Designing Price Incentives in a Network with Social Interactions

Maxime C. Cohen
NYU Stern School of Business, New York, NY 10012, maxime.cohen@stern.nyu.edu

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

Problem definition: The recent ubiquity of social networks allows firms to collect vast amount of data on their customers and on their social interactions. We consider a setting where a monopolist sells an indivisible good to consumers who are embedded in a social network.

Academic/Practical relevance: This an important problem as sellers can use available data to design and send targeted promotions that account for social externality effects and ultimately increase their profits.

Methodology: We capture the interactions among consumers using a broad class of non-linear utility models. This class extends the existing models by explicitly capturing externalities from subsets of agents (communities or groups) and includes several existing models as special cases (e.g., full information version of the triggering model). Assuming complete information about the interactions, we model the optimal pricing problem as a two-stage game. First, the firm designs prices to maximize profits and then consumers choose whether or not to purchase the item.

Results: Under positive network externalities, we show the existence of a pure Nash equilibrium that is preferred by both the seller and the buyers. Using duality theory and integer programming techniques, we reformulate the problem into a linear mixed-integer program (MIP). We derive efficient ways to optimally solve the MIP using its linear programming relaxation for two pricing strategies: discriminative and uniform. Finally, we propose two intuitive heuristic algorithms to solve the problem for which we derive worst-case parametric performance bounds.

Managerial implications: We draw interesting insights on the structure of the optimal prices and the seller's profit. In particular, we quantify the effect on prices when using a non-linear utility model relative to a linear model and identify settings when it is beneficial to offer a price below cost to influential agents. Finally, we extend our model and results to the case where the seller offers incentives (in addition to prices) to solicit actions so as to ensure network externality effects.

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

---

1. Introduction

The recent ubiquity of social networks has 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 on their customers, their buying behavior, and their social interactions. The challenge faced by every firm is how to process this data and turn it to actionable policies to improve their competitive advantage. In this paper, we focus on designing effective pricing strategies to enhance the profits of a 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. In recent times, word-of-mouth communication is just as likely to arise from social networks as from a neighbor across the fence. Consultants at The Conversation Group report that 65% of consumers who receive a recommendation from a contact on their social media have purchased the recommended product. In particular, personalized referrals from friends and family have been more effective in encouraging such purchases. Further, 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 are influenced by their reference groups (see e.g., Iyengar et al. 2011). The previous examples clearly indicate 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/her Facebook or Yelp page, s/he does influence his/her peers to purchase the same item and increases their valuation. The valuation increase can sometimes be non-linear. For example, if many other people have already recommended this item, the marginal effect of a new recommender can have a relatively small increase. Alternatively, sometimes having a single friend who buys and recommends the item may be sufficient to trigger a significant increase in valuation.

An important feature of the products or services we consider in this paper is the *local (non-linear) positive externalities*. This means that people *positively* influence each other's willingness-to-pay for an item (the setting with negative externalities is briefly discussed in Section 6.2). In addition, the item becomes more valuable to a person if many of his/her friends buy it even though there can be a decreasing (or increasing) marginal effect (submodular or supermodular). Examples of products with such effects include smartphones, tablets, fashion items, and cell phone plan subscriptions. Such positive externalities may be even more significant when a new generation of products is introduced to 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) 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 valuable for firms to identify these influential agents. As an example, 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 14.68M of fans who liked their Facebook page (February 2018). These fans can claim offers via the social platform and thus, directly influence their friends about purchasing. This interaction between the seller and the fan club allows the seller to keep its fan club engaged and to identify influencers. Ultimately, the seller can offer personalized prices or incentives to these influencers to increase the overall profitability.

In this paper, our goal is to develop a model that incorporates local non-linear externalities among potential buyers and design efficient algorithms to compute the optimal prices that maximize the seller's profit. We formulate the optimal pricing problem as a two-stage game between the seller and the agents in the network where the seller first offers prices and the agents then choose whether or not to purchase the (indivisible) item. The main contributions of the paper are as follows:

1. **Non-linear additive utility models with externalities.** We introduce and study a class of additive utility models. This class extends existing models by explicitly capturing externalities from subsets of buyers (communities or groups) and allowing a threshold on the number of agents needed to trigger the externality effect. Several commonly used models in the literature are included as special cases under full information (e.g., independent cascade model, linear threshold, and triggering model). In particular, the total value earned by an agent when purchasing the item is the sum of his/her own valuation and the valuation derived from externalities of all subsets of friends (see Section 2.1). This is a broad class of non-linear utility functions that can capture different influence structures, including special cases of supermodular and submodular, with respect to the number of neighbors.
2. **Reformulating the optimal pricing problem into an operational MIP.** The strategic complements nature of the second stage game guarantees the existence of a pure strategy Nash equilibrium (under non-negative externalities). Using duality theory, we derive equilibrium constraints and reformulate the two-stage problem into a non-convex integer program. We then transform it into an equivalent MIP using integer programming reformulation techniques. This resulting MIP holds under general externalities (positive or negative) and can be viewed as an operational pricing tool where one can easily incorporate business rules on prices.
3. **Efficient optimal algorithms for discriminative and uniform pricing strategies.** We develop efficient and scalable methods to optimally solve the MIP for two pricing strategies using the linear programming (LP) relaxation. We consider discriminative and uniform pricing strategies and present a solution method that is efficient (polynomial in the number of agents) and scalable to large networks. We also propose two heuristic algorithms (the *greedy expansion* and the *greedy removal*) that are intuitive and easy to implement; and we derive parametric bounds on their worst-case performance under non-negative externalities.

4. **Insights on the structure of the optimal pricing strategy.** Under discriminative prices, we show that the price of a buyer can be expressed as the sum of his/her own value and a markup term corresponding to the network externalities of agents who buy the item. Therefore, prices for influenced agents are higher (as expected). The seller's profit from network externalities comes from two types of agents: (i) high-valued customers who influence their neighbors, and (ii) low-valued customers who are highly influential and can sometimes be offered a price below cost. In addition, when comparing linear to non-linear utility models, we show that as we move from a linear model (with only pairwise interactions) to a non-linear one (that includes subsets of higher sizes), additional agents will buy, buyers will pay a higher price, and the profit increases. We also convey that a larger threshold on the minimum number of agents results in a smaller number of buyers and decreases the seller's profit.
5. **Price incentives that guarantee influence.** We extend our model and results to optimally design both prices and incentives to solicit actions so as to ensure externality effects. This new model we introduce is more realistic and allows the seller to ensure that network externalities among agents occur by offering both a price and a discount (incentive) to each buyer. The buyer can then decide between: (i) not buying the item, (ii) buying the item at full price, and (iii) buying the item while claiming the discount in exchange of influence actions specified by the seller (e.g., liking the product on social media or writing a review). Interestingly, the methods and results we develop extend to this richer model.

Literature review

Models that incorporate local network externalities find their origins in the papers by Farrell and Saloner (1985) and Katz and Shapiro (1985). These early papers assume that consumers are affected by the global consumption of all other consumers. 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 only interact with a subset of agents, also known as their neighbors (or friends). Although interactions are of a local nature, the utility of each player may still depend on the global structure of the network, given that each agent potentially interacts indirectly with a much larger set of agents. Models of local network externalities which explicitly account for the network structure have been proposed in various papers (e.g., Ballester et al. 2006 and Banerji and Dutta 2009). Several recent papers explicitly model the interactions among agents in social networks to study the network effects on marketing campaigns. The first among these are the papers on influence maximization (i.e., selecting the optimal set of nodes in a social network to maximize the spread) by Domingos and Richardson (2001) and Kempe et al. (2003) which aimed to identify influential agents in a network. Hartline et al. (2008), Akhlaghpour et al. (2010), and Arthur et al. (2009) extend this line of work to study optimal pricing strategies in networks. Hartline et al. (2008) focus on viral marketing strategies for revenue maximization where agents are offered the product in a sequential manner and show that a simple two-price strategies performs 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 at each stage. Arthur et al. (2009) allow agents to buy the product with a certain probability if the product is recommended by their friends who purchased the item. 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, however, we consider a simultaneous purchasing decision for all agents in the network, who are fully rational. Since the seminal work of Kempe et al. (2003) on the influence maximization problem, several papers followed up on this topic. In Chen et al. (2010), the authors propose a new heuristic algorithm that is easily scalable to millions of nodes. The work in Borodin et al. (2010) extends the influence maximization problem to a competitive setting. The authors show that the problem becomes more challenging and that greedy approaches cannot be used anymore. Another recent paper is the work by Gunnec and Raghavan (2017) that investigates social network influence in the context of product design (share-of-choice problem). Finally, Nakkas and Xu (2015) study bargaining in two-sided supply chain networks and examine how valuation heterogeneity among manufacturers influences the equilibrium prices and the trading pattern of the supply chain network.

Our paper is also related to several studies in the marketing literature, especially in the field of social marketing (see, e.g., Andreasen 1994). Since the introduction of the Bass model, the diffusion of innovation has been an active area of research (see, e.g., Mahajan et al. 1991, and the references therein). This model assumes that potential adopters of an innovative product are influenced by mass media (called “Innovators”) and word of mouth (called “Imitators”). A large number of marketing papers conduct empirical studies on the impact of word-of-mouth in the context of social networks (see, e.g., Goldenberg et al. 2001). Another stream of papers (both in economics and marketing) study models with *indirect network effects* (INE) which postulate that the utility of the primary product increases as more complements become available. Stremersch et al. (2007) provide a good marketing literature review on this topic. In Basu et al. (2003), the authors consider a model in which the utility of a product increases with the greater availability of complementary products and show that the INE can vary by product attributes. More recently, Lovett et al. (2013) present an empirical analysis on the relationship between brands and word-of-mouth. The authors argue that consumers spread the word on brands as a result of three drivers: social, emotional, and functional. Finally, several marketing and information systems papers consider the problem of running field experiments to identify causal estimates of social influence in networks or to asses the effectiveness of viral features in marketing campaigns (see, e.g., Aral and Walker 2011).

We model the pricing problem with simultaneous purchasing decisions as a two-stage game, where the seller first sets the prices and then, agents make their purchasing decisions. Rational behavior is captured by the Nash equilibrium. Three papers in this context are closely related to our work. Candogan et al. (2012) study optimal pricing strategies for a divisible good with linear utility functions under complete network information. The authors provide efficient algorithms to compute discriminative prices, 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 good with linear utility functions under incomplete information. Our work is in the similar light of the three aforementioned papers. In particular, we study the optimal pricing problem for indivisible items under a general class of non-linear network externality models. The class of models we propose can be seen as a deterministic generalization (due to the perfect information assumption) of previous utility models (e.g., the *independent cascade* in Kempe et al. (2003) and the *triggering* model in Seeman and Singer (2013)) and can capture submodular and supermodular functions. Our models allow to explicitly capture externalities of subsets of agents (communities or groups) and a threshold on the number of agents needed to trigger the externality effect. In addition, the techniques required to address the pricing problem under general utility models differ significantly from earlier papers. In particular, one cannot derive a closed form solution for the optimal pricing problem as in Candogan et al. (2012) and Bloch and Querou (2013). Instead, the equilibrium in our setting can be characterized by a system of non-convex constraints with integer variables. We use techniques from integer programming (IP) to reformulate the optimal pricing problem with the equilibrium constraints into a linear MIP. We refer the reader to the books by Nemhauser and Wolsey (1988) and Bertsimas and Weismantel (2005) for the IP techniques used in this paper.

Structure of the paper. In Section 2, we describe our model, assumptions, and dynamics of the two-stage game. In Section 3, we show the existence of a pure strategy Nash equilibrium for the purchasing game (under non-negative externalities). We use duality theory to formulate the problem as a MIP in Section 4 and derive efficient algorithms to optimally solve it for discriminative and uniform pricing strategies in Section 5. In Section 6, we propose two heuristic algorithms and discuss the setting with general externalities (positive or negative). Section 7 extends our results to the case where the seller designs both prices and incentives to guarantee network externality effects. In Section 8, we present computational experiments. Finally, our conclusions are reported in Section 9. The proofs of the technical results are relegated to the Appendix.

2. Problem setting

2.1. Utility model with network externalities

Consider a monopolist selling an indivisible product to N agents, denoted by the set $\mathcal{I} = \{1, ..., N\}$, embedded in a social network. We denote the set of *value interaction* of agent $i \in \mathcal{I}$ by $G_i = \{g_{S,i} | S \subset \mathcal{I} \backslash \{i\}\}$, where the element $g_{S,i}$ represents the marginal increase in value that agent i obtains by owning the product, when agents in S influence him/her.[1] This excludes the network effects due to subsets of S, each of which is modeled explicitly. Concrete examples will be presented in the sequel. Note that the network considered in this paper can be represented as a weighted bipartite graph, $G(\mathcal{I}, \Delta_{\mathcal{I}}, E)$, where $\mathcal{I}$ denotes the set of agents, $\Delta_{\mathcal{I}}$ corresponds to the set of subsets of $\mathcal{I}$, and E represents the links between the groups so that the weight of a link determines the externality effect from one subset to an agent. We next define the threshold $\Gamma \geq 1$. For any set S, a minimum number of $\min\{\Gamma, |S|\}$ buyers in S are required to influence agent i. Here $|.|$ refers to the cardinality of a set. The term $g_{\emptyset,i}$ (also denoted by g_i) is the marginal value that agent i derives by owning the product. If agent j does not influence agent i, then all terms $g_{S,i}$ where $j \in S$ are zero. On the other hand, if agent j influences agent i then at least one of the terms $g_{S,i}$ where $j \in S$ is non-zero. In this case, we refer to j as a *neighbor* (or friend) of i.

ASSUMPTION 1. *We make the following assumptions regarding the elements of G_i $\forall i \in \mathcal{I}$ and the corresponding utility model.*

a. The firm and the agents have perfect knowledge on externalities i.e., everyone knows G_i.

*b. **Positive externality:** The network externality of any set S on any agent i is non-negative. That is, $g_{S,i} \geq 0$, for any $S \subset \mathcal{I} \backslash \{i\}$ and $i \in \mathcal{I}$.*

*c. **Additive model:** The total value earned by an agent when purchasing the item is the sum of his/her own valuation and the valuation derived from the network externalities of all subsets of friends who own the item. Mathematically, agent i valuation is given by:*

$$v_i(\alpha_i, \boldsymbol{\alpha_{-i}}) = \alpha_i \Big[\sum_{S \subset \mathcal{I} \backslash \{i\}} g_{S,i} \alpha_S \Big], \tag{2.1}$$

where $\alpha_i \in \{0, 1\}$ is a binary variable that represents the purchasing decision of agent i, $\boldsymbol{\alpha_{-i}}$ represents the vector of purchasing decisions of all agents but i, and α_S is a binary function that maps the purchasing decisions of agents in S to a binary indicator which is 1 if agents in S influence i, and 0 otherwise. More specifically, we have:

$$\alpha_S = \max_{\substack{S' \subset S \\ |S'| = \min\{\Gamma, |S|\}}} \Big\{ \prod_{j \in S'} \alpha_j \Big\}. \tag{2.2}$$

[1] In the context of this paper, the term *influence* refers to exerting network externalities (or network effects).

Equation (2.2) captures the fact that agents in S influence i, if at least Γ of them purchase the product. Note that we consider a setting with perfect information in order to make the analysis tractable and to draw some insights on our problem. This is a common assumption in several previous papers, such as Candogan et al. (2012). The positive externality assumption is also a commonly used condition that allows to ensure the existence of an equilibrium, as we will show in Theorem 1. The setting in which externalities can be negative is discussed in Section 6.2.

The additive model in (2.1) together with (2.2) capture a broad class of linear and non-linear utility models with network externalities. In particular, it can be seen as a deterministic generalization (due to the perfect information assumption) of the *independent cascade* and the *linear threshold* models formulated in Kempe et al. (2003) and the *triggering* model studied in Seeman and Singer (2013). The triggering model generalizes the former two models by considering that agent i is influenced by a subset of its neighbors, S (called a trigger set), if any of the agents in S purchases the product. Note that this is a special case of our model when $\Gamma = 1$. For a general value of Γ, our model can be viewed as a *subset triggering model*. More precisely, agent i is influenced by its neighbors who belong to the trigger set S, if at least Γ of them purchase the item. We next present several examples from this class of models that capture linear and non-linear effects.

1. **Pairwise externality:** Suppose $g_{S,i} = 0$ for all subsets $|S| > 1$ and $g_i, g_{j,i} \geq 0$ $\forall i, j \in \mathcal{I}$ with $\Gamma = 1$. More specifically, we have:

$$v_i(\alpha_i, \boldsymbol{\alpha_{-i}}) = \alpha_i \Big[g_i + \sum_{j \in \mathcal{I} \setminus \{i\}} \alpha_j g_{j,i} \Big]. \tag{2.3}$$

   This case corresponds to a weighted bipartite graph $G(\mathcal{I}, \Delta_{\mathcal{I}}, E)$ with $\Delta_{\mathcal{I}} = \mathcal{I}$. The above valuation function captures only the marginal externality of each neighbor and is linear and additive across neighbors. This type of models has been the focus of several earlier papers such as Candogan et al. (2012), Bloch and Querou (2013), and Chen et al. (2011).

2. **Triple-wise externality:** Suppose $g_{S,i} = 0$ for all subsets $|S| > 2$ and $g_i, g_{j,i}, g_{\{j,k\},i} \geq 0$ $\forall i, j, k \in \mathcal{I}$. In particular, we have:

$$v_i(\alpha_i, \boldsymbol{\alpha_{-i}}) = \alpha_i \Big[g_i + \sum_{j \in \mathcal{I} \setminus \{i\}} \alpha_j g_{j,i} + \sum_{j,k \in \mathcal{I} \setminus \{i\};\ j \neq k} \alpha_{j,k} g_{\{j,k\},i} \Big]. \tag{2.4}$$

   In addition, we can have either of the following: (a) if $\Gamma = 1$, $\alpha_{j,k} = \max\{\alpha_j, \alpha_k\}$, (b) if $\Gamma = 2$, $\alpha_{j,k} = \alpha_j \alpha_k$. This valuation function is a special case of a supermodular utility model, where the marginal effect of an additional neighbor increases with the set of existing influencers. Observe that these effects are characterized by externality terms that do not decompose by agents and are hence non-linear with respect to subsets of neighbors.

   Observe that in the above model, the externality effects of each subset is modeled separately. For example, consider a small network with 3 agents and $\Gamma = 2$ and suppose that all 3 agents decide to buy the item. The total network externalities of agents 2 and 3 on agent 1 are given by: $g_{2,1} + g_{3,1} + g_{\{2,3\},1}$. The first two terms correspond to the pairwise effects (i.e., how agent j affects agent i on his own), whereas the term $g_{\{2,3\},1}$ represents the additional externality of agents 2 and 3 *together* on agent 1. In particular, $g_{\{2,3\},1}$ does not incorporate the impact of its subsets ($g_{2,1}$ and $g_{3,1}$). The weighted bipartite graph $G(\mathcal{I}, \Delta_{\mathcal{I}}, E)$ for this example can be found in Figure 9 (see Appendix A).

3. **Complete neighborhood triggering model:** Suppose $g_{S,i} = 0$ for all $S \subset \mathcal{I}\backslash\{i\}$ except a single set $\mathcal{N}_i$, which represents all the neighbors of i.

$$v_i(\alpha_i, \boldsymbol{\alpha_{-i}}) = \alpha_i g_{\mathcal{N}_i,i} \max_{\substack{S' \subset \mathcal{N}_i \\ |S'| = \min\{\Gamma, |\mathcal{N}_i|\}}} \Big\{ \prod_{j \in S'} \alpha_j \Big\}.$$

In this model, an agent is influenced if and only if at least Γ of his/her neighbors buy the item. By taking $\Gamma = 1$, we obtain a special case of a submodular influence model, where only the first purchasing neighbor triggers an externality effect (and thereafter, the function does not increase with the number of purchasing neighbors). For other values of Γ, the function is neither submodular nor supermodular.

2.2. Pricing and purchasing model

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

Our goal is to develop a general and efficient optimal pricing method. For a given set of prices chosen by the seller, the agents in the network simultaneously choose their actions to maximize their utility (i.e., we consider a simultaneous game). We assume that the utility model of an agent is the total value minus the price:

$$u_i(\alpha_i, \boldsymbol{\alpha_{-i}}, p_i) = v_i(\alpha_i, \boldsymbol{\alpha_{-i}}) - \alpha_i p_i = \alpha_i \Big[\sum_{S \subset \mathcal{I}\backslash\{i\}} g_{S,i}\alpha_S - p_i \Big]. \tag{2.5}$$

If $\alpha_i = 1$, agent i purchases the item and derives a utility equal to $\sum_{S \subset \mathcal{I}\backslash\{i\}} g_{S,i}\alpha_S - p_i$ and if $\alpha_i = 0$, the agent does not purchase the item and derives zero utility. Each agent can buy at most one unit of the item and either fully purchases the item or does not purchase it at all.

Each agent is assumed to be rational and utility maximizer. The utility maximization problem of agent i can be written as follows,

$$\max_{\alpha_i \in \{0,1\}} u_i(\alpha_i, \boldsymbol{\alpha_{-i}}, p_i). \tag{2.6}$$

If the utility of an agent is exactly zero, the tie is broken assuming this agent buys the item.

We assume that the seller is a profit maximizer with a linear manufacturing cost. The seller's problem is given by:

$$\max_{\mathbf{p} \in \mathbf{P}} \sum_{i \in \mathcal{I}} \alpha_i \left(p_i - c\right), \tag{2.7}$$

where the vector $\boldsymbol{\alpha}$ represents the purchasing decisions of the agents obtained from problem (2.6), and c is the unit manufacturing cost. If agent i decides to buy the product at the offered price p_i, α_i is equal to 1 and the seller earns a profit of $p_i - c$. If agent i decides not to purchase the item, it incurs zero profit to the seller. The profit is denoted by Π.

We view the entire problem, called the pricing-purchasing game, as a two-stage Stackelberg game. First, the seller leads by choosing the prices $\mathbf{p} \in \mathbf{P}$ to be offered to the agents. Second, the agents follow by deciding whether or not to purchase the item, α_i $\forall i \in \mathcal{I}$. We are interested in *subgame perfect equilibria* of this two-stage game (see, e.g., Fudenberg and Tirole 1991). For a fixed price vector, the second stage equilibria, referred to as the *purchasing equilibria*, are defined by:

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

We note that this definition is similar to the consumption equilibria of a divisible good 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. We also note that the two-stage problem is non-convex as it includes terms of the form $\alpha_i p_i$ in the seller's objective and $\alpha_i \alpha_S$ in the objective functions of the agents (which will be used as constraints in the seller's problem). In addition, the discrete nature of the purchasing decisions increases the complexity of the problem, as it yields a non-convex integer program.

3. Purchasing equilibria

In this section, we consider the second stage game and show the existence of a pure Nash equilibrium (PNE) strategy, given any price vector. We observe that there could be multiple pure Nash equilibria but we characterize all these equilibria via a system of constraints using duality theory. We also identify a mild condition which allows us to focus on the purchasing equilibrium that is preferred by both the seller and the network of agents. We later show that our optimization formulation naturally induces this preferred purchasing equilibrium.

Theorem 1. *Consider the second stage game played simultaneously by the network of agents.*

1. *The second stage game has at least one PNE for any given price vector* $\mathbf{p}$.

2. *There exists a small* $\epsilon \geq 0$ *such that a price perturbation* $p_i - \epsilon$ $\forall i \in \mathcal{I}$ *does not change any of the PNEs. In addition, it ensures that all agents in all PNEs strictly prefer one of the actions (buy or no-buy).*
3. *Among the multiple PNEs, there exists a unique Pareto optimal PNE in which each agent's utility is at least as large as in any other PNE and is strictly higher for at least one agent. This implies that all agents who buy in any PNE will also buy in the Pareto optimal PNE, while deriving a higher utility.*

The proof can be found in Appendix B. The existence of a PNE follows from the fact that for a given price vector, the second stage game is of strategic complements (see, e.g., Jackson and Zenou 2014). Note that the first part of Theorem 1 guarantees the existence of a PNE but not necessarily its uniqueness. Consider the following simple example in which two distinct PNEs arise. Assume a network with two symmetric agents that mutually influence one another: $g_1 = g_2 = 2$ and $g_{21} = g_{12} = 1$. Consider the price vector $p_1 = p_2 = 2.5$. In this case, we have two PNEs: 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. More precisely, for any price strictly larger than 3 or strictly smaller than 2, we have a unique equilibrium but for any price between 2 and 3, we have multiple equilibria. Nevertheless, for any price between 2 and $3 - \epsilon$, the purchasing equilibrium is preferred by the agents as they each derive a positive utility from buying. In particular, for any price but 3, ϵ can be set to 0 and for $p = 3$, any small positive number will work. As a result, even though there exist multiple equilibria, by reducing the price by ϵ, the purchasing equilibria is strictly preferred by both agents. The purpose of ϵ, as can be noted from the above example, is to ensure that agents with ties will buy without affecting other agents' decisions. Note that the value of ϵ can be taken very small so that it does not affect the revenue of the seller significantly.

In the last part of Theorem 1, we show that the seller's preferred equilibrium is unique and is also collectively preferred by all the agents. In particular, we show that among all PNEs, the preferred equilibrium has the property that any buyer in other PNEs will also buy in the preferred equilibrium. In Section 4, we demonstrate that the nature of the first stage game always induces the purchasing equilibrium of interest (in the above example, buy-buy) and hence, the rest of the paper focuses on this preferred equilibrium.

Characterization of the purchasing equilibria

The natural next step is to characterize the purchasing equilibria as a function of the prices, i.e., to derive 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 price vector. We begin by making the following observation regarding the utility maximization problem of each agent.

OBSERVATION 1. *Consider problem (2.6) for agent i under a given price vector $\mathbf{p}$. If other agents' decisions $\boldsymbol{\alpha_{-i}}$ are given, the problem of agent i has a tight linear programming (LP) relaxation.*

In fact, for fixed values of $\mathbf{p}$ and $\boldsymbol{\alpha_{-i}}$, the subproblem faced by agent i happens to be a simple assignment problem. If the quantity $\left(\sum_{S\subset\mathcal{I}\setminus\{i\}} g_{S,i}\alpha_S - p_i\right)$ is positive, $\alpha_i^* = 1$ and if this quantity is negative, $\alpha_i^* = 0$. Finally, if this 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. Therefore, the LP relaxation of the problem faced by agent i (for fixed values of $\mathbf{p}$ and $\boldsymbol{\alpha_{-i}}$) is integral.

Observation 1 allows us to transform problem (2.6) for agent i into a set of constraints by using duality theory. More specifically, this set of constraints consists of primal feasibility, dual feasibility, and strong duality conditions. For agent i, the constraints can be written as follows:

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

$$\text{Dual feasibility:} \quad y_i \ge \sum_{S\subset\mathcal{I}\setminus\{i\}} g_{S,i}\alpha_S - p_i. \tag{3.2}$$

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

$$\text{Strong duality:} \quad y_i = \alpha_i\Big(\sum_{S\subset\mathcal{I}\setminus\{i\}} g_{S,i}\alpha_S - p_i\Big). \tag{3.4}$$

Here, the variable y_i represents the dual variable of problem (2.6). Combining constraints (3.1–3.4) for all agents characterizes all the equilibria (mixed and pure) of the second stage game as a function of the prices. To restrict our attention to the pure Nash equilibria (whose 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 formulation. Note that the number of variables increases by N, as we add a dual continuous variable for each agent.

4. Optimal pricing: MIP formulation

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

We next formulate the pricing problem faced by the seller (denoted by Z) by incorporating the second stage PNE characterized by constraints (3.2–3.4) and $\alpha_i \in \{0,1\}$ for each agent. The binary function defining α_S in (2.2) is also included. 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 formulation is given by:

$$\max_{\substack{\mathbf{p}\in\mathbf{P}\\ \mathbf{y},\boldsymbol{\alpha}}} \sum_{i\in\mathcal{I}} \alpha_i(p_i - c) \tag{Z}$$

$$\text{s.t.} \quad \left.\begin{aligned} y_i &= \alpha_i\left(\textstyle\sum_{S\subset\mathcal{I}\setminus\{i\}} g_{S,i}\alpha_S - p_i\right)\\ y_i &\geq \textstyle\sum_{S\subset\mathcal{I}\setminus\{i\}} g_{S,i}\alpha_S - p_i\\ y_i &\geq 0\\ \alpha_i &\in \{0,1\}\end{aligned}\right\} \quad \forall\, i\in\mathcal{I}$$

$$\alpha_S = \max_{\substack{S'\subset S\\ |S'|=\min\{\Gamma,|S|\}}} \left\{\prod_{j\in S'}\alpha_j\right\} \quad \forall\, S\subset\mathcal{I}$$

In addition to the presence of binary variables and the binary functions defining α_S, one can see that problem Z is non-linear (and non-convex) as it includes terms of the form $\alpha_i\alpha_S$ and $\alpha_i p_i$. Therefore, problem Z is not easily solvable by commercial solvers. We next show that by introducing a few additional continuous variables, one can reformulate problem Z into an equivalent MIP with the same number of binary variables. To gain tractability, we consider valuation models in which the terms $g_{S,i}$ for large $|S|$ are set to zero. This motivates the following definition.

Definition 1. The K**-wise utility model** is a model where $g_{S,i} = 0$ for all subsets $|S| > K-1$.

Note that a larger value of K results in a more non-linear utility model when compared to a smaller K (e.g., for $K=2$, α_S is always linear while for $K=3$ it becomes quadratic). We next present the MIP for the K-wise utility model. We first define the following additional variables, while also redefining the variable α_S:

$$z_i = \alpha_i p_i \qquad \forall\, i\in\mathcal{I}, \tag{4.1}$$

$$\alpha_S = \prod_{j\in S}\alpha_j \qquad \forall\, 1<|S|<\Gamma+2, S\subset\mathcal{I}, \tag{4.2}$$

$$\beta_S = \max_{S'\subset S, |S'|=\Gamma}\{\alpha_{S'}\} \qquad \forall\, \Gamma<|S|<K, S\subset\mathcal{I}, \tag{4.3}$$

$$\eta_{S,i} = \beta_S\alpha_i \qquad \forall\, i\notin S, \Gamma<|S|<K, S\subset\mathcal{I}, i\in\mathcal{I} \tag{4.4}$$

The variables β_S and $\eta_{S,i}$ are defined for sets S satisfying $|S|>\Gamma$. The variables α_S are defined for sets S satisfying $|S|\leq\Gamma+1$. By using the binary nature of the variables and adding certain linear constraints, we can replace all non-linear terms in problem Z. This yields the following MIP denoted by Z-MIP. For simplicity of exposition, we present it for the case when $\Gamma = K-1$ (and hence the $\beta_S, \eta_{S,i}$ variables are absent). The more general formulation can be found in Appendix C.

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

s.t.

$$\left.\begin{aligned} y_i &= \sum_{\substack{|S|<K \\ S\subset I\setminus\{i\}}} g_{S,i}\alpha_{S\cup i} - z_i \\ y_i &\geq \sum_{\substack{|S|<K \\ S\subset I\setminus\{i\}}} g_{S,i}\alpha_S - p_i \\ y_i &\geq 0 \end{aligned}\right\} \quad \forall\, i\in\mathcal{I} \tag{4.5}$$

$$\left.\begin{aligned} z_i &\geq 0 \\ z_i &\leq p_i \\ z_i &\leq \alpha_i p^{max} \\ z_i &\geq p_i - (1-\alpha_i)p^{max} \end{aligned}\right\} \quad \forall\, i\in\mathcal{I} \tag{4.6}$$

$$\left.\begin{aligned} \alpha_S &\geq 0 \\ \alpha_S &\leq \alpha_{S\setminus\{i\}} \qquad \forall\, i\in S \end{aligned}\right\} \quad \forall\, 1<|S|<K+1, S\subset\mathcal{I} \tag{4.7}$$

$$\alpha_{S\cup\{i,j\}} \geq \alpha_{S\cup\{i\}} + \alpha_{S\cup\{j\}} - \alpha_S \qquad \forall\, |S|<K-1, S\subset\mathcal{I}\setminus\{i,j\}, \{i\}\neq\{j\}\subset\mathcal{I} \tag{4.8}$$

$$\alpha_i \in \{0,1\} \qquad \forall\, i\in\mathcal{I} \tag{4.9}$$

$$\alpha_\emptyset = 1 \tag{4.10}$$

In the above formulation, p^{max} denotes the maximal price allowed and is typically known from the context. For example, one can take $p^{max} = \max_i \left\{\sum_{S\subset I} g_{S,i}\right\}$ without affecting the problem at all, since no agent would ever pay a price beyond this value. The set of constraints (4.6) aims to linearize and guarantee the definition of the variable z_i. The sets of constraints (4.7) and (4.8) linearize and ensure the correctness of the variable α_S. For example, constraint (4.8) for agents i and j and $S=\emptyset$ is given by: $\alpha_{i,j} \geq \alpha_i + \alpha_j - 1$, which ensures $\alpha_{i,j} = \alpha_i\alpha_j$.

We note that in the Z-MIP formulation under the K-wise utility model, we have a total of (at most) $4N + \sum_{k=2}^{\Gamma+1}\binom{N}{k} + \sum_{k=\Gamma+1}^{K-1}\binom{N}{k}(1+N-k)$ variables ($4N$ for the α_i, $\mathbf{p}$, $\mathbf{y}$ and $\mathbf{z}$, the second term accounts for α_S, and the last term corresponds to β_S and $\eta_{S,i}$). However, only N variables are binary, while the remaining are all continuous. In particular, in the pairwise setting ($K=2$) we have $4N + \binom{N}{2}$ variables and in the triple-wise setting ($K=3$), we have $4N + \binom{N}{2} + \binom{N}{3}$ variables, if $\Gamma = 2$ and $4N + \binom{N}{2} + \binom{N}{3}(N-1)$ if $\Gamma=1$. In other words, for small values of K (e.g., 2 or 3), the number of variables is a small degree polynomial in N.

We conclude that our problem of designing prices for selling an indivisible good to agents embedded in a social network can be formulated as a MIP. This MIP is equivalent to the two-stage non-convex IP game we started with. This formulation can be viewed as an operational tool to solve the optimal pricing problem (as we will discuss in Section 6.2, the MIP formulation also holds when externalities can be negative). 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 always feasible. If the size of the network is not very large, one can still solve it efficiently using commercial MIP solvers. Moreover, it is possible to solve the problem offline (before launching a new product for example) so that the running time may not be of the highest consideration. Potentially, one can also consider network clustering methods to aggregate or coalesce several nodes to reduce the network size. If the network size is very large, one needs to find more efficient methods to solve Z-MIP. In the next section, we derive efficient optimal methods (polynomial in the number of agents) to solve the problem for two popular pricing strategies.

5. Efficient optimal algorithms

5.1. Discriminative prices

We next consider the general pricing strategy where the firm offers discriminative prices that potentially differ for each agent in the network, that is, $\mathbf{P} = \mathbb{R}_+^N$ in Z-MIP. This scenario is of interest in various practical settings where the seller gathers the purchasing history of each potential buyer, his/her geographical location as well as other attributes. It can also be used by the seller to understand who the influential agents in the network are and what is the maximal profit that can be achieved when using discriminative prices. The prices can then be implemented by setting the same ticket price for everyone, and sending out coupons with discriminative discounts. In fact, in practice, it often occurs that people receive different deals for the same item depending on the loyalty class, purchasing history, and geographical location. The method we propose aims to provide a systematic and automated way of finding the prices (equivalently, discounts) to offer to agents embedded in a social network based on their externalities to maximize the seller's profit.

As discussed, solving Z-MIP using an optimization solver may be impractical for large networks. We next show that solving the LP relaxation of Z-MIP yields the desired optimal integer solution. Consequently, one can solve the problem efficiently (polynomial in the number of agents and very fast in practice) and obtain an optimal solution even for large networks. Recall that the linearization of problem Z was possible due to the integrality of the decision variables. In other words, to reformulate problem Z into Z-MIP, the binary restriction was crucial. As a result, by introducing the new variables z_i, α_S, β_S, and $\eta_{S,i}$, one may potentially obtain fractional solutions that cannot be implemented in practice. However, the following theorem shows that the optimal solutions of Z-MIP can be identified using its continuous relaxation.

Theorem 2. *The optimal discriminative pricing solution of Z-MIP can be obtained efficiently (polynomial in the number of agents). In particular, problem Z-MIP admits a tight LP relaxation.*

The proof can be found in Appendix D. After solving the LP, we first order the agents in increasing order of α_i and then sequentially increase each agent's α_S value (where $i \in S$) to the next agent's α_i value (equal to 1 after the last iteration). We do so carefully by maintaining the feasibility of the LP relaxation and without affecting other buyers' decisions or decreasing the seller's profit. This process is repeated until all agents who bought a fractional amount fully purchase the item. One can use this constructive argument or a solution approach such as the simplex method to arrive at the optimal extreme points which are guaranteed to be integer.

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 applicable to large networks. Next, we derive some properties of the optimal solution. Interestingly, if we know the optimal set of buyers, the corresponding optimal prices can be obtained in closed form as summarized in the following observation.

OBSERVATION 2. *Suppose T^* is the optimal set of buyers, i.e., $T^* = \{i \in \mathcal{I} | \alpha_i^* = 1\}$. Let $S^*(T^*)$ be the sets for which $\alpha_S = 1$, $S \in S^*(T^*)$, obtained from (2.2). Then, the optimized prices are:*

$$p_i^* = \begin{cases} \sum_{S \in S^*(T^*) \setminus \{i\}} g_{S,i} & \forall i \in T^* \\ p^{max} & \textit{otherwise.} \end{cases} \tag{5.1}$$

In other words, the price of an agent who buys is the sum of his/her own value (g_i) and a markup term corresponding to the externalities of the "buying" network on this agent. Membership of agent i into the buying class depends on the self value g_i and on the externalities exerted by agent i on the network. The corresponding optimal profit is given by:

$$\begin{aligned} \Pi^* &= \sum_{i \in T^*} \Big[\sum_{S \in S^*(T^*) \setminus \{i\}} g_{S,i} - c \Big] \\ &= \underbrace{\sum_{\{i \in T^* | g_i \geq c\}} (g_i - c)}_{\textit{Profit in the absence of network effects}} + \underbrace{\sum_{\{i \in T^* | g_i < c\}} (g_i - c) + \sum_{\substack{i \in T^* \\ \emptyset \neq S \in S^*(T^*) \setminus \{i\}}} g_{S,i}}_{\textit{Incremental profit due to network effects}} . \end{aligned} \tag{5.2}$$

We separate the profit into two components. First, the profitable component of individual valuations, which the seller earns even in the absence of network externality effects. Second, the incremental profit due to network effects which balances the seller's profit from unprofitable individual valuations with the revenue gain from network externalities.

Note that prices for non-buyers are set to p^{max} to ensure that no agent will buy. The prices in (5.1) are optimal as they result in the maximal profit that can be extracted by the seller given that agents in T^* are buying. That is, the seller extracts the full consumer surplus as expected, since we consider a setting with complete information. One can see that T^* includes all agents in $\mathcal{I}$ whose individual valuations are profitable. In addition, some buyers may be offered a price below cost. One can view these agents as influencers who receive membership into the buying class due to their strong network externalities. On the flip side, the seller charges higher prices for strongly influenced agents. As a result, the seller taps into an additional source of profits by taking advantage of the network effects. The computational challenge lies in identifying the optimal set of buyers T^*. The LP approach presented in Theorem 2 allows to efficiently find T^* under the utility model (2.1).

We next use the closed form solutions in (5.1) and (5.2) to compare the prices and profit between the linear model in (2.3) and the non-linear models in (2.4)a and (2.4)b. More precisely, we consider a setting where $g_i \geq 0$ and $g_{j,i} \geq 0$ are the same across all three models. However, the $g_{\{j,k\},i}$ are set to zero in the linear model and are non-negative in the other models. We denote the problem and the optimal solutions of models (2.3), (2.4)a, and (2.4)b with subscripts $(K=2, \Gamma=1)$, $(K=3, \Gamma=1)$, and $(K=3, \Gamma=2)$ respectively, and make the following observation.

Observation 3. *The optimal solutions satisfy the following trends:*

- $T^*_{(K=2,\Gamma=1)} \subset T^*_{(K=3,\Gamma=2)} \subset T^*_{(K=3,\Gamma=1)}$, *where* $T^* = \{i \in \mathcal{I} | \alpha_i^* = 1\}$.
- $p^*_{i,(K=2,\Gamma=1)} \leq p^*_{i,(K=3,\Gamma=2)} \ \forall\ i \in T^*_{(K=2,\Gamma=1)}$ *and* $p^*_{i,(K=3,\Gamma=2)} \leq p^*_{i,(K=3,\Gamma=1)} \ \forall\ i \in T^*_{(K=3,\Gamma=2)}$.
- $\Pi^*_{(K=2,\Gamma=1)} \leq \Pi^*_{(K=3,\Gamma=2)} \leq \Pi^*_{(K=3,\Gamma=1)}$, *where* Π^* *denotes the optimal profit.*

As we move from the linear model in (2.3) to the non-linear model (2.4), additional agents will buy. In addition, buyers will pay a higher price, hence inducing a larger seller's profit. As the value of Γ decreases, additional agents will buy. In particular, as the $g_{\{j,k\},i}$ term increases, it results in a larger number of agents buying the item. The seller can also charge higher prices and consequently, earns a higher profit. The additional agents will further increase the profit. Interestingly, one can extend Observation 3 as a function of the degree of non-linearity K and the threshold value Γ.

Corollary 1. *Consider the non-linear utility model in* (2.1) *for given* K *and* Γ. *Then:*

- *The set of buyers, the optimal prices, and the seller's profit increase with* K.
- *The set of buyers, the optimal prices, and the seller's profit decrease with* Γ.

Corollary 1 allows us to understand the impact of both the degree of non-linearity and the threshold value of our utility model on the optimal outcomes. Together with (5.2), Corollary 1 highlights the importance of these parameters. In particular, the value of K affects the optimal prices and profit and the value of Γ affects the active influential sets S. This result suggests that incorporating non-linearity factors in the utility can significantly modify the pricing decisions.

5.2. Uniform price

In this section, we consider the case where the seller offers a uniform price across the network while incorporating externality effects. This scenario arise when the firm may not want to price discriminate due to fairness or ethical reasons and prefers to offer a uniform price. We observe that a similar result as in Theorem 2 for the setting with uniform pricing does not hold. In other words, by adding the (linear) uniform price constraint: $p_1 = p_2 = \ldots = p_N$ to Z-MIP, the corresponding

LP relaxation is no longer tight and we obtain fractional solutions that cannot be implemented in practice. Geometrically, it means that incorporating such a constraint in the Z-MIP formulation is equivalent to adding a cut that violates the integrality of the extreme points of the feasible region. Therefore, we propose an alternative approach to optimally solve the problem by using an efficient algorithm based on iteratively solving the relaxed Z-MIP, which is an LP.

THEOREM 3. *The optimal solution of Z-MIP 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

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 defined as $p_{min}^{(t)} = \max\left\{c, \min_{i\in\mathcal{I}} p_i^{(t)}\right\}$ 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 agents who receive prices less than or equal to the minimal discriminative price from the network (including all their edges). If there are no more agents in the network, go to step 5. Otherwise, go to step 4.
4. Re-solve the relaxed Z-MIP for the 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 $p_{min}^{(\hat{t})}$, where $\hat{t} = \arg\max \Pi^{(t)}$, i.e., the price that yields the highest profit.

Figure 1 Procedure for finding the uniform optimal price.

We show the termination of Algorithm 1 in finite time and prove its correctness in Appendix E. 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 such agent per iteration). As a result, it suffices to consider only a finite selection of prices (at least as high as cost) to identify the optimal uniform price.

6. Extensions

In this section, we consider two extensions of the models and results developed in Sections 3, 4, and 5. First, we present two heuristic methods to solve the problem under discriminative prices. Second, we discuss the setting where externalities among agents can be negative.

6.1. Heuristic methods

In Section 5, we developed efficient optimal algorithms to solve the problem for discriminative prices and uniform price. Our solution approach is based on solving a (continuous) LP. Although these algorithms are efficient for most practical instances, they require the use of an optimization solver. In this section, we present two intuitive heuristic methods which are motivated by exploiting the insights drawn in Observation 2. Such approaches are transparent and easy to interpret but will generally yield a sub-optimal solution. As noted in Observation 2, a key element of the optimal solution consists of identifying the optimal set of buyers T^*. Instead of solving an LP or enumerating all possible subsets (there are exponentially many), we propose two greedy approaches to construct the set of buyers. First, we present the *greedy expansion procedure*, where we iteratively add to the set of buyers agents who yield a positive marginal contribution to the profit. Second, we consider the *greedy removal procedure*, where the initial set of buyers includes all agents, and we iteratively remove the agent who decreases the seller's profit the most. The details of both procedures are reported in Algorithms 2 and 3 respectively. For simplicity of exposition, we focus on the utility model (2.3) that captures the individual network effects of neighbors only (i.e., $K = 2$ and $\Gamma = 1$). Note, however, that both heuristic methods easily extend to the more general utility model.

Algorithm 2 Greedy expansion procedure

Input: c, N, **and** G

Procedure

1. Assign all agents with $g_i \geq c$ to the set of buyers, which we denote by T_1^{GEP}. For all remaining agents $j \in N \setminus T_1^{GEP}$, update g_j to $\tilde{g}_j = g_j + \sum_{i \in T_1^{GEP}} g_{i,j}$.
2. Assign all agents j with $\tilde{g}_j \geq c$ to the set of buyers denoted by T_2^{GEP}.
3. Repeat steps 1-2 until convergence. After this step, we call the set of buyers T_3^{GEP}. We are now left only with agents such that $\tilde{g}_i < c$.
4. For all remaining agents $i \in N \setminus T_3^{GEP}$, compute the quantity $A_i = (\tilde{g}_i - c) + \sum_{k \in T_3^{GEP}} g_{i,k}$. If $A_i \geq 0$, add agent i to the set of buyers, which we denote by T_4^{GEP}. For all remaining agents $j \in N \setminus T_4^{GEP}$, update $\tilde{g}_j$ to $\tilde{g}_j + \sum_{i \in T_4^{GEP}} g_{i,j}$.
5. Repeat step 4 until convergence. After this step, we call the set of buyers T_5^{GEP}.
6. For the remaining agents, we can use one of the following three options: (i) solve a smaller scale LP with the remaining agents, (ii) test adding each pair (or higher subsets) of agents to the set of buyers, or (iii) simply label the remaining agents as non-buyers.

Figure 2 **Greedy expansion procedure.**

Both heuristic methods iteratively construct the set of buyers by exploiting the network externality structure among agents. In the greedy expansion procedure, we first ensure that all agents with a high self-value g_i purchase the item. We then use the network effects of such agents to identify new agents that will buy, and assign them to the set of buyers. Next, we iteratively include to the set of buyers, other agents who have a non-negative marginal increase in the total seller's profit (captured by A_i). At this point, we are left with a smaller set of agents for which it is harder to determine if they belong to the set of buyers. Those agents have a negative marginal contribution, when we add exactly one of them to the set of buyers. Nevertheless, it is possible that adding several of them simultaneously is profitable. To solve this sub-problem to optimality, one can solve a smaller LP or enumerate all possible subsets. Alternatively, one can simply assume

Algorithm 3 Greedy removal procedure

Input: c, N, **and** G

Procedure

1. Assign all agents to the set of buyers, i.e., $T_1^{GRP} = N$, and compute the vector of prices using equation (5.1) with T_1^{GRP} instead of T^*.
2. For each agent i, compute the quantity $B_i = (p_i - c) + \sum_{k \in T_1^{GRP}} g_{i,k}$. This quantity represents the contribution of including agent i to the set of buyers in the seller's profit.
3. If $B_i \geq 0, \forall i \in T_1^{GRP}$, the procedure terminates. Otherwise, remove the agent with the smallest (i.e., the most negative) value of B_i from the set of buyers. If several agents attain the smallest value, break ties randomly. Note that we can extend this step by considering simultaneously removing higher subsets of agents (e.g., pairs). After this step, we call the set of buyers T_3^{GRP}.
4. Repeat steps 1-3 (i.e., update the set of buyers to T_3^{GRP}, compute the prices and the quantity B_i for each i in T_3^{GRP}, and remove the agent with the smallest $B_i < 0$) until convergence. Note that in each step, we only need to update a small number of prices.

Figure 3 **Greedy removal procedure.**

that these agents are non-buyers. In the greedy removal procedure, we first assign all agents to the set of buyers. We then iteratively remove the agent with the most negative marginal decrease to the seller's profit (captured by B_i). We stop the procedure when removing a single agent from the set of buyers does not increase the profit anymore. We next present parametric bounds on the profit performance of these heuristic methods. We denote by C_I the set of agents such that $g_i < c$ and by $|C_I|$ the number of such agents. We call g_{max} the maximal value of $g_{i,j}$ for agents in C_I, i.e., $g_{max} = \max_{i \in C_I, j \in N} g_{i,j}$, and denote by $N_{I,max}$ the maximal number of neighbors for agents in C_I.

Proposition 1. *The greedy expansion and greedy removal procedures satisfy the following.*

1. *Assume that the remaining agents in step 6 of Algorithm 2 are labeled as non-buyers. In this case, the final set of buyers is T_5^{GEP} and satisfies $T_5^{GEP} \subseteq T^*$. In addition, the worst-case additive loss is $|C_I| N_{I,max} g_{max}$.*
2. *The greedy removal procedure admits a worst-case additive loss of $|C_I|c$.*

Interestingly, the performance of both heuristic methods depends on the agents from the set C_I. Specifically, the results of Proposition 1 can be shown by subtracting the contribution of all agents in C_I from the optimal profit (the formal proof is not reported for conciseness). Note that the first bound ($|C_I| N_{I,max} g_{max}$) represents the largest possible loss which is computed by assuming that all agents in C_I collectively yield a positive profit. In most instances, the loss will be much smaller (for example, one can replace C_I by the set of remaining agents in step 6 of Algorithm 2). Similarly, the second bound ($|C_I|c$) corresponds to the worst-case of mistakenly including all agents in C_I. Note that when $C_I = \emptyset$, both greedy procedures yield the optimal solution, as expected. The above parametric bounds admit the following interesting implication. For the greedy expansion procedure, in the worst-case, we miss all agents with $g_i < c$, whereas in the greedy removal procedure, in the worst-case, we mistakenly include all agents with $g_i < c$. This implies that when one heuristic method will not perform well, the other will, and vice-versa. Interestingly, one can show that the greedy removal procedure is optimal for the special case of supermodular valuation functions.

In summary, we have exploited the structure of the problem to propose two greedy heuristic procedures. As mentioned, one can also use the greedy expansion procedure to significantly reduce the size of the LP to compute the optimal solution, while this method can also be used as a sub-routine in Algorithm 1 instead of the LP.

6.2. Setting with negative externalities

In the previous sections, as well as in several papers in the literature, the focus has mainly been on settings with non-negative externalities. One exception is the recent work in Cao et al. (2017) where the authors take an algorithmic approach to solve the iterative pricing problem with negative externalities. They show that the problem is NP-hard and propose a 2-approximation algorithm. The problem considered in Cao et al. (2017) is different than ours as the seller can post an iterative list of uniform prices for all the agents (multiple selling rounds). In addition, their utility model is a special case of our model (they consider pairwise interactions terms which are all negative). In many practical applications, the externalities among agents in the network can be both positive and negative (e.g., negative reviews on a product sold via an online platform). In this section, we discuss how our results change for a setting where externalities can be either positive or negative.[2] First, the existence of a PNE for the second stage game is not always guaranteed when externalities can be negative (and therefore, Theorem 1 does not hold in general). Nevertheless, the characterization method of the purchasing equilibria still applies. In particular, the Z-MIP formulation presented in Section 4 remains the same. The main difference is that when externalities can take negative values, it is possible that the MIP is infeasible. In this case, it means that there is no PNE for the purchasing game under any price vector. By solving the Z-MIP problem, we can then determine whether a PNE exists, and if it does, we can compute the most profitable solution for the seller. We summarize this result in the following corollary.

Corollary 2. *Consider a setting with general externalities. If Z-MIP is infeasible, it means that there is no PNE under any price vector. Otherwise, its optimal solution leads to the optimal price vector, and to the resulting purchasing PNE. This optimal solution corresponds to the most profitable outcome for the seller.*

The implication of Corollary 2 is twofold: (i) the Z-MIP formulation allows us to determine whether a PNE exists for the network of agents under any price vector, and (ii) if it does, it can compute

[2] We refer to *general externalities* to describe the setting that includes both positive and negative externalities among agents in the network.

the optimal outcome for the seller. Computing an equilibrium under negative externalities is a hard problem (Cao et al. 2017 show that a similar version of this problem is NP-hard). As a result, the Z-MIP formulation we introduce in this paper allows us to solve the problem for a wide range of utility models under either positive or negative externalities. Unfortunately, the efficient optimal approach presented in Section 5.1 uses the non-negativity of the externalities to ensure the integrality of the formulation. We then propose to either solve Z-MIP directly or to use one of the heuristic methods presented in Section 6.1. Note, however, that the parametric bound in Proposition 1 holds only for the setting with non-negative externalities.

7. Price-incentives to guarantee influence

So far, we have assumed that network externality effects always materialize as long as agents purchase the item. This assumption is not realistic in many practical settings. After purchasing an item, it is sometimes not entirely natural to exert network externalities unless one takes some effort to do so. This, for example, could be by writing a review, endorsing the item on social media, or at the very least announcing the purchase. However, to the best of our knowledge, most previous work impose the assumption that purchasing is equivalent to influencing (i.e., exerting network externality effects) with the exception of Arthur et al. (2009). In the latter, the authors study a cash back setting where the seller offers an exogenous uniform cash reward to any recommender if s/he influenced at least one friend to purchase the item. We study a similar model in the context of purchasing equilibrium and the optimization framework proposed in this paper.

Consider a setting where the seller offers both a price and a discount (also referred to as incentive) to each agent. If the agent decides to purchase the item, s/he can claim the discount in return for influence actions such as liking the product in online platforms or writing a review. Using such a model, the seller can now ensure the externalities so that network effects are guaranteed to occur. In the previous setting where externalities were assumed to always occur, the actual profits may be far from the value predicted by the optimization. In fact, we show via a computational example in Section 8 that even if a a few agents do not exert his/her externality effects, it can significantly reduce the seller's profit. We next extend our model and results to this more general setting.

7.1. Model

For simplicity, we focus on the utility model (2.3) that captures individual externalities of neighbors only. We consider a model with a continuum of actions to influence neighbors. Let $t_i \geq 0$ denote the utility 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* of agent i and the variable γ_i as the *externality intensity* chosen by agent i. The parameter t_i can be estimated from historical data such as past purchases and number of reviews written. For a given price vector $\mathbf{p}$ and discount vector $\mathbf{d}$, we extend the utility function of agent i in (2.5) as follows:

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

where α_i is the binary purchasing decision of agent i and $\gamma_i \leq \alpha_i$. If agent i does not purchase the item, $\alpha_i = 0$ and $\gamma_i = 0$. In other words, the constraint $\gamma_i \leq \alpha_i$ captures the fact that only buyers can exert externalities on friends. However, 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 other agents but i. Similarly to problem (2.6), the utility maximization problem for agent i is given by:

$$\begin{aligned} \max_{\alpha_i, \gamma_i} \quad & u_i(\alpha_i, \gamma_i, \boldsymbol{\alpha_{-i}}, \boldsymbol{\gamma_{-i}}, p_i, d_i) \\ \text{s.t.} \quad & 0 \leq \gamma_i \leq \alpha_i \\ & \alpha_i \in \{0,1\}. \end{aligned} \tag{7.1}$$

In a similar way as problem (2.7), the seller's profit maximization problem can be written as:

$$\max_{\mathbf{p}, \mathbf{d} \in \mathbf{P}} \quad \sum_{i \in \mathcal{I}} \Big[\alpha_i \, (p_i - c) - \gamma_i d_i \Big]. \tag{7.2}$$

Here, the seller's decision variables are $\mathbf{p}$ and $\mathbf{d}$ which are two vectors of prices and discounts with a component potentially different for each agent. As before, these vectors can be chosen according to different strategies. For example, one can consider a fully discriminative, a uniform pricing strategy, or more generally, an hybrid model where the regular price is uniform across the network (i.e., $p_i = p_j$) but the discounts are tailored to the various agents. This hybrid setting corresponds to a common practice of online sellers who offer a standard posted price but design personalized discounts for different classes of customers (sent via targeted coupons). The variables α_i and γ_i are decided according to each agent's utility maximization problem given in (7.1). If agent i decides to buy the product, then the seller incurs a profit of $p_i - \gamma_i d_i - c$.

In the special case where $\alpha_i = \gamma_i$ and $t_i = 0 \ \forall i \in \mathcal{I}$, we recover our previous model. 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 an action to influence. Note that one can easily extend the model in this section to more than two prices to incorporate a finite set of different actions specified by the seller.

7.2. Results

Our goal is to extend our results to the above general setting. We next 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 vectors of prices* $\mathbf{p}$ *and discounts* $\mathbf{d}$ *chosen by the seller. A small perturbation in prices and discounts results in a Pareto optimal PNE that is preferred by both the seller and the network of agents.*

The proof of Theorem 4 is not presented for conciseness and is of similar nature as the proof of Theorem 1. In this case, a PNE is defined by restricting the purchasing decisions α_i to be 0 or 1. Nevertheless, we note that there always exists an equilibrium in which the variables γ_i are also all integer. More precisely, if $d_i - t_i > 0$ (recall that prices and discounts are given), γ_i can be set to 1, and otherwise $\gamma_i = 0$. As a result, there exists a PNE with γ_i integer. Note that a result similar to Observation 1 still holds, and hence 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 (7.1) to a set of constraints using duality theory:

$$\text{Primal feasibility:} \quad 0 \le \alpha_i \le 1, \tag{7.3}$$

$$0 \le \gamma_i \le \alpha_i. \tag{7.4}$$

$$\text{Dual feasibility:} \quad y_i - w_i \ge g_i + \sum_{j \in \mathcal{I}\setminus\{i\}} \gamma_j g_{ji} - p_i, \tag{7.5}$$

$$w_i \ge d_i - t_i, \tag{7.6}$$

$$y_i, w_i \ge 0. \tag{7.7}$$

$$\text{Strong duality:} \quad y_i = \alpha_i\Big(g_i + \sum_{j \in \mathcal{I}\setminus\{i\}} \gamma_j g_{ji} - p_i\Big) + \gamma_i(d_i - t_i). \tag{7.8}$$

In this case, we 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, we impose α_i to be binary for all $i \in \mathcal{I}$ to restrict to pure equilibria. We can then formulate the optimal pricing problem, similar to problem Z, that maximizes the profit given in (7.2) with the equilibrium constraints (7.3)-(7.8):

$$\begin{aligned} \max_{\substack{\mathbf{p},\mathbf{d}\in\mathbf{P} \\ \mathbf{y},\mathbf{w},\boldsymbol{\alpha},\boldsymbol{\gamma}}} \quad & \sum_{i\in\mathcal{I}} \Big[\alpha_i\,(p_i - c) - \gamma_i d_i\Big] \\ \text{s.t.} \quad & \text{constraints } (7.4) - (7.8), \alpha_i \in \{0,1\} \quad \forall\, i \in \mathcal{I}. \end{aligned} \tag{Zi}$$

We denote this problem by Zi, where i represents the model with incentives to guarantee influence of the present section. We make the following observation.

OBSERVATION 4. *Every optimal solution of problem Zi satisfies $d_i \le t_i$.*

This follows from the fact that the seller can reduce d_i to be equal to t_i, while maintaining feasibility and strictly increasing the objective. This implies that constraint (7.6) is redundant in the optimal pricing problem. By using constraints (7.5–7.7), one can always assign $w_i = 0$ while maintaining feasibility without altering the objective. This observation allows us to simplify problem Zi by removing all dual variables w_i $\forall i \in \mathcal{I}$.

PROPOSITION 2. *Problem Zi admits a tight continuous relaxation. Moreover, there always exists an optimal solution to problem Zi where all variables γs are also integer.*

The second result in Proposition 2 is interesting as 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 for a specific action to influence.

Problem Zi has non-linearities of the form $\alpha_i\gamma_j$, $\alpha_i p_i$, and $\gamma_i d_i$. Using the discreteness of α_i and γ_i from Proposition 2, one can transform problem Zi into the following MIP, denoted by Zi-MIP:

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

s.t.

$$\left.\begin{aligned} y_i &= \Big(\alpha_i g_i + \textstyle\sum_{j\in\mathcal{I}\setminus\{i\}} x_{ji} g_{ji} - z_i\Big) + (z_i^d - \gamma_i t_i) \\ y_i &\geq g_i + \textstyle\sum_{j\in\mathcal{I}\setminus\{i\}} \gamma_j g_{ji} - p_i \\ \gamma_i &\leq \alpha_i \\ y_i &\geq 0 \end{aligned}\right\} \quad \forall\, i\in\mathcal{I} \tag{7.9}$$

$$\left.\begin{aligned} z_i, z_i^d &\geq 0 \\ z_i &\leq p_i \\ z_i &\leq \alpha_i p^{max} \\ z_i &\geq p_i - (1-\alpha_i)p^{max} \\ z_i^d &\leq d_i \\ z_i^d &\leq \gamma_i p^{max} \\ z_i^d &\geq d_i - (1-\gamma_i)p^{max} \end{aligned}\right\} \quad \forall\, i\in\mathcal{I} \tag{7.10}$$

$$\left.\begin{aligned} x_{ji} &\geq 0 \\ x_{ji} &\leq \alpha_i \\ x_{ji} &\leq \gamma_j \\ x_{ji} &\geq \alpha_i + \gamma_j - 1 \end{aligned}\right\} \quad \forall\, i\neq j\in\mathcal{I} \tag{7.11}$$

$$\alpha_i, \gamma_i \in \{0,1\} \qquad \forall\, i\in\mathcal{I} \tag{7.12}$$

Note that we removed the dual variables w_i using Observation 4. We conclude that the problem of designing prices and incentives for selling an indivisible item to agents in a social network can be formulated as a MIP. For the case of discriminative prices and discounts, i.e., when $\mathbf{P} = \mathbb{R}_+^N \text{x}\ \mathbb{R}_+^N$, we next show a similar result as Theorem 2.

Theorem 5. *The optimal discriminative pricing solution of* Zi-MIP *can be obtained efficiently (polynomial in the number of agents). In particular, problem* Zi-MIP *admits a tight LP relaxation.*

The main idea behind the proof is composed of 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 construct a solution with α_i integer $\forall i\in\mathcal{I}$. Second, with the integer $\boldsymbol{\alpha}$ values obtained from the previous step, one can show that the objective does not decrease by modifying any component of $\boldsymbol{\gamma}$ to 1 by appropriately altering 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 a lower profit for the seller. However, this profit is guaranteed whereas in the previous case, the estimated profit can be far from the realized value if people fail to exert their externalities on neighbors (i.e., the model is misspecified). The difference in profits between both settings can be viewed as the price paid by the seller to guarantee network externalities and can be computed efficiently by solving both settings.

8. Computational experiments

In this section, we present computational experiments on simple networks to draw qualitative insights and to compare various pricing strategies including the richer model with incentives from Section 7. We consider a network with $N = 10$ agents and a utility model with $K = 2$ and $\Gamma = 1$.

Value of incorporating network externalities: In Fig. 4, we plot the optimal prices offered by the seller under discriminative and uniform pricing strategies, both with and without network externalities. 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 with network externalities). In this instance, each agent is connected to exactly three other agents, where we use $g_{j,i} = 1.25$ for any connected edge, $g_i = 3$, and $c = 2$.[3]

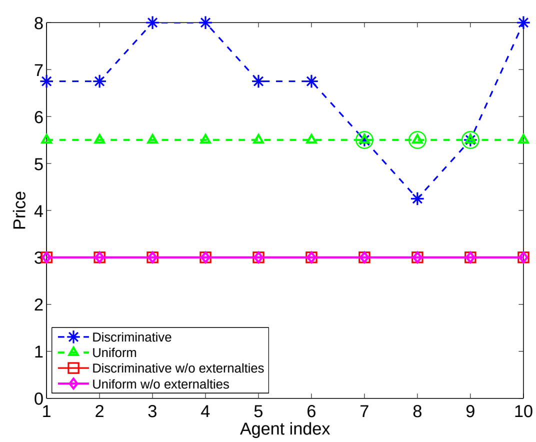

Figure 4 **Value of incorporating network externalities for discriminative and uniform pricing strategies.**

We observe that incorporating the positive externalities among the agents allows the seller to earn higher profits. In this particular example, the total profits are equal to 46.25 (discriminative prices) and 24.5 (uniform price) for the case with network externalities. In the case without network

[3] In Fig. 4, the exact network structure is such that agents 2, 3, and 10 have four influencers; agents 1, 2, 5, and 6 have three influencers; agents 7 and 9 have two influencers; and finally, agent 3 has one influencer.

externalities, the profits under both uniform and discriminative prices are equal to 10. This result is expected as each agent willingness-to-pay increases as their neighbors positively affect them. The seller can therefore charge higher prices and increase its profits. Fig. 4 also shows the added benefit of using a discriminative pricing strategy relative to a uniform price. When the firm has the additional flexibility to price discriminate and offers a different price to each agent in the network, the total profit can increase significantly.

Pricing an influencer: In Fig. 5, we present an example where it is beneficial for the seller to earn a negative profit ($p_i < c$) from an influential agent to extract significant positive profits from his/her neighbors. In particular, we consider a network where agent 5 is a very influential player with $g_{5,5}$ being very low (0.075) while $g_{5,j}$ is sufficiently high (1.38) for the four agents that s/he influences. Here, $g_{i,j} = 0.75$ for any other connected edge, $g_i = 1.5R\ \forall i \neq 5$ where $R = U[1,2]$ (a single i.i.d. instance is drawn) and $c = 2$. In this example, the optimal discriminative price for agent

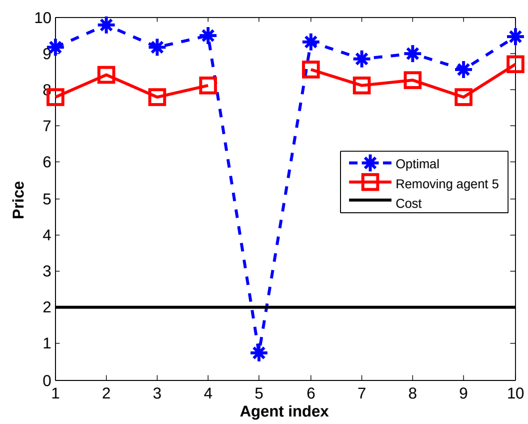

Figure 5 **Centrality effect: losing money on an influential agent.**

5 happens to be below cost. This illustrates the fact that agent 5 has an influential position in the network and therefore, the seller should strongly incentivize this agent. In particular, the optimal algorithm identifies this pattern and captures the fact that it is profitable to offer a low price to agent 5. This way, the seller loses a small amount of money from the influential agent but can extract higher profits from others. We now consider an alternative strategy where the seller decides to remove agent 5 from the network due to his/her low valuation (this is equivalent to offer a very large price to agent 5). In this case, all the optimal prices are decreased and the overall profit drops from 63.52 to 55.5. As a result, one can increase the profit by 14.5% by including agent 5.

Value of incorporating incentives that guarantee influence: In Fig. 6, we compare the optimal solution for discriminative prices to the more general model from Section 7 where the seller offers a uniform regular price ($p = 4$) and designs discriminative discounts in exchange for an action. As discussed, if the seller does not incentivize the agents to influence, it is possible that some of them would not exert their externalities. The goal of this experiment is to quantify the impact of having incentives. In this instance, each agent is connected to three other agents with $g_{j,i} = 0.75$ for any connected edge and $g_i = 1.5R$ where $R = U[1, 4.5]$ (a single i.i.d. instance is drawn). We assume that $t_i = U[0, 1]\ \forall\ i \neq 1$ (single i.i.d. drawn), $t_1 = 6.9$, and $c = 1$.

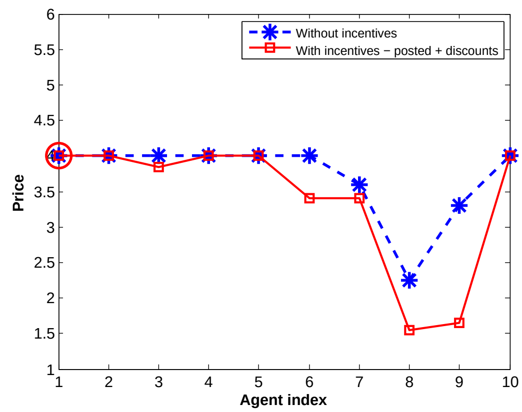

Figure 6 **Value of incorporating incentives that guarantee influence.**

We observe that the profit using the model without incentives is equal to 27.15. This profit is not guaranteed because some agents may not influence their peers (i.e., the model may be misspecified). In particular, in this example, suppose agents 5 and 10 who buy at the full price do not influence their neighbors. Agent 1 ends up not purchasing the item and consequently, does not influence his/her neighbors either. Finally, it happens that only agents 2, 5, and 10 buy the item yielding a profit of 9 as opposed to 27.15. As a result, the earlier model predicts a profit value which is significantly higher than the realized profit. On the other hand, in the model with incentives that guarantee influence, the total profit is equal to 20.85. In this case, agent 1 does not purchase the item and agents 5 and 10 do not influence anyone but other agents do. Observe that this is lower than 27.15 but significantly higher 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 profit value.

Symmetric agents with asymmetric incentives: In Fig. 7, we present a setting with symmetric agents who receive asymmetric incentives to influence their neighbors. In this instance, each agent has the same number of neighbors and the same self and cross valuations. In particular, we consider a complete graph with $g_i = 1.3$ and $g_{i,j} = 0.3$, a cost to influence $t_i = 2.2$, and $c = 0.2$. We compute the optimal discriminative prices which happen to be 3 for everyone and compare 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 agents are homogenous. Indeed, it is sufficient to incentivize any 6 out of 10 agents in the network (no matter which group of 6). These 6 agents receive a targeted discount to exert network effects on their peers that purchase at the full price.

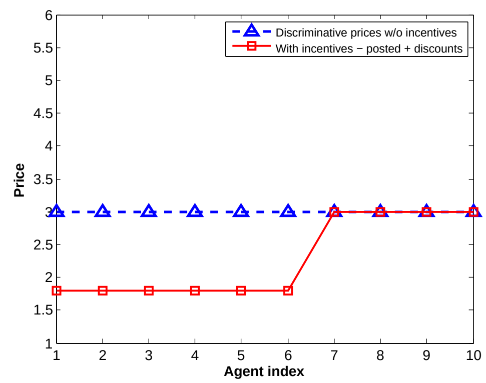

Figure 7 **Symmetric Graph with asymmetric incentives: with and without incentives.**

Effect of network topology on optimal prices: In Fig. 8, we consider different network topologies and compare the optimal discriminative prices and the corresponding profits. In all the scenarios, $g_i = 1.5R$ where $R = U[1,2]$ (a single i.i.d. instance is drawn), $g_{i,j} = 0.75$ when agent i influences agent j and 0 otherwise, and $c = 2$. For each network topology, we solve the optimal discriminative prices using the Z-MIP relaxation. We plot the optimal price vector for the different networks in Fig. 8. We observe that in our example, all agents purchase the item. In the complete graph, all nodes are connected to each other and hence, the profits are the highest (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. 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 other agents and in Star 2 agent 5 is influenced by all others. In both cases, the profits are equal to 8.2, as the total valuations in the system are identical. In Star 1, agent 5 receives a small discount to influence so that the prices offered to other buyers are slightly higher. In Star 2, the prices of all agents but 5 are slightly lower so that the seller can charge a high price to agent 5. As we can see, the prices and 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 connected the graph is, the higher the profit.

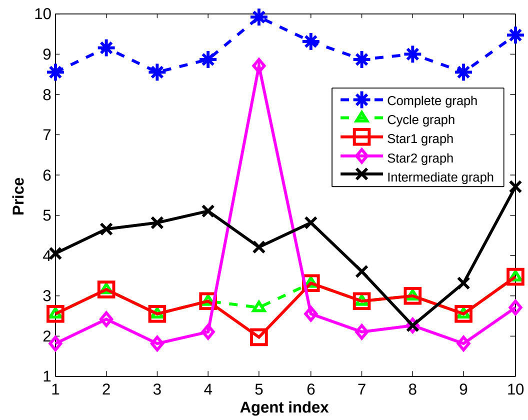

Figure 8 Optimal prices for various network topologies.

9. Conclusions

In this paper, we study an optimal pricing model for a firm that sells an indivisible good to agents embedded in a social network. We assume that agents interact and positively influence each others' purchasing decisions (via network externalities). We propose a broad class of non-linear utility models that explicitly capture externalities from subsets of agents (communities or groups) and allow a threshold on the number of agents needed to trigger the externality effect.

We model the problem as a two-stage game and reformulate it as a MIP with linear constraints. We view this MIP as an operational pricing tool that holds for general externalities (positive or negative) and that can incorporate various pricing business rules. For the case of discriminative and uniform pricing strategies (under positive externalities), we present efficient methods to optimally solve the MIP using its LP relaxation. We observe that the price of a buyer in the optimal discriminative solution can be expressed as the sum of its own value and a markup term corresponding to externalities from the network of agents who buy the item. The gain from network externalities comes from two types of customers: high-valued customers and low-valued customers who are influential and can sometimes be offered a price below cost. In addition, when comparing linear to non-linear utility models, we show that as we move from a linear model to a non-linear one, additional agents will buy, buyers will pay a higher price and hence, this induces higher profits. We also convey that a larger threshold on the minimum number of neighbors results in a smaller number of buyers and decreases the seller's profit. Hence, our analysis suggests that incorporating non-linearity effects in the utility model significantly modifies the pricing decisions.

We extend our pricing model and results to the case when the seller can design both prices and incentives to guarantee influence. This extension is important as in general, agents that buy do not necessarily exter network externalities on their peers. The seller can use incentives in exchange for an action such as a wall post or a review to guarantee network externality effects. Finally, we present computational experiments to highlight the benefits of incorporating network externalities and to compare the different pricing strategies.

References

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

Andreasen AR (1994) Social marketing: Its definition and domain. *Journal of public policy & marketing* 108–114.

Aral S, Walker D (2011) Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management science* 57(9):1623–1639.

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

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

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

Basu A, Mazumdar T, Raj S (2003) Indirect network externality effects on product attributes. *Marketing Science* 22(2):209–221.

Bertsimas D, Weismantel R (2005) *Optimization over integers*, volume 13 (Dynamic Ideas Belmont).

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

Borodin A, Filmus Y, Oren J (2010) Threshold models for competitive influence in social networks. *International Workshop on Internet and Network Economics*, 539–550 (Springer).

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

Cao Z, Chen X, Hu X, Wang C (2017) Approximation algorithms for pricing with negative network externalities. *Journal of Combinatorial Optimization* 33(2):681–712.

Chen W, Lu P, Sun X, Tang B, Wang Y, Zhu ZA (2011) Optimal pricing in social networks with incomplete information. *International Workshop on Internet and Network Economics*, 49–60 (Springer).

Chen W, Wang C, Wang Y (2010) Scalable influence maximization for prevalent viral marketing in large-scale social networks. *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1029–1038 (ACM).

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

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

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

Goldenberg J, Libai B, Muller E (2001) Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing letters* 12(3):211–223.

Gunnec D, Raghavan S (2017) Integrating social network effects in the share-of-choice problem. *Decision Sciences* 48(6):1098–1131.

Hartline J, Mirrokni V, Sundararajan M (2008) Optimal marketing strategies over social networks. *Proceedings of the 17th international conference on World Wide Web*, 189–198 (ACM).

Iyengar R, den Bulte CV, Valente TW (2011) Opinion leadership and social contagion in new product diffusion. *Marketing Science* 30(2):195–212.

Jackson MO, Zenou Y (2014) Games on networks. *Handbook of game theory* 4.

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

Kempe D, Kleinberg J, Tardos É (2003) Maximizing the spread of influence through a social network. *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 137–146 (ACM).

Lovett MJ, Peres R, Shachar R (2013) On brands and word of mouth. *Journal of marketing research* 50(4):427–444.

Luo ZQ, Pang JS, Ralph D (1996) *Mathematical programs with equilibrium constraints* (Cambridge University Press).

Mahajan V, Muller E, Bass FM (1991) New product diffusion models in marketing: A review and directions for research. *Diffusion of technologies and social behavior*, 125–177 (Springer).

Nakkas A, Xu Y (2015) The impact of valuation heterogeneity and network structure on equilibrium prices in supply chain networks. Technical report, Working paper.

Nemhauser GL, Wolsey LA (1988) *Integer and combinatorial optimization*, volume 18 (Wiley New York).

Seeman L, Singer Y (2013) Adaptive seeding in social networks. *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, 459–468.

Stremersch S, Tellis GJ, Franses PH, Binken JL (2007) Indirect network effects in new product growth. *Journal of Marketing* 71(3):52–74.

Topkis DM (1979) Equilibrium points in nonzero-sum n-person submodular games. *Siam Journal on control and optimization* 17(6):773–787.

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 Itai Ashlagi, Adam Elmachtoub, Markus Ettl, Georgia Perakis, and Velibor V. Mišić for their valuable feedback that helped improve the paper.

Appendix A: Weighted bipartite graph for the example in Section 2.1

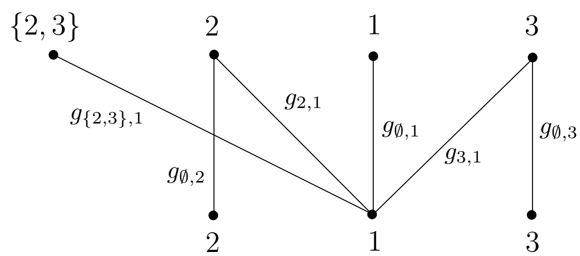

Figure 9 Example of the weighted bipartite graph.

Appendix B: Proof of Theorem 1

1. The existence of a PNE for any given price vector follows from the fact that the second stage game is of strategic complements. In other words, for any $\alpha_i \geq \alpha_i'$ and $\boldsymbol{\alpha_{-i}} \geq \boldsymbol{\alpha'_{-i}}$ (componentwise), we have increasing differences, that is, $u_i(\alpha_i, \boldsymbol{\alpha_{-i}}, p_i) - u_i(\alpha_i', \boldsymbol{\alpha_{-i}}, p_i) \geq u_i(\alpha_i, \boldsymbol{\alpha'_{-i}}, p_i) - u_i(\alpha_i', \boldsymbol{\alpha'_{-i}}, p_i)$. This follows from the positive externality assumption (see Assumption 1.b). Consequently, using the result from Theorem 1 in Section 3.2 of Jackson and Zenou (2014) (the same result can also be found in Topkis (1979)), we conclude the existence of a PNE.

2. If there are no ties in any of the PNEs, one can take $\epsilon = 0$. Consider the case where there are ties for some agents in one or more PNEs. In this case, one can choose $\epsilon > 0$ to be very small such that (i) agents that were buying in any PNE are still buying (in particular, their utility strictly increases and they become better-off), (ii) agents that were not buying (i.e., have negative utility) continue not to buy, and (iii) agents who were indifferent (i.e., exactly at zero utility), become strictly better-off as they derive a positive utility when the price is reduced by ϵ. Consequently, all the ties are eliminated for all PNEs.

3. In the case of a unique PNE, this is by definition a Pareto optimal PNE for the agents. We next consider a setting where there are multiple equilibria. Consider any agent whose actions differ in the different PNEs. As there are two actions, one of the actions is to buy in one of the PNEs. Note that when the agent buys, s/he derives a (strictly) positive utility with the perturbed prices and zero utility otherwise. As a result, this agent prefers to buy in the Pareto optimal solution. By using the non-negativity assumption, this agent can only non-negatively impact other agents' valuations and increase their

utility. This implies that all other agents also prefer the buying decision of the focal agent. Similarly, one can argue that all agents who buy in one of the equilibria will buy (and prefer) the Pareto solution. The only agents who remain are the ones that do not buy in any equilibria. Note that they will not buy in the Pareto optimal solution either. This Pareto optimal solution is also a PNE: the buyers have no incentive to deviate as they derive a strictly positive utility and the non-buyers should not deviate either due to their negative utility from buying. □

Appendix C: General Z-MIP formulation

In this section, we present the generalization of Z for any $\Gamma < K$. Recall that in Section 4 we presented the formulation for the case when $\Gamma = K - 1$.

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

s.t.

$$\left.\begin{aligned} y_i &= \sum_{\substack{|S|\le\Gamma\\ S\subset\mathcal{I}\setminus\{i\}}} g_{S,i}\alpha_{S\cup i} + \sum_{\substack{\Gamma<|S|<K\\ S\subset\mathcal{I}\setminus\{i\}}} g_{S,i}\eta_{S,i} - z_i\\ y_i &\ge \sum_{\substack{|S|\le\Gamma\\ S\subset\mathcal{I}\setminus\{i\}}} g_{S,i}\alpha_S + \sum_{\substack{\Gamma<|S|<K\\ S\subset\mathcal{I}\setminus\{i\}}} g_{S,i}\beta_S - p_i\\ y_i &\ge 0 \end{aligned}\right\} \quad \forall\, i\in\mathcal{I} \tag{C.1}$$

$$\left.\begin{aligned} z_i &\ge 0\\ z_i &\le p_i\\ z_i &\le \alpha_i p^{max}\\ z_i &\ge p_i - (1-\alpha_i)p^{max} \end{aligned}\right\} \quad \forall\, i\in\mathcal{I} \tag{C.2}$$

$$\left.\begin{aligned} \alpha_S &\ge 0\\ \alpha_S &\le \alpha_{S\setminus\{i\}} \qquad \forall\, i\in S \end{aligned}\right\} \quad \forall\, 1<|S|<\Gamma+2, S\subset\mathcal{I} \tag{C.3}$$

$$\alpha_{S\cup\{i,j\}} \ge \alpha_{S\cup\{i\}} + \alpha_{S\cup\{j\}} - \alpha_S \qquad \forall\, |S|<\Gamma, S\subset\mathcal{I}\setminus\{i,j\}, \{i\}\ne\{j\}\subset\mathcal{I} \tag{C.4}$$

$$\alpha_{S'} \le \beta_S \le 1 \qquad \forall\, |S'|=\Gamma, \Gamma<|S|<K, S'\subset S\subset\mathcal{I} \tag{C.5}$$

$$\beta_S \le \sum_{|S|'=\Gamma, S'\subset S} \alpha_{S'} \qquad \forall\, \Gamma<|S|<K, S\subset\mathcal{I} \tag{C.6}$$

$$\left.\begin{aligned} \eta_{S,i} &\ge 0\\ \eta_{S,i} &\le \alpha_i\\ \eta_{S,i} &\le \beta_S\\ \eta_{S,i} &\ge \beta_S + \alpha_i - 1 \end{aligned}\right\} \quad \forall\, \Gamma<|S|<K, S\subset\mathcal{I}\setminus\{i\}, i\in\mathcal{I} \tag{C.7}$$

$$\alpha_i \in \{0,1\} \qquad \forall\, i\in\mathcal{I} \tag{C.8}$$

$$\alpha_\emptyset = 1 \tag{C.9}$$

The sets of constraints (C.3), (C.4), (C.5), and (C.6) linearize and ensure the correctness of the variables α_S and β_S. For example, constraint (C.4) for agents i and j and $S=\emptyset$ is given by:

$\alpha_{i,j} \geq \alpha_i + \alpha_j - 1$, which along with other constraints ensures $\alpha_{i,j} = \alpha_i \alpha_j$. On the other hand, constraint (C.6) for $S = \{i, j\}$ and $\Gamma = 1$ is given by: $\beta_{i,j} \leq \alpha_i + \alpha_j$, which again along with other constraints ensures $\beta_{i,j} = \max\{\alpha_i, \alpha_j\}$. Finally, constraints (C.7) linearize the $\eta_{S,i}$ variable in a similar fashion as in constraints (C.3) and (C.4).

Appendix D: Proof of Theorem 2

Before proving the main theorem, we state and prove the next Lemma that identifies the optimal values of all variables, given α_i $\forall i \in \mathcal{I}$ (discrete or fractional).

Lemma 1. *For given (discrete or fractional) α_i $\forall i \in \mathcal{I}$, the revenue maximizing solution (and hence profit maximizing because α_i values are fixed) is given by:*

$$p_i^* = z_i^* + (1-\alpha_i)p^{\max} \qquad \forall i \in \mathcal{I} \tag{D.1}$$

$$z_i^* = \sum_{\substack{S \subset I \setminus \{i\} \\ |S| \leq \Gamma}} g_{S,i} \alpha^*_{S \cup i} + \sum_{\substack{S \subset I \setminus \{i\} \\ |S| > \Gamma}} g_{S,i} \eta^*_{S,i} \qquad \forall i \in \mathcal{I} \tag{D.2}$$

$$\eta^*_{S,i} = \min\{\beta^*_S, \alpha_i\} \qquad \forall\, i \notin S, \Gamma < |S| < K, S \subset \mathcal{I}, i \in \mathcal{I} \tag{D.3}$$

$$\beta^*_S = \min\left\{1, \sum_{S' \subset S, |S'| = \Gamma} \alpha^*_{S'}\right\} \qquad \forall\, \Gamma < |S| < K, S \subset \mathcal{I}, \tag{D.4}$$

$$\alpha^*_S = \min_{i \in S}\{\alpha_i\} \qquad \forall\, 1 < |S| \leq \Gamma + 1, S \subset \mathcal{I}. \tag{D.5}$$

We begin by showing (D.1–D.2) first assuming that all remaining variables $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\eta}$ are given. Consider the remainder of the feasibility constraints (4.6) for each agent. Eliminating y_i reduces them to:

$$z_i \geq \max\{0, p_i - (1-\alpha_i)p^{\max}\} \tag{D.6}$$

$$z_i \leq \min\left\{p_i, \alpha_i p^{\max}, \sum_{\substack{S \subset I \setminus \{i\} \\ |S| \leq \Gamma}} g_{S,i} \alpha_{S \cup i} + \sum_{\substack{S \subset I \setminus \{i\} \\ |S| > \Gamma}} g_{S,i} \eta_{S,i}, \sum_{\substack{S \subset I \setminus \{i\} \\ |S| \leq \Gamma}} g_{S,i} \left[\alpha_{S \cup i} - \alpha_S\right] + \sum_{\substack{S \subset I \setminus \{i\} \\ |S| > \Gamma}} g_{S,i} \left[\eta_{S,i} - \beta_S\right] + p_i\right\}. \tag{D.7}$$

We know that $g_{S,i}\alpha_{S \cup i} \leq g_{S,i}\alpha_i$ and $g_{S,i}\eta_{S,i} \leq g_{S,i}\alpha_i$. Therefore, $\sum_{\substack{S \subset I \setminus \{i\} \\ |S| \leq \Gamma}} g_{S,i}\alpha_{S \cup i} + \sum_{\substack{S \subset I \setminus \{i\} \\ |S| > \Gamma}} g_{S,i}\eta_{S,i} \leq \sum_{S \subset I \setminus \{i\}} g_{S,i}\alpha_i \leq p^{\max}\alpha_i$, where the last inequality follows from the definition of $p^{\max}$. Note that the objective aims to maximize z_i. Since p_i is also a decision variable and increasing p_i increases the value for z_i, one can set $p_i^* = z_i^* + (1-\alpha_i)p^{\max}$. Therefore, we obtain:

$$0 \leq z_i \leq \sum_{\substack{S \subset I \setminus \{i\} \\ |S| \leq \Gamma}} g_{S,i}\alpha_{S \cup i} + \sum_{\substack{S \subset I \setminus \{i\} \\ |S| > \Gamma}} g_{S,i}\eta_{S,i} \tag{D.8}$$

$$0 \leq \sum_{\substack{S \subset I \setminus \{i\} \\ |S| \leq \Gamma}} g_{S,i}\left[\alpha_{S \cup i} - \alpha_S\right] + \sum_{\substack{S \subset I \setminus \{i\} \\ |S| > \Gamma}} g_{S,i}\left[\eta_{S,i} - \beta_S\right] + (1-\alpha_i)p^{\max}. \tag{D.9}$$

We next show that constraint (D.9), which is independent of z_i, always holds allowing us to identify z^*. Observe that $[\alpha_S - \alpha_{S \cup i}] \leq [1 - \alpha_i]$. The inequality follows from constraint (C.4) when $\{j\} = \emptyset$. Similarly, $[\beta_S - \eta_{S,i}] \leq (1 - \alpha_i)$ using the last constraint of (C.7). These inequalities along with the non-negative influence terms and the definition of $p^{\max}$ prove (D.9). Therefore, maximizing z_i results in $z_i^* = \sum_{\substack{S\subset I\setminus\{i\}\\|S|\le\Gamma}} g_{S,i}\alpha_{S\cup i} + \sum_{\substack{S\subset I\setminus\{i\}\\|S|>\Gamma}} g_{S,i}\eta_{S,i}$, hence proving our claim about (D.2).

Given the form of z^* and the fact that we are maximizing it, α_S and $\eta_{S,i}$ should be set at their maximum values as the influence terms are always non-negative. Assuming that all $\boldsymbol{\alpha}$ values are given, (D.3) and (D.4) follow from constraints (C.7) and constraints (C.5–C.6) respectively. We next show that the solution in (D.5) is also the largest feasible solution for every α_S variable. One can see that for $|S|=2$ the result holds (similar to (D.3) above). We next show inductively that (D.5) holds for larger $|S|$ values. Assume that it is true for some $|S|=k$. Constraints (C.3–C.4) are as follows:

$$\alpha_S \le \min_{k\in S\setminus\{i\}} \alpha_k \qquad \forall\, i\in S$$
$$\alpha_S \ge \min_{k\in S\setminus\{i\}} \alpha_k + \min_{k\in S\setminus\{j\}} \alpha_k - \min_{k\in S\setminus\{i,j\}} \alpha_k \qquad \forall\, S\ni\{i\}, S\ni\{j\}, \{i\}\ne\{j\}\subset\mathcal{I}.$$

The largest feasible value of α_S from the first set of constraints is $\min_{i\in S}\left[\min_{k\in S\setminus\{i\}}\alpha_k\right] = \min_{k\in S}\alpha_k$. We next show that this is feasible to the second constraint. Observe that $\min_{k\in S}\alpha_k$ is the same as $\min_{k\in S\setminus\{i\}}\alpha_k$ or $\min_{k\in S\setminus\{j\}}\alpha_k$ (or both if they are equal). In addition, $\min_{k\in S\setminus\{i,j\}}\alpha_k$ is always greater than either of those terms and therefore, the constraint holds. Note that by setting all variables at their largest values, we arrived at a feasible and revenue maximizing solution. □

Consider solving the relaxed version of Z-MIP, where the binary constraints for each α_i $\forall\, i\in\mathcal{I}$ are replaced by: $0\le\alpha_i\le 1$. Let $V^* = \left(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \boldsymbol{\eta}^*, \mathbf{p}^*, \mathbf{y}^*, \mathbf{z}^*\right)$ be a fractional optimal solution to the relaxed problem with a corresponding objective Π^*. We construct a new solution with all α_is being integer and show its feasibility to problem Z-MIP (hence relaxed Z-MIP as well) with an objective at least as good as V^*. We denote the solution we construct by $\tilde{V} = \left(\tilde{\boldsymbol{\alpha}}, \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\eta}}, \tilde{\mathbf{p}}, \tilde{\mathbf{y}}, \tilde{\mathbf{z}}\right)$ and its corresponding objective by $\tilde{\Pi}$. We next construct this new solution.

For any agent i, $\tilde{\alpha}_i = \lceil\alpha_i^*\rceil$ where $\lceil.\rceil$ refers to the ceiling function that maps a real number to the smallest following integer. Let T be the subset of agents who buy in $\tilde{V}$, i.e., $\tilde{\alpha}_i = 1$ which also refers to those with $\alpha_i^* > 0$. Then, $\tilde{\alpha}_S = 1$ for any $S\subset T$ and $|S|\le\Gamma+1$, and 0 otherwise. In addition, $\tilde{\beta}_S = 1$ for all sets $S\subset\mathcal{I}$ such that $|S|>\Gamma$ and $|S\cap T|\ge\Gamma$. Also, $\tilde{\eta}_{S,i}=1$ if both $\tilde{\beta}_S=1$ and $\tilde{\alpha}_i=1$. Finally, if $\tilde{\alpha}_i=1$, $\tilde{z}_i=\tilde{p}_i$ (described below) and $\tilde{y}_i=0$. Otherwise, $\tilde{z}_i=\tilde{y}_i=0$ and $\tilde{p}_i=p^{\max}$.

From Lemma 1, we know that:

$$\Pi^* = \sum_{i\in T}\Big[\sum_{\substack{S\subset I\setminus\{i\}\\|S|\le\Gamma}} g_{S,i}\alpha^*_{S\cup i} + \sum_{\substack{S\subset I\setminus\{i\}\\|S|>\Gamma}} g_{S,i}\eta^*_{S,i}\Big] - c\sum_{i\in T}\alpha_i^* = \sum_{i\in T}\Big[\sum_{\substack{S\subset T\setminus\{i\}\\|S|\le\Gamma}} g_{S,i}\alpha^*_{S\cup i} + \sum_{\substack{S\subset I\setminus\{i\}\\|S|>\Gamma,|S\cap T|\ge\Gamma}} g_{S,i}\eta^*_{S,i}\Big] - c\sum_{i\in T}\alpha_i^*$$

$$\tilde{\Pi} = \sum_{i\in T}\Big[\sum_{\substack{S\subset T\setminus\{i\}\\|S|\le\Gamma}} g_{S,i} + \sum_{\substack{S\subset I\setminus\{i\}\\|S|>\Gamma,|S\cap T|\ge\Gamma}} g_{S,i}\Big] - c|T|.$$

We next reorder the agents, rewrite the objective function, and argue that a certain property holds for a partial sum. We then introduce a sequence of iterative steps to show that $\tilde{\Pi}\ge\Pi^*$.

- Order the agents in the set $T=\{k_1,k_2,\dots,k_{|T|}\}$ such that $\alpha^*_{k_1}\geq\alpha^*_{k_2}\geq\cdots\geq\alpha^*_{k_{|T|}}$.
- Create the nested sets of agents: $T_1\subset T_2\subset\cdots\subset T_{|T|}=T$, where $\{k_1\}=T_1, T_1\cup\{k_2\}=T_2,\dots,T_m\cup\{k_{m+1}\}=T_{m+1},\dots,T_{|T|}=T$.
- Rewrite Π^* as follows:

$$\Pi^*=\sum_{m=1}^{|T|}\Big[\sum_{i\in T_{m-1}}\Big(\sum_{\substack{S\ni k_m\\ S\subset T_m\setminus\{i\},|S|\leq\Gamma}} g_{S,i}\alpha^*_{S\cup i}+\sum_{\substack{S\ni k_m,S\subset I\setminus\{i\}\\ |S|>\Gamma,|S\cap T_m|=\Gamma}} g_{S,i}\eta^*_{S,i}\Big)+\sum_{\substack{S\subset T_{m-1}\\ |S|\leq\Gamma}} g_{S,k_m}\alpha^*_{S\cup k_m}+\sum_{\substack{S\subset I\setminus\{k_m\}\\ |S|>\Gamma,|S\cap T_m|\geq\Gamma}} g_{S,k_m}\eta^*_{S,k_m}-c\alpha^*_{k_m}\Big].$$

  We build the above objective by considering the marginal terms obtained by adding one agent at a time starting from k_1 to $k_{|T|}$. For every agent that we add, say k_m, related value terms and cost terms are included. The value term corresponds to all influence terms related to all sets S that consists of k_m. In particular, they consist of terms where k_m is the influencer for all agents added so far (i.e., $i\neq k_m, i\in T_{m-1}$) and terms where k_m is influenced by previous agents (i.e., by agents in T_{m-1} on $i=k_m$).
- Substituting (D.3–D.5) from Lemma 1 in Π^*, we obtain:

$$\Pi^*=\sum_{m=1}^{|T|}\Big[\sum_{i\in T_{m-1}}\Big(\sum_{\substack{S\ni k_m\\ S\subset T_m\setminus\{i\},|S|\leq\Gamma}} g_{S,i}+\sum_{\substack{S\ni k_m,S\subset I\setminus\{i\}\\ |S|>\Gamma,|S\cap T_m|=\Gamma}} g_{S,i}\Big)+\sum_{\substack{S\subset T_{m-1}\\ |S|\leq\Gamma}} g_{S,k_m}+\sum_{\substack{S\subset I\setminus\{k_m\}\\ |S|>\Gamma,|S\cap T_m|\geq\Gamma}} g_{S,k_m}-c\Big]\alpha^*_{k_m}. \tag{D.10}$$

We know that $\Pi^*\geq 0$ because a no-buy for everyone results in 0 profit and is a feasible solution. Moreover, any cumulative sum starting from the last term in (D.10) has to be non-negative:

$$\Pi^*_{[l,|T|]}\geq 0\quad\forall l=1,\dots,|T|. \tag{D.11}$$

Otherwise, it would have been beneficial to set α_{k_m} values corresponding to these agents $(\alpha_{k_l},\dots\alpha_{k_{|T|}})$ to 0, and restrict T to just T_{l-1} (where $T_0=\emptyset$).

We now introduce a sequence of iterative steps to show that $\tilde{\Pi}\geq\Pi^*$. We choose a decreasing iteration counter, l, and show that the above property holds in each step.

1. Let $l=|T|$, $V^l=V^*$ and $\Pi^l=\Pi^*$.
2. If $\exists\ i\in\mathcal{I}$ s.t. $0<\alpha^l_i<1$, go to step 3. Otherwise, set $\tilde{\Pi}=\Pi^l$ and the procedure terminates.
3. For all $i\in\{k_l,\dots,k_{|T|}\}$, increase $\alpha^{l-1}_i=\alpha^l_{k_{l-1}}$ and no change otherwise, i.e., $\alpha^{l-1}_i=\alpha^l_i$.

   This results in $\Pi^{l-1}_{[l,|T|]}\geq\Pi^l_{[l,|T|]}$ as $\Pi^l_{[l,|T|]}\geq 0$ and $\Pi^{l-1}_{[i,i]}=\Pi^l_{[i,i]}\ \forall i=1,\dots,l-1$. Therefore, we obtain $\Pi^{l-1}_{[1,|T|]}\geq\Pi^l_{[1,|T|]}$. Also, from the previous step (or (D.11) when $l=|T|$), we know that $\Pi^l_{[i,|T|]}\geq 0\ \forall\ i=1,\dots,|T|$. Therefore, $\Pi^{l-1}_{[i,|T|]}\geq 0\ \forall\ i=1,\dots,|T|$.
4. Proceed back to step 2 after setting $l:=l-1$.

As a result, the algorithm terminates in at most $|T|$ steps and the final solution is such that all α_i values are integer as $\alpha_{k_0}=1$. Note that in the above steps, we did not discuss the feasibility of the solution in each step. It can be shown that the solution remains feasible to the relaxed Z-MIP in each iteration, although it is not relevant from the perspective of the proof and hence omitted. The final integer solution $\tilde{\alpha}$ is feasible to the relaxed Z-MIP and also to Z-MIP with an objective $\tilde{\Pi}\geq\Pi^*$, hence concluding the proof. □

Appendix E: 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. We denote by $I_T(\leq N)$ the total number of iterations and by $N^{(t)}$ the number of agents in the network at iteration $t=1,2,\ldots,I_T$.

Next, we show that the only candidates for the optimal uniform price are $p_{min}^{(t)}$ $\forall$ $t\in\{1,\ldots,I_T\}$ (see the definition in Algorithm 1). First, observe that the uniform optimal price cannot be smaller than $p_{min}^{(1)}$. Indeed, for any price $p\leq p_{min}^{(1)}$, all agents that bought in the discriminative case will still buy at this smaller price. However, a lower price than $p_{min}^{(1)}$ will result in lower profit (per buyer) for the seller. It is possible though that some new agents would buy the item at the lower price inducing an overall higher profit. Nevertheless, one can see that the new lower price in a uniform pricing scheme certainly will not be less than c. Therefore, it suffices to consider prices that are at least c but lower than $p_{min}^{(1)}$, if any. If this is the case, it would have been profitable to offer this price (which is higher than c) to those agents in the discriminative pricing scheme as well. Since it was not optimal to offer a lower price than $p_{min}^{(1)}$ to the non-buyers, it is not profitable to decrease the uniform price lower than $p_{min}^{(1)}$. As a result, 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 lose the buyers with $p_i\leq p_{min}^{(1)}$ from the discriminative pricing scheme. Otherwise, in the discriminative case one would offer a higher price. 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 I_T candidates, we just need to evaluate the corresponding profits (denoted by $\Pi^{(t)}$ $\forall t=1,2,\ldots,I_T$) and choose the one that yields the maximal profit. One can do so by using the following relation:

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

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

Appendix F: 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 4). 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 variables α_i^* $\forall i\in\mathcal{I}$ are integer. Second, we construct from the latter solution a new solution with all 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:

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

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

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

Here, 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 $\left(g_{jj} + \sum_i \gamma_i^* g_{ij} - p_j^*\right) = 0$ 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 hence, $y_j^* = 0$. 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 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, purchasing decisions of agents $i \in S_j$ remain the same. Third, since 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.

We 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 $p_j^* - c \leq 0$ as 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 $0 < \alpha_j^* < 1$, both $\overline{\Pi}$ and $\underline{\Pi}$ are lower or equal than Π^*. By requiring $\Pi^* - \overline{\Pi} \geq 0$ together with $\Pi^* - \underline{\Pi} \geq 0$ and using the fact that $p_j^* - c \leq 0$, we obtain the condition: $\alpha_j \geq \gamma_j$. However, from feasibility, we know that $\alpha_j \leq \gamma_j$ and thus $\alpha_j = \gamma_j$. 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 also optimal solutions as they are feasible and yield the same objective as 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 with an objective as good as the initial fractional solution. Since the number of agents is finite, this step is repeated at most N times.

Then, we 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.

At this point, we know 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 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 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 the following feasible solutions (denoted by $\overline{V}$ and $\underline{V}$) to the relaxed problem Zi:

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

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

As before, 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. First, we note that we construct the dual variables for agent j to satisfy all 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 neighbors of agent j exactly by the change in the level of influence from agent j on them and hence, the purchasing decisions of agents $i \in S_j$ remain the same. Third, since 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 also remains feasible for them.

We 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 $0 < \gamma_j^* < 1$, $d_j^* - \sum_{i \in S_j} \alpha_i^* g_{ji} = 0$ should hold. Otherwise, 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 also optimal solutions as they are feasible and yield the same objective as 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 problem with an objective 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 with integer purchasing decisions and integer γ variables. $\square$