Theoretical Computer Science 476 (2013) 1–15

Contents lists available at SciVerse ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

Equilibrium pricing with positive externalities[☆]

Nima AhmadiPourAnari[a], Shayan Ehsani[b], Mohammad Ghodsi[c,d], Nima Haghpanah[e], Nicole Immorlica[e], Hamid Mahini[c,*], Vahab Mirrokni[f]

[a] *University of California at Berkeley, EECS Department, United States*
[b] *Stanford University, Management Science and Engineering Department, United States*
[c] *Sharif University of Technology, Computer Engineering Department, Iran*
[d] *Institute for Research in Fundamental Sciences (IPM), P.O. Box: 19395-5746, Tehran, Iran*
[e] *Northwestern University, EECS Department, United States*
[f] *Google Research NYC, 76 9th Ave, New York, NY 10011, United States*

ARTICLE INFO

Article history:
Received 14 August 2011
Received in revised form 10 April 2012
Accepted 23 January 2013
Communicated by P.G. Spirakis

Keywords:
Social networks
Game
Revenue maximization
Market
Equilibrium
Historical externalities
Pricing

ABSTRACT

We study the problem of selling an item to strategic buyers in the presence of positive *historical externalities*, where the value of a product increases as more people buy and use it. This increase in the value of the product is the result of resolving bugs or security holes after more usage. We consider a continuum of buyers that are partitioned into *types* where each type has a valuation function based on the actions of other buyers. Given a fixed sequence of prices, or *price trajectory*, buyers choose a day on which to purchase the product, i.e. they have to decide whether to purchase the product early in the game or later after more people already own it. We model this strategic setting as a game, study existence and uniqueness of the equilibria, and design an FPTAS to compute an approximately revenue-maximizing pricing trajectory for the seller in two special cases: the *symmetric* settings in which there is just a single buyer type, and the *linear* settings that are characterized by an initial type-independent bias and a linear type-dependent influenceability coefficient.

1. Introduction

Many products like softwares, electronics, or automobiles evolve over time. When a consumer considers buying such a product, he faces a tradeoff between buying a possibly sub-par early version versus waiting for a fully functional later version. Consider, for example, the dilemma faced by a consumer who wishes to purchase the latest Windows operating system. By buying early, the consumer takes full advantage of all the new features. However, operating systems may have more bugs and security holes at the beginning, and hence a consumer may prefer to wait with the rationale that, if more people already own the operating system, then more bugs will have already been uncovered and corrected. The key observation is, the more people that use the operating system, or any product for that matter, the more inherent value it accrues. In other words, the product exhibits a particular type of externality, a so-called *historical externality*.[1]

☆ A preliminary version of this paper appeared with the same name in WINE 2010 AhmadiPourAnari et al. (2010) [1].

* Corresponding author. Tel.: +98 9124660470.
E-mail addresses: nima.ahmadi.pour.anari@gmail.com (N. AhmadiPourAnari), shayan.ehsani@gmail.com (S. Ehsani), ghodsi@sharif.edu (M. Ghodsi), nima.haghpanah@gmail.com (N. Haghpanah), nicimm@gmail.com (N. Immorlica), hamid.mahini@gmail.com, mahini@ce.sharif.edu (H. Mahini), mirrokni@google.com (V. Mirrokni).

1 Note that this is different from the more well-studied notion of externalities in the computer science literature where a product (e.g., a cell phone) accrues value as more consumers buy it simply because the product is used in conjunction with other consumers.

0304-3975/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
doi:10.1016/j.tcs.2013.01.014

How should a company price a product in the presence of historical externalities? A low introductory price may attract early adopters and hence help the company extract greater revenue from future customers. On the other hand, too low a price will result in significant revenue loss from the initial sales. Often, when faced with such a dilemma, a company will offer an initial promotional price at the product's release in a limited-time offer, and then raise the price after some time. For example, when releasing Windows 7, Microsoft announced a two-week pre-order option for the Home Premium Upgrade version at a discounted price of $50; thereafter the price rose to $120, where it has remained since the pre-sale ended on July 11th, 2009. Additionally, beta testers, who can be interpreted as consumers who "bought" the product even prior to release, received the release version of Windows 7 for free (as is often the case with software beta-testers).[2]

We study this phenomenon in the following stylized model: a monopolistic seller wishes to derive a pricing and marketing plan for a product with historical externalities. To this end, she commits to a price trajectory. Potential consumers observe the price trajectory of the seller and make simultaneous decisions regarding the day on which they will buy the product (and whether to buy at all). The payoff of a consumer is a function of the day on which he bought the product, the price on that day, and the set of consumers who bought before him. We compute the equilibria of the resulting sequential game and observe that the revenue-maximizing price trajectories for the seller are increasing, as in the Windows 7 example above.

A few words are in order about our model. First, we focus on settings in which the seller has the ability to commit to a price trajectory. Such commitments are observed in many settings especially at the outset of a new product (see the Windows 7 example described above) and have been assumed in prior models in the economics literature on pricing as well as in other games in the form of Stackelberg strategies (see Section 1.2). Further, commitment increases revenue: clearly a seller who commits to price trajectories can extract at least as much revenue as a seller who does not (or cannot). This is because he can commit to the outcome that would have happened in the absence of commitment. Notice that in our game of complete information in the presence of rational agents, the seller (in fact, all the players) can precisely predict the outcome of the game without commitment. We further observe via example in Section 2 that in fact commitment enables a seller to extract unboundedly higher revenue than in settings without commitment. Second, we assume a consumer's payoff is only a function of past purchases; i.e. consumers have no utility for future purchases. We motivate this in the Windows 7 example by arguing that bugs are resolved in proportion to usage rates. Of course, strictly speaking, consumers of Windows 7 benefit from future purchases as well via software updates and the like. However, this forward-looking benefit is substantially dampened in comparison to past benefits by safety and security risks, and time commitments involved in updates. Another justification for our payoff model comes from consumers' uncertainties regarding products. In many settings, consumers have signals regarding the value of a product (say an electronic gadget like the iPad for example), but do not observe its precise value until the time of purchase. Past purchases and the ensuing online reviews may help consumers improve their estimates of their values prior to purchase, an especially important factor for risk-adverse buyers.

1.1. Our results

We focus on the non-atomic setting in which we have a continuum of consumers so that each consumer is infinitesimally small and therefore his own action has a negligible effect on the actions of others. Consumers are drawn from a (possibly infinite) set of types. These types capture varying behavior among consumer groups. We study a sequential game in which the seller first commits to a price trajectory and then the consumers simultaneously choose when and whether to buy in the induced normal-form game among them. We study subgame perfect equilibria. We first observe in Section 3 that equilibria exist due to a slight generalization of a paper of [18].

We then turn to the question of uniqueness in Section 4. We focus on *well-behaved equilibria* in which consumers with non-negative utility always purchase the product (thus indifferent consumers purchase the product). In general multiple such equilibria may exist. However, in an aggregate model in which the value function of each consumer type depends only on the *aggregate* behavior of the population (i.e. the total fraction of potential consumers that have bought the product and not the total fraction of various types), then we are able to show that when they exist, well-behaved equilibria of this game are unique in the sense that the fraction of purchases per-type-per-day is fixed among all equilibria. This enables us to search for the revenue-maximizing price trajectory. In Section 5, we address this question in settings in which we either have just one type or there are multiple types whose valuation functions are linear in the aggregate, both of which are special cases of the aggregate model discussed above. For each price trajectory, we define its revenue to be the amount of money consumers spend on the product. We first show that given oracle access to the valuation functions, an optimal solution can be found only with infinitely many queries. We then design an FPTAS to find the revenue-maximizing price trajectory for a monopolistic seller in these settings. We do this via a reduction to a novel rectangular covering problem in which we must find the discounted area-maximizing set of rectangles that fit underneath a given curve. We study the rectangular covering problem in Section 6.

In summary, we get the following main result for the settings described above: first, every price trajectory has well-behaved equilibria that are revenue-unique and revenue-maximal among all possible consumer equilibria. Second, it is

[2] Historical prices, announced upon the press release, can be found in archived versions of various technology news websites such as Ars Technica [22] and the Microsoft blog [16]. The current prices were accessed on Microsoft's website at the time of submission.

possible to (approximately) compute the revenue-maximizing price trajectory. Hence, the strategy tuple in which the seller announces this (approximately) revenue-maximizing price trajectory and the consumers respond by playing a well-behaved equilibrium is an (approximately) subgame perfect equilibrium of our game.[3]

1.2. Related work

Our work falls in the long line of literature investigating pricing and marketing of products that exhibit externalities [2–5,7,10–12,14,15,19,20,24]. Among these, the paper of [5] is most closely related to our own work. [5] analyzed the two and infinite period pricing problems in the presence of linear historical externalities and the study equilibria of the induced games both with and without commitment. They observe, as we do, that optimal price trajectories are increasing. The historical externalities that we study generalize the externalities of [5], and in this more general model, we solve for the optimal price sequence for any fixed number of price periods. Most of the remaining externalities literature studies externalities in which consumers care about the total population of users of a product and hence their utility is affected by future sales as well as past sales. Although the phenomenon studied is different from ours, some of the modeling assumptions in these papers are similar to ours. For example, in the economics literature [7], also consider a seller that commits to a price trajectory and then observe that the revenue-maximizing price sequence with fully rational consumers (playing a Bayesian equilibrium) is increasing. Similar to our model, they study the pricing problem in the presence of a continuum of consumers.

In the computer science literature [2,12,6] study algorithmic questions regarding revenue maximization over social networks for products with externalities. However, their models assume naive behavior for consumers. Namely, they assume consumers act myopically, buying the product on the first day in which it offers them positive utility without reasoning about future prices and sales that could affect optimal buying behavior and long-term utility. Furthermore, [12] allow the seller to use adaptive price discrimination. In contrast, we model consumers as fully rational agents that strategically choose the day on which to buy based on full information regarding all future states of the world and a sequence of public posted prices. While the correct model of pricing and consumer behavior probably lies somewhere between these two extremes, we believe studying fully rational consumers is an important first step in relaxing myopic assumptions. [6] study such a model in which players act strategically. They study a two stage pricing problem in which the seller first determines the prices for a *divisible* good, and then the buyers simultaneously decide about the amount of consumption. In contrast to our model and similar to [12], the seller is allowed to use price discrimination. Also, the strategy of buyers in their setting is the amount of consumption, not the time in which consumption happens. They assume that each player's utility is a non-decreasing function of the consumption of other buyers. Similar to our model, they study the subgame perfect equilibria of the two-stage game. They study the structural properties of the optimal prices, as well as the algorithmic problem of finding the optimal price sequence.

Our work can be viewed as an analysis of the optimal commitment strategy of a seller in the presence of historical externalities. [23] first studied optimal commitment strategies, called Stackelberg strategies, in the Cournot's duopoly model [9], proving that a firm who can commit has strategic advantage. The computational aspects of Stackelberg strategies have been studied in generic normal form [8,21] and extensive form games [17]. [17] suggested polynomial-time algorithms for various models of perfect information extensive form games and proved NP-hardness of the problem with imperfect information even with two players. A perfect information extensive form game is one in which at every stage of the game, every player is exactly aware of what has happened earlier in the game. This does not happen in games with simultaneous move, as in our game.

2. Model

We wish to study the sale of a good by a monopolistic seller over k days to a set of potential consumers or buyers. We model our setting as a sequential game whose players consist of the monopolistic seller and a continuum of potential consumers or buyers $b \in [0, 1]$. In our game, the seller moves first, selecting a *price trajectory* $p = (p_1, \ldots, p_k)$ where $p_i \in \mathbb{R}$ assigns a (possibly negative) price p_i to each day i. The buyers move next, selecting a day on which to buy the product given the complete price trajectory, as described below.

The buyers are partitioned into n types $T_1, \ldots, T_n$ where for $t \in \{1, \ldots, n\}$, T_t is a subinterval of $[0, 1]$.[4] We say that a buyer b has type t if $b \in T_t$. The strategy set $A = \{1, \ldots, k\} \cup \{\emptyset\}$ indicates the day on which the product is bought ($\emptyset$ is used to indicate that the product was not purchased). Hence the strategy profile of the buyer population can be represented by a $(k+1) \times n$ matrix $X = \{X_{i,t}\}_{i=1,\ldots,k+1;t=1,\ldots,n}$ where entry $X_{i,t}$ indicates the fraction of buyers that are of type t and buy the product *before* day i, and we define $X_{1,t} = 0$ for all t. Note that by normalization $\sum_t X_{k+1,t} \leq 1$ and $1 - \sum_t X_{k+1,t}$ is the fraction of buyers that do not buy the product at any time. Corresponding to this matrix X we also define the *marginal strategy profile* matrix $x = \{x_{i,t}\}_{i=1,\ldots,k;t=1,\ldots,n}$ where $x_{i,t} = X_{i+1,t} - X_{i,t}$ is the fraction of buyers who are of type t and buy on

[3] We must also define what happens off the equilibrium path in subgame perfect equilibria; here we assume that for any price trajectory announced by the seller, the consumer respond by playing a well-behaved equilibrium.

[4] Later, we will generalize this to infinitely many types.

day i. In the special case when there is only 1 type, we use X_i as a scalar to denote the fraction of buyers who buy before day i and x_i as a scalar to denote the fraction of buyers who buy on day i.

Given a strategy profile X, we define the value of buyers of type t buying on day i by a value function $F_i^t(X_i)$ where X_i is the i'th row of X (hence buyers are indifferent to future buying decisions). Note the explicit dependence of F on time, which allows $F_i^t(X_i)$ to be different than $F_j^t(X_j)$, for $i \neq j$. The revenue-maximization results in Section 5 further assume that the dependence of $F_i^t(X_i)$ on i is of the form $F_i^t(X_i) = \beta^i F^t(X_i)$ for $\beta \in [0, 1]$. This special case is of particular interest as the β factor models settings in which the value degrades over time due to, for example, a reduction in the novelty of the product.

Given a strategy profile X, the payoff of buyers of type t who buy on day i is defined to be $F_i^t(X_i) - p_i$. We additionally allow buyers to have a discount factor α such that their payoff is $(1-\alpha)^i(F_i^t(X_i) - p_i)$. Thus α represents the way in which agents discount future payoffs with respect to present payoffs. We say that a strategy profile X is a Nash equilibrium of the induced subgame given by price trajectory p, or equivalently $X \in NE(p)$, if for any buyer of type t who buys on day i we have $i \in \arg\max_j (F_j^t(X_j) - p_j)(1-\alpha)^j$, and the strategy is $\emptyset$ whenever the maximum is not positive (in which case the buyer's payoff is zero). We call an equilibrium *well-behaved* if all indifferent buyers buy, i.e. a buyer does not buy if and only if his payoff $(1-\alpha)^i(F_i^t(X_i) - p_i)$ is negative on all days $1 \leq i \leq k$. We say that (p, X) is a (well-behaved) equilibrium if the profile X is a (well-behaved) Nash equilibrium for the subgame of price trajectory p. Equivalently, a marginal strategy profile x of a strategy profile X is a (well-behaved) Nash equilibrium for the subgame of price trajectory p if the profile X is a (well-behaved) Nash equilibrium for the subgame of price trajectory p.

Given a price trajectory p and a marginal strategy profile x that arises in the subgame induced by p, we define the payoff of the seller to be the *revenue* of x for p, which is $R(p, x) = \sum_{i=1}^{k}\sum_{t=1}^{n} x_{i,t} p_i (1-\alpha)^i$. A *subgame perfect equilibrium* of the sequential game is then a price trajectory p^* and a set of marginal strategy profiles x_p for each possible price trajectory p such that: (1) x_p is a Nash equilibrium of the subgame induced by p, and (2) p^* maximizes $R(p, x_p)$. The *outcome* of this subgame perfect equilibrium is (p^*, x_{p^*}) and its revenue is $R(p^*, x_{p^*})$.

We are interested in computing the outcome in a revenue-maximizing subgame perfect equilibrium. To do so, we must compute a price trajectory which maximizes the revenue of the seller in equilibrium. Note that this is equal to finding the best response of the seller given the strategies $\{x_p\}$ of the buyers. We solve this problem for special settings in which there exist revenue-maximizing well-behaved equilibria in $NE(p)$ for any price trajectory p, allowing us to maximize over them. These settings are as follows. For the purpose of these definitions, we allow each buyer to have a unique type and hence there are infinitely many types. We will use $b \in [0, 1]$ to denote the type of buyer b.

Definition 1. The aggregate model: The value function of each type in this model is a function of the *aggregate* behavior of the population and is invariant with respect to the behavior of each separate type. That is, the value function of buyer b is a function of X_i only, where X_i is a scalar indicating the total fraction of all buyers who buy before day i. In this instance, we overload the notation for the value function and let $F_i^b(X_i)$ indicate the value of buyer b (hence $F_i^b(\cdot)$ now maps the unit interval to the non-negative reals).

Definition 2. The linear model: This is a special case of the aggregate model which is defined by a function F_i, an initial bias $I \in \mathbb{R}^+$, and a function $C : [0, 1] \to \mathbb{R}^+$ so that the value of buyer b is $F_i^b(X_i) = I + C(b) \cdot F_i(X_i)$. We further define the commonly-known distribution $\mathcal{C} : \mathbb{R} \to [0, 1]$ such that $\mathcal{C}(c^*)$ indicates the fraction of buyers b with $C(b) \leq c^*$.

Definition 3. The symmetric model: In this version we only have one type, that is, $F_i^b = F_i$ for all b.

We note that alternatively, one could model this pricing game as a sequential game with multiple stages where in each day i the seller selects a price p_i and then buyers simultaneously choose whether to buy or not. Such a model is appropriate when it is not possible for a seller to commit to a price trajectory in advance. Again, in this setting, one could study the subgame perfect equilibria and analyze the resulting revenue. The revenue with commitment is at least as high as that without commitment, since the seller can commit to the outcome of the game without commitment. The example below shows that the revenue without commitment can actually be unboundedly less.

Example. We study the outcome of the sequential game without commitment. We consider buyers $b \in [0, 1]$ with utility function $F^b(X) = (1 + a|X|)b$ for a non-negative constant a, and focus on a setting with $k = 2$ days. In a strategy profile, the seller picks a price p_1 for day 1 and a price $p_2(S)$ for any subset $S \subseteq 2^{[0,1]}$ of buyers that remain in day 2. Each buyer decides whether to buy on day 1 given any possible announcement p_1 of the seller, and also whether to buy on day 2 given any possible announcement p_2 of the seller and decisions S of other buyers. A strategy profile is a subgame perfect equilibrium if each player is playing a best response given the strategies of others. In particular, this implies that the seller must set $p_2(S)$ equal to the revenue-maximizing price for buyers S.

Assume a subgame perfect equilibrium in which the prices on path of play are p_1 and p_2, and subsets $S_1, S_2 \subseteq 2^{[0,1]}$ buy at days 1 and 2, respectively. We claim $p_1 \geq p_2$. Suppose for contradiction that $p_1 < p_2$. First note that if $|S_1| = 0$, then buying on day 1 gives higher utility that buying on day 2, and therefore $S_2 = \emptyset$. But this gives the seller zero revenue at day 2 which is not his best strategy as he can lower the price of day 2 and gain positive revenue. So assume from now on that $|S_1| > 0$. The set of buyers who buy at the first day are the ones who have non-negative utility of buying on day 1, and also more utility of buying on day 1 than day 2. The first constraint is equal to having $b \geq p_1$, and the second to having b such that $b(1+0) - p_1 > b(1+a|S_1|) - p_2$, or equivalently $b < \frac{p_2 - p_1}{a|S_1|}$ (remember that we assumed $|S_1| > 0$). Let $b^* = \frac{p_2 - p_1}{a|S_1|}$. The set of buyers buying on day 1 is equal to the buyers with b satisfying $p_1 \leq b$ and also $b < b^*$. Our assumption that $|S_1| > 0$ now implies that $p_1 < b^*$. We observe that $p_2 \in (p_1(1+a|S_1|), b^*(1+a|S_1|))$, as can be seen from the equality $p_2 = a|S_1|b^* + p_1$ and the inequality $b^* > p_1$. By the properties of subgame-perfect equilibria, the seller has to play his best response in the subgame induced at day 2. In day 2 the utilities of players are outside the interval $(p_1(1+a|S_1|), b^*(1+a|S_1|))$. In fact the utility of player b at day 2 is in the interval $(p_1(1+a|S_1|), b^*(1+a|S_1|))$ if and only if $p_1 < b$ and $b < b^*$ which implies the player b bought in day 1. So $p_2 \in (p_1(1+a|S_1|), b^*(1+a|S_1|))$ cannot be a revenue-maximizing price. We conclude that in any subgame perfect equilibrium, the seller has to set the first price at least equal to second price, and no player will buy in the first day. The highest revenue for the seller is to set the price of the second day equal to the monopoly price (and any weakly higher price for the first day), which gives revenue 1/4.

Now consider the game with commitment in which the seller selects a price trajectory (p_1, p_2) and then buyers choose their actions. Now the profile of strategies $p_1 = 1/2$, $p_2 = \frac{8+3a}{16}$, $S_1 = [1/2, 3/4]$, $S_2 = [3/4, 1]$ is an equilibrium, and its revenue is $1/4 + 3a/64$. We conclude by observing that the ratio of the revenue with commitment to that without goes to infinity as a grows.

An important observation is that without externalities (when $a = 0$), both equilibria would have been the same, with revenue equal to the optimal revenue of the 1-stage game. This is in fact true for any form of utility functions.

3. Existence

In this section we study the existence of equilibria in the subgame induced by a given price trajectory p. For the aggregate model, the existence of equilibria follows from a result of [18]. Here we use Kakutani's fixed point theorem [13] to prove a slight generalization of this result applicable to models with finitely many types and continuous valuation functions. We also give examples of games with a non-continuous valuation function in which no equilibria exist.

Note that while the above theorems prove existence of equilibria, they do not guarantee that the equilibria are necessarily well-behaved (i.e. all indifferent buyers buy at some point). Well-behaved equilibria may not exist. Furthermore, even when they exist they do not necessarily maximize revenue. However, for the linear and symmetric models that we focus on, we can show that well-behaved equilibria do in fact exist and maximize revenue. The existence results are presented in Section 3.1; the results in Section 5 imply the revenue-maximizing property.

We show the pricing game has an equilibrium whenever buyers have finitely many types and continuous valuation functions F_i^t, $1 \leq t \leq n$, $1 \leq i \leq k$. To do so, we define a set-valued function on the space of marginal strategy profile matrices whose fixed point is an equilibrium of our game and show the existence of a fixed point using Kakutani's fixed point theorem (KFPT).

We start by defining the notation used in the statement of KFPT. For arbitrary sets S_1 and S_2, let $\phi : S_1 \to 2^{S_2}$ be a set-valued function, i.e. a function from S_1 to the power set of S_2. We say that ϕ has a closed graph if the set $\{(s_1, s_2) | s_1 \in S_1, s_2 \in \phi(s_1)\}$ is a closed subset of $S_1 \times S_2$ in the product topology.

Theorem 1. *Kakutani Fixed Point Theorem (KFPT) [13]: Let S be a non-empty, compact and convex subset of Euclidean space R^n for some n. Let $\phi : S \to 2^S$ be a set-valued function on S with a closed graph and the property that $\phi(s)$ is non-empty and convex for all $s \in S$. Then ϕ has a fixed point s such that $s \in \phi(s)$.*

We are now ready to prove the equilibrium's existence. We define ϕ to be the correspondence that maps strategy profile matrices to the set of best-response matrices. That is, for a strategy profile x, $\phi(x)$ is a strategy profile in which each buyer chooses the action that gives him maximum utility, assuming that other buyers are using strategy x. After showing that the mapping ϕ has the desired properties, we invoke KFPT to show that this mapping has a fixed point. This implies that existence of a strategy profile in which each buyer is playing best response to the actions of others, and is therefore by definition an equilibrium.

Theorem 2. *If valuation functions are increasing and continuous, then the subgame induced by any price trajectory p has an equilibrium in the general model with finitely many types.*

Proof. Let S be a subset of the Euclidean space $\mathbb{R}^{k \times n}$ consisting of all valid marginal profile matrices x. Also let $\mu(t) = |T_t|$ be the length of T_t for each type t (recall that T_t is a subinterval of $[0, 1]$). Each $x \in S$ is a marginal strategy profile matrix $x = (x_{1,1}, \ldots, x_{k,n})$, where $x_{i,t}$ is the fraction of buyers who are of type t and choose to buy on day i, with the constraint that for each $1 \leq t \leq n$, the inequality $\sum_i x_{i,t} \leq \mu(t)$ holds. Define $\phi : S \to 2^S$ to be the function assigning each $x \in S$ the set of all marginal strategy profile matrices $y \in \phi(x)$ which are simultaneous best-responses to the profile x. Formally, $\phi(x)$ consists of all y satisfying the following conditions:

1. A buyer buys in y only if they get non-negative utility in x: $\sum_i y_{i,t} > 0$ only if there exists j such that $F_j^t(X_j) - p_j \geq 0$,
2. If some type has a positive utility in x, then they all buy in y: if $F_j^t(X_j) - p_j > 0$ then $\sum_i y_{i,t} \geq \mu(t)$,
3. If a buyer buys in y, then he does so on a day which gives him maximum utility in x i.e. for every day i and type t, if $y_{i,t} > 0$ then $i \in \arg\max_j F_j^t(X_j) - p_j$.

If the conditions of the KFPT hold, we get a fixed point $x \in S$; i.e. a point x for which $x \in \phi(x)$. As discussed above, any such fixed point is an equilibrium of our game. Now let us prove that the set S and the function ϕ satisfy the conditions of KFPT. Set S can be defined as the set of points $x \in \mathbb{R}^{k\times n}$ satisfying the following linear inequalities, i.e. $\forall i, t : x_{i,t} \geq 0$, and $\forall t : \sum_i x_{i,t} \leq \mu(t)$. As a result, S, being the intersection of half-spaces, is a polyhedron, and clearly is closed and convex. The set S is also bounded, because each $x_{i,t}$ lies in the interval $[0, \mu(t)]$. So S is a compact and convex subset of $\mathbb{R}^{k\times n}$.

Let x be an arbitrary point in S. The set $\phi(x)$ can be defined as the intersection of S, and a set of (possibly open) half-spaces defined by linear inequalities listed in the conditions above. Thus, $\phi(x)$ is a convex set. It is also nonempty as each buyer of each type t has a well-defined set of best-responses to X, the cumulative corresponding to marginal profile x, which is either some day j if $F_j^t(X_j) - p_j \geq 0$ or the empty strategy (not buying) otherwise.

It only remains to show that the graph of ϕ is a closed subset of $\mathbb{R}^{2(k\times n)}$. We will show that each (x, y) lying outside the graph is contained in an open neighborhood which also lies outside the graph. This neighborhood will be of the form $A \times B$, where A is an open neighborhood of x and B is an open neighborhood of y. Since (x, y) is not in the graph, either (x, y) is not in $S \times S$ or y does not satisfy one of the conditions defining $\phi(x)$. In the former case, since $S \times S$ is closed, we can find a suitable neighborhood of (x, y) having no intersection with $S \times S$ and by extension the graph of ϕ.

So assume that y does not satisfy at least one of the constraints defining $\phi(x)$. Let $U : \mathbb{R}^{k\times n} \to \mathbb{R}^{k\times n}$ be the function assigning each $x \in S$ the matrix of utilities $\{u_{i,t}\}$, where $u_{i,t} = F_i^t(X_i) - p_i$ denotes the utility of buying on day i for a buyer of type t. Assuming the valuation functions are continuous and increasing, U and hence U^{-1} is continuous and invertible. As $y \notin \phi(x)$, there is some type t such that either:

1. $\sum_j y_{j,t} > 0$ and for all days j, $u_{j,t} < 0$: Let $A = \{U^{-1}(\{u_{i,t}\}) \mid \forall j : u_{j,t} < 0\}$, $B = \{\{y_{i,t}\} \mid \sum_j y_{j,t} > 0\}$,
2. Or $\max_j u_{j,t} > 0$ and $\sum_i y_{i,t} < \mu(t)$: Let $A = \{U^{-1}(\{u_{i,t}\}) \mid \exists i : u_{i,t} > 0\}$, $B = \{\{y_{i,t}\} \mid \sum_i y_{i,t} < \mu(t)\}$,
3. Or there exists $j^* \notin \arg\max_j u_{j,t}$ and $y_{j^*,t} > 0$: Let $A = \{U^{-1}(\{u_{i,t}\}) \mid \exists i : u_{i,t} > u_{j^*,t}\}$, $B = \{\{y_{i,t}\} \mid y_{j^*,t} > 0\}$.

Then $A \times B$ is an open neighborhood containing (x, y) since U^{-1} is continuous and invertible and we are applying it to an open subset of the domain in each case. Also, $A \times B$ has no intersection with the graph of ϕ, which proves that the graph of ϕ is closed. Hence all assumptions of the KFPT hold and we have an equilibrium. □

We next show via an example that the game may have no equilibria when valuation functions are not continuous.

Example 3. *Suppose there are two days $k = 2$ and only one type in the market. Let $F_i(X) = 1$ if $X \leq \frac{1}{2}$ and $F_i(X) = \frac{3}{2} + X$ otherwise, for $i = 1, 2$. We claim price trajectory $p = (0, 1)$ has no equilibrium. Consider any strategy profile X. If $x_1 \leq 1/2$, then the payoff of day 1 is 1 and day 2 is zero, so buyers who do not buy in day one are not playing a best-response, and X is not an equilibrium. On the other hand, if $x_1 > 1/2$ then the payoff on day 1 is still 1 but the payoff on day 2 is now strictly greater than $3/2 + 1/2 - 1 \geq 1$. Hence the buyers who buy in day one are not playing a best-response and so X is not an equilibrium.*

While the previous example is enough to show that continuity is necessary for existence of equilibria, one might notice that we can resolve the issue by changing the function to $F_i(X) = 1$ if $X < \frac{1}{2}$, and $F_i(X) = \frac{3}{2} + X$ otherwise, for $i = 1, 2$. In this case $x = (1/2, 1/2)$ is an equilibrium for price trajectory $p = (0, 1)$. The following example is more robust to the changes in the valuation function.

Example 4. *Assume there are three types and the fraction of each type is $\frac{1}{3}$. The valuation function $F_i^t(X) = 2$ if $X_{t'} < \frac{1}{3}$ and $F_i^t(X) = 4$, where $t' = 2, 3, 1$ respectively for $t = 1, 2, 3$. Then similar reasoning shows price trajectory $p = (1, 2)$ has no equilibrium.*

3.1. Existence of well-behaved equilibria

Our revenue results assume the existence of revenue-maximizing well-behaved equilibria for all price sequences. Unfortunately, this does not hold even for the aggregate model, as the following example shows:

Example 5. *Suppose there are three days $k = 3$ and two types in the market, each of which contains half of the total population. Let $F_i^1(X_i)$ and $F_i^2(X_i)$ be as shown in Fig. 1, for $i = 1, 2, 3$. Note that these are aggregate valuation functions. Consider the price trajectory $p = (1, 2, 3)$.*

If we assume that a buyer may decide not to buy the product when the utility of buying is 0, then we have an equilibrium in this example. The vector strategy profile

$$x = \begin{bmatrix} 0 & 0.25 \\ 0.35 & 0 \\ 0 & 0.25 \end{bmatrix}$$

is an equilibrium and a 0.15 fraction of type 1 will not buy the product. However a case analysis shows there are no well-behaved equilibria.

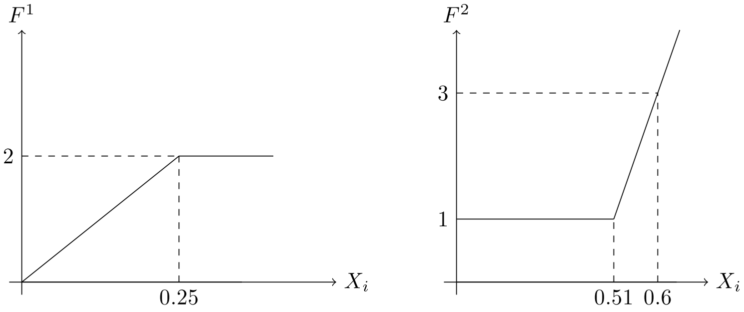

Fig. 1. The valuation function of type 1 and 2.

Fortunately, we can show that for the linear and symmetric models, well-behaved equilibria exist. We actually show something a bit more general; namely we derive a condition on the utilities that is sufficient to guarantee the existence of well-behaved equilibria in which either *all* or *no* buyers buy. We will present the proof for finitely many types for simplicity. The proof extends to infinitely many types.

Theorem 6. *If either* $\min_{t,i} F_i^t(0) \geq \min_i p_i$ *or* $\max_{t,i} F_i^t(0) < \min_i p_i$, *then there exists a well-behaved equilibrium in the subgame induced by p.*

Proof. If $\max_{t,i} F_i^t(0) < \min_i p_i$ then the strategy profile in which no one buys is an equilibrium, because when no one else is buying, a player of type t would receive utility $F_i^t(0) - p_i$ from buying at day i, which is negative. Hence everyone would receive negative utility from buying and the strategy profile is a well-behaved equilibrium in which no buyer buys.

So assume that $\min_{t,i} F_i^t(0) \geq \min_i p_i$. Consider the modified price trajectory $(p_1 - \epsilon, \ldots, p_k - \epsilon)$ for some $\epsilon > 0$ (recall that our model permits negative prices). Using this price sequence, all buyers buy on some day since the utility of buying on the day with the minimum price is strictly positive. Hence the new price trajectory has a well-behaved equilibrium. We will show that this well-behaved equilibrium is also an equilibrium for the original price trajectory. When prices are all decreased by the same amount all utilities are also decreased by that same amount. Hence the relative ordering of utilities is the same for both price trajectories. So after changing the prices back to the original ones, everyone is still buying an optimal day. The only equilibrium condition that might not hold anymore is the one asserting that buyers only buy when the optimum utility is non-negative. However, since $\min_{t,i} F_i^t(0) \geq \min_i p_i$, everyone has non-negative utility on the day with the minimum price and so has non-negative optimal utility. Therefore this is a well-behaved equilibrium for the original price trajectory in which all buyers buy. □

If $F_i^t(0)$ is the same for all i, t, then $\max_{t,i} F_i^t(0) = \min_{t,i} F_i^t(0)$. So at least one of the conditions of Theorem 6 holds and we can conclude Corollary 7. Note that for both the symmetric and linear models, $F_i^t(0)$ is the same for all i, t.

Corollary 7. *If* $F_i^t(0)$ *is the same for all* i, t *then there is always a well-behaved equilibrium.*

4. Uniqueness of equilibria

The following example shows that if we allow the valuation functions to be sensitive to the behavior of each type separately, our game might have more than one equilibrium, with different revenues for the seller. This holds even when all the valuation functions are continuous.

Example 8. *Assume that there are two types and* $F_i^t(Y) = Y_{t'} + 2$ *where* $t = 1, 2$ *and* $t' \neq t$. *In other words a valuation functions depends only on the behavior of buyers of the other type. The population of type* 1 *buyers are* 0.3 *and the population of type* 2 *buyers are* 0.7. *Suppose that the seller wants to sell the product in two days and* $p_1 = 1$ *and* $p_2 = 1.2$.

Two strategy profiles $x = \begin{bmatrix} 0.3 & 0 \\ 0 & 0.7 \end{bmatrix}$ *and* $x' = \begin{bmatrix} 0 & 0.7 \\ 0.3 & 0 \end{bmatrix}$ *are equilibria of this game. In the first equilibrium all the type* 1 *(respectively* 2*) buyers buy the product on day* 1 *(respectively* 2*). The revenue is* $0.3 \times 1 + 0.7 \times 1.2 = 1.14$ *in this equilibrium. In the second one all the type* 1 *(respectively* 2*) buyers buy the product on day* 2 *(respectively* 1*). The revenue is* $0.7 \times 1 + 0.3 \times 1.2 = 1.06$ *in this equilibrium.*

Here we show that well-behaved equilibria are revenue-unique, that is, they all have the same revenue. We show this for an infinite number of types in the aggregate model which generalizes both the linear and symmetric models. We will show that in all of the well-behaved equilibria the fraction of people buying on a certain day is the same. In turn, it implies that the revenues of all well-behaved equilibria are the same and hence the well-behaved equilibria are revenue-unique. In what follows, we consider the equilibria of a fixed price sequence p. We start with a definition. Consider two well-behaved equilibria x and y. Partition the set of k days to two sets as follows. Call a day i a *level 1* day, denoted $i \in D_1(x, y)$, if $X_i < Y_i$. Otherwise, if $X_i \geq Y_i$, we call i a *level 2* day, denoted $i \in D_2(x, y)$.

Lemma 9. *Assume that there exist two distinct well-behaved equilibria x and y in the subgame induced by p. Then there exists a buyer whose strategy in x is a day in $D_1(x, y)$ and whose strategy in y is a day in $D_2(x, y)$.*

Proof. Since $x \neq y$, assume without loss of generality that there exists a day $\hat{i}$ such that $X_{\hat{i}} < Y_{\hat{i}}$. Let S_1 be the set of buyers buy on a level 1 day in x and S_2 be the set of buyers who buy on a level 2 day in y. We want to show that $S_1 \cap S_2$ is a non-empty set. Define $S = S_1 \cup S_2$. We show that $|S_1| + |S_2| > |S|$. Therefore S_1 and S_2 have a nonempty intersection whose elements are the buyers we are looking for. We do this by showing that $|S_1| + |S_2| > min(X_{k+1}, Y_{k+1}) \geq |S|$.

First observe that $|S| \leq min(X_{k+1}, Y_{k+1})$ since

- If $b \in S_1$ then b buys on some level 1 day $i \in D_1(x, y)$ in x. Therefore his utility on day i in y is non-negative. As y is a well-behaved equilibrium, this means that b must buy on some day in y. Thus b buys in both x and y.
- If $b \in S_2$ then b buys on some level 2 day in y and so similar reasoning shows that b must also buy on some day in x. Thus again b buys in both x and y.

Hence $|S| \leq min(X_{k+1}, Y_{k+1})$ as claimed.

Next let z_i be equal to x_i for $i \in D_1(x, y)$, and equal to y_i for $i \in D_2(x, y)$.[5] Note that $|S_1| + |S_2| = \sum_{i \in D_1(x,y)} x_i + \sum_{i \in D_2(x,y)} y_i = \sum_{i=1}^{k} z_i$. Thus it suffices to show that $\sum_{i=1}^{k} z_i > min(X_{k+1}, Y_{k+1}$. Let $Z_i = z_1 + \cdots + z_{i-1}$. We show that $Z_{k+1} > min(X_{k+1}, Y_{k+1})$. To do so we use the following claim:

Claim 10. *For each day i if $Z_i \geq$ (respectively $>$) $\min(X_i, Y_i)$, then we have $Z_{i+1} \geq$ (respectively $>$) $\min(X_{i+1}, Y_{i+1})$.*

Proof. For day i, there are four possibilities: If i and $i-1$ are level 1 days, then $Z_{i+1} = Z_i + z_i = Z_i + x_i \geq (>) X_i + x_i = X_{i+1}$. If i is a level 1 day and $i-1$ is a level 2 day, then $Z_{i+1} = Z_i + z_i = Z_i + x_i \geq (>) Y_i + x_i > X_i + x_i = X_{i+1}$. Note that in this case we get strict inequality even assuming weak inequality for i. If i is a level 2 day and $i-1$ is a level 1 day, then $Z_{i+1} = Z_i + z_i = Z_i + y_i \geq (>) X_i + y_i \geq Y_i + y_i = Y_{i+1}$. Finally, if i and $i-1$ are both level 2 days, then $Z_{i+1} = Z_i + z_i = Z_i + y_i \geq (>) Y_i + y_i = Y_{i+1}$. □

Now we complete the proof of the lemma by making the following two observations: First, for $i = 1$, we have $X_1 = Y_1 = Z_1 = 0$ so by induction for all i we have $Z_i \geq \min(X_i, Y_i)$. Second, recall that there exists a day $\hat{i}$ of level 1, and since day 1 is level 2, there must be a day i that falls in the second case of the claim for which we must have $Z_i > \min(X_i, Y_i)$. Then by induction we will have $Z_{k+1} > min(X_{k+1}, Y_{k+1})$. □

Theorem 11. *Let $F_i^b(X)$ be a strictly increasing function for each buyer b and day i. For a price sequence p and two well-behaved equilibria x and y in the induced subgame, we have $X_i = Y_i$ for all i.*

Proof. Assume for contradiction that we have two well-behaved equilibria x and y and a day i for which $X_i \neq Y_i$. Again assume without loss of generality that $X_i < Y_i$. By Lemma 9 we know that there exists a buyer b who buys on a level 1 day in x and buys on a level 2 day in y. Assume that b buys on day i in x and on day j in y. Then $F_i^b(X_i) - p_i \geq F_j^b(X_j) - p_j$ and $F_j^b(Y_j) - p_j \geq F_i^b(Y_i) - p_i$. Adding the two inequalities we get: $F_i^b(X_i) + F_j^b(Y_j) \geq F_j^b(X_j) + F_i^b(Y_i)$. On the other hand since i is a level 1 day, $X_i < Y_i$; hence by monotonicity $F_i^b(X_i) < F_i^b(Y_i)$. Since j is a level 2 day, $X_j \geq Y_j$; hence $F_j^b(Y_j) \leq F_j^b(X_j)$. The addition of these two inequalities contradicts the previous one. □

5. Revenue maximization

In this section, we solve the revenue-maximizing problem in two special cases: the discounted version of the symmetric model, and the general linear model without discount factors. In both cases, we provide an FPTAS to compute the revenue-maximizing price sequence. Before proceeding, however, we show with simple examples that given oracle access to the value functions, it is impossible to find the optimal solution with finite number of queries. This implies that we only hope for approximation schemes for the problem.

Example 12. *Let us first consider the symmetric model. Fix a parameter $a > 1/2$. Assume that for $X \geq a$, $F(X) = X$, and $F(X) = 0$ otherwise. Assume also that $k = 2$. In this instance, the unique optimum solution is the price sequence $(0, a)$, with the well-behaved equilibrium $(1/2, 1/2)$. The optimal revenue is therefore $a/2$. In order to compute the optimal solution, an algorithm needs to find a, and therefore to make infinite number of queries to the function F.*

We can modify the above instance to make it also an instance of the linear model. Set $I_b = 0$. Also set $C_b = 0$ if $b \leq 1/2$, and $C_b = 1$ otherwise. Again, the unique optimal solution is to set prices $(0, a)$, and an optimal algorithm needs to make infinitely many queries to the function F in order to find a.

[5] Strictly speaking z is not a valid marginal strategy profile as some buyers will buy in two different days in z.

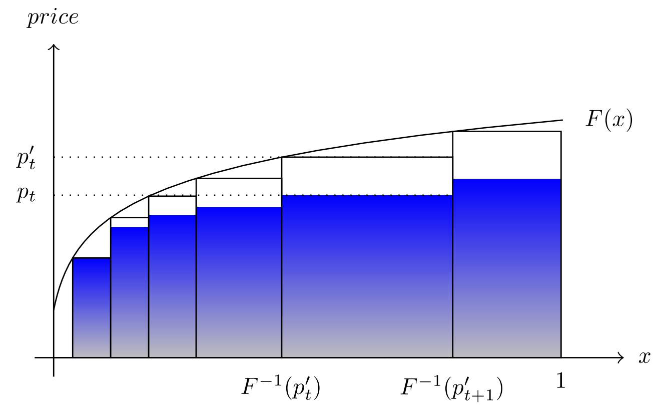

Fig. 2. The discounted RCP problem. The blue area is the total area covered by rectangles, discounted by the index of each rectangle.

5.1. Symmetric model

Since all players in this model have the same valuation function F, the marginal strategy profile matrix will reduce to the vector $x = (x_1, \ldots, x_k)$. Also, fixing p and x, the utility of buyer b for the item on day i is $F_i^b(X_i) = F(X_i)\beta^i(1-\alpha)^i - p_i(1-\alpha)^i$, and the revenue $R(p, x) = \sum_i x_i p_i (1-\alpha)^i$. By renaming $q_i = p_i(1-\alpha)^i$ and $\gamma = \beta(1-\alpha)$, the utility of buyer b for the item on day i will be $F(X_i)\gamma^i - q_i$, and the revenue becomes $\sum_i x_i q_i$. Using this new notation, we may assume without loss of generality that the only discount factor is γ. For convenience, we use p for the discounted prices q.

Since we only have one type in this model, we know that the utility of buying in day i is equal among all players. We use the term *utility of a day i*, denoted by u_i, for $u_i = F(X_i)\gamma^i - p_i$. Define $u(p, x) = \max_i u_i$. Consider a price sequence and its equilibrium strategy profile x. We get the following properties immediately from the facts that players are utility maximizers: (i) players are allowed to choose inaction and have utility zero, (ii) they choose to buy if there is a day with a strictly positive utility. First, if there is a day i with $x_i > 0$, then $u(p, x) \geq 0$ and $u_i = u(p, x)$. Second, if there is a day i with $u_i > 0$, then $\sum_{i=1}^k x_i = 1$.

First, we observe the following lemma:

Lemma 13. *Let $\hat{p}$ be the revenue-maximizing price vector that results in equilibrium $\hat{x}$. Then $u(\hat{p}, \hat{x}) = 0$.*

Proof. Assume for contradiction that $u(\hat{p}, \hat{x}) = w > 0$. Let $\hat{w} = (w, \ldots, w)$ be a vector of length k with all of its elements equal to w, and consider the price sequence $\hat{p} + \hat{w}$ and vector $\hat{x}$. The utility of each day decreases by the same amount $\hat{w}$, and the set of maximizers have positive utility, that is, $u(\hat{p} + \hat{w}, \hat{x}) \geq 0$. Therefore, each day i with $\hat{x}_i > 0$ is still a maximizer with $u_i \geq 0$. We conclude that $\hat{x}$ is an equilibrium for $\hat{p} + \hat{w}$. Since $\sum_i \hat{x}_i = 1$, this price sequence has strictly greater revenue, contradicting the optimality of $\hat{p}$. $\square$

Assume that there is a price sequence p with equilibrium x and $u(p, x) = 0$ such that for some day i, we have $x_i = 0$ and $x_{i+1} > 0$. Then we can define a new price sequence $\tilde{p}$ which is equal to p except that $\tilde{p}_j = p_{j+1}/\gamma$ for each $j \geq i$. Also define the vector $\tilde{x}$ to be equal to x except that $\tilde{x}_j = x_{j+1}$ for each $j \geq i$, and $\tilde{x}_k = 0$. One can observe that the pair $(\tilde{p}, \tilde{x})$ is an equilibrium with no less revenue. So we can assume without loss of generality that for a revenue maximizing price sequence $\hat{p}$ associated with $\hat{x}$, there exists a $k' \leq k$ such that $x_i \neq 0$ if and only if $i \leq k'$. For such a price sequence, Lemma 13 shows that $F(X_i)\gamma^i - p_i = 0$ for each $1 \leq i \leq k'$. As a result, we have $X_i = F^{-1}(p_i/\gamma^i)$, which is well-defined as F is increasing. Now set $p'_t = p_t/\gamma^t$. The fraction of people buying on day i and paying price p_i is equal to $x_i = F^{-1}(p'_{i+1}) - F^{-1}(p'_i)$. So the revenue is $\sum_i x_i p_i = \sum_i (F^{-1}(p'_{i+1}) - F^{-1}(p'_i))p'_i\gamma^i$. This sum is equal to the sum of the areas of a number of rectangles, discounted by γ, that are fit under the graph of F (see Fig. 2). So the revenue maximization problem reduces to the following *rectangular covering* problem.

Definition 4 (*Rectangular Covering Problem (RCP)*)**.** Given an increasing function F and an integer k, find a sequence p of size at most k that maximizes the discounted total area of the rectangles fit under the graph of F, that is,

$$p \in \arg\max_{p'} \sum_t (F^{-1}(p'_{t+1}) - F^{-1}(p'_t))p'_t\gamma^t.$$

In Section 6.1, we present an FPTAS to solve the RCP for concave valuation function and then show how to generalize the proof to non-concave functions in Section 6.2. See Section 6 for a general treatment of different versions of the RCP.

5.2. Linear model

To solve the problem in the linear version, we first study the properties of any equilibrium in Section 5.2.1. We also show how to derive prices and the total revenue from vector x. Then we study the revenue maximization problem in the linear version in Section 5.2.2. We prove that we can approximate the maximum revenue by solving the RCP for a specific curve. This curve is obtained from the function F and the distribution function $\mathcal{C}$.

5.2.1. Equilibrium properties:

Fix a price sequence p with an equilibrium x. Since we do not have discounts, we can assume without loss that for some $k' \leq k$, a purchase happens in day i if and only if $i \leq k'$ (just remove the days with no purchase to the end). So we can assume that for all $i < j \leq k'$, we have $X_i < X_j$. The utility of a buyer b for buying on day i is $I_b + C_b F(X_i) - p_i$. In order to concentrate on network externalities, for now we restrict the model and assume that I_b is always equal to a fixed constant I for all buyers. Then the utility can be written as $I + C_b F(X_i) - p_i$.

Consider the set of points $q_i = (F(X_i), p_i)$, $1 \leq i \leq k'$. Define $\hat{s}(i,j)$ to be equal to $(p_i - p_j)/(F(X_i) - F(X_j))$, which is the slope of the line between points q_i and q_j. Also let $s(i) = \hat{s}(i, i+1)$. In Lemma 14, we prove that s is non-decreasing in i. Lemma 15 then shows that the utility of buyer b will be maximized on day i if and only if $C_b \in [s(i-1), s(i)]$. Finally, we use these lemmas in Lemma 16 to show how to find a price vector given x. We use these properties in the next section in order to find the desirable equilibrium.

For a fixed b, let $i > j$ be two distinct days. The player prefers day i to j if $I + C_b F(X_i) - p_i \geq I + C_b F(X_j) - p_j$. The above inequality can be written as (recall that we know $X_i > X_j$, and therefore $F(X_i) > F(X_j)$):$C_b \geq \frac{p_i - p_j}{F(X_i) - F(X_j)}$. The converse is also true. If C_b is less than $(p_i - p_j)/(F(X_i) - F(X_j))$, then day j will be preferred to day i.

Lemma 14. *For the function s defined above, $s(i)$ is non-decreasing in i.*

Proof. Let $i, i+1, i+2$ be three consecutive days, and let b be a buyer who chooses to buy on day $i+1$. For b, day $i+1$ is at least as good as days $i, i+2$. Hence C_b must be greater than or equal to $s(i)$ and less than or equal to $s(i+1)$. We conclude that $s(i+1) \geq s(i)$. □

Lemma 15. *If $C_b \in [s(i-1), s(i)]$ then b will have the maximum utility buying on day i.*

Proof. Assume that $C_b \in [s(i-1), s(i)]$. For each $j > i$, the utility of day i is at least as good as that of day j, because $C_b \leq s(i) \leq \hat{s}(i,j)$. Similarly for each $j < i$, the utility of day i is at least as good as that of day j, because $C_b \geq s(i-1) \geq \hat{s}(j,i)$. The two special cases $C_b \in [0, s(1)]$ and $C_b \in [s(k-1), \infty)$ are dealt with using the same arguments. □

This lemma enables us to find the key relation between prices and vector x.

Lemma 16. *In an equilibrium, the following holds for each $2 \leq i \leq k$: $p_i - p_{i-1} = (F(X_i) - F(X_{i-1}))\mathcal{C}^{-1}(X_i)$.*

Proof. The fraction of people who buy on day i is exactly the fraction whose C_b's lie inside interval $[s(i-1), s(i)]$. So we have $x_i = X_{i+1} - X_i = \mathcal{C}(s(i) - s(i-1))$. The two sequences $\{X_1, \ldots, X_k\}$ and $\{\mathcal{C}(s(0)), \mathcal{C}(s(1)), \ldots, \mathcal{C}(s(k-1))\}$ have identical differences of consecutive terms. They also have identical initial elements $X_1 = \mathcal{C}(s(0)) = 0$. Hence they are identical and we have $X_i = \mathcal{C}(s(i-1))$.

On the other hand $s(i-1)$ is equal to $\frac{p_i - p_{i-1}}{F(X_i) - F(X_{i-1})}$ by definition. Therefore we can conclude the desirable result. □

5.2.2. Revenue maximization:

We have analyzed the properties of any equilibrium. In this part, we study properties of optimal equilibria. We show in Lemma 17 that in the revenue-maximizing equilibrium, the first price p_1 will be equal to I. Using this result, we express the total revenue in the revenue-maximizing equilibrium as a function of vector x. Details are in Lemma 18. Finally, in Lemma 19, we prove that the revenue maximization problem can be reduced to the RCP, and therefore there exists an FPTAS to solve the revenue maximization problem (Theorem 23).

Lemma 17. *In a revenue-maximizing equilibrium, $p_1 = I$.*

Proof. Let x, p be the defining vectors of an equilibrium. Obviously $p_1 \leq I$, because people who buy on the first day have nonnegative utility. Now raise all elements of p by $I - p_1$ to get p'. One can verify that x and p' still define an equilibrium. The new equilibrium's revenue is greater than the original one's. Hence $p_1 = I$ in the revenue-maximizing equilibrium. □

Note that by calculating all X_i with respect to vector x, we know the values $p_i - p_{i-1}$ using Lemma 16. On the other hand, $p_1 = I$ in the revenue-maximizing equilibrium. So we can calculate all prices, and therefore x uniquely determines prices. Let Price(x) be the optimal price vector of x. Vectors x and Price(x) define an equilibrium. Hence it suffices to view everything as functions of the free variable x. Now we express the revenue in an optimal equilibrium in terms of vector x.

Lemma 18. *If x and p correspond to the revenue-maximizing equilibrium, the total revenue can be expressed by the following formula:*

$$R(p, x) = I + \sum_{i=2}^{k} (1 - X_i)\mathcal{C}^{-1}(X_i)(F(X_i) - F(X_{i-1})).$$

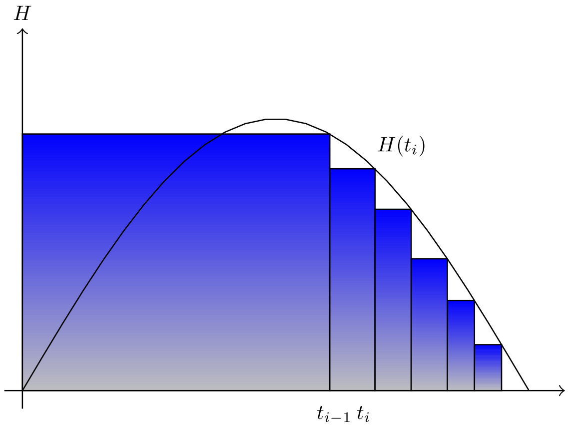

Fig. 3. Revenue with respect to non-monotone function H.

Proof. Since the utility of buying on the first day is nonnegative for everybody, all buyers would choose to buy, i.e. $X_{k+1} = 1$. The total revenue can be written as:

$$R(p,x) = \sum_{i=1}^{k} p_i x_i = p_1 \left(\sum_{j=1}^{k} x_j\right) + \sum_{i=2}^{k} (p_i - p_{i-1}) \left(\sum_{j=i}^{k} x_j\right).$$

To interpret the above formula, note that the price change from day $i-1$ to day i is exerted on all buyers who have bought on day i or later. Since $\sum_{j=1}^{k} x_j = 1$, we re-order summations and write $R(p,x) = p_1 + \sum_{i=2}^{k}(p_i - p_{i-1})(1 - X_i)$. Now we can substitute for p_1 and $p_i - p_{i-1}$ from Lemmas 16 and 17 and write

$$R(p,x) = I + \sum_{i=2}^{k} (1 - X_i)\mathcal{C}^{-1}(X_i)(F(X_i) - F(X_{i-1})). \quad \square$$

The next step is to reduce the problem of maximizing revenue to the RCP. Note that I is a constant in Lemma 18 which does not affect revenue maximization.

Lemma 19. *The problem of maximizing* $\sum_{i=2}^{k}(1 - X_i)\mathcal{C}^{-1}(X_i)(F(X_i) - F(X_{i-1}))$ *can be reduced to the RCP.*

Proof. Since F is a monotone and hence a bijective function, we can write the revenue as

$$R = \sum_{i=2}^{k} (1 - F^{-1}(F(X_i)))\mathcal{C}^{-1}(F^{-1}(F(X_i))(F(X_i) - F(X_{i-1})).$$

The above formula which only depends on $F(X_i)$'s can be interpreted as an RCP instance. Let $F_{min} = F(0), F_{max} = F(1)$. We define $H : [F_{min}, F_{max}] \to \mathbb{R}^{\geq 0}$ as follows : $H(x) = (1 - F^{-1}(x))\mathcal{C}^{-1}(F^{-1}(x))$. Note that the revenue is exactly $R = \sum_{i=2}^{k} H(F(X_i))(F(X_i) - F(X_{i-1}))$. Therefore we want to find values $t_i = F(X_i)$'s in order to maximize R. Revenue R can be shown as the total area of some rectangles with one corner on curve H (See Fig. 3).

The problem of finding these rectangles is very similar to the RCP. If we define $H'(x) = H(-x)$ and $t'_i = -t_{k-i}$ then we can write the revenue as:

$$\begin{aligned} R &= \sum H(t_i)(t_i - t_{i-1}) = \sum H'(-t_i)(t_i - t_{i-1}) \\ &= \sum H'(t'_{k-i})(t'_{k-i+1} - t'_{k-i}) = \sum H'(t'_j)(t'_{j+1} - t'_j)) \end{aligned}$$

So we should solve the RCP for function H'. $\square$

We have proved that the revenue maximization problem can be solved using the RCP. Therefore, there exists an FPTAS to solve this problem (Theorem 23 in Section 6.2).

6. Rectangular covering problem

Recall that in the RCP we want to place k rectangles under the curve F, and maximize the discounted covering area (see Fig. 2). We propose an FPTAS for this problem for concave functions in Section 6.1 and for bounded slope functions in Section 6.2. The algorithm described in Lemma 20 can be used to solve the RCP for functions F with $F(0) > 0$. The running-time of the algorithm is $poly(k, \log_{1+\epsilon} F(1)/F(0))$ in this case. In the remaining part of paper we assume that F is non-decreasing and $F(1) = 1$. We prove that these assumptions do not hurt the generality in Section 6.3.

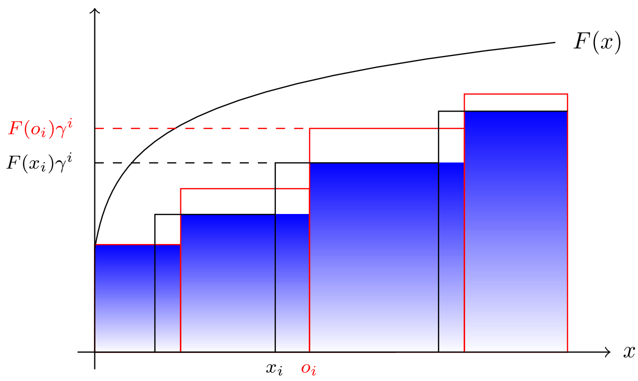

Fig. 4. The region with area A.

Now we define some restricted versions of the RCP.

Definition 5. δ**-RCP**: it is an instance of the RCP in which every $F(x_i)$ should be greater than or equal to δ.

Definition 6. (δ_2, δ_1)**-RCP**: it is an instance of the RCP in which every $F(x_i)$ should be outside of a given interval (δ_2, δ_1).

6.1. Concave functions

In this section we propose an FPTAS for the RCP when the function F is concave and non-decreasing.

Lemma 20. *For every $\epsilon > 0$ and $\delta \geq 0$, the δ-rectangular covering problem can be solved in poly$(k, \log_{1+\epsilon}(F(1)/\delta))$ time, with approximation factor $1+\epsilon$.*

Proof. First, we define a sequence $S = (s_1, s_2, \ldots, s_m)$ and then prove that if we choose indices $x_1, x_2, \ldots, x_k$ from the sequence S, we could approximate the optimum revenue. Let $s_i = F^{-1}((1+\epsilon)^{i-1}\delta)$. In fact, $F(s_{i+1}) = (1+\epsilon)F(s_i)$ for every $i < m$, $s_0 = 0$ and $s_m \leq 1 < s_{m+1}$.

Assume that the optimum solution in the RCPis the sequence $o_1, o_2, \ldots, o_k$. Therefore the value of the optimum solution is $OPT = \sum_{i=0}^{k}(o_{i+1} - o_i)F(o_i)\gamma^i$. Let x_i be the maximum index of sequence S with value no more than o_i. So $F(o_i) \leq F(x_i)(1+\epsilon)$. We can bound OPT as follows (in all equations assume that $x_0 = o_0 = 0$ and $o_{k+1} = x_{k+1} = 1$).

$$OPT = \sum_{i=0}^{k}(o_{i+1} - o_i)F(o_i)\gamma^i \leq \sum_{i=0}^{k}(o_{i+1} - o_i)(1+\epsilon)F(x_i)\gamma^i = (1+\epsilon)A.$$

The region with area A has been shown in Fig. 4. First observe A is less than or equal to $\sum_{i=0}^{k}(x_{i+1} - x_i)F(x_i)\gamma^i$. This is because

$$\begin{aligned}
A &= \sum_{i=0}^{k}(o_{i+1} - o_i)F(x_i)\gamma^i \\
&= \sum_{i=1}^{k} o_i(F(x_{i-1}) - F(x_i))\gamma^i + o_{k+1}F(x_k)\gamma^{k+1} \\
&\leq \sum_{i=1}^{k} x_i(F(x_{i-1}) - F(x_i))\gamma^i + x_{k+1}F(x_k)\gamma^{k+1} \\
&= \sum_{i=0}^{k}(x_{i+1} - x_i)F(x_i)\gamma^i,
\end{aligned}$$

where the inequality followed since $o_i \geq x_i$ for $1 \leq i \leq k$, and F is a non-decreasing function.

As a result, if we choose indices $x_1, x_2, \ldots, x_k$ optimally from the sequence S we can approximate the optimum value within a factor $1+\epsilon$.

Now we design a dynamic programming algorithm to find indices $x_1, x_2, \ldots, x_k$ optimally from S. For $n \leq m$ and $r \leq k$, let $A[n, r]$ be equal to the optimum solution when we select r indices from the subsequence $(s_1, s_2, \ldots, s_n)$ of S and also s_n has been selected. We have

$$A[n, r] = max_{n' < n}\{A[n', r-1] + (s_n - s_{n'})F(s_{n'})\gamma^r\}.$$

To see the above equality, first assume that $n' < n$ is the highest index which is less than n and is selected by the optimum solution. Then the area between n' and n is going to be $(s_n - s_{n'})F(s_{n'})\gamma^r$. Furthermore, the optimal solution will find the optimum covering when $r-1$ indices are chosen from the subsequence $(s_1, s_2, \dots, s_{n'})$, whose value is by definition $A[n', r-1]$. The above equality then states that the optimum solution picks $n' < n$ in order to maximize value.

For $m = \Theta(\log_{1+\epsilon}(F(1)/\delta))$, the running time of the algorithm is $\Theta(poly(m, k))$. □

Now we are ready to propose an FPTAS for the RCP for concave functions.

Theorem 21. *For every $\epsilon > 0$, the rectangular covering problem with concave functions can be solved in $poly(k, 1/\log(1+\epsilon), \log(1/\epsilon))$ time with approximation factor $1+\epsilon$.*

Proof. Let $\epsilon' = \epsilon/2$. Assume we want to solve the δ-RCP for a concave function F with $\delta = F(1)(\frac{1}{1+\epsilon'})^n$, where $n = \log_{1+\epsilon'} \frac{4(1+2\epsilon')}{\epsilon'}$. Since F is concave, the optimum solution for the RCP OPT is at least $F(1)/4$. Let the optimum solution for the δ-RCP be OPT_δ and the solution found by Lemma 20 be A_δ. We have proved that $A_\delta \geq \frac{OPT_\delta}{1+\epsilon'}$. On the other hand, the optimum solution for the RCP is at most $OPT_\delta + \delta$. So $A_\delta \geq \frac{OPT-\delta}{1+\epsilon'} = \frac{OPT}{1+\epsilon'} - \frac{F(1)}{(1+\epsilon')^{n+1}} \geq \frac{OPT}{1+\epsilon'} - \frac{4OPT}{(1+\epsilon')^{n+1}}$. If we replace $(1+\epsilon')^n$ by $\frac{4(1+2\epsilon')}{\epsilon'}$, we can conclude that $A_\delta \geq \frac{OPT}{1+\epsilon}$.

We have used Lemma 20 with ϵ' and $\delta = F(1)(\frac{1}{1+\epsilon'})^n$. So the algorithm runs in $\Theta(k, \log_{1+\epsilon'} F(1)/\delta) = \Theta(k, n)$ time, where $n = \Theta(\frac{\log(1/\epsilon)}{\log(1+\epsilon)})$. □

6.2. Functions with bounded slope

In this section we propose an FPTAS to solve the RCP for a function F when $\gamma = 1$ and $F'(x) \leq L$.

Lemma 22. *For every $\delta_1, \delta_2, \epsilon > 0$, the (δ_2, δ_1)-rectangular covering problem can be solved in $poly(k, \log_{1+\epsilon}(\frac{1}{1-F^{-1}(\delta_2)}), \log_{1+\epsilon}(\frac{1-\delta_2}{\delta_1-\delta_2}))$ time with approximation factor $1+\epsilon$.*

Proof. First, we define a sequence $S = (s_1, s_2, \dots, s_m)$ and then prove that if we choose indices $x_1, x_2, \dots, x_k$ from sequence S, we could approximate the optimum. The sequence S is the union of two sequences S' and S'', which are constructed as follows:

- Sequence $S' = (s'_1, s'_2, \dots, s'_{m'})$ consists of indices with $F(s'_i) \geq \delta_1$. Let $s'_1 = F^{-1}(\delta_1)$ and $s'_{i+1} = F^{-1}((1+\epsilon)(F(s'_i) - \delta_2) + \delta_2)$ and $s'_{m'} = 1$. In fact we have $F(s'_{i+1}) - \delta_2 \leq (1+\epsilon)(F(s'_i) - \delta_2)$ for every $i < m'$. The length of the sequence S' is $m' = \log_{1+\epsilon}(\frac{1-\delta_2}{\delta_1-\delta_2})$.
- Sequence $S'' = (s''_1, s''_2, \dots, s''_{m''})$ consists of indices with $F(s''_i) \leq \delta_2$. Let $s''_1 = 0$ and $1 - s''_i = (\frac{1}{1+\epsilon})^{i-1}$ and $s''_{m''} = F^{-1}(\delta_2)$. In fact we have $1 - s''_i \leq (1+\epsilon)(1 - s''_{i+1})$. The length of sequence S'' is $m'' = \log_{1+\epsilon}(\frac{1}{1-F^{-1}(\delta_2)})$.

Assume optimum indices are $o_1, o_2, \dots, o_k$ in the (δ_2, δ_1)-RCP. Let the value of the optimum solution be $O = \sum_{i=1}^{k}(o_{i+1} - o_i)F(o_i)$. We know that every $F(o_i)$ is outside of the interval (δ_2, δ_1). For every $o_i \leq F^{-1}(\delta_2)$, Let x_i be the minimum index in sequence S (S'') with value not less than o_i, and for every $o_i \geq F^{-1}(\delta_1)$, let x_i be the maximum index in sequence S (S') with value no more than o_i. Assume that $x_i \in S''$ for every index $i < k'$ and $x_i \in S'$ for every index $i \geq k'$. We can now bound OPT as follows. (In all equations assume that $x_0 = o_0 = 0$ and $o_{k+1} = x_{k+1} = 1$.)

$$OPT = \sum_{i=0}^{k}(o_{i+1} - o_i)F(o_i) = \sum_{i=0}^{k'-1}(o_{i+1} - o_i)F(o_i) + \sum_{i=k'}^{k}(o_{i+1} - o_i)F(o_i). \tag{1}$$

The right hand side of above equation can be written as:

$$\sum_{i=0}^{k'-1}(o_{i+1} - o_i)F(o_i) = \sum_{i=0}^{k'-1}(1 - o_i)(F(o_i) - F(o_{i-1})) - F(o_{k'-1})(1 - o_{k'}). \tag{2}$$

And the left hand side can be written as:

$$\sum_{i=k'}^{k}(o_{i+1} - o_i)F(o_i) = \sum_{i=k'}^{k}(o_{i+1} - o_i)(F(o_i) - F(o_{k'-1})) + F(o_{k'-1})(1 - o_{k'}). \tag{3}$$

Rewrite equality (1) with respect to equality (2) and (3).

$$OPT = \sum_{i=0}^{k'-1}(1 - o_i)(F(o_i) - F(o_{i-1})) + \sum_{i=k'}^{k}(o_{i+1} - o_i)(F(o_i) - F(o_{k'-1})). \tag{4}$$

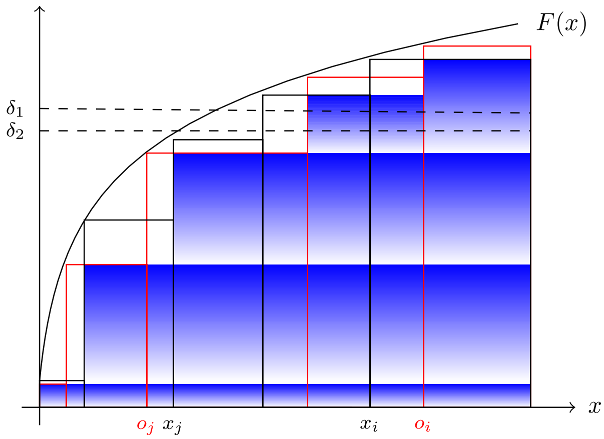

Fig. 5. The region with area A.

For every $i < k'$, we have $(1-o_i) \le (1+\epsilon)(1-x_i)$. On the other hand for every $i \ge k'$, we have $F(o_i)-\delta_2 \le (1+\epsilon)(F(x_i)-\delta_2)$. Since $F(o_{k'-1}) \le \delta_2$, we have $F(o_i) - F(o_{k'-1}) \le (1+\epsilon)(F(x_i) - F(o_{k'-1}))$ for every $i \ge k'$. With respect to these facts we can bound the optimum as follows:

$$\begin{aligned} OPT &\le (1+\epsilon)\sum_{i=0}^{k'-1}(1-x_i)(F(o_i)-F(o_{i-1})) + (1+\epsilon)\sum_{i=k'}^{k}(o_{i+1}-o_i)(F(x_i)-F(o_{k'-1})) \\ &= (1+\epsilon)A. \end{aligned} \tag{5}$$

The region with area A has been shown in Fig. 5. It is clear that A is less than or equal to $\sum_{i=1}^{k}(x_{i+1}-x_i)F(x_i)$. So if we choose the indices $x_1, x_2, \dots, x_k$ from the sequence S, we can approximate the optimum within a factor $1+\epsilon$. Now we design a dynamic programming algorithm to find the indices $x_1, x_2, \dots, x_k$ from S. For $n \le m$ and $r \le k$, let $A[n, r]$ be equal to the best solution when we select indices from the subsequence $(s_1, s_2, \dots, s_n)$ of S and also s_n has been selected. We have:

$$A[n, r] = max_{n'<n}\{A[n', r-1] + (s_n - s_{n'})F(s_{n'})\}. \quad \square$$

To see the above equality, first assume that $n' < n$ is the highest index which is less than n and is selected by the optimum solution. Then the area between n' and n is going to be $(s_n - s_{n'})F(s_{n'})$. Furthermore, the optimal solution will find the optimum covering when $r-1$ indices are chosen from the subsequence $(s_1, s_2, \dots, s_{n'})$, whose value is by definition $A[n', r-1]$. The above equality then states that the optimum solution picks $n' < n$ in order to maximize value.

Consider an instance of the RCP with optimum value OPT. Construct k^2+1 instances of the (δ_2, δ_1)-RCP from the RCP instance. In the i-th instance we set $\delta_2 = \frac{i-1}{k^2+1}$ and $\delta_1 = \frac{i}{k^2+1}$ and assume that the value of the goal function in the optimum solution of this instance is OPT_i. Assume that $o_1, o_2, \dots, o_k$ are the optimum indices in the RCP instance. It is clear that there exist $1 \le j \le k^2+1$ such that every o_i is outside of the interval $(\frac{j-1}{k^2+1}, \frac{j}{k^2+1})$. Therefore we have $OPT = OPT_j$. Assume that we solve all of the k^2+1 instances of the (δ_2, δ_1)-RCP using Lemma 22. Let A_i be the output of the algorithm for i-th instance. We proved in Lemma 22 that $OPT_i \le (1+\epsilon)A_i$. So if we return $max_i A_i$ as the output for the RCP instance, we can approximate the optimum within a factor $1+\epsilon$. Now, we prove that this algorithm runs in polynomial time. To see this, we show that if we set $\delta_2 = \frac{i-1}{k^2+1}$ and $\delta_1 = \frac{i}{k^2+1}$, then $\log_{1+\epsilon}(\frac{1}{1-F^{-1}(\delta_2)})$ and $\log_{1+\epsilon}(\frac{1-\delta_2}{\delta_1-\delta_2})$ are polynomial. First we have $\frac{1-\delta_2}{\delta_1-\delta_2} = k^2+2-i$ which is polynomial with respect to k. On the other hand we have $\delta_2 + \int_{F^{-1}(\delta_2)}^{1} F'(x)dx = F(1) = 1$. Therefore $\int_{F^{-1}(\delta_2)}^{1} F'(x)dx = 1-\delta_2 \ge \frac{1}{k^2+1}$. If we assume that $F'(x) \le L$ we can conclude that $1-F^{-1}(\delta_2) \ge \frac{1}{L(k^+1)}$. So $\frac{1}{1-F^{-1}(\delta_2)}$ is at most $L(k^2+1)$ which is polynomial with respect to k and L. Hence, we get the following:

Theorem 23. *For every $\epsilon > 0$, the rectangular covering problem with $F'(x) \le L$ can be solved in $poly(k, \log_{1+\epsilon} k, \log_{1+\epsilon} L)$ time, with approximation factor $1+\epsilon$.*

6.3. Assumptions about the Function F

Here, we prove that the following assumptions can be made for the RCP without loss of generality.

- *F is non-decreasing*: Let $G(x) = \max_{0\le y\le x} F(y)$. We prove that the best rectangular covering of G and F are exactly the same. We prove this statement by showing that every optimum solution of the RCP for F is a solution for G with same objective value and vice versa. First assume that $x_1, x_2, \ldots, x_k$ are optimum indices corresponding to some rectangular covering of F. For every $1 \leq i \leq k$, we prove that $F(x_i) = G(x_i)$. If it is not the case, let j be the smallest index which $F(x_j) < G(x_j)$. Let $x_{j'}$ be the smallest index for which $G(x_{j'}) = G(x_j)$. It is clear that $F(x_{j'})$ is the biggest value in interval $[0, x_j]$. Now if we replace x_j by $x_{j'}$ in the optimum indices, the change in the objective function will be:

$$(x_{j'} - x_{j-1})F(x_{j-1})\gamma^{j-1} + (x_{j+1} - x_{j'})F(x_{j'})\gamma^{j} - (x_j - x_{j-1})F(x_{j-1})\gamma^{j-1} + (x_{j+1} - x_j)F(x_j)\gamma^{j}.$$

Note that $F(x_{j'})\gamma^j$ is greater than $F(x_j)\gamma^j$ and $F(x_{j-1})\gamma^{j-1}$. Therefore, the amount of change in the objective function is positive, which is a contradiction with optimality of indices $x_1, x_2, \ldots, x_k$. So the sequence $(x_1, x_2, \ldots, x_k)$ is a solution for curve G with the same objective value.

On the other hand, assume that $(x_1, x_2, \ldots, x_k)$ is an optimum sequence for some rectangular covering of curve G. Note that G is non-decreasing. If there are two indices $x_{j'}$ and x_j such that $x_{j'} < x_j$ and $F(x_{j'}) = F(x_j)$, then x_j will not appear in any optimum solution of G. If we remove indices with this property from the search space, only indices x_i with $G(x_i) = F(x_i)$ remain. So the sequence $(x_1, x_2, \ldots, x_k)$ is a solution for curve F with the same objective value.

- $F(1) = 1$: Let $F(1) = c \neq 1$. Define $G(x) = \frac{F(x)}{c}$ and solve the problem for G. Assume $x_1, x_2, \ldots, x_k$ are indices corresponding to some rectangular covering for curve G. It is clear that:

$$\sum_{i=0}^{k}(x_{i+1} - x_i)G(x_i)\gamma^i = \frac{1}{c}\sum_{i=0}^{k}(x_{i+1} - x_i)F(x_i)\gamma^i.$$

So we can convert every solution for covering the area under the curve G to a solution for covering area under curve F and vice versa. So by solving the RCP for curve G we can solve the problem for curve F.

7. Conclusion

In this paper, we studied the problem of optimal item pricing in the presence of historical network externalities and strategic buyers. We obtained necessary and sufficient conditions on the valuation functions for the existence and uniqueness of equilibria. We defined the revenue maximization problem in settings in which the equilibrium exists and is unique, and presented nearly-optimal algorithms for the two special cases of linear and symmetric valuation functions. The problem of designing algorithms or proving hardness (e.g. inapproximability) for the more general aggregate model remains open.

We studied the historical externalities setting in which the valuation for a purchase on a given day may depend only on the purchases on previous purchases. However, in many settings, users may benefit from purchases that happen in the future. Characterizing equilibria and studying the optimal pricing problems in these settings are interesting subjects of study for future research.

Finally, one can define and study the externalities model in non-monopolistic settings with competing sellers, and compare the behavior of equilibria with and without such competing sellers. It would also be interesting to study the extensions of the classical models of duopoly from algorithmic perspective.

References

[1] N. AhmadiPourAnari, S. Ehsani, M. Ghodsi, N. Haghpanah, N. Immorlica, H. Mahini, V.S. Mirrokni, 2010, Equilibrium pricing with positive externalities, in: WINE: Workshop on Internet and Network Economics.

[2] H. Akhlaghpour, M. Ghodsi, N. Haghpanah, H. Mahini, V.S. Mirrokni, A. Nikzad, 2010, Optimal iterative pricing over social networks, in: WINE: Workshop on Internet and Network Economics.

[3] W.B. Arthur, Competing technologies, increasing returns, and lock-in by historical events, Economic Journal 99 (1989) 116–131.

[4] N. Bansal, N. Chen, N. Cherniavsky, A. Rudra, B. Schieber, M. Sviridenko, Dynamic pricing for impatient bidders, in: SODA, 2007, pp. 726–735.

[5] B. Bensaid, J.P. Lesne, Dynamic monopoly pricing with network externalities, International Journal of Industrial Organization 14 (1996) 837–855.

[6] K. Bimpikis, O. Candogan, A. Ozdaglar, 2010, Optimal pricing in the presence of local network effects, in: WINE: Workshop on Internet and Network Economics.

[7] L. Cabral, D. Salant, G. Woroch, 1994, Monopoly pricing with network externalities, Industrial Organization 9411003. EconWPA.

[8] V. Conitzer, T. Sandholm, 2006, Computing the optimal strategy to commit to, in: ACM Conference on Electronic Commerce, pp. 82–90.

[9] A.A. Cournot, 1838, Recherches sur les principes mathmatiques de la theorie des richesses (researches into the mathematical principles of the theory of wealth). Hachette, Paris.

[10] P. Domingos, M. Richardson, Mining the network value of customers, in: KDD'01, ACM, New York, NY, USA, 2001, pp. 57–66.

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

[12] J. Hartline, V.S. Mirrokni, M. Sundararajan, 2008, Optimal marketing strategies over social networks, in: WWW, pp. 189–198.

[13] S. Kakutani, A generalization of Brouwers' fixed point theorem, Duke Mathematical Journal 8 (1941) 457–459.

[14] M.L. Katz, C. Shapiro, Network externalities, competition, and compatibility, American Economic Review 75 (1985) 424–440.

[15] D. Kempe, J. Kleinberg, Tardos Éva, Influential nodes in a diffusion model for social networks, in: ICALP, Springer Verlag, 2005, pp. 1127–1138.

[16] B. LeBlanc, 2009, Showing our thanks to Windows 7 beta testers. The Windows Blog.

[17] J. Letchford, V. Conitzer, 2010, Computing optimal strategies to commit to in extensive-form games, in: ACM Conference on Electronic Commerce, pp. 83–92.

[18] A. Mas-Colell, On a theorem of Schmeidler, Journal of Mathematical Economics 13 (1984) 201–206.

[19] R. Mason, Network externalities and the coarse conjecture, European Economic Review 44 (2000) 1981–1992.

[20] E. Mossel, S. Roch, On the submodularity of influence in social networks, in: STOC'07, ACM, New York, NY, USA, 2007, pp. 128–134.

[21] P. Paruchuri, J.P. Pearce, J. Marecki, M. Tambe, F. Ordóñez, S. Kraus, Efficient algorithms to solve Bayesian Stackelberg games for security applications, in: AAAI, 2008, pp. 1559–1562.

[22] E. Protalinski, 2009, Windows 7 pricing announced. Ars Technica.

[23] H.R. von Stackelberg, Marktform und gleichgewicht, Springer-Verlag, 1934.

[24] P. Sskilahti, 2007, Monopoly pricing of social goods. MPRA Paper 3526. University Library of Munich, Germany.