RC25575 (WAT1511-038) November 12, 2015
Mathematics

IBM Research Report

Designing Price Incentives in a Network with Social Interactions

Maxime Cohen
MIT
Cambridge, MA 02139

Pavithra Harsha
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598 USA

Designing Price Incentives in a Network with Social Interactions

Maxime Cohen

Operations Research Center, MIT, Cambridge, MA 02139, maxcohen@mit.edu

Pavithra Harsha

IBM Research Thomas J. Watson Research Center, Yorktown Heights, NY, pharsha@us.ibm.com

The recent ubiquity of social networks allows firms to collect vast amount of data about their customers and including their social interactions with other customers. We consider a setting where a monopolist sells an indivisible good to consumers that are embedded in a social network. Assuming complete information about the interactions between consumers, we model the optimal pricing problem of the firm as a two-stage game. First, the firm designs prices to maximize his profits and second the consumers choose whether to purchase the item or not at the given prices so as to optimize their own payoffs. Assuming *positive* network externalities, we show the existence of a pure strategy Nash equilibrium for the second stage game. Using duality theory and integer programming techniques, we reformulate the problem into a single stage linear mixed-integer program (MIP). This formulation can be used by any firm as an operational pricing tool as it can easily incorporate business rules on prices and constraints on network segmentation. We derive efficient and scalable ways of optimally solving the MIP using its linear programming relaxation for two different pricing strategies chosen by the seller: discriminative pricing where the monopolist offers prices that may differ by agent and the uniform pricing where the prices are the same for all agents. Further we extend our model and results to the case when the seller offers incentives in addition to prices to solicit actions such as buyer reviews to ensure influence on other agents. Finally, we present computational insights by comparing the various pricing strategies and highlighting the benefits of incorporating social interactions. In particular, we provide instances where it is beneficial for the seller to earn negative profit on an influential agent in order to extract significant positive profits on others.

Key words: Pricing, social networks, network externalities, influence, integer program, game theory

1. Introduction

The recent ubiquity of social networks have revolutionized the way people interact and influence each other. The overwhelming success of social networking platforms such as Facebook and Twitter allows firms to collect unprecedented volumes of data about their customers, their buying behavior including their social interactions with other customers. The challenge that confronts every firm, from big to small, is how to process this data and turn it into actionable policies so as to improve their competitive advantage. In this paper, we focus on the design of effective pricing strategies to improve profitability of a monopolistic firm that sells indivisible goods or services to agents embedded in a social network.

Word-of-mouth communication between agents has always been an effective marketing tool for several businesses. In recent times, word of mouth communication is just as likely to arise from social networks or smart phone applications as from a neighbor across the fence. According to a recent study by SproutSocial, 74% of consumers rely on social networks to guide their purchase decisions. Consultants at The Conversation Group report that 65% of consumers who receive a recommendation from a contact on their social media sites have purchased a product that was recommended to them. In particular, personalized referrals from friends and family have been more effective in encouraging such purchases. Finally, nearly 93% of social media users have either made or received a recommendation for a product or service. Academic research on consumer behavior shows that consumers' purchasing decisions and product evaluations are influenced by their reference groups (see e.g. Iyengar et al. (2011)). All this clearly indicates that people influence their connections. They not only guide their purchasing behavior but more importantly alter their willingness-to-pay for various items. For example, when an individual buys a product and posts a positive review on his Facebook page, he not only influences his peers to purchase the same item or service but also increases their valuation for the item.

An important feature of the products or services we consider in this paper is that they exhibit local *positive externalities*. This means people *positively* influence each other willingness-to-pay for an item which gets more valuable to a person if many more of his friends buy it. Examples of products with such effects include smart phones, tablets, certain fashion items and cell phone plans subscriptions just to name a few. Such positive externalities are even more significant when new generation of products are introduced in the market and people use social networks as a way to accelerate their friends' awareness about the item.

It is common practice that a very small number of highly influential people (e.g., certified bloggers on certain websites) receive the item nearly for free to increase the awareness of the remaining population. Mark W. Schaefer, the author of: "Return on Influence" reports that: "For the first time, companies large and small can find these passionate influencers (using social networks), connect to them and turn them into brand advocates". Therefore, it can be very valuable for firms to localize these influential agents. These days many online sellers let consumers sign in with their Facebook account. Consequently, they have access to their personal information such as age, gender, geographical location, number of friends and more importantly, their network. Various sellers even build a Facebook page to advertise their firm through social platforms. For example, the large US corporation Macy's has more than 11.5M of fans that liked their Facebook page (June 2013). These fans can then claim offers using the social platform that allows for a certain degree of publicity and thus directly influence their friends about purchasing. This interaction between the seller and the fan club not only allows the seller to keep their fan club engaged and interested in the brand, but also enables the seller to localize the influencers, provide them with personalized prices or incentives, for example, and turn them into influencers, thereby increasing overall profitability.

In this paper, we consider a setting where a monopolist sells an indivisible good to consumers embedded in a social network. Our goal is to develop a model that incorporates the positive social externalities between potential buyers and design efficient algorithms to compute the optimal prices for the item to maximize the seller's profitability. In particular, the method we propose aims to provide a systematic and automated way of finding the prices to offer to the agents based on their influence level so as to maximize the total profit of the seller. We view the interaction between seller and buyers as a two-stage game where the seller first offers prices and the agents then choose whether to purchase the item or not at the offered prices. We capture the utility of an agent using a linear additive valuation model wherein the total value for an agent in owning the item as the sum of the agent's own value as well as the (positive) influence from the agent's friends who own the item. Our first result shows the existence of a pure strategy Nash equilibrium for the second stage game given any vector of prices. Using duality theory, we derive equilibrium constraints and reformulate the two-stage problem faced by the seller into a single stage non-convex integer programming problem. We then transform it into an equivalent mixed-integer program (MIP) using reformulation techniques from integer programming. This resulting MIP can be viewed as an operational pricing tool as any firm can easily incorporate business rules on prices and constraints on network segmentation. For example, the seller can identify bounded and tiered prices for the members in its fan club based on their loyalty class such as platinum or gold.

Because the pricing model is formulated as a MIP, on one hand it is not scalable to large networks but on the other hand it lends itself into the traditional optimization framework where one can explore and exploit various advancements in the field of optimization. In this paper, we derive efficient and scalable methods to solve the MIP to optimality for two distinct pricing strategies using linear programming (LP) relaxation of the MIP. We consider the discriminative and the uniform pricing strategies and present a solution method that is efficient (polynomial in the number of agents) and scalable to large networks.

In the above model, we assume that anyone who buys the items influences their peers as long as they purchase the item. This assumption is not realistic in many practical settings. Indeed, after purchasing an item, it is sometimes not entirely natural to influence friends about the product unless one takes some effort to do so. This for example could be writing a review, endorsing the item on their wall, blogging about the item or at the very least announcing that they have purchased the item. In practice, it is popular for firms to offer cash rewards for recommending or referring friends. For example, Groupon, a popular deal-of-the-day website that features discounted gift certificates usable at local or national companies offers $10 for referring a friend. Similarly, American Express provides a bonus for each approved referral. Interestingly, online booking service sites like HolidayCheck.com provides incentives for people to just share their experience on the booking site.

We therefore extend our models to the case when the seller can design both prices and incentives so as to maximize profitability while aiming to guarantee influence by soliciting influence actions in return for the incentives. In our extended model, the seller ensures the influence among the agents by offering a price and a discount (incentive) to each buyer. The buyer can then decide between several options: (i) not buy the item, (ii) buy the item at the full price or (iii) buy the item and claim a portion of the discount in return for influence actions specified by the seller. In this last alternative, the agent receives a small discount in exchange for a simple action such as liking the product and a more significant discount by taking a time-consuming action such as writing a detailed review thereby influencing friends by varying degrees in the respective cases. The utility model of an agent requires an additional input parameter in this setting. We refer to this as the *influence cost* for an agent which is the utility value of the effort it takes for an agent to influence his neighbors. This parameter can be estimated from historical data using the intensity of online activity for past purchases, the number of reviews written, the corresponding incentives needed and other data from cookies. With this more general pricing setting, the seller can ensure the influence among the agents so that the network externalities effects are guaranteed to occur. Interestingly, the methods and results we develop for our previous model extend entirely to this model.

In summary, the main contributions of the paper are as follows:

1. We formulate the optimal pricing problem for an indivisible item as a two-stage game between seller and the buyers so as to maximize the seller's profitability and show the existence of a pure strategy equilibrium for the second stage game.
2. We reformulate the game as a single stage MIP. This MIP can be viewed as an operational pricing model as a firm can easily incorporate business rules on prices and constraints on network segmentation.
3. We develop efficient and scalable LP based methods that are polynomial in the number of agents to solve the MIP optimally for discriminative and uniform pricing strategies.
4. We extend our pricing model and the corresponding results to optimally design both prices and incentives so as to guarantee influence among agents with incentives that solicit influence actions while maximizing the profitability of the seller.

5. We present computational results on example networks to draw qualitative insights about the incorporation of social interactions, different pricing strategies and the richer model with incentives, pricing influencers and pricing with different network topologies.

Literature review

Models that incorporate local network externalities find their origins in the works of Farrell and Saloner (1985) and Katz and Shapiro (1985). These early papers assume that consumers are affected by the global consumption of all other players. In other words, the network effects are of global nature, i.e., the utility of a consumer depends directly on the behavior of the entire set of agents in the network. In our model, consumers directly interact only with a subset of agents, also known as their neighbors. Although interactions are of a local nature, the utility of each player may still depend on the global structure of the network, because each agent potentially interacts indirectly with a much larger set of agents than just his neighbors.

Models of local network externalities which explicitly take into account the network structure have been proposed in several papers including Ballester et al. (2006), Banerji and Dutta (2009) and Sundararajan (2007). Ballester et al. (2006) propose a model in which individuals located in a network choose actions (criminal activities) which affect the payoffs of other individuals in the network using linear influence models which we adopt as well. Banerji and Dutta (2009) study a setting where firms sell to consumers located on a network with local adoption externalities. Their characterize networks which can sustain different technologies in equilibrium and show that even if rival firms engage in Bertrand competition, this form of network externalities permits strong market segmentation. The paper by Sundararajan (2007) presents a model of local network effects in which agents connected in a social network value the adoption of a product by a heterogeneous subset of other agents in their neighborhood and have incomplete information about the adoption complementarities between all other agents. Another related area of research is network games. Our second stage problem takes the form of a network game where agents in a network interact with each other. A recent series of papers that study network games include Galeotti et al. (2010) and the references therein.

Several recent papers explicitly include the interactions among agents in social networks to study social network's effects on various marketing problems. The first among these are the works on influence maximization by Domingos and Richardson (2001) and Kempe et al. (2003) which aimed to identify influential agents in a network. Several recent papers such as those by Hartline et al. (2008), Akhlaghpour et al. (2010) and Arthur et al. (2009) extend these works to study optimal pricing strategies in networks. Hartline et al. (2008) focus on viral marketing strategies for revenue maximization where the agents are offered the product in a sequential manner and show simple two-price strategies (the *Influence-and-Exploit* strategy where the seller chooses a set of consumers who get the product for free and use the optimal myopic pricing for the rest) performs very well relative to the optimal strategy which is NP-hard. Akhlaghpour et al. (2010) extend this approach to a multi-stage model where the seller sets different prices for each stage. Arthur et al. (2009) allow buyers to buy the product with a certain probability if the product is recommended by their friends who purchased the item. The main difference in the approaches in these papers and ours stems from the timing of purchasing decisions. These papers consider sequential purchases where myopic consumers base their consumption decisions on the number of consumers who have already bought the product. In our paper, we consider a simultaneous purchasing decision for all the agents in the network, who are fully rational.

In the literature and in our paper, pricing with simultaneous purchasing decisions is set up as a two-stage game, where the seller designs the prices in the first stage and agents respond by playing their purchasing decisions. Agents rational behavior in this case is captured by the Nash equilibrium (or Bayesian Nash equilibrium if the information is incomplete). Three papers in this context closely related to our work are Candogan et al. (2012), Bloch and Querou (2013) and Chen et al. (2011). Candogan et al. (2012) study optimal pricing strategies for a divisible good with linear social influence functions so as to maximize profits of the seller who has complete information of the network. They provide efficient algorithms to compute fully discriminative prices as well as the uniform optimal price and show that the problem is NP-hard when the monopolist is restricted to two pre-specified prices. Bloch and Querou (2013) and Chen et al. (2011) study the optimal pricing problem of an indivisible item with linear utility functions under incomplete information. Bloch and Querou (2013) study externalities resulting either from local network interactions or from prices and distinguish between a single global monopoly and several local monopolies. Chen et al. (2011) assume a uniform prior and propose an efficient uniform pricing algorithm for revenue maximization but show that the general discriminative pricing problem is NP-hard. Our work is in the similar light of the three aforementioned papers for the case of an indivisible item under complete information. But the model and the techniques required to address our setting are quite different from these papers. In particular, the second-stage equilibrium unlike the private valuation case or the divisible good case, can be characterized only with a system of equilibrium constraints that happen to be highly non-linear and non-convex with integer variables and this entirely alters the simple quadratic form of optimal pricing problem that could be solved in closed form as observed by Candogan et al. (2012), Bloch and Querou (2013). In fact, in our setting the optimal pricing problem is cast as a mathematical program with equilibrium constraints (MPEC) (see Luo et al. (1996)). We refer the reader to the books by Nemhauser and Wolsey (1988) and Bertsimas and

Weismantel (2005) for the integer programming reformulation techniques that we use in this paper to address the non-convexities and arrive at a MIP.

Because of the modeling flexibility of our approach, to the best of our knowledge, we are the first paper that provides an explicit optimization formulation for the pricing problem that can incorporate business rules on prices and constraints on market segmentation. Moreover, we are also able to extend the model and results of the paper to a practical setting where potential buyers are explicitly given incentives, potentially different for each agent, and a choice to influence their neighbors in addition to a price.

Finally, another recent paper on incorporating the effects of social network influence but unrelated to pricing is by Gunnec and Raghavan (2012). The latter investigates social network influence in the context of product design, in particular, in the share-of-choice problem, and construct a genetic algorithm to solve the problem.

Structure of the paper: The rest of the paper is organized as follows. In Section 2, we describe the model and our assumptions as well as the dynamics of the two-stage game. In Section 3, we show the existence of a pure strategy Nash equilibrium for the second-stage purchasing game. We use duality theory to formulate the problem as a MIP in Section 4. In Section 5, we derive efficient algorithms to solve the MIP for the discriminative and uniform pricing strategies. In Section 6, we extend our model to the case when the seller can design both price and incentives to guarantee the influence among agents in the network. In Section 7, we present computational experiments to draw some qualitative insights. Finally, we present our conclusions in Section 8. The proofs of the different propositions and theorems are relegated to the Appendix.

2. Model

Consider a monopolistic firm selling an indivisible product to N agents denoted by the set $\mathcal{I} = \{1, ..., N\}$ embedded in a social network. We denote the *value interaction* matrix for this product by G, where the element g_{ji} represents the marginal increase in value that agent i obtains by owning the product when agent j owns also the product. In particular, g_{ii} is the marginal value that agent i derives from himself by owning the product. If agent j does not influence agent i, then $g_{ji} = 0$.

Assumption 1. We make the following assumptions about G:

a. Only positive influences occur among the agents in the network, i.e., that $g_{ji} \geq 0$ for all i, j.
b. The firm and the agents have perfect knowledge of the network externalities i.e., everyone knows G.

Let the vector $\mathbf{p} \in \mathbf{P}$ denote the prices offered by the firm for the indivisible product. In particular, $p_i \in \mathbb{R}$ is the i^{th} element of the vector $\mathbf{p}$ and represents the price offered to agent i by the seller. Here,

$\mathbf{P}$ is assumed to be a polyhedral set that represents the feasible pricing strategies of the firm, which possibly includes several business constraints on prices and on network segmentation. For example, the firm can adopt a discriminative pricing strategy where each agent may potentially receive a different price. In this case, $\mathbf{P} = \mathbb{R}^N$. In addition, one can restrict the values of these prices to lie between p_L and $p_U (\geq p_L)$, i.e., $\mathbf{P} = \{\mathbf{p} \in \mathbb{R}^N | p_L \leq p_i \leq p_U \ \forall i\}$. A common pricing strategy is to adopt a single uniform price for all the agents across the network. Here, $\mathbf{P} = \{\mathbf{p} \in \mathbb{R}^N | p_i = \bar{p} \ \forall i, \ \bar{p} \in \mathbb{R}\}$. In a similar fashion, depending on the application, the firm can select some appropriate business constraints to impose on the pricing strategy. Finally, $\mathbf{P}$ can also incorporate specific constraints on the network segmentation. For example, motivated by business practices, a particular segment of agents should be offered the same price or special members (loyal customers) need to receive a lower price than regular customers.

Our goal is to develop a general optimal pricing method for the firm that incorporates the different business rules as constraints. Before we mathematically formulate the problems of the potential buyers and the firm, we summarize our assumptions about them below.

Assumption 2. We assume the following about each agent $i \in \mathcal{I}$ in the network:

a. Each agent has a linear additive form of the utility as described below in (2.1).
b. Each agent is assumed to be rational and a utility maximizer.
c. Each agent can buy at most one unit of the item and either fully purchases the item or does not purchase it at all.
d. If the utility of an agent is zero, the tie is broken assuming this agent buys the item.

Assumption 3. We assume that the seller is a profit maximizer as described below in (2.3) and has a linear manufacturing cost.

For a given set of prices chosen by the seller, the agents in the network aim to collectively maximize their utility from purchasing the item. We capture the linear additive valuation model of an agent by assuming that the total value for owning the item is the sum of the agent's own valuation and the valuation derived from the (positive) influences of the agent's friends who own the item. In particular, the utility of agent i is given by:

$$u_i(\alpha_i, \boldsymbol{\alpha_{-i}}, p_i) = \alpha_i \Big[g_{ii} + \sum_{j \in \mathcal{I} \setminus i} \alpha_j g_{ji} - p_i \Big], \tag{2.1}$$

where $\alpha_i \in \{0,1\}$ is a binary variable that represents the purchasing decision of agent i and $\boldsymbol{\alpha_{-i}}$ represents the vector of purchasing decisions of all the agents but i in the network. If $\alpha_i = 1$, agent i purchases the item and derives a utility equal to $g_{ii} + \sum_{j \in \mathcal{I} \setminus i} \alpha_j g_{ji} - p_i$ from owning the item and if, on the other hand, $\alpha_i = 0$, the agent does not purchase the item and derives zero utility. The utility maximization problem of agent i can then be written as follows:

$$\max_{\alpha_i \in \{0,1\}} \; u_i(\alpha_i, \boldsymbol{\alpha_{-i}}, p_i). \tag{2.2}$$

The profit maximizing problem of the seller is given by:

$$\max_{\mathbf{p} \in \mathbf{P}} \sum_{i \in \mathcal{I}} \alpha_i \left(p_i - c\right), \tag{2.3}$$

where α_i's are the purchasing decisions of the agents obtained from the utility maximization subproblem 2.2 and c is the unit manufacturing cost of the item. If agent i decides to buy the product at the offered price p_i, α_i is equal to 1 and the firm incurs a profit of $p_i - c$. If agent i decides not to purchase the item, it incurs zero profit to the seller. The firm designs the prices to offer to the different agents depending on the pricing strategy $\mathbf{P}$ employed.

We view the entire problem as a two-stage Stackelberg game, referred also to as the pricing-purchasing game. First, the seller leads the game by choosing the vector of prices $\mathbf{p}$ to be offered to the potential buyers. Second, the agents follow by deciding whether or not to purchase the item at the offered prices. In other words, the firm sets the prices $\mathbf{p} \in \mathbf{P}$ and the network of agents collectively follow with their decisions, $\alpha_i \; \forall i \in \mathcal{I}$. We are interested in *subgame perfect equilibria* of this two-stage pricing-purchasing game (e.g. see more details in Fudenberg and Tirole (1991)). For a fixed vector of prices offered by the seller, the equilibria of the second stage game, referred to as the *purchasing equilibria* are defined as follows:

$$\alpha_i^* \in \arg\max_{\alpha_i \in \{0,1\}} u_i\big(\alpha_i, \boldsymbol{\alpha^*_{-i}}, p_i\big) \quad \forall i \in \mathcal{I}. \tag{2.4}$$

We note that this definition is similar to the consumption equilibria for a divisible item (or service) in Candogan et al. (2012). However, in our case the decision variables α_i are restricted to be binary so that agents cannot buy fractional amounts of the item and have to either buy it fully or not to buy it at all.

We also note that the overall two-stage problem is non-linear and non-convex as it includes terms of the form $\alpha_i p_i$ in the seller's objective function and $\alpha_i \alpha_j$ in the objective functions of the agents which we will see soon appears as constraints in the seller's problem. In addition, the discrete nature of the purchasing decisions makes it even more complicated in that we are working with a non-convex integer program. Therefore, one cannot directly apply tractable convex optimization methods to solve the problem to optimality. In the next section, we start by considering the second stage purchasing game and show the existence of an equilibrium such as in Eq. (2.4), for any given vector of prices. We then characterize the equilibria by a set of constraints for any price vector. In Section 4, we use this characterization to formulate the optimal pricing problem as a MIP.

3. Purchasing equilibria

In this section, we consider the second stage purchasing game and show the existence of a pure Nash equilibrium (PNE) strategy, given any vector of prices $\mathbf{p}$ specified by the seller. We observe that there could be multiple pure Nash equilibria for this game but we characterize all these equilibria through a system of constraints using duality theory.

3.1. Existence of the purchasing equilibria

The existence of a PNE for the second stage game is summarized in the following theorem.

Theorem 1. The second stage game has at least one pure Nash equilibrium for any given vector of prices $\mathbf{p}$ chosen by the seller.

The proof can be found in the Appendix A. We note that theorem 1 guarantees the existence of a PNE but not necessarily its uniqueness. Consider the following example in which two distinct PNE's arise. Assume a network with two symmetric agents that mutually influence one another: $g_{11} = g_{22} = 2$ and $g_{21} = g_{12} = 1$. Consider the given price vector: $p_1 = p_2 = 2.5$. In this case, we have two PNE's: buy-buy and no buy-no buy. In other words, if player 1 buys, player 2 should buy but if player 1 does not buy, player 2 will not either. Therefore, uniqueness is not guaranteed.

We note that the existence of a PNE for the second stage game is not always guaranteed for the case with negative externalities. Nevertheless, it may be possible to find sufficient conditions on the valuation functions to ensure the existence. This direction is not pursued in this paper as we focus on positive externalities.

A common assumption in games with multiple equilibria is that the Nash equilibrium that is actually played relies on the presence of some mechanism or process that leads the agents to play this particular outcome (see Fudenberg and Tirole (1991) for more details). We impose a similar assumption in our setting and assume that the seller can identify some simple (low cost) strategies to guide the players to his preferred Nash equilibrium. In the above example, reducing the price for one of the players to $p = 2 - \epsilon$ for a small $\epsilon > 0$ is enough to guarantee the preferred buy-buy equilibrium and discard the undesired no buy-no buy equilibrium. We refer the interested readers to the paper by Gunnec and Raghavan (2012) where the authors encounter a similar setting and propose a secondary seeding algorithm called the least cost influence problem that minimizes the total cost of incentives offered to all the players in order to achieve the preferred solution in their setting. We will see in Section 4 that the nature of the first stage game induces the preferred equilibrium buy-buy and hence a similar secondary mechanism is potentially required.

3.2. Characterization of the purchasing equilibria

The natural next step is to characterize the purchasing equilibria as a function of the prices. In other words, we would like to characterize the functions $\alpha_i(\mathbf{p})$ $\forall i \in \mathcal{I}$. This will allow us to reduce the two-stage problem to a single optimization formulation, where the only variables are the prices. In our setting, a closed form expression for $\alpha_i(\mathbf{p})$ is not straightforward. Instead, by using duality theory, we characterize the set of constraints the equilibria should satisfy for any given vector of prices. We begin by making the following observation regarding the utility maximization problem of any agent.

Observation 1. Given a vector of prices $\mathbf{p}$, let us consider the subproblem 2.2 for agent i. If the decisions of the other agents $\boldsymbol{\alpha_{-i}}$ are given, the problem of agent i has a tight linear programming (LP) relaxation.

In fact, the sub-problem faced by agent i happens to be an assignment problem for fixed values of $\mathbf{p}$ and $\boldsymbol{\alpha_{-i}}$. More specifically, let us consider the LP obtained by the continuous relaxation of the binary constraint $\alpha_i \in \{0,1\}$ to $0 \le \alpha_i \le 1$. One can view this LP as a relaxation purchasing game where agents can purchase fractional amounts of the item and therefore adopt mixed strategies. If the quantity $\left(g_{ii} + \sum_{j\in\mathcal{I}\setminus i} \alpha_j g_{ji} - p_i\right)$ (which is exactly known since $\mathbf{p}$ and $\boldsymbol{\alpha_{-i}}$ are given) is positive, $\alpha_i^* = 1$ and if the quantity is negative, $\alpha_i^* = 0$. Finally, if the quantity is equal to zero, α_i^* can be any number in $[0,1]$ so that the agent is indifferent between buying and not buying the item. Therefore, the LP relaxation of the subproblem of agent i for fixed values of $\mathbf{p}$ and $\boldsymbol{\alpha_{-i}}$ is tight, meaning that all the extreme points are integer. Equivalently, for any feasible fractional solution, one can find an integral solution with at least the same objective.

Observation 1 allows us to transform the relaxation of subproblem 2.2 for agent i into a set of constraints by using duality theory of linear programming. More specifically, these constraints comprise of primal feasibility, dual feasibility and strong duality conditions. In the case of subproblem 2.2 for agent i, the constraints can be written as follows:

$$\text{Primal feasibility:} \quad 0 \le \alpha_i \le 1 \tag{3.1}$$

$$\text{Dual feasibility:} \quad y_i \ge g_{ii} + \sum_{j\in\mathcal{I}\setminus i} \alpha_j g_{ji} - p_i \tag{3.2}$$

$$y_i \ge 0 \tag{3.3}$$

$$\text{Strong duality:} \quad y_i = \alpha_i\Big(g_{ii} + \sum_{j\in\mathcal{I}\setminus i} \alpha_j g_{ji} - p_i\Big) \tag{3.4}$$

Here, the variable y_i represents the dual variable of subproblem 2.2 for agent i. Combining the above constraints (3.1–3.4) for all the agents $i \in \mathcal{I}$ characterizes all the equilibria (mixed and pure) of the second stage game as a function of the prices. In order to restrict our attention to the pure

Nash equilibria (that the existence is guaranteed by theorem 1), one can impose $\alpha_i \in \{0,1\}\ \forall\ i$. Observe that this characterization has reduced $N+1$ interconnected optimization problems to be compactly written as a single optimization formulation. We note that the number of variables increases by N as we add a dual continuous variable for each agent's subproblem.

4. Optimal pricing: MIP formulation

In this section, we use the existence and characterization of PNE to transform the two-stage optimal pricing problem into a single optimization formulation. This formulation happens to be a non-convex integer program but depicts some interesting properties. We then reformulate the problem to arrive at a MIP with linear constraints.

We next formulate the optimal pricing problem faced by the seller (denoted by problem Z) by incorporating the second stage PNE characterized by the set of constraints (3.1–3.4) for each agent. The class of optimization problems with equilibrium constraints is referred to as MPEC (Mathematical Program with Equilibrium Constraints) and is well studied in the literature (see e.g. Luo et al. (1996)). The equilibrium constraints in our case include constraints (3.2–3.4) and $\alpha_i \in \{0,1\}$ instead of constraint (3.1) for all agents to restrict to the pure Nash equilibria. The formulation is given by:

$$\begin{aligned}
\max_{\substack{\mathbf{p}\in\mathbf{P}\\ \mathbf{y},\boldsymbol{\alpha}}} \quad & \sum_{i\in\mathcal{I}} \alpha_i (p_i - c) \\
\text{s.t.} \quad & \left.\begin{aligned}
y_i &= \alpha_i \big(g_{ii} + \textstyle\sum_{j\in\mathcal{I}\setminus i} \alpha_j g_{ji} - p_i\big) \\
y_i &\geq g_{ii} + \textstyle\sum_{j\in\mathcal{I}\setminus i} \alpha_j g_{ji} - p_i \\
y_i &\geq 0 \\
\alpha_i &\in \{0,1\}
\end{aligned}\right\} \ \forall\ i \in \mathcal{I}
\end{aligned} \tag{Z}$$

In addition to the presence of binary variables, one can see that the above optimization problem is non-linear and non-convex as it includes terms of the form $\alpha_i\alpha_j$ and $\alpha_i p_i$. Therefore, problem Z is not easily solvable by commercially available solvers. We next prove the following interesting tightness result of problem Z that allows us to view the problem as a non-convex continuous, instead of an integer problem. This is not of immediate consequence in this section but provides insight to one of our main results presented in theorem 2.

Proposition 1. Problem Z admits a tight continuous relaxation.

The proof can be found in the Appendix B. We next show that by introducing a few additional continuous variables, one can reformulate problem Z into an equivalent MIP formulation that has a linear objective with linear constraints and the same number of binary variables. We first define the following additional variables:

$$
\begin{aligned}
z_i &= \alpha_i p_i && \forall\ i \in \mathcal{I} \\
x_{ij} &= \alpha_i \alpha_j && \forall\ j > i \text{ where } i, j \in \mathcal{I}.
\end{aligned}
$$

By using the binary nature of the variables and adding certain linear constraints, one can replace all the non-linear terms in problem Z. that is now equivalent to the following MIP formulation denoted by Z-MIP:

$$
\max_{\substack{\mathbf{p}\in\mathbf{P} \\ \mathbf{y},\mathbf{z},\mathbf{x},\boldsymbol{\alpha}}} \quad \sum_{i\in\mathcal{I}} \left(z_i - c\alpha_i\right) \tag{Z-MIP}
$$

s.t.

$$
\left.\begin{aligned}
y_i &= \alpha_i g_{ii} + \textstyle\sum_{j\in\mathcal{I}\setminus i} g_{ji} x_{ji} - z_i \\
y_i &\geq g_{ii} + \textstyle\sum_{j\in\mathcal{I}\setminus i} \alpha_j g_{ji} - p_i \\
y_i &\geq 0
\end{aligned}\right\} \ \forall\ i \in \mathcal{I} \tag{4.1}
$$

$$
\left.\begin{aligned}
z_i &\geq 0 \\
z_i &\leq p_i \\
z_i &\leq \alpha_i p^{max} \\
z_i &\geq p_i - (1-\alpha_i) p^{max}
\end{aligned}\right\} \ \forall\ i \in \mathcal{I} \tag{4.2}
$$

$$
\left.\begin{aligned}
x_{ij} &\geq 0 \\
x_{ij} &\leq \alpha_i \\
x_{ij} &\leq \alpha_j \\
x_{ij} &\geq \alpha_i + \alpha_j - 1 \\
x_{ij} &= x_{ji}
\end{aligned}\right\} \ \forall\ j > i \text{ where } i, j \in \mathcal{I} \tag{4.3}
$$

$$
\alpha_i \in \{0, 1\} \qquad \forall\ i \in \mathcal{I} \tag{4.4}
$$

In the above formulation, p^{max} denotes the maximal price allowed and is typically known from the context. For example, one can take: $p^{max} = \max_i \left\{g_{ii} + \sum_{j\neq i} g_{ji}\right\}$ without affecting the problem at all, since no agent would ever pay a price beyond that value. The set of constraints (4.2) aims to guarantee the definition of the variable z_i, whereas the set of constraints (4.3) ensures the correctness of the variable x_{ij}. One can note that in the above Z-MIP formulation, we have a total of $\frac{N^2}{2} + 3.5N$ variables ($4N$ for $\boldsymbol{\alpha}$, $\mathbf{p}$, $\mathbf{y}$ and $\mathbf{z}$ and $\frac{N(N-1)}{2}$ for $\mathbf{x}$) but only N of them are binary, while the remaining are all continuous.

We conclude that the problem of designing prices for selling an indivisible item to agents embedded in a social network can be formulated as a MIP that is equivalent to the two-stage non-convex

IP game we started with. As a result, one can easily incorporate various business constraints such as pricing policies, market segmentation, inter-buyers price constraints, just to name a few. In other words, this formulation can be viewed as an operational tool to solve the optimal pricing problem of the seller. This is in contrast to previous approaches that proposed tailored algorithms for the problem where one cannot easily incorporate business rules. However, solving a MIP may not be very scalable. If the size of the network is not very large, one can still solve it using commercially available MIP solvers. Moreover, it is possible to solve the problem off-line (before launching a new product for example) so that the running time might not be of first importance. Potentially, one can also consider network clustering methods to aggregate or coalesce several nodes to a single virtual agent in order to reduce the scale of the network. If the size of the network is very large, one needs to find more efficient methods to solve the Z-MIP problem. In the next section, we derive efficient methods (polynomial in the number of agents) to solve it to optimality for two different but popular pricing strategies.

5. Efficient algorithms

5.1. Discriminative Prices

In this section, we consider the most general pricing strategy where the firm offers discriminative prices that potentially differ for each agent, depending on his influence in the network. In particular, $\mathbf{P} = \mathbb{R}^N$ in problem Z-MIP. This scenario is of interest in various practical settings where the seller gathers the purchasing history of each potential buyer, his geographical location as well as other attributes or features. It can also be used by the seller to understand who are the influential agents in the network and what is the maximal profit he can potentially achieve if he were to discriminate prices at the individual level. The prices can then be implemented by setting a ticket price that is the same for all the agents and sending out coupons with discriminative discounts to the potential buyers in the network. In fact, in practice, it often occurs that people receive different deals for the same item depending on the loyalty class, purchase history and store location. It is also common that a very small number of highly influential people (e.g., certified bloggers) receive the item for free or at a very low price. The method we propose aims to provide a systematic and automated way of finding the prices (equivalently, the discounts) to offer to the agents embedded in a social network based on their influence so as to maximize the total profit of the seller.

Solving the Z-MIP problem presented in the previous section using an optimization solver may be not practical for a very large scale network. We next show that solving the LP relaxation of the Z-MIP problem yields the desired optimal integer solution. Consequently, one can solve the problem efficiently (polynomial in the number of agents) and obtain an optimal solution even for large scale networks. This result is very interesting because the linearization of problem Z was possible only under the assumption of integrality of the decision variables. In other words, in order to reformulate problem Z into problem Z-MIP, the binary restriction was crucial. It is therefore possible that because of variables z_i and x_{ij}, new fractional solutions that cannot be practically implemented are introduced. However, the following theorem shows that the optimal solutions of Z-MIP can be identified using its relaxation.

THEOREM 2. The optimal discriminative pricing solution of the Z-MIP problem can be obtained efficiently (polynomial in the number of agents). In particular, problem Z-MIP with $\mathbf{P} = \mathbb{R}^N$ admits a tight LP relaxation.

The proof can be found in the Appendix C. We not only show that the LP relaxation is tight but also provide a constructive method of rounding the fractional LP solution to obtain an integer solution that is as good in terms of the profit. One can employ this constructive method or use a method like simplex to arrive at the optimal extreme points which we know are integer. Here forth, when we refer to the solution of Z-MIP, we refer to its integer optimal solution only.

The result of theorem 2 suggests an efficient method to solve the problem that we formulated as a two stage non-convex integer program. The LP based method inherits all the complexity properties of linear programming and is thus scalable and applicable to large scale networks. In the next section, we consider adding constraints on the pricing strategy by investigating the case of designing a single uniform price across the network.

5.2. Uniform Price

In this section, we consider the case where the seller offers a uniform price across the network while incorporating the effects of social interactions. This scenario may arise when the firm may not want to price discriminate due to fairness or ethical reasons and prefers to offer a uniform price. It is also interesting to compare the total profits achieved by this pricing strategy to the case where discriminative prices are used. In particular, one can quantify how much the seller is losing by working with a uniform pricing strategy. We observe that a similar result to theorem 2 for the uniform pricing case does not hold. In other words, by adding the linear uniform price constraint: $p_1 = p_2 = \ldots = p_N$ to the Z-MIP formulation as an additional business rule, the corresponding LP relaxation is no longer tight and we obtain fractional solutions that cannot be implemented in practice. Geometrically, this means that incorporating such a constraint in the Z-MIP formulation is equivalent to add a cut that violates the integrality of the extreme points of the feasible region. Therefore, we propose an alternative approach that solves the problem optimally by an efficient algorithm (polynomial in the number of agents) that is based on iteratively solving the relaxed Z-MIP, which is an LP. We summarize this result in the following theorem.

THEOREM 3. The optimal solution of the Z-MIP problem for the case of a single uniform price can be obtained efficiently (polynomial in the number of agents) by applying algorithm 1.

Algorithm 1 Procedure for finding the uniform optimal price

Input: c, N **and** G

Assumption: $g_{ji} \geq 0 \quad \forall i, j \in \mathcal{I}$

Procedure

1. Set the iteration number to $t = 1$, solve the relaxed Z-MIP (an LP) and obtain the vector of discriminative prices $\mathbf{p}^{(1)}$.
2. Find the minimal discriminative price $p_{min}^{(t)} = \min_{i \in \mathcal{I}} p_i^{(t)}$ and evaluate the objective function $\Pi^{(t)}$ with $p_i = p_{min}^{(t)} \ \forall \ i \in \mathcal{I}$ using formula (D.1).
3. Remove all the nodes with the minimal discriminative price from the network (including all their edges). If there are no more agents in the network, go to step 5. If not, go to step 4.
4. Re-solve the relaxed Z-MIP for the new reduced network and denote the output by $\mathbf{p}^{(t+1)}$. Set $t := t + 1$ and go to step 2.
5. The optimal uniform price is equal to $p_{min}^{(\hat{t})}$, where $\hat{t} = \arg\max \Pi^{(t)}$ i.e., the price that yields the larger profits.

We propose a method that iteratively solves the LP relaxation for discriminative prices to arrive at the optimal uniform price. The details of this procedure are summarized in algorithm 1. We show its termination in finite time and prove its correctness by showing that it yields the optimal solution of the uniform pricing problem (in polynomial complexity). The proof of correctness of algorithm 1 can be found in the Appendix D. At a high level, the procedure in Algorithm 1 iteratively reduces the size of the network by eliminating agents with low valuations (at least one per iteration). As a result, it suffices for one to consider only a finite selection of price values to identify the optimal uniform price.

6. Price-incentives to guarantee influence

This far, we have assumed that consumers always influence their peers as long as they purchase the item. This assumption is not realistic in many practical settings. Indeed, after purchasing an item, it is sometimes not entirely natural to influence friends about the product unless one takes some effort to do so. This, for example, could be by writing a review, endorsing the item on their wall, blogging about the item or at the very least announcing the purchase. However, to the best of our knowledge, most previous works impose such an assumption with the exception of Arthur et al. (2009). In the latter, the authors study a cash back setting where the seller offers an exogenous specified uniform cash reward to any recommender if he influenced at most one of his neighbors and they buy the item. We consider a variant of this model but study it in the context of purchasing equilibrium and the optimization framework we have proposed in this paper.

Consider a setting where the seller offers both a price and a discount (also referred to as an incentive) to each agent in the network. Each agent can then decide whether to buy the item or not. If the agent decides to purchase the item, he can claim a fraction of the discount offered by the seller in return for influence actions. These can include liking the product or a wall post in an online social platform such as Facebook or writing a review so that these actions can be digitally tracked by the seller. The agent receives a small discount in exchange for a simple action such as liking the product and a more significant discount by taking a time-consuming action such as writing a detailed review. Such incentive mechanisms are commonplace these days. For example, online booking agencies request reviews of booked hotels on their website in return for certain loyalty benefits. Using such a model, the seller can now ensure the influence among the agents so that the network externalities effects are guaranteed to occur. In particular, the profits obtained through the optimization are guaranteed for the seller since each agent claims the discounted price as soon as the influence action is taken. In the previous setting where externalities are assumed to always occur, the actual profits may be far from the value predicted by the optimization. In fact, we show using a computational example in Section 7 that even if a single agent does not impart his influence, it can significantly reduce the total profits of the seller. We now extend our model and results to this more general setting where the seller can design price-incentives to guarantee social influence.

We consider a model with a continuum of actions to influence ones' neighbors. Let $t_i \geq 0$ denote the utility equivalent of the maximal effort needed by agent i to claim the entire discount offered by the seller. If agent i decides to purchase the item, we assume that $\gamma_i t_i$ is the effort required by agent i to claim a fraction γ_i of the discount, where $0 \leq \gamma_i \leq 1$. We view t_i as the *influence cost* for agent i and the variable γ_i as the *influence intensity* chosen by agent i. The parameter t_i can be estimated from historical data using the intensity of online activity for past purchases, the number of reviews written, the corresponding incentives needed and data from cookies. For a given set of prices $\mathbf{p}$ and discounts $\mathbf{d}$ chosen by the seller, we extend the utility function of agent i in Eq. (2.1) as follows:

$$u_i(\alpha_i, \gamma_i, \boldsymbol{\alpha_{-i}}, \boldsymbol{\gamma_{-i}}, p_i, d_i) = \alpha_i \Big(g_{ii} + \sum_{j \in \mathcal{I} \setminus i} \gamma_j g_{ji} - p_i \Big) + \gamma_i (d_i - t_i), \tag{6.1}$$

where $\gamma_i \leq \alpha_i$ and α_i is the binary purchasing decision of agent i. So, if agent i does not purchase the item, $\alpha_i = 0$ and $\gamma_i = 0$ as well. But if agent i purchases the item, then $\alpha_i = 1$ and γ_i can be any number in $[0,1]$ as chosen by agent i. Here, $\boldsymbol{\alpha_{-i}}$ and $\boldsymbol{\gamma_{-i}}$ are the decisions of all the other agents but i. Similarly to problem 2.2, the utility maximization problem for agent i can be written as follows:

$$\begin{aligned} \max_{\alpha_i,\gamma_i} \quad & u_i(\alpha_i,\gamma_i,\boldsymbol{\alpha_{-i}},\boldsymbol{\gamma_{-i}},p_i,d_i) \\ \text{s.t.} \quad & 0 \leq \gamma_i \leq \alpha_i \\ & \alpha_i \in \{0,1\} \end{aligned} \tag{6.2}$$

In a similar way as problem 2.3, the seller's profit maximization problem can be written as:

$$\max_{\mathbf{p},\mathbf{d}\in\mathbf{P}} \quad \sum_{i\in\mathcal{I}} \Big[\alpha_i\left(p_i - c\right) - \gamma_i d_i\Big]. \tag{6.3}$$

Here, the decision variables of the seller are $\mathbf{p}$ and $\mathbf{d}$ which are two vector of prices and discounts with an element for each agent in the network. These vectors can be chosen according to different pricing strategies. For example, one can consider a fully discriminative or a fully uniform pricing strategy or more generally, an hybrid model where the regular price is uniform across the network $(p_i = p_j)$ but the discounts are tailored to the various agents. This hybrid setting corresponds to a common practice of online sellers that offer a standard posted price for the item but design personalized discounts for different classes of customers that are sent via e-mail coupons. Finally, similarly to the previous setting, one can incorporate various polyhedral business rules on prices, discounts and constraints on network segmentation. The variables α_i and γ_i are decided according to each agent's utility maximization problem given in (6.2). If agent i decides to buy the product, then the seller incurs a profit of $p_i - \gamma_i d_i - c$.

We note that in the special case where $\alpha_i = \gamma_i$ and $t_i = 0$ $\forall i \in \mathcal{I}$, we recover the previous model where the seller offers a single price to each agent and any buyer is assumed to always influence his peers. In addition, by adding the constraint $\gamma_i \in \{0,1\}$ we have an interesting setting where each agent can only buy at two different prices: a full price p_i that does not require any action and a discounted price $p_i - d_i$ that requires some action to influence. Note that one can easily extend the model in this section to more than two prices so as to incorporate a finite but discrete set of different actions secified by the seller.

Our goal is to extend the results from previous sections for this new setting with incentives to guarantee influence. We begin by studying the purchasing equilibria of the second stage game. By using a similar methodology as in section 3, one can show that for any given prices and discounts there exists a PNE for the second stage game.

Theorem 4. The second stage game has at least one pure Nash equilibrium for any given vector of prices $\mathbf{p}$ and discounts $\mathbf{d}$ chosen by the seller.

The proof of theorem 4 is in a similar spirit as theorem 1. In particular, if for some agent $0 < \alpha_i^* < 1$, α_i^* is increased to 1 while keeping the exact same value for γ_i^*. Therefore, by using a similar construction procedure as in theorem 1, one can obtain a PNE (the complete proof is not provided for conciseness). Note that in this case, a PNE is defined such that the binary purchasing decisions α_i are all integer. However, one can also note that there always exists an equilibrium for which the variables γ_i are all integer as well. More precisely, if $d_i - t_i > 0$ (remember that the prices and discounts are given), γ_i can be increased to 1 and otherwise $\gamma_i = 0$. We therefore have the existence of a PNE with γ_i integer as well.

One can see that a result similar to Observation 1 still holds and therefore one can characterize the equilibria (mixed and pure) as a set of constraints where the binary variables are relaxed to be continuous. In this case, one can transform subproblem 6.2 of agent i to a set of feasibility constraints using duality theory as follows:

$$\text{Primal feasibility:} \quad 0 \le \alpha_i \le 1 \tag{6.4}$$

$$0 \le \gamma_i \le \alpha_i \tag{6.5}$$

$$\text{Dual feasibility:} \quad y_i - w_i \ge g_{ii} + \sum_{j \in \mathcal{I} \setminus i} \gamma_j g_{ji} - p_i \tag{6.6}$$

$$w_i \ge d_i - t_i \tag{6.7}$$

$$y_i, w_i \ge 0 \tag{6.8}$$

$$\text{Strong duality:} \quad y_i = \alpha_i \Big(g_{ii} + \sum_{j \in \mathcal{I} \setminus i} \gamma_j g_{ji} - p_i \Big) + \gamma_i (d_i - t_i) \tag{6.9}$$

We now have two continuous dual variables y_i and w_i, together with two dual feasibility constraints for each agent i. Similar to the earlier setting, in order to restrict to the pure Nash equilibria (that is necessary for the problem of optimal pricing), we need to impose α_i to be binary variables for all agents $i \in \mathcal{I}$. We can then formulate the optimal pricing problem faced by the seller, similar to problem Z, that maximizes the profits given in (6.3) with the equilibrium constraints (6.4)-(6.9), where the constraints on α_i are replaced by the binary versions as follows:

$$\begin{aligned} \max_{\substack{\mathbf{p},\mathbf{d} \in \mathbf{P} \\ \mathbf{y},\mathbf{w},\boldsymbol{\alpha},\boldsymbol{\gamma}}} \quad & \sum_{i \in \mathcal{I}} \Big[\alpha_i \left(p_i - c \right) - \gamma_i d_i \Big] \\ \text{s.t.} \quad & \text{constraints } (6.5) - (6.9), \alpha_i \in \{0,1\} \quad \forall\, i \in \mathcal{I} \end{aligned} \tag{Zi}$$

We denote this problem by Zi where i represents the model with incentives to guarantee influence of this present section. We impose the following assumption on the agents to address the ties in utilities.

Assumption 4. If the discount offered to agent i is such that $d_i = t_i$, then agent i decides to influence i.e., $\gamma_i > 0$.

The seller can always ensure such a condition by increasing the discount by a small factor $\epsilon > 0$. In addition, the nature of the first stage problem guarantees this condition at optimality. One can then make the following Observation.

OBSERVATION 2. Every optimal solution of problem Z*i* satisfies $d_i \leq t_i$.

Indeed, the seller can always reduce d_i to be equal to t_i while maintaining feasibility and strictly increasing the objective function. This implies that the constraint (6.7) is redundant in the optimal pricing problem. Consequently and by using the constraints (6.6–6.8), one can always assign $w_i = 0$ in the pricing problem while maintaining feasibility and without altering the objective function. This observation allows us to simplify problem Z*i* by removing all the dual variables w_i $\forall i \in \mathcal{I}$. We next extend proposition 1 for this setting.

PROPOSITION 2. Problem Z*i* admits a tight continuous relaxation. Moreover, there always exists an optimal solution to problem Z*i* where all the variables γ's are integer as well.

The proof of this proposition is provided in the Appendix E. The second result in this proposition is interesting because it implies that even though the seller allows for a continuum of influence actions, the buyer would either fully influence or not influence at all. As a result, this is equivalent to the setting where the seller offers only two options: a full price p_i and a discounted price $p_i - d_i$ in exchange of a specific action to influence.

Problem Z*i* has non-linearities of the form $\alpha_i\gamma_j$, $\alpha_i p_i$ and $\gamma_i d_i$. Using the discrete nature of the variables α_i and γ_i from proposition 2, one can transform problem Z*i* to the following MIP formulation, denoted by Z*i*-MIP:

$$\max_{\substack{\mathbf{p},\mathbf{d}\in\mathbf{P} \\ \mathbf{y},\mathbf{z},\mathbf{z}^{\mathbf{d}},\mathbf{x},\boldsymbol{\alpha},\boldsymbol{\gamma}}} \quad \sum_{i\in\mathcal{I}} \left(z_i - z_i^d - c\alpha_i\right) \tag{Zi-MIP}$$

s.t.

$$\left.\begin{aligned} & y_i = \left(\alpha_i g_{ii} + \textstyle\sum_{j\in\mathcal{I}\setminus i} x_{ji} g_{ji} - z_i\right) + \left(z_i^d - \gamma_i t_i\right) \\ & y_i \geq g_{ii} + \textstyle\sum_{j\in\mathcal{I}\setminus i} \gamma_j g_{ji} - p_i \\ & \gamma_i \leq \alpha_i \\ & y_i \geq 0 \end{aligned}\right\} \forall\, i \in \mathcal{I} \tag{6.10}$$

$$\left.\begin{aligned} z_i, z_i^d &\geq 0 \\ z_i &\leq p_i \\ z_i &\leq \alpha_i p^{max} \\ z_i &\geq p_i - (1-\alpha_i) p^{max} \\ z_i^d &\leq d_i \\ z_i^d &\leq \gamma_i p^{max} \\ z_i^d &\geq d_i - (1-\gamma_i) p^{max} \end{aligned}\right\} \forall\, i \in \mathcal{I} \tag{6.11}$$

$$\left.\begin{aligned} x_{ji} &\geq 0 \\ x_{ji} &\leq \alpha_i \\ x_{ji} &\leq \gamma_j \\ x_{ji} &\geq \alpha_i + \gamma_j - 1 \end{aligned}\right\} \forall\, i \neq j \in \mathcal{I} \tag{6.12}$$

$$\alpha_i, \gamma_i \in \{0, 1\} \qquad \forall\, i \in \mathcal{I} \tag{6.13}$$

where p^{max} is the maximum price allowed. Note that we removed the dual variables w_i by using Observation 2. We conclude that the problem of designing prices and incentives for selling an indivisible item to agents embedded in a social network can be formulated as a MIP where one can incorporate business rules on prices and on constraints on network segmentation. However, solving a MIP may not be very scalable. For the case of discriminative prices and discounts, i.e., when $\mathbf{P} = \mathbb{R}^N$x $\mathbb{R}^N$, we are able to retrieve a similar result as theorem 2. The result is summarized in the following theorem.

Theorem 5. The optimal discriminative pricing solution of the Zi-MIP problem can be obtained efficiently (polynomial in the number of agents). In particular, problem Zi-MIP with $\mathbf{P} = \mathbb{R}^N$x $\mathbb{R}^N$ admits a tight LP relaxation.

The proof is in a similar spirit as theorem 2 and is not reported entirely due to space limitations. However, the main idea can be folded into the following two steps. First, fix the values of γ_i, z_i^d and proceed in the same fashion as in theorem 2 to show how to construct a solution with α_i integer $\forall i \in \mathcal{I}$. Next, with the integer values of $\boldsymbol{\alpha}$ obtained from the previous step, one can show that the objective does not change when we modify any component of $\boldsymbol{\gamma}$ to 0 or 1 by appropriately modifying the prices of the neighbors so that their actions do not change as in proposition 2.

In comparison to problem Z-MIP with a single price for each agent, problem Zi-MIP yields potentially lower profits for the seller. However, these profits are guaranteed whereas in the previous case, the estimated profits can be far from the actual values if people fail to influence their neighbors. The difference in profits between both settings can be viewed as the price the seller has to pay to guarantee the influence between agents in the network and can be computed efficiently by solving both settings.

An interesting observation that we see throughout this section is that even though our model allows a continuum of influence actions, the optimal prices can be designed in such a way that only two price options suffice. More specifically, the two options are a full price with no action required and a discounted price which requires an influence action in return. In the next section, we present numerical experiments to highlight several key insights including a case that illustrates the benefit of the model with incentives that guarantees influence.

7. Computational experiments

In this section, we present computational experiments on simple example networks to draw qualitative insights about incorporating social interactions, comparing various pricing strategies including the richer model with incentives developed in this paper. We consider a network with $N = 10$ agents.

Value of incorporating network externalities: In Fig. 1, we plot the optimal prices offered by the seller to the different agents under the discriminative and uniform pricing strategies with and without social interactions. The circles around the markers, whenever present, depict the fact that the agent decided not to purchase the item at the offered price (agents 7, 8 and 9 for uniform price and agent 8 for uniform without externalities). In this instance, each agent is randomly connected to three other agents with $g_{ji} = 1.25$ for any connected edge, $g_{ii} = 2.5R$ where R is a uniform random variable in $[1,2]$ (denoted by $U[1,2]$) and $c = 2$.

We observe that by incorporating the positive externalities between the agents, the seller earns higher profits. In this particular example, the total profits are equal to 50.75 (discriminative prices) and 24.5 (uniform price) for the case with network externalities compared to 14.5 and 9 for the case without network externalities. This result is expected because every agent's willingness-to-pay increases as their neighbors positively influence them. The seller can therefore charge even higher prices and increase his profits. Fig. 1 also shows the added benefit from using a discriminative pricing strategy compared to a uniform single price. When the firm has the additional flexibility to price discriminate and offers a different price to each agent in the network, the total profits can increase significantly. In the example above, only one agent is offered a price that is lower than the optimal uniform price.

Pricing an influencer: In Fig. 2, we present an example where it is beneficial for the seller to earn negative profit ($p_i < c$) on some influential agent i in order to extract significant positive profits on his neighbors. In particular, we consider a network where agent 5 is a very influential

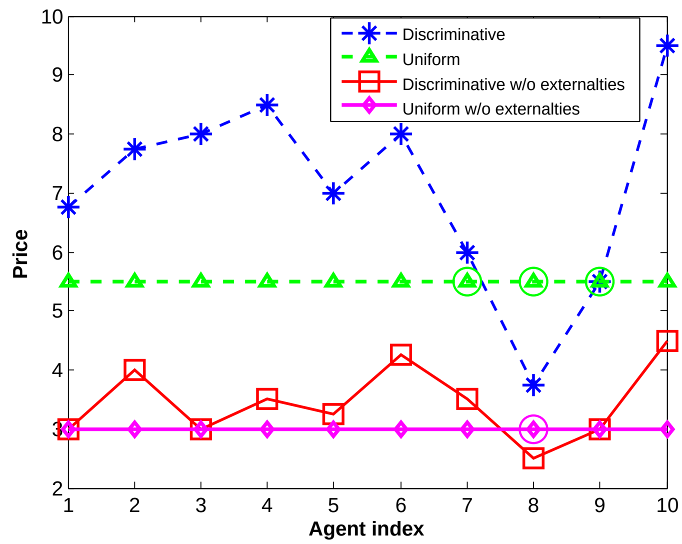

Figure 1 Value of incorporating network externalities for the discriminative and uniform pricing strategies.

player with g_{55} being very low (0.075) while g_{5j} is sufficiently high (1.38) for the four agents that he influences. Here, $g_{ij} = 0.75$ for any other connected edge, $g_{ii} = 1.5R\ \forall i \neq 5$ where $R = U[1,2]$ and $c = 2$. The optimal discriminative price vector includes a price for agent 5 that happens to be lower

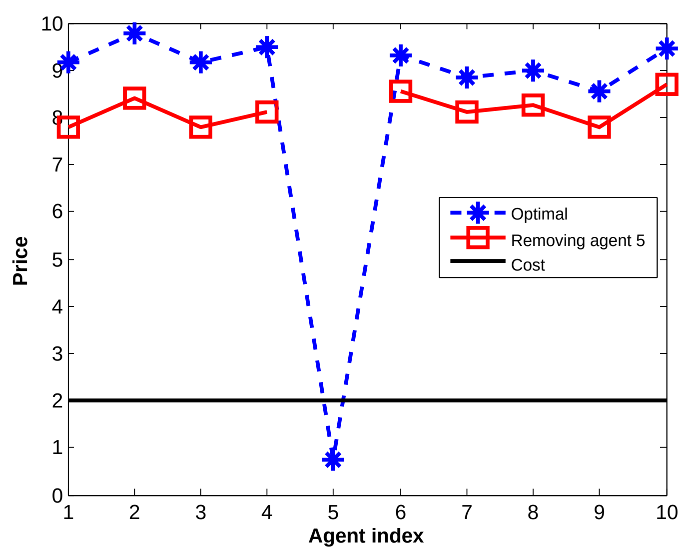

Figure 2 Centrality effect: losing money on an influential agent.

than the cost. This illustrates the fact that agent 5 has a central and influential position in the network and therefore, the seller should strongly incentivize this player. In particular, the optimal algorithm identifies this feature and captures the fact that it is profitable to offer a very low price to this person so that he can influence other people about the product. This way, the seller loses some small amount of money on the influential agent but is able to extract higher profits on his neighbors. We now compare this to an alternate strategy where the seller decides to remove agent 5 from the network due to his low valuation. In this case, all the optimal prices are decreased and the overall profit drops from 63.52 to 55.5 units so that one can increase profits by about 14.5% by including player 5.

Value of incorporating incentives that guarantee influence: In Fig. 3, we compare the optimal solution for discriminative prices to the extended model introduced in Section 6 where the seller offers a uniform regular price ($p=4$) and designs discriminative discounts in exchange of some action to influence. In this instance, every agent is randomly connected to three other agents with $g_{ji}=0.75$ for any connected edge and $g_{ii}=1.5R$ where $R=U[1,4.5]$. We assume $t_i=U[0,1]\ \forall\ i\neq 1$, $t_1=6.9$ and $c=1$.

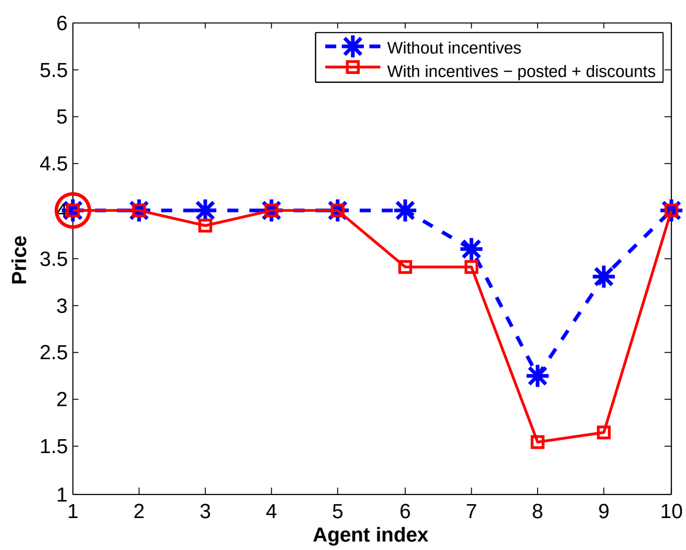

Figure 3 Value of incorporating incentives that guarantee influence.

We observe that the total profit using the earlier model (without incentives to influence i.e., $t_i=0$) is equal to 27.15. This profit is not guaranteed because some agents may not influence their peers. In particular, in this example, suppose agents 5 and 10 who buy at full price do not influence their neighbors which includes agent 1. Agent 1 ends up not purchasing the item and consequently does not influence his neighbors either. Finally, it so happens that only agents 2, 5 and 10 buy the item yielding a profit of 9 as opposed to 27.15. Consequently, the earlier model predicts a value for the profits that is much higher than the realized one even if a few agents do not influence. On the other hand, in the model with incentives that guarantee influence (t_i is taken into account), the total profits are equal to 20.85 and agent 1 does not purchase the item. Observe that this is lower than 27.15 but way larger than 9. Therefore, the model with incentives provides the seller with the flexibility of using prices together with incentives that result in a higher degree of confidence on the predicted profits.

Symmetric agents with asymmetric incentives: In Fig. 4, we present a setting with symmetric agents who receive asymmetric incentives to influence their neighbors. In this instance, every agent not only has the same number of neighbors but also the same self and cross valuations. In particular, we consider a complete graph with $g_{ii} = 1.3$ and $g_{ij} = 0.3$, a cost to influence $t_i = 2.2$ and $c = 0.2$. We also assume that the item has a posted price equal to 3. We compute the optimal discriminative prices which happen to be at 3 for everyone and compare them to the case where the seller designs incentives to guarantee influence by offering two prices using problem Zi-MIP. Interestingly, the optimal solution for the model with incentives is not symmetric despite the fact that all the agents are homogenous. Indeed, it is sufficient for the seller to incentivize 6 out of the 10 agents in the network (no matter which group of 6). These 6 agents receive a targeted discount to influence their peers that purchase at the full posted price.

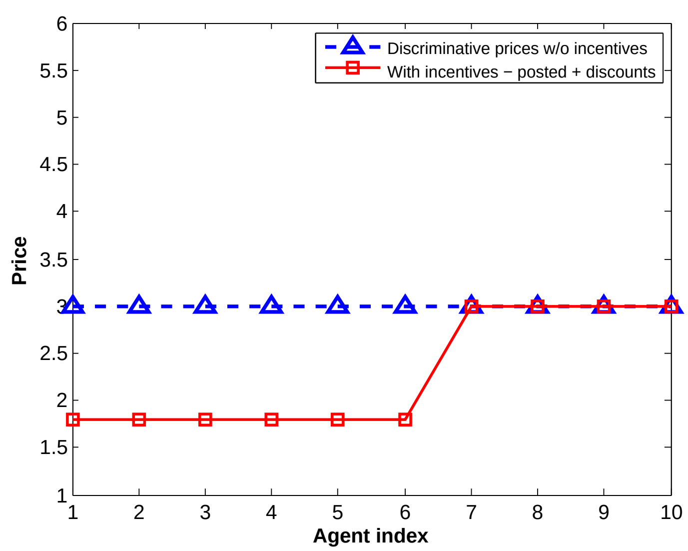

Figure 4 Symmetric Graph with asymmetric incentives: with and without incentives.

Effect of network topology on optimal prices: In Fig. 5, we consider different network topologies and compare the optimal discriminative prices as well as the corresponding profits. In all the scenarios, $g_{ii} = 1.5R$ where $R = U[1,2]$, $g_{ij} = 0.75$ when agent i influences agent j and 0 otherwise and $c = 2$. For each network topology, we solve the optimal discriminative prices using the relaxation of Z-MIP. We plot the optimal price vector for the different networks in Fig. 5. We observe that in our example, all the agents always decide to purchase the item. In the complete graph, all the nodes are connected to each other and therefore the profits are the highest and equal to 70.15. In the intermediate topology where each agent has three neighbors, the total profits are equal to 22.45. The cycle graph is a network where the nodes are connected in a circular fashion, where each agent has one ingoing and one outgoing edge (influences one agent and influenced by one). In this case, the total profits are equal to 8.95. Star 1 and star 2 are star graphs with a central agent being agent 5. In star 1, agent 5 influences all the other agents and in star 2 agent 5 is influenced by all the others. In both cases, the profits are equal to 8.2. This is interesting to observe that both star networks yield the same profits as the total valuations in the system are the same. In star1, agent 5 receives a small discount to influence so that the prices of the others are slightly higher. In star 2, the prices of all the agents but 5 are slightly lower so that the seller can charge a high price to agent 5. As we observe the prices for the different network topologies, we note that the value of the prices and the profits increase with the number of edges in the graph. Indeed, each additional edge corresponds to an agent increasing another agent's willingness-to-pay and therefore the more the graph is connected, the larger are the profits.

8. Conclusions

In this paper, we present an optimal pricing model for a profit maximizing firm that sells an indivisible item to agents embedded in a social network who interact with each other and positively influence each others' purchasing decisions. We model the problem as a two stage game where the seller first offers prices and the agents collectively follow with their purchasing decisions by taking into account their neighbors influences. Using equilibrium existential properties that we prove, duality theory and techniques from integer programming, we reformulate the two stage pricing problem as a MIP formulation with linear constraints. We view this MIP as an operational pricing tool that any firm can use by incorporating various business rules on prices and constraints on network segmentation. The main advantage of this formulation is that it allows us to cast the problem into the traditional optimization framework, where one can explore and exploit various advancements in optimization techniques. For the case of discriminative and uniform pricing strategies, we present efficient methods to optimally solve the MIP that are polynomial in the number of agents using its LP relaxation.

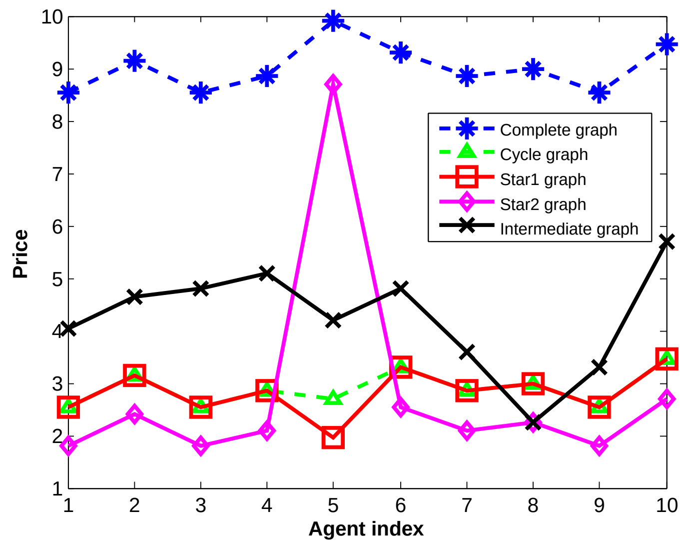

Figure 5 Optimal prices for various network topologies.

We extend our proposed model and results to the case when the seller can design both prices and incentives to guarantee influence amongst agents. This is because, in general, agents that buy need not necessarily influence their peers. The seller can use incentives in exchange for an action such as an endorsement, a wall post or a review to guarantee influence. Finally using computational experiments, we highlight the benefits of incorporating network externalities, compare the different pricing strategies and the more general model with incentives. In particular, we also show that sometimes it is beneficial for the seller to earn negative profit on an influential agent in order to extract significant positive profits on others.

As a part of future work, it would be interesting to study (both analytically and computationally) the difference in profits between the models with and without incentives depending on the input parameters. The optimization framework for optimal pricing presented in this paper allows one to explore decomposition techniques for other complex pricing strategies, and stochastic and robust optimization methods to handle partially observable noisy social network data.

Appendix A: Proof of Theorem 1

We use a constructive argument to show the existence of a PNE strategy. We construct the equilibrium from a mixed strategy Nash equilibrium which we know exists by appealing to the Nash's Existence Theorem for finite games. Note that the second stage purchasing game is a finite game because both the number of players and the strategy space are finite. A mixed Nash equilibrium in our setting refers to agents mixing between two pure strategies: buying and not buying. Mathematically, this is equivalent to the case where the variable α_i can take fractional values i.e., $0 \leq \alpha_i \leq 1$

in problem (2.4). We next show how one can construct a PNE from any given mixed Nash equilibrium. The following steps provide our proposed iterative constructive method, where t denotes the iteration number:

1. Let $t=1$ and $D_t=\Phi$ (the empty set).
2. Fix the strategy of the players in the set D_t to fully purchase the item i.e., $\alpha_i^t=1\ \forall i\in D_t$.
3. From the Nash's Existence Theorem for finite games, we know that the purchasing game with the N players imposing $\alpha_i^t=1\ \forall i\in D_t$ has a mixed Nash equilibrium strategy, say $\alpha_i^{t,*}\ \forall i\in\mathcal{I}$.
4. If $\alpha_i^{t,*}$ are integer for all $i\in\mathcal{I}$, a PNE strategy with $\alpha_i^{t,*}$ is identified and the procedure terminates. If not, let $S_t=\{i\in\mathcal{I}\backslash D_t|\alpha_i^{t,*}>0\}$ (note that $S_t\neq\Phi$).
5. Update the set D_t as follows: $D_{t+1}:=D_t\cup S_t$ and repeat from step 2 by setting $t:=t+1$.

We first note that at each iteration of the procedure, at least one agent belongs to the set S_t and is allocated to the set D_{t+1}. Therefore, the algorithm terminates in finite time, more precisely, at most after N iterations, N being the number of agents in the network. Let $T\le N$ denote the total number of iterations. We also note that at termination, all the α's are integer. What remains to be shown is that the output of the above procedure is indeed a PNE i.e., a best response for every player $i\in\mathcal{I}$. Consider all the agents in the set $\mathcal{I}\backslash D_T$ (if any). Their actions $\alpha_i^{T,*}$ are best responses to all the other players because (i) they are obtained from a mixed Nash equilibrium strategy in step 3 and (ii) are integer (at termination). Now, let us consider all the agents in the set $D_T=\{i\in\mathcal{I}|\alpha_i^{T-1,*}>0\}$. By increasing the value of $\alpha_i^{T-1,*}\ \forall\ i\in D_T$ to 1, it does not decrease the payoff of agent i. More precisely, the agents with fractional values in the set D_T have a zero utility so that they are indifferent between buying and not buying. In addition, since g_{ji} are non-negative, this does not decrease the payoffs of any other player in the network. Therefore, setting $\alpha_i^{T,*}\ \forall\ i\in D_T$ to 1 is a best response for all the agents in D_T. This completes the proof of existence of a PNE strategy for the second stage game. □

Appendix B: Proof of Proposition 1

Consider the continuous relaxation of problem Z that replaces the binary constraint $\alpha_i\in\{0,1\}$ by $0\le\alpha_i\le 1\ \forall\ i\in\mathcal{I}$. Let $V^*=\left(p_i^*,y_i^*,\alpha_i^*\right)\ \forall i\in\mathcal{I}$ be an optimal solution for the relaxed problem with the corresponding objective Π^*. Assume that the latter optimal solution has at least one fractional component i.e., $\exists\ j\in\mathcal{I}$ s.t. $0<\alpha_j^*<1$. Now, consider two alternative feasible solutions, denoted by $\overline{V}$ and $\underline{V}$ to the relaxed problem Z given by:

$$\overline{V}=\left(\overline{p}_i,\overline{y}_i,\overline{\alpha}_i\right)=\begin{cases}\left(p_i^*,y_i^*,1\right) & \text{if } i=j\\ \left(p_i^*+(1-\alpha_j^*)g_{ji},y_i^*,\alpha_i^*\right) & \forall i\in S_j\\ \left(p_i^*,y_i^*,\alpha_i^*\right) & \text{otherwise}\end{cases}$$

$$\underline{V} = \left(\underline{p}_i, \underline{y}_i, \underline{\alpha}_i\right) = \begin{cases} \left(p_i^*, y_i^*, 0\right) & \text{if } i = j \\ \left(p_i^* - \alpha_j^* g_{ji}, y_i^*, \alpha_i^*\right) & \forall i \in S_j \\ \left(p_i^*, y_i^*, \alpha_i^*\right) & \text{otherwise} \end{cases}$$

where S_j denotes the set of neighbors of agent j (excluding j). We observe that both solutions are feasible to the problem for the three following reasons. First, since $0 < \alpha_j^* < 1$ it implies that $\left(g_{jj} + \sum_i \alpha_i^* g_{ij} - p_j^*\right) = 0$ otherwise it cannot be a best response for agent j and cannot satisfy the equilibrium constraints. Therefore, changing α_j^* to 1 or 0 does not affect the best response of agent j. Second, we have modified the prices of the neighbors of agent j exactly by the change in the level of influence from agent j on them. Therefore, it yields the same profit for agent $i \in S_j$ and their purchasing decisions remain the same. Third, since the agents in $\mathcal{I} \backslash (\{j\} \cup S_j)$ are unaffected by the change in α_j^* or p_i^* $\forall i \in S_j$, the solution remains feasible for them as well.

Let us denote the objective corresponding to these new solutions by $\overline{\Pi}$ and $\underline{\Pi}$ respectively. We observe that $\Pi^* - \overline{\Pi} = -(1 - \alpha_j^*)\left[(p_j^* - c) + \sum_{i \in S_j} \alpha_i^* g_{ji}\right]$ and $\Pi^* - \underline{\Pi} = \alpha_j^* \left[(p_j^* - c) + \sum_{i \in S_j} \alpha_i^* g_{ji}\right]$. Since Π^* is the optimal value of the objective and $0 < \alpha_j^* < 1$, it should be the case that $\left[(p_j^* - c) + \sum_{i \in S_j} \alpha_i^* g_{ji}\right] = 0$. Because if not, one of the solutions we constructed is strictly better than the optimal solution and this is a contradiction. Consequently, one can see that both $\overline{V}$ and $\underline{V}$ are optimal solutions as well since they are feasible and yield the same objective than V^*. In the process, we have therefore reduced the number of fractional components by one. One can now repeat the same procedure for each fractional value α_j^* to derive a constructive way of identifying a feasible integral solution to the original problem with an objective function that is as good as the initial fractional solution. Note that since the number of agents is finite, this step is repeated at most N times. One can conclude that the continuous relaxation of problem Z is tight, meaning that for any feasible fractional solution, one can find an integral solution with at least the same objective if not better. $\square$

Appendix C: Proof of Theorem 2

Consider solving the relaxed version of the problem Z-MIP where the binary constraints for each α_i $\forall$ $i \in \mathcal{I}$ are replaced by the constraint: $0 \leq \alpha_i \leq 1$ $\forall$ $i \in \mathcal{I}$. Let $V^* = \left(\alpha^*, \mathbf{p}^*, \mathbf{y}^*, \mathbf{z}^*, \mathbf{x}^*\right)$ be a fractional optimal solution to the relaxed problem with a corresponding objective Π^*. We construct a new solution with all the α_i's being integer, show its feasibility to the problem Z-MIP (hence relaxed Z-MIP as well) with an objective that is at least as good as V^*. Let us denote the solution we construct by $\tilde{V} = \left(\tilde{\alpha}, \tilde{\mathbf{p}}, \tilde{\mathbf{y}}, \tilde{\mathbf{z}}, \tilde{\mathbf{x}}\right)$ and its corresponding objective by $\tilde{\Pi}$. We construct this new solution iteratively where t denotes the iteration number:

1. Let $t = 1$, $V_t = V^*$ and $\Pi_t = \Pi^*$.

2. If V_t is fractional i.e., $\exists\ i \in \mathcal{I}$ s.t. $0 < \alpha_i^t < 1$ then proceed to step 3. Otherwise, set $\tilde{V} = V_t$, $\tilde{\Pi} = \Pi_t$ and the procedure terminates.
3. Construct a new solution V_{t+1} as follows:
   - For agent i: $\alpha_i^{t+1} = 1,\ p_i^{t+1} = \min\{p_i^t, g_{ii} + \sum_{j \in \mathcal{I}\setminus i} \alpha_j^t g_{ji}\},\ z_i^{t+1} = p_i^{t+1},\ y_i^{t+1} = \max\{0, g_{ii} + \sum_{j \in \mathcal{I}\setminus i} \alpha_j^t g_{ji} - p_i^t\}$.
   - For the agents j that are influenced by i ($g_{ij} > 0$): $\alpha_j^{t+1} = \alpha_j^t,\ p_j^{t+1} = p_j^t + (1-\alpha_i^t) g_{ij},\ y_j^{t+1} = y_j^t,\ z_j^{t+1} = z_j^t + (\alpha_j^t - x_{ij}^t) g_{ij},\ x_{ij}^{t+1} = \alpha_j^t$.
   - For the remaining agents, keep all the variables at the same values as in V_t.

   Let us denote the objective value corresponding to this solution by Π_t.
4. Proceed back to step 2 after setting $t := t+1$.

Observe that the above procedure updates exactly one fractional α_i from the initial solution V^* at each step until the point where all the α_i's are integer. Therefore, the algorithm terminates in at most N steps (N being the number of agents) and the final solution is such that all the α_i's are integer.

We next show that the solution $\tilde{V}$, obtained at the end of the above procedure, is feasible to the problem Z-MIP while the objective $\tilde{\Pi} \geq \Pi^*$. We proceed by showing the following 2 steps. First at each iteration of the procedure, we have a solution that remains feasible to a portion of the relaxed Z-MIP problem (and therefore may be infeasible to the relaxed Z-MIP), yet it increases the objective function i.e., $\Pi_{t+1} \geq \Pi_t$ at each iteration. Second, at termination we have a solution denoted by $\tilde{V}$ that is feasible to the relaxed Z-MIP. Note that at termination, the solution $\tilde{V}$ integer and feasible (shown below) to the relaxed Z-MIP and therefore also feasible to Z-MIP.

Let us now show the feasibility condition. We use an induction argument to show that at any iteration t, given that V_t satisfies constraints (4.1), constraints (4.3), constraints (4.2) and not including the last constraint of (4.2), the new constructed solution V_{t+1} also satisfies this set of constraints. We know that $V_1 = V^*$ satisfies all the constraints of the relaxed Z-MIP. Consider any iteration $t > 1$ and the step 3 where the solution V_{t+1} is constructed. For all the agents that were not influenced by agent i, we did not modify any of the variables and hence all their constraints remain feasible. Now let us consider agent i and agents j influenced by i and show feasibility. One can see that $x_{ji}^{t+1} = x_{ij}^{t+1} = \alpha_j^{t+1} = \alpha_j^t$ satisfies all the constraints (4.3) on the variables x_{ij}. Because $\alpha_i^{t+1} = 1$ and $z_i^{t+1} = p_i^{t+1}$, the constraints (4.2) on z_i are satisfied as well. Observe that in the two cases for agent i, y_i^{t+1} is chosen so that it is feasible to constraints (4.1). Now, let us consider any agent j that is influenced by i (i.e., $g_{ij} > 0$). In this case as well, the way we constructed z_j^{t+1} and p_j^{t+1} guarantees the feasibility of constraints (4.1) because the right hand sides of the constraints remain the same. We only need now to check the feasibility for the first two constraints of (4.2) for z_j. Since $(\alpha_j^t - x_{ij}^t) g_{ij} \geq 0$, we obtain: $z_j^{t+1} \geq 0$. Also, we have $z_j^{t+1} \leq p_j^{t+1}$ as $x_{ij}^t \geq \alpha_i^t + \alpha_j^t - 1$.

Finally, since $y_j^{t+1} \geq 0$, we know from the first equation of (4.1) that $z_j^{t+1} \leq \alpha_j^{t+1} g_{jj} + \sum_{k \in \mathcal{I} \setminus k} g_{kj} x_{kj}^{t+1}$ which is in turn less than $\alpha_j^{t+1}(g_{jj} + \sum_{k \in \mathcal{I} \setminus k} g_{kj}) \leq \alpha_j^{t+1} p^{max}$ by the definition of p^{max}. Therefor, this completes the induction argument.

Note that we are not guaranteed to satisfy the last constraint of (4.2) for z_j in the above induction argument. However, we can make the following two observations that guarantee its feasibility at termination. First, whenever $\alpha_j^t = 0$, $z_j^{t+1} = z_j^t = 0$ because $x_{ij}^t = 0$. This implies for all the agents such that $\alpha_j^t = 0$, the last constraint of (4.2) for z_j is always feasible. Second, whenever $\alpha_j^t = 1$, $z_j^t = p_j^t$ and $x_{ij}^t = \alpha_i^t$. Therefore, in the solution V_{t+1}, $z_j^{t+1} = z_j^t + (1 - x_{ij}^t) g_{ij} = p_j^t + (1 - \alpha_i^t) g_{ij} = p_j^{t+1}$ and consequently, the last constraint of (4.2) for z_j is always feasible whenever $\alpha_j^t = 1$ as well. Since all the α_j's are integer at termination, we have a solution that satisfies the last constraint of (4.2) for z_j. Combining this with the induction argument above, we constructed a new solution $\tilde{V}$ that is feasible for the relaxed Z-MIP problem. Interestingly, we started with an initial feasible solution V^* and went to multiple solution update steps (at most N) where the solution is feasible only to a portion of the relaxed Z-MIP problem (and therefore may be infeasible to the relaxed Z-MIP) but terminated at a solution $\tilde{V}$ that is feasible to the entire relaxed Z-MIP.

What remains to be shown is that every update step results in an objective satisfying $\Pi_{t+1} \geq \Pi_t$. For simplicity, we present the proof for the case with $c = 0$ and we will briefly explain how to extend it for the more general case with $c > 0$ at the end. The difference of profits in two successive iterations is given by: $\Pi_{t+1} - \Pi_t = z_i^{t+1} - z_i^t + \sum_{j \in S_i} (\alpha_j^t - x_{ij}^t) g_{ij}$. Indeed, the profit induced by the agents not influenced by i are the same. We know from Step 3 of the procedure that: $z_i^{t+1} = \min\{p_i^t, g_{ii} + \sum_{j \in \mathcal{I} \setminus i} \alpha_j^t g_{ji}\}$. In addition, from the feasibility of z_i^t, we have: $z_i^t \leq \min\{p_i^t, \alpha_i^t p^{max}, \alpha_i^t g_{ii} + \sum_{j \in \mathcal{I} \setminus i} x_{ji}^t g_{ji}\}$. Note that $\alpha_i^t p^{max} \geq \alpha_i^t (g_{ii} + \sum_{j \in \mathcal{I} \setminus i} g_{ji})$ by the definition of p^{max} and the latter is greater than $\alpha_i^t g_{ii} + \sum_{j \in \mathcal{I} \setminus i} x_{ji}^t g_{ji}$. Therefore we obtain, $z_i^t \leq \min\{p_i^t, \alpha_i^t g_{ii} + \sum_{j \in \mathcal{I} \setminus i} x_{ji}^t g_{ji}\}$. Observe also that $\alpha_i^t g_{ii} + \sum_{j \in \mathcal{I} \setminus i} x_{ji}^t g_{ji} < g_{ii} + \sum_{j \in \mathcal{I} \setminus i} \alpha_j^t g_{ji}$ because $0 < \alpha_i^t < 1$. This implies that $z_i^{t+1} - z_i^t \geq 0$. Since we also have: $(\alpha_j^t - x_{ij}^t) \geq 0 \ \forall \ j \in \mathcal{I}$ one can conclude that $\Pi_{t+1} - \Pi_t \geq 0$ and this completes the proof.

We now briefly discuss how to extend the proof for the case where $c > 0$. First, if $c < \min_i g_{ii}$, then it is never profitable to earn negative profits on any agent. In this case, the exact same proof applies by the following transformation of variables: reduce all the g_{ii}'s by c and assume $c = 0$. Second, let us consider the more general case where the cost c can be larger than some of the self-valuations g_{ii}. In this case, it may be profitable for the seller to lose some money (earning negative profits) on influential agents in order to extract some significant positive profits on others. Therefore, by increasing α_i to 1 for such an agent, the same proof cannot apply in a straightforward manner. However, if we start with a fractional solution where it was profitable to lose some money on agent i, it means that the profits earned from his neighbors due to the contribution of agent i are larger than the loss incurred by agent i. Consequently, by increasing further α_i to 1, it becomes even more profitable and therefore it is optimal to have $\alpha_i = 1$ in this case too. We note that the formal above proof for the case $c = 0$ can be slightly modified to handle the case $c > 0$ at the expense of a more cumbersome notation and is in a similar spirit than the proof of proposition 1. $\square$

Appendix D: Proof of correctness of Algorithm 1

First, we note that after each iteration of the procedure, at least one agent is removed from the network. Therefore, the algorithm clearly terminates in finite time, more precisely, at most after N iterations. Let us denote by $T(\leq N)$ the total number of iterations and by $N^{(t)}$ the number of agents in the network at iteration $t = 1, 2, \ldots, T$.

Next, we show that the only candidates for the optimal uniform price are $p_{min}^{(t)}$ $\forall\ t \in \{1, \ldots, T\}$. The first observation is that the uniform optimal price cannot be smaller than $p_{min}^{(1)}$. Indeed, for any price $p \leq p_{min}^{(1)}$, all the agents that bought in the discriminative case will still buy at this smaller price. But a lower price than $p_{min}^{(1)}$ will result in a lower profit to the seller from these buyers. It is possible though that some new agents buy the item at the lower price that can result in an overall higher profit. However, if this is the case it would have been profitable to offer this price to those agents in the discriminative case as well. But because it was not optimal to offer a low price than $p_{min}^{(1)}$ to the non-buyers, it is clearly not profitable to decrease the uniform price lower than $p_{min}^{(1)}$. Therefore, we conclude that the optimal uniform price cannot be smaller than $p_{min}^{(1)}$. We now consider the case where the uniform price is larger than $p_{min}^{(1)}$. In this case, we clearly lose the buyers with $p_i = p_{min}^{(1)}$ from the discriminative case. Otherwise, in the discriminative case one would offer them higher prices. We can therefore remove those agents from the network. Now applying the same argument, it is the case that the uniform optimal price cannot be equal to a value that is strictly between $p_{min}^{(1)}$ and $p_{min}^{(2)}$. By repeating this procedure, we conclude that the optimal uniform price has to be equal to one of the $p_{min}^{(t)}$ prices. In order to select the best uniform price among these T candidates, we just need to evaluate the corresponding profits (denoted by $\Pi^{(t)}$ $\forall t = 1, 2, \ldots, T$) and choose the one that yields the maximal profits. One can do so by using the following relation:

$$\Pi^{(t)} = \left(p_{min}^{(t)} - c\right) \sum_{i=1}^{N^{(t)}} \alpha_i^{(t)}, \tag{D.1}$$

where $N^{(t)}$ is the remaining number of agents in the network at iteration t. $\square$

Appendix E: Proof of Proposition 2

Consider the continuous relaxation of problem Zi that replaces the binary constraint $\alpha_i \in \{0, 1\}$ by $0 \leq \alpha_i \leq 1$ $\forall\ i \in \mathcal{I}$. We consider the version of problem Zi without the dual variables w_i (see

Observation 2). Let $V^* = \left(p_i^*, d_i^*, y_i^*, \alpha_i^*, \gamma_i^*\right)$ $\forall i \in \mathcal{I}$ be an optimal solution for the relaxed problem with the corresponding objective Π^*. We divide the proof into two parts. First, we show that given any optimal solution, one can construct a new optimal solution for which all the variables α_i^* $\forall i \in \mathcal{I}$ are integer. Second, we construct from the latter solution a new solution with all the variables γ_i^* $\forall i \in \mathcal{I}$ integer as well. Assume that the initial optimal solution has at least one fractional component i.e., $\exists\, j \in \mathcal{I}$ s.t. $0 < \alpha_j^* < 1$. Now, consider three other feasible solutions $\tilde{V}, \overline{V}$ and $\underline{V}$ to the relaxed problem as follows: *Zi* given by:

$$\tilde{V} = \left(\tilde{p}_i, \tilde{d}_i, \tilde{y}_i, \tilde{\alpha}_i, \tilde{\gamma}_i\right) = \begin{cases} \left(p_i^*, d_i^*, y_i^*, 1, \gamma_i^*\right) & \text{if } i = j \\ V_i^* & \forall i \in S_j \\ V_i^* & \text{otherwise} \end{cases}$$

$$\overline{V} = \left(\overline{p}_i, \overline{d}_i, \overline{y}_i, \overline{\alpha}_i, \overline{\gamma}_i\right) = \begin{cases} \left(p_i^*, d_i^*, 0, 1, 1\right) & \text{if } i = j \\ \left(p_i^* + (1-\gamma_j^*) g_{ji}, d_i^*, y_i^*, \alpha_i^*, \gamma_i^*\right) & \forall i \in S_j \\ V_i^* & \text{otherwise} \end{cases}$$

$$\underline{V} = \left(\underline{p}_i, \underline{d}_i, \underline{y}_i, \underline{\alpha}_i, \underline{\gamma}_i\right) = \begin{cases} \left(p_i^*, d_i^*, 0, 1, 0\right) & \text{if } i = j \\ \left(p_i^* - \gamma_j^* g_{ji}, d_i^*, y_i^*, \alpha_i^*, \gamma_i^*\right) & \forall i \in S_j \\ V_i^* & \text{otherwise} \end{cases}$$

where S_j denotes the set of neighbors of agent j (excluding j). We observe that all three solutions are feasible to the problem for the following reasons. First, since $0 < \alpha_j^* < 1$ it implies that $\big(g_{jj} + \sum_i \gamma_i^* g_{ij} - p_j^*\big) = 0$ as otherwise it cannot be a best response for agent j and cannot satisfy the equilibrium constraints. In addition, to ensure feasibility, we should have either $\gamma_j^* = 0$ or $d_j^* = t_j$ and therefore $y_j^* = 0$. Therefore, changing α_j^* to 1 or 0 does not affect the best response of agent j as far as α_j is concerned. Note that we construct the dual variable for agent j to satisfy all the feasibility constraints. Second, we have modified the prices of the neighbors of agent j exactly by the change in the level of influence from agent j on them and therefore the purchasing decisions for agents $i \in S_j$ remain the same. Third, since the agents in $\mathcal{I} \backslash (\{j\} \cup S_j)$ are unaffected by the change in α_j^* or p_i^* $\forall i \in S_j$, the solution remains feasible for them as well.

Let us denote the objective corresponding to these new solutions by $\tilde{\Pi}$, $\overline{\Pi}$ and $\underline{\Pi}$ respectively. We observe that $\Pi^* - \tilde{\Pi} = -(1-\alpha_j^*)(p_j^* - c)$. Since, V^* is an optimal solution, it has to be the case that $p_j^* - c \le 0$ because otherwise $\tilde{V}$ is a better solution. In addition, we observe that $\Pi^* - \overline{\Pi} = -(1-\alpha_j^*)(p_j^* - c) + (1-\gamma_j^*) d_j^* - \sum_{i \in S_j} \alpha_i^* g_{ji} (1-\gamma_j^*)$ and $\Pi^* - \underline{\Pi} = \alpha_j^*(p_j^* - c) - \gamma_j^* d_j^* + \sum_{i \in S_j} \alpha_i^* g_{ji} \gamma_j^*$. Since Π^* is the optimal value of the objective and $0 < \alpha_j^* < 1$, it should be the case that both $\overline{\Pi}$ and $\underline{\Pi}$ are lower or equal than Π^*. By requiring $\Pi^* - \overline{\Pi} \ge 0$ together with $\Pi^* - \underline{\Pi} \ge 0$ and using the fact that $p_j^* - c \le 0$, we obtain the condition: $\alpha_j \ge \gamma_j$. However, from the feasibility constraint, we know that $\alpha_j \le \gamma_j$ and therefore $\alpha_j = \gamma_j$. By using this fact, we obtain: $\Pi^* - \overline{\Pi} =$

$-(1-\alpha_j^*)\left(p_j^*-c-d_j^*+\sum_{i\in S_j}\alpha_i^* g_{ji}\right)$ and $\Pi^*-\underline{\Pi}=\alpha_j^*\left(p_j^*-c-d_j^*+\sum_{i\in S_j}\alpha_i^* g_{ji}\right)$. Since $0<\alpha_j^*<1$, it has to be the case that both $\overline{V}$ and $\underline{V}$ are optimal solutions as well since they are feasible and yield the same objective than V^*. We therefore have reduced the number of fractional components by one. One can now repeat the same procedure for each fractional value α_j^* to derive a constructive way of identifying a feasible integral solution to the original problem with an objective function that is as good as the initial fractional solution. Note that since the number of agents is finite, this step is repeated at most N times. One can conclude that the continuous relaxation of problem Zi is tight, meaning that for any feasible fractional solution, one can find an integral solution with at least the same objective if not better.

We now know that there exists an optimal solution with α_i^* integer $\forall i\in\mathcal{I}$. We next show the following result that allows to guarantee the integrality of γ_i^* $\forall i\in\mathcal{I}$ at optimality. In other words, it is optimal for each buyer to either fully influence (i.e., $\alpha_i^*=\gamma_i^*=1$) and receive the full discount or not to influence at all (i.e., $\gamma_i^*=0$) and pay the full price. Consider the optimal integer purchasing decisions α_i^* $\forall i\in\mathcal{I}$. For all the agents k with $\alpha_k^*=0$, it is clear from feasibility that $\gamma_k^*=0$. Consider a given optimal solution denoted by V^* with $\alpha_j^*=1$ and assume by contradiction that $0<\gamma_j^*<1$. Consider two alternative feasible solutions $\overline{V}$ and $\underline{V}$ to the relaxed problem Zi given by:

$$\overline{V}=\left(\overline{p}_i,\overline{d}_i,\overline{y}_i,\overline{\alpha}_i,\overline{\gamma}_i\right)=\begin{cases}\left(p_i^*,d_i^*,y_i^*,1,1\right) & \text{if } i=j\\ \left(p_i^*+(1-\gamma_j^*)g_{ji},d_i^*,y_i^*,\alpha_i^*,\gamma_i^*\right) & \forall i\in S_j\\ V_i^* & \text{otherwise}\end{cases}$$

$$\underline{V}=\left(\underline{p}_i,\underline{d}_i,\underline{y}_i,\underline{\alpha}_i,\underline{\gamma}_i\right)=\begin{cases}\left(p_i^*,d_i^*,y_i^*,1,0\right) & \text{if } i=j\\ \left(p_i^*-\gamma_j^* g_{ji},d_i^*,y_i^*,\alpha_i^*,\gamma_i^*\right) & \forall i\in S_j\\ V_i^* & \text{otherwise}\end{cases}$$

where S_j denotes the set of neighbors of agent j (excluding j). We observe that both solutions are feasible to the problem for the following reasons. We first note that we construct the dual variables for agent j to satisfy all the feasibility constraints. Indeed, since $0<\gamma_j^*<1$, it has to be the case that $d_j^*=t_j$ as otherwise it cannot be a best response for agent j and cannot satisfy the equilibrium constraints. Second, we have modified the prices of the neighbors of agent j exactly by the change in the level of influence from agent j on them and therefore the purchasing decisions for agents $i\in S_j$ remain the same. Third, since the agents in $\mathcal{I}\backslash(\{j\}\cup S_j)$ are unaffected by the change in α_j^* or p_i^* $\forall i\in S_j$, the solution remains feasible for them as well.

Let us denote the objective corresponding to these new solutions by $\overline{\Pi}$ and $\underline{\Pi}$ respectively. We observe that $\Pi^*-\overline{\Pi}=(1-\gamma_j^*)\left[d_j^*-\sum_{i\in S_j}\alpha_i^* g_{ji}\right]$ and $\Pi^*-\underline{\Pi}=-\gamma_j^*\left[d_j^*-\sum_{i\in S_j}\alpha_i^* g_{ji}\right]$. Since Π^* is the optimal value of the objective and $0<\gamma_j^*<1$, it should be the case that $d_j^*-\sum_{i\in S_j}\alpha_i^* g_{ji}=0$. Because if not, one of the solutions we constructed is strictly better than the optimal solution and this is a contradiction. Consequently, one can see that both $\overline{V}$ and $\underline{V}$ are optimal solutions as well since they are feasible and yield the same objective than V^*. In the process, we have therefore reduced the number of fractional components by one. One can now repeat the same procedure for each fractional value γ_j^* to derive a constructive way of identifying a feasible integral solution to the original problem with an objective function that is as good as the fractional solution. Note that since the number of agents is finite, this step is repeated at most N times. In conclusion, the continuous relaxation of problem Zi always has an optimal solution such that not only the purchasing decisions are integer but the variables γ are integer too. $\square$

Acknowledgments

The first author would like to thank IBM Research as this work would not have been possible without an internship at the IBM T. J Watson Research Center during the summer of 2012. The authors would also like to thank Georgia Perakis, Markus Ettl, Adam Elmachtoub and Itai Ashlagi for their valuable feedback that has helped us improve the paper.

References

Akhlaghpour, H., M. Ghodsi, N. Haghpanah, V. Mirrokni, H. Mahini, A. Nikzad. 2010. Optimal iterative pricing over social networks. *WINE. Lecture Notes in Computer Science* **6484** 415–423.

Arthur, D., R. Motwani, S. Aneesh, Y. Xu. 2009. Pricing strategies for viral marketing on social networks. *The 5th International Workshop on Internet and Network Economics*. 101–112.

Ballester, C., A. Calvo-Armengol, Y. Zenou. 2006. Who's who in networks. wanted: The key player. *Econometrica* **74**(5) 1403–1417.

Banerji, A., B. Dutta. 2009. Local network externalities and market segmentation. *International Journal of Industrial Organization* **27** 605–614.

Bertsimas, D., R. Weismantel. 2005. *Optimization over integers*, vol. 13. Dynamic Ideas Belmont.

Bloch, F., N. Querou. 2013. Pricing in social networks. *Games and economic behavior* **80** 243–261.

Candogan, O., K. Bimpikis, A. Ozdaglar. 2012. Optimal pricing in networks with externalities. *Operations Research* **60** 883–905.

Chen, W., P. Lu, X. Sun, B. Tang, Y. Wang, Z. A. Zhu. 2011. Optimal pricing in social networks with incomplete information. *The 7th Workshop on Internet & Network Economics* .

Domingos, P., M. Richardson. 2001. Mining the network value of customers. *Proceedings of the 7th Conference on Knowledge Discovery and Data Mining* 57–66.

Farrell, J., G. Saloner. 1985. Standardization, compatibility and innovation. *RAND Journal of Economics* **16** 70–83.

Fudenberg, D., J. Tirole. 1991. *Game theory*. MIT press.

Galeotti, A., S. Goyal, M. Jackson, F. Vega-Redondo, L. Yariv. 2010. Network games. *Review of Economics Studies* **77**(1) 218–244.

Gunnec, D., S. Raghavan. 2012. Integrating social network effects in product design. *Submitted to Management Science* .

Hartline, J., V. S. Mirrokni, M. Sundararajan. 2008. Optimal marketing strategies over social networks. *Proceedings of the 17th International Conference on World Wide Web* .

Iyengar, R., C. Van den Bulte, T. W. Valente. 2011. Opinion leadership and social contagion in new product diffusion. *Marketing Science* **30**(2) 195–212.

Katz, M., C. Shapiro. 1985. Network externalities, competition and compatibility. *American Economic Review* **75** 424–440.

Kempe, D., J. Kleinberg, E. Tardos. 2003. Maximizing the spread of influence through a social network. *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* .

Luo, Z. Q., J. S. Pang, D. Ralph. 1996. *Mathematical programs with equilibrium constraints*. Cambridge University Press.

Nemhauser, G. L., L. A. Wolsey. 1988. *Integer and combinatorial optimization*, vol. 18. Wiley New York.

Sundararajan, A. 2007. Local network effects and complex network structure. *The BE Journal of Theoretical Economics* **7**(1) Article 46.