A Theory of Loss-Leaders: Making Money by Pricing Below Cost

Maria-Florina Balcan, Avrim Blum, T-H. Hubert Chan, and MohammadTaghi Hajiaghayi

Computer Science Department, Carnegie Mellon University
{ninamf,avrim,hubert,hajiagha}@cs.cmu.edu

Abstract. We consider the problem of assigning prices to goods of fixed marginal cost in order to maximize revenue in the presence of single-minded customers. We focus in particular on the question of how pricing certain items below their marginal costs can lead to an improvement in overall profit, even when customers behave in a fully rational manner. We develop two frameworks for analyzing this issue that we call the *discount* and the *coupon* models, and examine both fundamental "profitability gaps" (to what extent can pricing below cost help to improve profit) as well as algorithms for pricing in these models in a number of settings considered previously in the literature.

1 Introduction

The notion of *loss-leaders*, namely pricing certain items below cost in a way that increases profit overall from the sales of other items, is a common technique in marketing. For example, a hamburger chain might price its burgers below production cost but then have a large profit margin on sodas. Grocery stores often give discounts that reduce the cost of certain items even to zero, making money from other items the customers will buy while in the store.

Such "loss leaders" are often viewed as motivated by psychology: producing extra profit from the emotional behavior of customers who are attracted by the good deals and then do not fully account for their total spending. Alternatively, they are also often discussed in the context of selling goods of decreasing marginal cost (so the loss-leader of today will be a profit center tomorrow once sales have risen). However, even for items of fixed marginal cost, with fully rational customers who have valuations on different bundles of items and act to maximize utility, pricing certain items below cost can produce an increase in profit. For example, DeGraba [5] analyzes equilibria in a 2-firm, 2-good Hotelling market, and argues that the power of loss leaders is that they provide a method for focusing on high-profit customers: "a product could be priced as a loss leader if, in a market in which some customers purchase bundles of products that are more profitable than bundles purchased by others, the product is purchased primarily by customers that purchase more profitable bundles." Balcan and Blum [1] give an example, in the context of pricing n items of fixed marginal cost to a set

X. Deng and F.C. Graham (Eds.): WINE 2007, LNCS 4858, pp. 293–299, 2007.

of single-minded customers, where allowing items to be priced below cost can produce an $\Omega(\log n)$ factor more than the maximum possible profit obtained by pricing all items above cost. However, the problem of developing algorithms taking advantage of this idea was left as an open question.

In this paper we consider this problem more formally, introducing two theoretical models which we call the *coupon model* and the *discount model* for analyzing the profit that can be obtained by pricing below cost. These models are motivated by two different types of settings in which such pricing schemes can naturally arise. We then develop algorithms for several problems studied in the literature, including the "highway problem" [8] and problems of pricing vertices in graphs, as well as analyze fundamental gaps between the profit obtainable under the different models. It is worth noting that the algorithmic problem becomes much more difficult in these settings than in the setting where pricing below cost is not allowed.

The two models we introduce are motivated by two types of scenarios. In the *discount model*, we imagine a retailer (say a supermarket or a hamburger chain) selling n different types of items, where each item i has some fixed marginal (production) cost c_i to the retailer. The retailer needs to assign a sales price s_i to each item, which could potentially be less than c_i. That is, the profit margin $p_i = s_i - c_i$ for item i could be positive or negative. The goal of the retailer is to assign these prices so as make as much profit as possible from the customers. We will be considering the case of single-minded customers, meaning that each customer j has some set S_j of items he is interested in and will purchase the entire set (one unit of each item $i \in S_j$) if its total cost is at most his valuation v_j, else nothing. As an example, suppose we have two items $\{1, 2\}$, each with production cost $c_i = 10$ and two customers, one interested in item 1 only and willing to pay 20, and the other interested in both and willing to pay 25. In this case, by setting $s_1 = 20$ and $s_2 = 5$ (which correspond to profit margins $p_1 = 10$ and $p_2 = -5$, and hence the second item is priced below cost) the retailer can make a total profit of 15. This is greater than the maximum profit (10) obtainable from these customers if pricing below cost were not allowed.

One thing that makes the discount model especially challenging is that profit is not necessarily monotone in the customers' valuations. For instance, in the above example, if we add a new customer with $S_j = \{2\}$ and $v_j = 3$ then the solution above still yields profit 15 (because the new customer does not buy), but if we increase v_j to 10, then any solution will make profit at most 10.

The second model we introduce, the *coupon model*, is designed to at least satisfy monotonicity. This model is motivated by the case of goods with zero marginal cost (such as airport taxes or highway tolls). However, rather than setting actual negative prices, we instead will allow the retailer to give credit that can be used towards other purchases. Formally, each item i has marginal cost $c_i = 0$ and is assigned a sales price p_i which can be positive or negative, and the price of a bundle S is $\max(\sum_{i \in S} p_i, 0)$, which is also the profit for selling this bundle. We again consider single-minded customers. A customer j will purchase his desired bundle S_j iff its price is at most his valuation v_j. Note that in this model we are assuming *no free disposal*: the customer is only interested in a particular set of items and will not purchase a superset even if cheaper (e.g., in the case of highway tolls, we assume a driver would either use the highway to go from his source to his destination or not, but would not travel additional stretches of highway just to save on tolls). As an example of the coupon model, consider a highway with three toll portions (items) 1, 2, and 3. Assume there are four drivers (customers) A, B, C, and D as follows: A, B, and C each only use portions 1, 2, and 3 respectively, but D uses all three portions. Assume that A, C, and D each are willing to pay 10 while B is wiling to pay only 1. In this case, by setting $p_1 = p_3 = 10$ and $p_2 = -10$, we have a solution with profit of 30 (driver B gets to travel for free, but is not actually paid for using the highway). This is larger than the maximum profit possible (21) in the discount model or if we are not allowed to assign negative prices. Note that unlike the discount model, the coupon model does satisfy monotonicity.

We can make the discount model look syntactically more like the coupon model by subtracting production costs from the valuations. In this view, $w_j := v_j - \sum_{i \in S_j} c_i$ represents the amount *above production cost* that customer j is willing to pay for S_j, and our goal is to assign positive or negative profit margins p_i to each item i to maximize the total profit $\sum_{j: w_j \geq p(S_j)} p(S_j)$ where $p(S_j) = \sum_{i \in S_j} p_i$. It is interesting in this context to consider two versions: in the *unbounded discount* model we allow the p_i to be as large or as small as desired, ignoring the implicit constraint that $p_i \geq -c_i$, whereas in the *bounded discount* model we impose those constraints. Note that in this view, the only difference between the unbounded discount model and the coupon model is that in the coupon model we redefine $p(S_j)$ as $\max(\sum_{i \in S_j} p_i, 0)$.

We primarily focus on two well-studied problems first introduced formally by Guruswami et al. [8]: the *highway tollbooth* problem and the *graph vertex pricing* problem. In the highway tollbooth problem, we have n items (highway segments) $1, \ldots, n$, and each customer (driver) has a desired bundle that consists of some interval $[i, i']$ of items (consecutive segments of the highway). The seller is the owner of the highway system, and would like to choose tolls on the segments (and possibly also coupons in the coupon model) so as to maximize profits. Even if all customers have the same valuation for their desired bundles, we show that there are $\log(n)$-sized gaps between the profit obtainable in the different models. In the graph vertex pricing problem, we instead have the constraint that all desired bundles S_j have size at most 2. Thus, we can consider the input as a multi-graph whose vertex set represent the set of items and whose edges represent the costumers who want end-points of the edges. We show that if this graph is *planar* then one can in fact achieve a PTAS for profit in each model.

It is worth mentioning we do not focus on *incentive-compatibility* aspects in this paper since one can use the generic reductions in [3] to convert our approximations algorithms into good truthful mechanisms. In this version, we only state the results and the reader is referred to the full version [2] for the proofs.

2 Notation, Definitions, and Gaps Between the Models

We assume we have m customers and n items (or "products"). We are in an *unlimited supply* setting, meaning that the seller is able to sell any number of units of each item. We consider *single-minded customers*, which means that each customer is interested in only a single bundle of items and has valuation 0 on all other bundles. Therefore, valuations can be summarized by a set of pairs (e, v_e) indicating that a customer is interested in bundle (hyperedge) e and values it at v_e. Given the hyperedges e and valuations v_e, we wish to compute a pricing of the items that maximizes the seller's profit. We assume that if the total price of the items in e is at most v_e, then the customer (e, v_e) will purchase all of the items in e, and otherwise the customer will purchase nothing. Given a price vector $\mathbf{p}$ over the n items, it will be convenient to define $p(e) = \sum_{i \in e} p_i$.

Let us denote by E the set of customers, and V the set of items, and let h be $\max_{e \in E} v_e$. Let $G = (V, E, v)$ be the induced hypergraph, whose vertices represent the set of items, and whose hyperedges represent the customers. Notice that G might contain self-loops (since a customer might be interested in only a single item) and multi-edges (several customers might want the same subset of items). The special case that all customers want at most two items, so G is a graph, is known as the *graph vertex pricing* problem [1]. Another interesting case considered in previous work [1,8] is the *highway* problem. In this problem we think of the items as segments of a highway, and each desired subset e is required to be an interval $[i, j]$ of the highway.

Reduced Instance: In many of our algorithms, it is convenient to think about the *reduced instance* $\tilde{G} = (V, E, w)$ of the problem which is defined as follows. Let b_i denote the marginal cost of item i. Suppose customer e has valuation v_e. Then, in the reduced instance, its valuation becomes $w_e := v_e - \sum_{i \in e} b_i$. Now, if we give item i a price p_i in the reduced instance, then its real selling price would be $s_i := p_i + b_i$. In previous work [1,8,4], the focus was on pricing above cost, which in our notation, corresponds to the case where $p_i \geq 0$, for every item i. However, as mentioned in the introduction, in many natural cases, we can potentially extract more profit by pricing certain items below cost (which corresponds to the case where $p_i < 0$).

From now on, we always think in terms of the reduced instance. We formally define all the *pricing models* we consider as follows:

Positive Price Model: In this model, we require the selling price of an item to be at or above its production cost. Hence, in the reduced instance, we want the price vector $\mathbf{p}$ with positive components $p_i \geq 0$ that maximizes $\text{Profit}_{pos}(\mathbf{p}) = \sum_{e: w_e \geq p(e)} p(e)$. Let $\mathbf{p}^*_{pos}$ be the price vector with the maximum profit under positive prices and let $\text{OPT}_{pos} = \text{Profit}_{pos}(\mathbf{p}^*_{pos})$.

Discount Model: In this model, the selling price of an item can be arbitrary. In particular, the price can be below the cost, or even below zero. We want the price vector $\mathbf{p}$ that maximizes $\text{Profit}_{disc}(\mathbf{p}) = \sum_{e: w_e \geq p(e)} p(e)$. Let $\mathbf{p}^*_{disc}$ be the price vector with the maximum profit and let $\text{OPT}_{disc} = \text{Profit}_{disc}(\mathbf{p}^*_{disc})$.

B-Bounded Discount Model: In this model, the selling price of an item i can be below its production cost b_i, but cannot be below zero. This corresponds to a negative price in the reduced instance, but it is bounded below by $-b_i$. For simplicity, we assume that the production costs of all items are each B. We want the price vector $\mathbf{p}$ with components $p_i \geq -B$ that maximizes $\text{Profit}_B(\mathbf{p}) = \sum_{e:w_e \geq p(e)} p(e)$. Let $\mathbf{p}_B^*$ be the price vector with the maximum profit and let $\text{OPT}_B = \text{Profit}_B(\mathbf{p}_B^*)$. Observe that $\text{OPT}_{pos} \leq \text{OPT}_B \leq \text{OPT}_{disc}$.

Coupon Model: This model makes most sense in which the items have zero marginal cost, such as airport taxes or highway tolls. In this model, the selling price of an item can actually be negative. However, we impose the condition that the seller not make a loss in any transaction with any customer. We want the price vector $\mathbf{p}$ that maximizes $\text{Profit}_{coup}(\mathbf{p}) = \sum_{e:w_e \geq p(e)} \max(p(e), 0)$. Let $\mathbf{p}_{coup}^*$ be the price vector with the maximum coupon profit and let $\text{OPT}_{coup} = \text{Profit}_{coup}(\mathbf{p}_{coup}^*)$. From the definition, it is immediate that $\text{OPT}_{pos} \leq \text{OPT}_{coup}$.

Gaps between the Models. We state below a few fundamental gaps between the profits obtainable in these models.

Theorem 1. *For the highway problem, there exists an $\Omega(\log n)$ gap between the positive price model and the (B-bounded) discount model, even for $B = 1$. Moreover, there exists an $\Omega(\log n)$ gap between the coupon model and the (B-bounded) discount model.*

Theorem 2. *For the graph vertex pricing problem*[1]*, there exists an $\Omega(\log B)$ gap between the positive price model and the B-bounded discount model, even for a bipartite graph.*

3 Main Tools and Main Results

We describe now the main tools used in the paper. These tools allow us to give bounds on the prices of items in an optimal solution in each of the pricing models.

DAG Representation of the Highway Problem: We describe here an alternative representation of the Highway Problem. This representation proves to be extremely convenient both for the analysis and for the design of algorithms.

Suppose the n items are in the order $l_1, l_2, \ldots, l_n$, with corresponding prices $p_1, p_2, \ldots, p_n$. Then, for each $0 \leq i \leq n$, we have a node v_i labelled with the partial sum $s_i := \sum_{j=1}^{i} p_i$, where $s_0 = 0$. A customer corresponds to a subset of the form $\{l_i, \ldots, l_j\}$, which is represented by a directed arc from v_{i-1} to v_j.

Lemma 1. *Under all pricing models (positive price model, (bounded) discount model, coupon model), there is always an optimal solution such that $s_{max} - s_{min} \leq nh$, where $s_M := \max\{s_i : 0 \leq i \leq n\}$ and $s_m := \min\{s_i : 0 \leq i \leq n\}$, and h is the maximum valuation.*

[1] The graph vertex pricing problem is APX-hard under all our models. One can easily extend the result in [8] to our setting too.

Existence of Bounded Solution for Graph Vertex Pricing: Recall that in the graph setting, we denote the set of items by V, and each customer is interested in at most two items. We represent the set of customers interested in exactly two items by the set of (multi) edges E, and the set of customers interested in exactly one item by the (multi) set N, where for each $e \in E \cup N$, $w_e \in \mathbb{Z}$ is customer e's valuation.

Lemma 2. *Under all the pricing models (the coupon model and (bounded) discount model), there is an optimal price vector $\mathbf{p}^* \in \mathbb{R}^V$ that is half-integral if all customers' valuations are integral. Moreover, if all valuations are at most h, then $\mathbf{p}^*$ can be chosen to be bounded in the sense that for all $v \in V$, $|\mathbf{p}^*(v)| \leq 2nh$.*

3.1 Coupon Model

The main feature of the coupon model is that even when the sum of the prices for the items that a customer wants is negative, the net profit obtained from that customer is zero.

A Constant Factor Approximation for the Highway Problem: We show here a constant factor approximation algorithm for the highway problem under the coupon model, in the case where all the customers' valuations are identical.

Theorem 3. *There is a 2.33-approximation algorithm under the coupon model for the highway problem in the case when all all customers' valuations are all 1.*

Proof. First, we represent the problem as a DAG as described above: each node corresponds to a partial sum and each customer is represented as a directed edge from its left node to its right node. We then use the approximation algorithm presented in [7] for the MAX DICUT problem to get a $\frac{1}{0.859}$-approximation for OPT that uses no more than two levels, i.e., the partial sums are either 0 or 1. Hence, in order to show the result, it suffices[2] to show that there exists a solution in which the partial sums are either 0 or 1 and has profit at least $\frac{1}{2}\text{OPT}_{coup}$. Consider the partial sums in an optimal solution. Observe that for each customer from which we get a profit (of 1), we still obtain a profit for that customer after modifying the solution in exactly one of the following ways: If a partial sum is even, set it to 0, otherwise set it to 1. If a partial sum is even, set it to 1, otherwise set it to 0. Hence, by choosing the modification that yields higher profit, the claim follows. □

Theorem 4. *Under the coupon model we have a fully polynomial time approximation scheme for the case that the desired subsets of different customers form a hierarchy.*

Planar and Minor-free Graph Vertex Pricing Problem: We give a PTAS that uses negative prices to obtain $(1+\epsilon)$-approximation, using decomposition techniques for H-minor-free graphs by Demaine et al. [6]

[2] If all the valuations are integral, then there exists an optimal solution with all prices integral, under all our models (positive, coupon, and (B-bounded) discount models).

Theorem 5. *There exists a PTAS for minor-free instances of the graph vertex pricing problem under the coupon model.*

3.2 *B*-Bounded Discount Model

The main feature is that the net profit we obtain from a customer is exactly the sum of the prices of the items in the bundle of that customer, and hence can be negative. As explained in the introduction, the extra condition that the price of an item must be at least $-B$ corresponds to the real life situation in which the selling price of an item can be below its cost, but not negative.

Theorem 6. *There exists an $O(B)$ approximation algorithm for the vertex pricing problem under the B-bounded discount model.*

There exists an PTAS for minor-free instances of the graph vertex pricing problem under the B-bounded discount model for fixed B under either one *of the following assumptions: (1) All customers have valuations at least 1, or (2) There is no multi-edge in the graph.*

Theorem 7. *There exists an FPTAS for the case that the desired subsets of different customers form a hierarchy under both the discount and the B-bounded discount models.*

References

1. Balcan, M.-F., Blum, A.: Approximation Algorithms and Online Mechanisms for Item Pricing. In: ACM Conference on Electronic Commerce (2006)
2. Balcan, M.-F., Blum, A., Chan, T.-H.H., Hajiaghai, M.T.: A theory of loss-leaders: Making money by pricing below cost. Technical Report, CMU-CS-07-142 (July 2007)
3. Balcan, M.-F., Blum, A., Hartline, J., Mansour, Y.: Mechanism Design via Machine Learning. In: 46th Annual IEEE Symposium on Foundations of Computer Science, pp. 605–614. IEEE Computer Society Press, Los Alamitos (2005)
4. Briest, P., Krysta, P.: Single-Minded Unlimited Supply Pricing on Sparse Instances. In: Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms, ACM Press, New York (2006)
5. DeGraba, P.: Volume discounts, loss leaders, and competition for more profitable customers. Federal Trade Commission Bureau of Economics, Working Paper 260 (February 2003)
6. Demaine, E.D., Hajiaghayi, M., Kawarabayashi, K.: Algorithmic Graph Minor Theory: Decomposition, Approximation, and Coloring. In: 46th Annual IEEE Symposium on Foundations of Computer Science, pp. 637–646. IEEE Computer Society Press, Los Alamitos (2005)
7. Feige, U., Goemans, M.X.: Approximating the Value of Two Prover Proof Systems, with Applications to MaxtwoSat and MaxDicut. In: Proceedings of the Third Israel Symposium on Theory of Computing and Systems, pp. 182–189 (1995)
8. Guruswami, V., Hartline, J., Karlin, A., Kempe, D., Kenyon, C., McSherry, F.: On Profit-Maximizing Envy-Free Pricing. In: Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1164–1173. ACM Press, New York (2005)