R Garey and D S Johnson. *Computers and intractability. A guide to the theory of NP-completeness. A Series of Books in the Mathematical Sciences*. WH Freeman and Company, San Francisco, CA, 1979.
12. M Goemans and D Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995.
13. J Hartline, V Mirrokni, and M Sundararajan. Optimal marketing strategies over social networks. *Proceedings of the 17th international conference on World Wide Web*, 2008.
14. R Johari and S Kumar. Congestible services and network effects. *Working paper*, 2010.
15. M Katz and C Shapiro. Technology adoption in the presence of network externalities. *Journal of Political Economy*, 94(4):822–841, 1986.
16. D Kempe, J Kleinberg, and É Tardos. Maximizing the spread of influence through a social network. *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2003.
17. J MacKie-Mason and H Varian. Pricing congestible network resources. *IEEE Journal of Selected Areas in Communications*, 1995.
18. A Sundararajan. Local network effects and complex network structure. *The BE Journal of Theoretical Economics*, 7(1), 2007.
19. D M Topkis. *Supermodularity and complementarity*. Princeton University Press, Princeton, NJ, 1998.