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

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

July 2007
CMU-CS-07-142

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

* School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.
{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 *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. To design our algorithms, we use several tools including a particular DAG representation and graph decomposition techniques which may be of independent interest.

Research supported in part by NSF grants CCR-0105488, CCR-0122581, and IIS-0121678.

Keywords: Pricing Below Cost, Unlimited Supply, Combinatorial Auctions, Single Minded Bidders, Mechanism Design, Profit Maximization

1 Introduction

The notion of *loss-leaders*, namely pricing certain items below cost in a way that increases profit overall from 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. Video game makers often sell game consoles below cost and make up their profit on the games themselves.

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 [7] 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 [2] give an example, in the context of pricing n items of fixed marginal cost to a set of single-minded customers, where allowing items to be priced below cost can produce an $\Omega(\log n)$ factor more profit than possible by pricing all items above cost. However, the issue 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 *discount model* and the *coupon 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" [14] and problems of pricing vertices in graphs, as well as analyze fundamental gaps between the profit obtainable under the different models.

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 her 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 at least satisfies monotonicity: because no individual customer produces a loss, increasing the valuation of a customer cannot reduce overall profit.

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. [14]: 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. In this setting, the case that all customers have the same valuation v is not interesting (you can simply price all items at $v/2$), but if customers have different valuations then there again are gaps between the models. We show that if this graph is *planar* then one can in fact achieve a PTAS for profit in each model.

Related work: Revenue maximizing auctions and algorithmic pricing problems have generated a great deal of interest recently; for a survey see the chapter "Profit Maximization in Mechanism Design" of Hartline and Karlin in [16]. Here we briefly describe related algorithmic work in the context of combinatorial auctions with single-minded customers.

Algorithmic pricing problems of this form were first posed by Guruswami et al. [14] though item-pricing for unit-demand consumers with several alternative payment rules (i.e., rules that do not represent quasi-linear utility maximization) were independently considered by Aggarwal et al. [1]. Guruswami et al. [14] show an $O(\log m + \log n)$-approximation for the general *hypergraph* problem (where customers have valuations over larger subsets,), where n is the number of items (vertices) and m is the number of customers (hyperedges). They also show that even the *graph* vertex pricing problem is APX-hard — and this is true even when all valuations are identical (if self-loops are allowed) or all valuations are either 1 or 2 (if self-loops are not allowed). In related work, Hartline and Koltun [15] give a $(1+\epsilon)$-approximation that runs in time exponential in the number of vertices, but that is near-linear time when the total number of vertices in the hypergraph is constant. Recently, Demaine, Feige, Hajiaghayi, and Salavatipour [8] have shown that it is hard to approximate the hypergraph vertex pricing problem within a factor of $\log^{\delta} n$, for some $\delta > 0$, assuming that NP $\not\subseteq$ BPTIME$\left(2^{n^{\epsilon}}\right)$ for some $\epsilon > 0$. In the case that the cardinality of the desired bundles is bounded by k, Briest and Krysta [4] give an $O(k^2)$ approximation algorithm, which is improved to $O(k)$ by Balcan and Blum [2]. Finally, both Briest and Krysta [4] and Grigoriev et al. [13] proved that optimal pricing is weakly NP-hard for the special case of the *highway* problem. For the same problem [2] improved the approximation of Guruswami et al. [14] from $O(\log m + \log n)$ to $O(\log n)$. We note that all previous work focused just on the case that pricing below production cost is not allowed; pricing below cost was posed as an intriguing open question in [2].

Our results: The pricing problems as defined above are very general. In this paper, we focus on designing efficient algorithms for several important cases of this problem with guaranteed approximation factors.

We begin by proving fundamental profitability gaps between the different models under consideration. We then present several structural properties of our models that will be used in our algorithms. One of these is that for the graph and highway problems, we can polynomially bound the absolute values of positive and negative prices in an optimal solution in the coupon and unbounded discount models. Note that this is not trivial since one could imagine very large (e.g., exponential) positive and negative prices which cancel each other out. In fact, it is not even true for the general hypergraph problem even for edges of size 3. We also introduce a DAG representation for dealing with the highway tollbooth problem, that is especially convenient for algorithms. Roughly speaking, in this representation, the problem becomes one of partitioning a DAG into levels to optimize different objectives depending on the models. Here the vertices of the DAG are partial sums for initial segments of the highway tolls and the edges correspond to the customers who want to use this highway partially. We use the DAG representation both for designing our algorithms, and for proving the existence of bounded optimal solutions under all models (see above).

Next, we present two constant factor approximation algorithms for the highway problem under the coupon model when the customers valuations are all equal. The first has approximation factor 2.33 and is based on a semidefinite programming approach. The second approximation is a simpler 4-approximation *randomized* algorithm which has the advantage of being *oblivious*; here, we post prices for the items independently from the customers' valuations (even before seeing any customers) and in this sense our algorithms are even stronger than online pricing in this model. Our 4-approximation extends to the tollbooth problem on tree-networks as well (not just the line). Moreover, these algorithms generalize to the case when customers have a small number of distinct valuations. Furthermore, with extra assumptions, we obtain a QPTAS for the highway problem under general customer valuations with a bounded number of negatively priced items.

We show two polynomial-time approximation schemes (PTAS's), one for vertex pricing of planar graphs (more generally graphs excluding a fixed minor) and one for the special case of the highway problem when the customers form a hierarchy, under both coupon and discount models. In the former PTAS, we use the recent result of Demaine et al. [9] for decomposing H-minor-free graphs into graphs of bounded treewidth, and for the latter PTAS, we use bounded negative prices (as mentioned above) together with dynamic programming.

It is worth mentioning we do not focus on *incentive-compatibility* issues in this paper, since that issue is very similar to the one considered in [3, 2].

2 Notation and Definitions

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 [2, 5, 6]. Another interesting case considered in previous work [2, 5, 6, 14] 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 [2, 5, 6, 14], 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 $\mathrm{Profit}_{pos}(\mathbf{p}) = \sum_{e: w_e \geq p(e)} p(e)$. Let $\mathbf{p}^*_{pos}$ be the price vector with the maximum profit under positive prices and let $\mathrm{OPT}_{pos} = \mathrm{Profit}_{pos}(\mathbf{p}^*_{pos})$.

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 $\mathrm{Profit}_{disc}(\mathbf{p}) = \sum_{e: w_e \geq p(e)} p(e)$. Let $\mathbf{p}^*_{disc}$ be the price vector with the maximum profit and let $\mathrm{OPT}_{disc} = \mathrm{Profit}_{disc}(\mathbf{p}^*_{disc})$.

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 $\mathrm{Profit}_B(\mathbf{p}) = \sum_{e: w_e \geq p(e)} p(e)$. Let

$\mathbf{p}_B^*$ be the price vector with the maximum profit and let $\mathrm{OPT}_B = \mathrm{Profit}_B(\mathbf{p}_B^*)$. Observe that from the definitions $\mathrm{OPT}_{pos} \leq \mathrm{OPT}_B \leq \mathrm{OPT}_{disc}$.

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 $\mathrm{Profit}_{coup}(\mathbf{p}) = \sum_{e:w_e \geq p(e)} \max(p(e), 0)$. Let $\mathbf{p}_{coup}^*$ be the price vector with the maximum coupon profit and let $\mathrm{OPT}_{coup} = \mathrm{Profit}_{coup}(\mathbf{p}_{coup}^*)$. From the definition, it is immediate that $\mathrm{OPT}_{pos} \leq \mathrm{OPT}_{coup}$.

2.1 Gaps between the Models

We state below a few fundamental gaps between the models. In proving these gaps we use the following note which follows from a proof in [14].

Note 1 *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 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 $B = 1$.*

Proof: In this example, each customer e has valuation $w_e = 1$. Let S_0 be the instance that consists of one item i_0 and one customer of the form $\{i_0\}$. We then define S_r recursively for $r \geq 1$. In order to construct S_r, we first construct two copies of S_{r-1}, placed side by side, with one new item placed between the two copies; finally, we include 2^r customers, each of whose set includes all items.

It follows that the number of items in S_r is $n_r := 2^{r+1} - 1$. Moreover, by setting the prices of the items to (1, -1, 1, -1, $\cdots$, -1, 1), we can collect one unit of profit from each customer, and hence we have $\mathrm{OPT}_B(S_r) = 2^r r$.

Next, we show by induction on r that $\mathrm{OPT}_{pos}(S_r) = 2^{r+1} - 1$. The claim is trivial for $r = 0$ as there is only 1 customer. Assume the result is true for some $r \geq 0$. Consider the instance S_{r+1}. Observe that we can assume that the optimal solution has integral prices, by Note 1. If we collect some profit from customers whose sets include all items, there can be exactly one item priced 1 and the rest priced 0. Hence, in this case, the profit is $1 + 2 + \cdots + 2^{r+1} = 2^{r+2} - 1$. If we do not collect any profit those customers, then we have two independent copies of S_r and hence by induction hypothesis, the maximum profit collected is $2(2^{r+1} - 1) = 2^{r+1} - 2$. Hence, we have $\mathrm{OPT}_B/\mathrm{OPT}_{pos} = \Omega(r) = \Omega(\log n_r)$ and the claim follows. ■

Theorem 2 *For the highway problem, there exists an $\Omega(\log n)$ gap between the coupon model and the (B-bounded) discount model.*

Proof: This proof is a bit more involved than for Theorem 1 and appears in Appendix A. ■

Theorem 3 *For the graph vertex pricing problem, there exists an $\Omega(\log B)$ gap between the positive price model and the B-bounded discount model, even for a bipartite graph.*

Proof: Analogous to the example in [2]. ■

Note 2 *The graph vertex pricing problem is APX-hard under all our models.*[1]

[1] One can easily extend the result in [14] to our setting too.

3 Main Tools

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

3.1 DAG Representation of the Highway Problem

We describe here an alternative, very convenient representation of the Highway Problem. This representation proves to be extremely convenient both for the analysis (as seen in the following lemma) and for the design of algorithms (as seen in the algorithms we present in Section 4.1).

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.*

Proof: In each of the models, consider the optimal solution that minimizes $s_M - s_m$. Suppose in this case, we still have $s_M - s_m > nh$. Then, it follows that there must exists an open interval I in the real line of length L strictly greater than h such that no s_i lies in I. We show that it is possible to reduce all those s_i's that is to the right of I by $\delta := L - h$ without decreasing the profit. This would contradict the minimality of $s_M - s_m$.

Note that we only have to consider arcs between the left and right of I. If there are no such arcs, then of course we can reduce all those s_i's to the right of I by δ without changing the profit. If there is an arc going from s_i to s_j where $s_i < I$ and $s_j > I$, then no profit is generated from the customer corresponding to this arc anyway. After reducing s_j by δ, the valuation of this arc is still at least h and so the profit due to this arc cannot decrease.

It remains to consider an arc going from s_j to s_i, where $s_j > I$ and $s_i < I$. There can be no such arc in the positive price model. In this case, we have to suffer a loss of $s_j - s_i$ in the (bounded) discount model and just gain nothing in the coupon model. Hence, if we reduce s_j by δ, the loss suffered would be reduced by δ in the discount model due to this arc, and no change in the coupon model. In either case, the total profit does not decrease. ■

3.2 The Existence of a Bounded Solution for Graph Vertex Pricing

Remember 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 (in particular 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 $v \in V$, $|\mathbf{p}^*(v)| \leq 2nh$.*

Proof: Our description works for all models and distinctions would be made where necessary. We first consider B-bounded discount model. Observe that prices must be in the range $[-B, h+B]$. Hence, if B is small, then the result is trivial. If B is large, then we actually show that it is not necessary to price too low by resorting to the proof for the (unbounded) discount model.

Observe that the optimal price vector $\mathbf{p}^*$ is an extreme point of some polytope which we describe as follows. Suppose we know $A \subseteq E$ and $B \subseteq N$ are the customers from whom we obtain a positive profit in an optimal solution. Consider the following polytope:

$$p_i + p_j \quad \leq w_e, \quad \text{for } e = \{i, j\} \in A \tag{1}$$

$$p_i \quad \leq w_e, \quad \text{for } e = \{i\} \in B \tag{2}$$

Note that the objective function we are maximizing is different under different models.

Coupon model: $\sum_{\{i,j\} \in A}(p_i + p_j) + \sum_{\{i\} \in B} p_i$

discount model: $\sum_{\{i,j\} \in E}(p_i + p_j) + \sum_{\{i\} \in N} p_i$

Hence, an optimal solution $\mathbf{p}^*$ corresponds to some basic solution specified by a system of equalities indexed by $J \subseteq E$ and $I \subseteq N$.

$$p_i + p_j \quad = w_e, \quad \text{for } e = \{i, j\} \in J$$

$$p_i \quad = w_e, \quad \text{for } e = \{i\} \in I$$

We first show that there exists a half-integral $\mathbf{p}$ such that for all $e = \{i, j\} \in J$, $\mathbf{p}(i) + \mathbf{p}(j) = w_e$. Observe that if a component of the graph (V, J) contains some i such that $\{i\} \in I$, then for all j in that component, $\mathbf{p}(j)$ is integral. Hence, without loss of generality, we assume that I is empty. First consider the case when the graph (V, J) is bipartite with bipartition (A, B). Observe that for any $\delta > 0$, we can increase the price of all items in A by δ and decrease the price of all items in B by δ without violating any equalities in J. By considering each connected component of (V, J) and observing that w_e's are integral, we conclude there exists an integer $\mathbf{p}$ satisfying all equalities in J.

For general (V, J), we use a standard technique. Consider two copies V^1, V^2 of V such that for each $e = \{i, j\} \in J$, we have two equalities: $p_i^1 + p_j^2 = w_e$ and $p_i^2 + p_j^1 = w_e$. It follows that $\{p_i = \frac{1}{2}(p_i^1 + p_i^2)\}$ is a solution to the original system *iff* $\{p_i^1, p_i^2\}$ is a solution to the new system. Since the new system corresponds to a bipartite graph, it must have an integral solution. Hence, it follows that the original system has a half-integral solution.

We next show that there exists a bounded optimal price vector $\mathbf{p}^*$. We consider each connected component C in (V, J). Consider any two vertices x and y in C and any $\mathbf{p}$ satisfying the equalities in J. Then, there exists $x = x_0, x_1, \ldots, x_k = y$, where $k \leq |C| - 1$, such that for $1 \leq i \leq k$, $e_i = \{x_{i-1}, x_i\} \in J$. Hence, it follows that $\mathbf{p}(x_{i-1}) + \mathbf{p}(x_i) = c_i$. Multiplying each equation by $(-1)^{i+1}$ and summing up over i, we have $|\mathbf{p}(x_0) + (-1)^{k+1}\mathbf{p}(x_k)| = |\sum_{i=1}^{k}(-1)^{i+1}c_i| \leq kh \leq (n-1)h$. Hence, if we can show that there exists $x \in C$ such that $|\mathbf{p}(x)| \leq nh$, then it follows that for all y in C, $|\mathbf{p}(y)| \leq 2nh$. This is certainly achieved if the component contains some i such that $\{i\} \in I$. Hence, without loss of generality, we assume this is not the case. If the component C is bipartite, then using a similar argument as above, one can show that there exists a solution $\mathbf{p}$ such that for some x in C, $\mathbf{p}(x) = 0$. If the component C is non-bipartite, then there must be an odd cycle $x_0, x_1, \ldots, x_k, x_{k+1} = x_0$, where k is even. As before, we have $2x_0 = x_0 + (-1)^{k+2}x_{k+1} = \sum_{i=1}^{k+1} c_i$. Hence, it follows that $|x_0| \leq \frac{1}{2}nh$. ■

4 Coupon Model

We next consider the 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 still zero.

4.1 Constant Factor Approximation Algorithms for the Highway Problem

We show in the following constant factor approximation algorithms for the highway problem under the coupon model, in the case when all the customers' valuations are identical.[2]

We start by presenting our first 4-approximation algorithm. As mentioned before, this algorithm is oblivious, i.e., we post prices for the items *before* seeing any customers and in this sense our guarantees are even stronger than those for the online pricing in this model. We also show how we can extend this algorithm to tollbooth problem on tree-networks.

Since all valuations w_e are equal, for simplicity and without loss of generality we assume all are equal to 1.

Theorem 4 *There is a* 4*-approximation algorithm under the coupon model for the highway problem in the case when all all customers' valuations are precisely* 1.

Proof: First, we represent the problem as a DAG as described in Section 3.1: each node corresponds to a partial sum and each customer is represented as a directed edge from its left node to right node. Then, we simply randomly assign 0 or 1 to each partial sum independently. (This then clearly corresponds to a pricing vector in the original instance that uses only prices 0, 1 and -1.)

To see that this algorithm is a 4-approximation algorithm, it's enough to simply notice that with probability $\frac{1}{4}$ we get a profit of 1 from each customer; this clearly implies a 4-approximation algorithm for the coupon model. ■

Markov Chain interpretation: The randomized algorithm described in Theorem 4 has a Markov Chain interpretation. The system has a state that corresponds to the partial sum of the prices of items beginning from the left. If the state is 0, then with probability $\frac{1}{2}$, the state remains the same and the price of the next item is 0, and with probability $\frac{1}{2}$, the state changes to 1 and the price of the next item is 1. Similarly, if the state is 1, then with probability $\frac{1}{2}$, the state remains the same and the price of the next item is 0, and with probability $\frac{1}{2}$, the state changes to 0 and the price of the next item is -1.

Note: Our randomized algorithm can actually be generalized for the Highway Problem on Trees. In this setting, we are given a tree and each edge corresponds to an item, while a customer's bundle corresponds to a (simple) path in the tree, and as before we assume that each customer's valuation is 1. For a precise algorithm and proof see Theorem15 in Appendix A.

Note that the algorithms in Theorems 4 and 15 are extremely simple and they use *posted* prices, which directly implies we can adapt them to the online setting. We present in the following a more refined approximation algorithm that achieves a better 2.33-approximation guarantee for the basic highway problem.

Theorem 5 *There is a* 2.33*-approximation algorithm under the coupon model for the highway problem in the case when all all customers' valuations are precisely* 1.

Proof: First, we represent the problem as a DAG as described in Section 3.1: each node corresponds to a partial sum and each customer is represented as a directed edge from its left node to right node. We then use the algorithm 0.859 approximation algorithm presented in [12] for the MAX DICUT problem to get $\frac{1}{0.859}$ 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 to use Note 1 and show that there exists a solution in which the partial sums are either 0 or 1 and has profit at least $\frac{1}{2}\mathrm{OPT}_{coup}$. Consider the partial sums in an optimal

[2] For this special case the problem was shown to be solvable in polynomial time under the positive price model; specifically, [14] shows an exact $O(m^3)$-time dynamic programming algorithm.

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:
1. If a partial sum is even, set it to 0, otherwise set it to 1.
2. 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. ■

Note that Lemma 1 implies here a fully polynomial time approximation scheme for the case that the desired subsets of different (single-minded) customers form a hierarchy. Specifically:

Theorem 6 *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.*

Proof: We can extend the algorithm presented in [2] for the positive price model. The correctness follows from the analysis in [2] by additionally using the result in Lemma 1. ■

Moreover, under additional assumptions, we can obtain $(1+\epsilon)$-approximation with a QPTAS, which is based on the construction by Elbassioni et. al. [11]. We sketch the proof of the following theorem in Appendix B.

Theorem 7 *There exists a QPTAS for the highway problem under the coupon model with general valuations if there exists an optimal pricing that has only a bounded number of negatively priced items, and the corresponding revenue obtained from positively priced items is at least twice that of the revenue from negatively priced items.*

4.2 Planar and Minor-free Graph Vertex Pricing Problem

Recall that in the graph vertex pricing problem, E is the multi-set of customers who are interested in exactly two items, and N is the multi-set of customers who are interested in exactly one item. For any graph H, the instance is H-minor free if the graph (V, E) is H-minor-free. For example planar graphs are both $K_{3,3}$-minor-free and K_5-minor-free. In this subsection, we consider minor-free instances of the problem under the coupon model. We first show that in for this type of restricted input, the optimal profit obtained under the positive price model is at least a constant fraction of that obtained under the coupon Model.

Constant Gap between Positive and Coupon Models for Planar and minor-free Graphs

We use the following deep result from DeVos et al. [10] that states that a minor-free graph can be edge-partitioned into a constant number of trees.

Theorem 8 ([10]) *For any graph H, there exists an integer J_H such that any H-minor free graph can be edge-partitioned into J_H trees.*

Theorem 9 *Consider the graph vertex pricing problem on a H-minor free graph. Then, there exists a constant C_H depending only on H such that the optimal profit under the positive price model is at least $1/C_H$ times the optimal profit under the coupon model.*

Proof: First, suppose we have the vertex price problem on a tree. Pick any vertex as the root and consider the odd and the even level edges. An odd level edge is one in which the further end point from the root is at odd distance from the root. Then, by obtaining profit from either the odd or the even level edges using positive prices, we can obtain at least half the optimal profit obtained under the coupon model. By Theorem 8, any H-minor free graph can be partitioned into J_H trees. Hence, by considering profit from positive pricing for each of the trees, we conclude that be setting $C_H := 2J_H$, there exists a positive pricing that achieves at least $1/C_H$ fraction of the optimal profit under coupon model. ■

PTAS for Planar and Minor-free Graph Instance

Theorem 10 *There exists a PTAS for minor-free instances of the graph vertex pricing problem under the coupon model.*

We will use the following decomposition procedure for H-minor-free graphs by Demaine et al. [9]- see Theorem 3.1.

Theorem 11 (PARTITION INTO BOUNDED-TREEWIDTH GRAPHS)*For any graph H, there is a constant C_H, such that for any integer $k \geq 2$ and for any H-minor-free graph G, the edges of G can be partitioned into k sets such that any $k-1$ of the sets induce a graph of treewidth at most $C_H k$. Furthermore, such a partition can be found in polynomial time.*

Proof of Theorem 10 We next describe how to use Theorem 11 to solve the minor-free graph vertex pricing problem.

Algorithm Outline We use the decomposition procedure in Theorem 11 to partition the edges of (V, E) into $k = \frac{1}{\epsilon}$ sets $E = \cup_{i=1}^{k} E_i$. We discard the set E_i whose sum of valuations is the smallest, and let E' to be the union of the remaining $k-1$ parts. Since (V, E') has bounded treewidth, we can use dynamic programming to solve the instance restricted to E' and N exactly, thereby achieving a solution whose value is at least $(1-\epsilon)$ that of the optimal solution.

We next describe a dynamic program to solve the problem restricted to the customers in $E' \cup N$. Let $\mathcal{T} = (T, E_T)$ be a tree decomposition of (V, E'). In particular, for each $t \in T$, there exists $V_t \subseteq V$ such that $V = \cup_t V_t$, and for each $\{i,j\} \in E'$, there exists $t \in T$ such that $\{i,j\} \subseteq V_t$. Moreover, if $\{t_1, t_2\}, \{t_2, t_3\} \in E_T$, then $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$. Since the treewidth of (V, E') is at most $C_H k$, we can assume for all $t \in T$, $|V_t| = \kappa := C_H k + 1$.

Pick any node in T to be the root r. The terms *child, parent, ancestor* and *descendant* are used in their usual sense. For any $t \in T$, let $\mathcal{T}(t)$ be the subtree rooted at t. For any $t \in T$, let $A(t)$ be the set of customers whose items are contained in $\cup_{q \in \mathcal{T}(t)} V_q$.

By Lemma 2, we only need to consider prices in $U := \{s : |s| \leq 2nh, 2s \in \mathbb{Z}\}$. The running time would have hence have a dependence on h. If h is large, then we can use a standard rounding argument. Let $\delta := \epsilon h/2n$ and consider prices in $U' := \{l\delta : |l\delta| \leq 2nh, l \in \mathbb{Z}\}$. Note that $|U'| = O(mn^2/\epsilon)$. Let S be the one among U and U' that has smaller cardinality. Hence, we consider prices from a set of size at most $|S| = O(n) \cdot \min\{mn/\epsilon, h\}$.

For each $t \in T$, we fix some ordering on the r vertices in V_t such that any $x \in S^\kappa$ represents a price assignment for the items in V_t.

Suppose $u, v \in T$ such that u is the parent of v. Given $x_u, x_v \in S^\kappa$, let $B(u, v, x_u, x_v)$ be the profit obtained from customers of the form $e = \{i, j\}$ such that $i \in V_u$, $j \in V_v$ and $e \notin A(v)$, where items in V_u are priced according to x_u and those in V_v according to x_v. If x_u and x_v are inconsistent because of common items in V_u and V_v, let $B(u, v, x_u, x_v)$ be $-\infty$.

Given $t \in T$ and $x \in S^\kappa$, let $C(t, x)$ be the maximum profit obtained from customers in $A(t)$ such that items in V_t are priced according to x. If t is a leaf, then it is trivial to compute $C(t, x)$. Otherwise, $C(t, x)$ can be computed by the following equation: $C(t, x) = \sum_u \max_{x' \in S^r}(B(t, u, x, x') + C(u, x'))$, where the summation is over all children u of t. The optimal solution to the problem is given by $\max_{x \in S^\kappa} C(r, x)$. To compute each $C(t, x)$, we need to try each x' for each child of x. Hence, this takes time $n \cdot |S|^\kappa$. The total number of such entries is at most $n \cdot |S|^\kappa$. Hence, the total time is $n^2 |S|^{2\kappa} \leq n^2 \cdot |S|^{O(C_H/\epsilon)}$, where $|S| = O(n) \cdot \min\{mn/\epsilon, h\}$.

5 B-Bounded Discount Model

We next consider the 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. We state below our main results in this model and we prove them in the Appendix A.

Theorem 12 *There exists an $O(B)$ approximation algorithm for the vertex pricing problem under the B-bounded discount model.*

Theorem 13 *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.
(2) There is no multi-edge in the graph.

Theorem 14 *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.*

6 Discussion and Open Problems

In this paper, we have formally introduced the problem of assigning positive or negative markups to goods of fixed marginal cost in order to maximize revenue in the presence of single-minded customers, and we have shown that by pricing certain items below their marginal costs, we can obtain a substantial improvement in the overall profit. We have also developed approximation algorithms for the *discount* and *coupon* models for several interesting special cases considered previously in the literature.

Obtaining constant factor appropriation algorithms in the coupon model for general graph vertex pricing problem[3] and the highway problem with arbitrary valuations seems believable but very challenging. In the light of the proof construction of the recent almost logarithmic hardness result by Demaine et al. [8], it is conceivable that in the coupon model, finding an optimal pricing for general single-minded customers (even all with valuations one) is hard to approximate better than a logarithmic factor.

We have also studied the *discount without loss* model, another mathematically interesting model for pricing below cost. This is a model "in between" the *discount* model and the *coupon* model. Here our goal is to compete with an optimum which prices the items in a way that for any customer, the sum of the selling prices of items in his desired set is at least the sum of the marginal costs. We present some results in Appendix C and it would be interesting to explore this model further.

[3]Note that obtaining a constant factor approximation for bipartite graphs implies a constant factor approximation for general graphs in the coupon model.

References

[1] G. Aggarwal, T. Feder, R. Motwani, and A. Zhu. Algorithms for multi-product pricing. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*, pages 72–83, 2004.

[2] M.-F. Balcan and A. Blum. Approximation Algorithms and Online Mechanisms for Item Pricing. In *ACM Conference on Electronic Commerce* , 2006.

[3] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism Design via Machine Learning. In *46th Annual IEEE Symposium on Foundations of Computer Science*, pages 605 – 614, 2005.

[4] P. Briest and P. Krysta. Single-Minded Unlimited Supply Pricing on Sparse Instances. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[5] P. Briest and P. Krysta. Single-Minded Unlimited Supply Pricing on Sparse Instances. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1093 – 1102, 2006.

[6] P. Briest and P. Krysta. Buying Cheap is Expensive: Hardness of Non-Parametric Multi-Product Pricing. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2007.

[7] P. DeGraba. Volume discounts, loss leaders, and competition for more profitable customers. Federal Trade Commission Bureau of Economics, Working Paper 260, February 2003.

[8] E. D. Demaine, U. Feige, M. Hajiaghayi, and M. R. Salavatipour. Combination Can Be Hard: Approximability of the Unique Coverage Problem. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 162 – 171, 2006.

[9] E. D. Demaine, M. Hajiaghayi, and K. Kawarabayashi. Algorithmic Graph Minor Theory: Decomposition, Approximation, and Coloring. In *46th Annual IEEE Symposium on Foundations of Computer Science*, pages 637–646, 2005.

[10] M. DeVos, G. Ding, B. Oporowski, D. P. Sanders, B. Reed, P. Seymour, and D. Vertigan. Excluding any graph as a minor allows a low tree-width 2-coloring. In *Journal of Combinatorial Theory Series B. Volume 91 , Issue 1*, pages 25–41, 2004.

[11] K. Elbassioni, R. Sitters, and Y. Zhang. A quasi-ptas for envy-free pricing on line graphs. In *Proceedings of the 15th Annual European Symposium on Algorithms*, 2007.

[12] U. Feige and M. X. Goemans. 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*, pages 182–189, 1995.

[13] A. Grigoriev, J. van Loon, R. Sitters, and M. Uetz. How to Sell a Graph: Guideliness for Graph Retailers. Meteor Research Memorandum RM/06/001, Maastricht University, 2005.

[14] V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On Profit-Maximizing Envy-Free Pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1164 – 1173, 2005.

[15] J. Hartline and V. Koltun. Near-Optimal Pricing in Near-Linear Time . In *9th Workshop on Algorithms and Data Structures*, pages 422 – 431, 2005.

[16] N. Nisan, T. Rougharden, E. Tardos, and V. V. (Eds.). *Algorithmic Game Theory*. 2006. To appear.

A Proofs

Theorem 2 For the highway problem, there exists an $\Omega(\log n)$ gap between the coupon model and the (B-bounded) discount model.
Proof: Again, we set $B := 1$. We construct a similar example T_r. For $r = 0$, T_0 is the same as S_0. But for $r \geq 1$, T_r consists of two copies of T_{r-1}, but *without* an extra item between them. We also have 2^r customers, each of which has valuation 1 and corresponds to a set containing all items. In this case, the number of items is $n_r := 2^r$.

Using a randomized argument shown later in Theorem 4 from Section 4, $\text{OPT}_{coup} = \Omega(2^r r)$. We next consider the profit of T_r under the B-bounded discount model. First, observe that the optimal solution can be attained by integer prices. We first show that if the sum of the prices of all items is λ, then the profit under the discount model is at most $\lambda(2^{r+1} - 1)$. We show this by induction on r. For $r = 0$, the result is trivial. Consider T_{r+1} and suppose the sum of the prices of all items is λ. Recall that T_{r+1} consists of two copies of T_r. Let λ_1 and λ_2 be the sum of the prices of items in each copy. Then, we have $\lambda = \lambda_1 + \lambda_2$. We consider separate cases.

(1) $\lambda > 1$ In this case, the customers that are interested in all items will not buy. Hence, the maximum profit obtained is, by induction hypothesis, $\lambda_1(2^{r+1} - 1) + \lambda_2(2^{r+1} - 1) = \lambda(2^{r+1} - 1) \leq \lambda(2^{r+2} - 1)$.

(2) $\lambda \leq 1$ In this case, the net profit obtained from the 2^{r+1} customers who want all items is $\lambda \cdot 2^{r+1}$. By the induction hypothesis, the net profit by all other customers is at most $\lambda_1(2^{r+1} - 1) + \lambda_2(2^{r+1} - 1) = \lambda(2^{r+1} - 1)$. Hence, the total net profit is at most $\lambda(2^{r+1} + 2^{r+1} - 1) = \lambda(2^{r+2} - 1)$, as required.

We next show that for the instance T_r, the maximum net profit is achieved when the sum of the prices of all items is 1, and hence by the result we just show, the maximum net profit obtained from T_r is at most $2^{r+1} - 1$. Again, we show this by induction on r. First, the result is trivial for $r = 0$. Next, consider T_{r+1}. Observe that if we just set the price of one item to be 1 and the rest to 0, then the net profit obtained is $2^{r+2} - 1$. By the previous result, we know if the sum of the prices of all items is non-positive, then the net profit obtained is also non-positive. Hence, it suffices to consider the case when the sum of the prices of all items is greater than 1. In this case, the customers that are interested in all items will not buy. By induction hypothesis on the two copies of T_r, the maximum net profit obtained is at most $2(2^{r+1} - 1) < 2^{r+2} - 1$, as required.

Hence, $\text{OPT}_{coup}/\text{OPT}_B = \Omega(\log n_r)$. ■

Theorem 15 *There is a* 4*-approximation Algorithm under the coupon model for the highway problem on trees in the case when all all customers' valuations are precisely* 1 *and each customer's bundle corresponds to a simple path in the tree.*

Proof: Pick any node v_0 as the root. We assign a label to each node v to the tree, which is the sum of the prices of the items corresponding to the edges in the path from root v_0 to node v. Again, for each node, we assign a label from $\{0, 1\}$ uniformly at random. This correspond to a pricing. In particular, the price of an item corresponding to an edge (p_v, v), where p_v is the parent, is the label of v minus that of p_v.

We show that for each customer, the probability that we obtain a profit from that customer is $\frac{1}{4}$. Suppose the bundle of the customer corresponds to some path P. Observe that there is a unique node u that is an ancestor of all nodes on path P. We consider two cases.

(1) Path P is a sub-path of the path from the node u to some leaf. In this case, there is a node w that is a descendant of all nodes on path P. Hence, we can obtain profit from that customer *iff* the label of u is 0 and the label of w is 1. This happens with probability $\frac{1}{4}$.

(2) Path P consists of two such paths as in (1). In this case, there are two nodes w_1 and w_2 such that every node on path P is either on the path from u to w_1 or the path from u to w_2. It follows that we can obtain a profit from the customer *iff* the label of u is 0, and exactly one of the labels of w_1 and w_2 is 1 and the other 0. This also happens with probability $\frac{1}{4}$. ■

Theorem 12 There exists an $O(B)$ approximation algorithm for the vertex pricing problem under the B-bounded discount model.

Proof: Suppose the optimal net profit of for an instance of the vertex price problem under the B-bounded discount model is OPT. Let m be the number of customers whose valuations are at least 1. We show that there exists a solution consisting of non-negative prices that has net profit at least $\frac{OPT}{2B+1}$. Since there is a $O(1)$-approximation algorithm [2] for positive price model, it follows that there is a $O(B)$-approximation algorithm under the B-bounded discount model.

We consider the following two solutions with non-negative prices.

(1) Set the price of every item to be $\frac{1}{2}$. Then, it follows the total profit obtained is m, the total number of customers.

(2) Consider the solution under the B-bounded discount model that attains OPT. We modify the solution in the following way. If the price of an item is p, we change it to $\max\{0, p-B\}$. Observe that if a customer is buying a bundle, then after the modification, the customer is still going to buy the bundle. However, for every customer, we can potentially lose at most $2B$. Hence, the net profit for this non-negative prices is at least $OPT - 2mB$.

Observe that if $m \geq \frac{OPT}{2B+1}$, then we can just use the solution in (1). Otherwise, $OPT - 2mB \geq \frac{OPT}{2B+1}$, and so we can use the solution in (2). ■

Theorem 13 There exists a 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.

(2) There is no multi-edge in the graph.

Proof: The proof is very similar to that of Theorem 10. We still decompose the edges of the graphs into k groups, and after discarding one group the resulting graph has low treewidth, and so we can use dynamic program. However, since we are working under the B-bounded discount model, we could potentially lose a lot of profit from the customers that correspond to the discarded edges. We consider cases under either of the following extra assumptions.

(1) All customers have valuations at least 1. In this case, we set $k = \frac{2B+1}{\epsilon}$. Let m be the total number of customers. Then, we know $OPT \geq m$, because we can set each item to have price $\frac{1}{2}$. Suppose we pick any one of the k groups of edges to discard. Then, it follows the expected sum of valuations of the remaining customers is exactly $(1-\frac{1}{k})$ of the original sum. Moreover, the expected loss due to the customers discarded can be at most $\frac{2mB}{k} \leq \frac{2B}{k} \cdot OPT$. Hence, it follows that the expected total profit is at least $(1-\epsilon)OPT$.

(2) There is no multi-edge in the graph. This means that for any two items, there is exactly one customer that is interested in both items. We do a pre-processing step. If there is an item corresponding to a vertex such that all edges incident on it has non-positive valuations, then we set the price of that item to $+\infty$. Essentially, we have removed that vertex, together with the edges that incident on it. Let n be the number of remaining items and m be the number of remaining edges. It follows that if we set the price of each remaining item to $\frac{1}{2}$. Then, the profit obtained is at least $\frac{n}{2}$. Hence, $OPT \geq n/2$. Now the graph is planar and hence, $m \leq 3n$. It follows $m \leq 6\,OPT$. Hence, setting $k := \frac{B+3}{3\epsilon}$ and using the same argument as in the first case, the expected total profit is at least $(1-\epsilon)OPT$. ■

Theorem 14 There exists a fully polynomial time approximation scheme for the case that the desired subsets

of different customers form a hierarchy under the both the discount and the B-bounded discount models.
Proof: We can extend the algorithm presented in [2] for the positive price model. The correctness follows from the analysis in [2], and in the case of discount model, by additionally using the result in Lemma 1. ■

B QPTAS for the Highway Problem

We describe a sketch construction of a QPTAS for obtaining $(1+\epsilon)$-approximation the highway problem for general customer valuations, with additional assumptions. This approach works for both the coupon model and the bounded discount model. The method we use is based on the dynamic program by Elbassioni et. al. [11] for solving the highway problem with only non-negative prices. We would only highlight the modifications we make and the reader is advised to refer to [11] for the complete construction. We first state the result.

Theorem 7 There exists a QPTAS for the highway problem under the coupon model if there exists an optimal pricing that has only a bounded number C of negatively priced items, and the corresponding revenue obtained from positively priced items is at least twice that of the revenue from negatively priced items. The running time of the algorithm is $(\frac{m}{\epsilon})^{O(\frac{C}{\epsilon}(\log\frac{m}{\epsilon})^2)}$.

The main idea is as follows. Suppose the items are in the order $\{1,2,\ldots,n\}$, and there are m bidders. We pick the mid-point $\widehat{i} := \lfloor n/2 \rfloor$, and recursively solve for both the left $\{1,\ldots,\widehat{i}\}$ and right $\{\widehat{i}+1,\ldots,n\}$ halves. The difficulty is how to handle bidders whose bundles contain both items $\widehat{i}$ and $\widehat{i}+1$. The idea in [11] is to guess $1 = l_k < l_{k-1} < \cdots < l_1 \leq \widehat{i} < \widehat{i}+1 \leq r_1 < r_2 < \cdots < r_k = n$ and partial sums $f_1, f_2, \ldots, f_k$ and $g_1, g_2, \ldots, g_k$ such that the price function p satisfies the following.

1. For $1 \leq j \leq k$, $f_j = \sum_{i=r_j}^{\widehat{i}} p(i)$ and $g_j = \sum_{i=\widehat{i}+1}^{r_j} p(i)$.
2. For $1 \leq j < k$, $f_{j+1} \geq (1+\epsilon) f_j$ and $g_{j+1} \geq (1+\epsilon) g_j$.
3. For $1 \leq j < k$ $\sum_{i=r_{j+1}+1}^{\widehat{i}} p(i) < (1+\epsilon) f_j$ and $\sum_{i=\widehat{i}+1}^{r_{j+1}-1} p(i) < (1+\epsilon) g_j$.

The choice of r_j's, l_j's, f_j's and g_j's forms a configuration. Note any pricing q has a configuration consistent with it. Now, given a bidder interested in items $I = \{l, l+1, \ldots, r-1, r\}$, where $l \leq \widehat{i}$ and $r \geq \widehat{i}+1$, there exists s and t such that $l_{s+1} < l \leq l_s$ and $r_t \leq r < r_{t+1}$. Observe that the following property is used crucially in [11] for $(1+\epsilon)$-approximation.

$$f_s + g_t \leq q(I) \leq (1+\epsilon)(f_s + g_t) \tag{3}$$

Note that using standard rounding techniques, we can assume $k = O(\log P / \log(1+\epsilon))$, where $P := mn^2/\epsilon$. Hence, the number of configurations is at most $L = (nP)^{2k}$, and we have the following recursion for running time: $T(n) \leq poly(m) + 2L \cdot T(n/2)$. Hence, we get $T(n) \leq (2L)^{\log n} \cdot poly(m)$.

Observe that the inequality (3) holds only when non-negative prices are assigned. In particular, if $q(\{l,\ldots,\widehat{i}\}) > 0$ and $q(\{\widehat{i}+1,\ldots,r) < 0$, then $f_s + g_t$ is no longer a good approximation to $q(I)$ when f_s and g_t have similar magnitude. Although we cannot make guarantees about local rounding, we can ensure the rounding gives a good approximation globally if we make the following assumption.

There exists an optimal pricing q such that $\sum_{e\in E}\sum_{i\in e:q(i)>0} q(i) \geq 2\sum_{e\in E}\sum_{i\in e:q(i)<0}(-q(i))$.

Intuitively, this assumption states that the positive revenue of the optimal solution has to be at least twice that of the negative revenue.

Another technical issue that one encounters is that with negative prices, the f'_j and g_j's are no longer monotone. Hence, we would have to bound the number C of negatively priced items as well. Hence, at the beginning, we have to guess which items are negatively priced and their corresponding negative prices. This introduce an overall multiplicative factor of $(nP)^C$ in the running time.

Moreover, there would be C monotone segments and hence we would have to guess kC indices j's on both the left and right hand sides, where $k = O(\log P / \log(1+\epsilon))$. Hence, the number of configurations is $L = (nP)^{2Ck}$. As noted in [11], we can always assume $n \leq 2m$. Hence, the running time for the dynamic program with both positive and negative prices is at most $(nP)^C \cdot (2L)^{\log n} \cdot poly(m) = (\frac{m}{\epsilon})^{O(\frac{C}{\epsilon}(\log \frac{m}{\epsilon})^2)}$, which is quasi-polynomial in m.

C Algorithms for Discount without Loss on Highway Problem

As mentioned in Section 6, in the *discount without loss* model hte goal is to compete with an optimum which prices the items in a way that for any customer, the sum of the selling prices of items in his desired set is at least the sum of the marginal costs. More formally, we want the price vector $\mathbf{p}$ that maximizes $\text{Profit}_{rwl}(\mathbf{p}) = \sum_{e: w_e \geq \sum_{i \in e} p_i} \sum_{i \in e} p_i$, subject to the extra condition that for each customer e, $\sum_{i \in e} p_i \geq 0$. Let $\mathbf{p}^*_{rwl}$ be the price vector with the maximum profit and let $\text{OPT}_{rwl} = \text{Profit}_{rwl}(\mathbf{p}^*_{rwl})$.

This model has exactly the same DAG representation for the highway problem as before except there is no backward edges in the corresponding leveled graph. In particular, if there are n items, then for each $0 \leq i \leq n$, there is a node v_i that represents the partial sum $s_i := \sum_{l=1}^{i} p_l$, where p_l is the price of the lth item. Moreover, if a customer is interested in items $i+1$ through j (inclusive), then the customer is represented by the (directed) arc (v_i, v_j).

Below we present two theorems in this model for the highway problem.

Theorem 16 *For the discount without loss model on the highway problem in which all valuations are 1, there exists a randomized algorithm that gives $O(\log n)$ approximation, where n is the number of items.*

Proof: We consider the DAG representation of the instance, in which each node corresponds to a partial sum. In the discount without loss model, each customer corresponds to an arc (v_i, v_j) and hence the label of v_i cannot be larger than that of v_j. Therefore, all the customers impose a partial ordering on the nodes. We pick any topological order that satisfy the partial ordering. If no such topological order exists, then the instance is infeasible. We relabel the nodes such that v_0 appears first in the ordering and v_n the last.

We group the arcs corresponding to customers into $\log n$ groups such that each group is defined to be $E_k := \{(v_i, v_j) : 2^k \leq j - i < 2^{k+1}\}$, for $0 \leq k \leq \lfloor \log_2 n \rfloor$.

Let k^* be the group that contains the largest number of customers. Set $p := 1/2^{k^*+1}$. We set the labels for the nodes in the following way. We set the label of v_0 to be 0. In general, given that the label of v_i is set to l, with probability p the label of v_{i+1} is set to $l+1$, and with probability $1-p$ it is set to l.

We next show that each customer (v_i, v_j) in group E_{k^*} contributes to a profit of 1 with probability at least $\frac{1}{8}$. Let $r := j - i$ and hence $2^{k^*} \leq r < 2^{k^*+1}$. Note that we receive profit from the customer *iff* the label increases exactly once going from v_i to v_j. This happens with probability $rp(1-p)^{r-1} \geq \frac{1}{2}(1-p)^{1/p} \geq \frac{1}{8}$, since $p \leq \frac{1}{2}$.

Hence, it follows that the expected profit returned is an $O(\log n)$ approximation. ■

We can solve a restricted version of the problem exactly. Recall that in the DAG representation, each node has a label that corresponds to a partial sum.

Theorem 17 *For the discount without loss model on the highway problem, the restricted version in which each node in the DAG representation must have a label from* $\{0, 1\}$ *can be solved exactly by linear programming.*

Proof: Recall that we consider the highway problem in which every customer has valuation 1. Moreover, we add an extra condition such that when items are traversed from one end to another in the linear order, if we ignore the items that are priced 0, then the the prices of the remaining items must alternate between 1 and -1. Hence, when we see an item priced 1, we must first see an item priced -1, before we see an item priced 1 again.

Observe that under the alternate-$\{-1, 1\}$ condition, each s_i can only take one of two values. Hence, we have the following LP relaxation for the restricted version of the highway problem under the discount without loss model.

$$\begin{array}{lll} \max & \sum_{(i,j)\in E}(s_j - s_i) & \\ \text{s.t.} & s_j - s_i \geq 0 & \forall (i,j) \in E \\ & 0 \leq s_i \leq 1 & \forall 0 \leq i \leq n \end{array}$$

It is not difficult to show that this LP relaxation actually gives the optimal integer solution as well. We give the following rounding algorithm. Consider picking a random s uniformly from $(0, 1)$. If $s_i < s$, then round it down to 0; otherwise, round it up to 1. It follows that for every customer $(i, j) \in E$, the probability that we get one unit of profit from it is exactly $s_j - s_i$. Hence, it follows that the expected total profit we obtain is $\sum_{(i,j)\in E}(s_j - s_i)$, i.e., the expected value of an integer solution obtained from rounding is equal to the value of the optimal fractional solution. Since every integer solution has value at most that of the optimal fractional solution, it follows that every integer solution obtained from rounding is optimal. ■