

A Nonparametric Approach to Multiproduct Pricing

Paat Rusmevichientong

School of Operations Research and Industrial Engineering, 221 Rhodes Hall, Cornell University, Ithaca, New York 14853,
paatrus@cornell.edu

Benjamin Van Roy, Peter W. Glynn

Department of Management Science and Engineering, Stanford University, Stanford, California 94305
{bvr@stanford.edu, glynn@stanford.edu}

Developed by General Motors (GM), the Auto Choice Advisor website (<http://www.autochoiceadvisor.com>) recommends vehicles to consumers based on their requirements and budget constraints. Through the website, GM has access to large quantities of data that reflect consumer preferences. Motivated by the availability of such data, we formulate a nonparametric approach to multiproduct pricing.

We consider a class of models of consumer purchasing behavior, each of which relates observed data on a consumer's requirements and budget constraint to subsequent purchasing tendencies. To price products, we aim at optimizing prices with respect to a sample of consumer data. We offer a bound on the sample size required for the resulting prices to be near-optimal with respect to the true distribution of consumers. The bound exhibits a dependence of $O(n \log n)$ on the number n of products being priced, showing that—in terms of sample complexity—the approach is scalable to large numbers of products.

With regards to computational complexity, we establish that computing optimal prices with respect to a sample of consumer data is NP-complete in the strong sense. However, when prices are constrained by a price ladder—an ordering of prices defined prior to price determination—the problem becomes one of maximizing a supermodular function with real-valued variables. It is not yet known whether this problem is NP-hard. We provide a heuristic for our price-ladder-constrained problem, together with encouraging computational results.

Finally, we apply our approach to a data set from the Auto Choice Advisor website. Our analysis provides insights into the current pricing policy at GM and suggests enhancements that may lead to a more effective pricing strategy.

Subject classifications: pricing: multiproduct pricing; marketing: buyer behavior, choice models, estimation/statistical techniques; industries: automotive.

Area of review: Manufacturing, Service, and Supply Chain Operations.

History: Received February 2003; revision received April 2004; accepted December 2004.

1. Introduction

Availability of data on consumer preferences, together with sophisticated analytical tools, enables increases in profit through optimization of prices. In addition to airlines and hotels, which have been traditional users of revenue management (McGill and Van Ryzin 1999, Weatherford and Bodily 1992), many companies in retail and manufacturing have started to employ advanced pricing strategies to increase their bottom lines.

By identifying product features for which consumers are willing to pay a premium, Ford has developed a pricing strategy that encourages consumers to purchase more expensive vehicles, resulting in a significant increase in profit (Coy 2000). Dell uses a sophisticated pricing strategy, quoting different prices to different market segments for the same product. This strategy enables Dell to increase its market share and profitability (McWilliams 2001). Using pricing optimization software, ShopKo Stores identifies an optimal markdown strategy for selling its line of nylon track pants, enabling it to clear out its inventory by the end of the selling season while maintaining a high profit

margin. The company has successfully applied the same technology to other products as well (Merrick 2001).

In this paper, we develop a new approach to price optimization that leverages data collected through operation of a website that hosts a “product recommendation engine.” Our work is motivated by the availability of data from the Auto Choice Advisor (ACA) website (<http://www.autochoiceadvisor.com>) and an aim towards optimization of General Motors (GM) vehicle prices. Developed by GM, the ACA website recommends vehicles to consumers based on their requirements and budget constraints. Through the website, GM has access to large quantities of data that reflect consumer preferences.

Table 1 provides an example of a record generated during a visit to the ACA website. The system records a budget and a list of recommended vehicles for each visitor. The visitor specifies her budget explicitly. The list of vehicles is generated based on requirements elicited by the website. The vehicles in the list are ranked based on how well each matches the visitor's requirements, with the Honda Accord Value Sedan being the best match in this example. Note that

Table 1. An example of consumer preference data collected by the ACA website.

Budget	Recommended vehicles	Rank
\$18,000	Honda Accord Value Sedan	1st
	Saturn L100 (GM)	2nd
	Dodge Stratus SXT Sedan	3rd
	Chevrolet Malibu Sedan (GM)	4th
	Mitsubishi Lancer LS	5th

the website recommends vehicles with prices exceeding the budget, provided that the vehicles meet other requirements.

The large quantities of data being generated by the ACA website may offer an improved understanding of consumer preferences, and thereby create an opportunity for GM to increase profit by adjusting vehicle prices. In this paper, we explore an approach to tapping this value.

1.1. Problem Formulation

We consider a market with N competing products. Our objective is to set prices of $n \leq N$ of the products, given that the other $N - n$ product prices remain fixed. We index the former set of products by $A = \{1, \dots, n\}$ and the latter by $\bar{A} = \{n + 1, \dots, N\}$. The set A represents our products, whereas \bar{A} represents products offered by competitors. We assume that the prices of competitors' products are fixed at d_{n+1}, \dots, d_N .

Information about a consumer is represented by a profile $\omega = (b, Z)$, where $b \in \mathfrak{R}_+$ denotes the consumer's budget, and $Z = (z_1, \dots, z_\ell)$ denotes the ordered list of recommended products. We denote by Ω the set of all possible consumer profiles and by \mathbb{F} the Borel σ -algebra generated by this set. Let \mathcal{P} be the probability distribution such that if a consumer is selected uniformly at random from the set of all consumers in the world, the probability that her profile ω is in a set $X \in \mathbb{F}$ is $\mathcal{P}(X)$. We make the following assumption on the distribution \mathcal{P} .

ASSUMPTION 1.1. *There exists $B < \infty$ such that $\mathcal{P}\{\omega \in \Omega: b \leq B\} = 1$.*

We denote the vector of our product prices by $p \in \mathfrak{R}_+^n$. We will model each consumer's purchasing decision as depending only on her profile ω and the product prices p . This purchasing behavior is captured by a choice function $C: \mathfrak{R}_+^n \times \Omega \rightarrow \{0\} \cup A \cup \bar{A}$, where 0 denotes a decision to purchase no product. In particular, given prices p , a consumer with profile ω is assumed to purchase product $C(p, \omega)$. For convenience, let $p_0 = 0$, representing the price of no purchase. Thus, the price paid by the consumer can be written as $p_{C(p, \omega)}$ if $C(p, \omega) \in \{0\} \cup A$, or $d_{C(p, \omega)}$ if $C(p, \omega) \in \bar{A}$.

The revenue we receive from a consumer with profile ω is denoted by

$$R(p, \omega) = \begin{cases} p_{C(p, \omega)} & \text{if } C(p, \omega) \in A, \\ 0 & \text{otherwise.} \end{cases}$$

Our objective is to set prices of our products to maximize expected revenue

$$p^* \in \arg \max_{p \in \mathfrak{R}_+^n} E[R(p, \omega)], \quad (1.1)$$

where the expectation is taken with respect to \mathcal{P} . We will assume knowledge of the revenue function R , but not the distribution \mathcal{P} . As such, we cannot compute the expectation. Instead, we will make use of a set of sample customer profiles $\omega_1, \dots, \omega_M$, and aim at setting prices to maximize a sample estimate of the expectation:

$$\hat{p}^M \in \arg \max_{p \in \mathfrak{R}_+^n} \frac{1}{M} \sum_{j=1}^M R(p, \omega_j). \quad (1.2)$$

In the following sections, we analyze the sample and computational complexity of the proposed approach. In particular, we place a bound on the number M of samples required, under certain assumptions, for \hat{p}^M to nearly optimize $E[R(p, \omega)]$ and the computation required to accomplish this.

We impose three requirements on the choice function. First, we assume that if a consumer makes a purchase, the product is from her recommended list and priced under her budget.

ASSUMPTION 1.2. *For all $p \in \mathfrak{R}_+^n$ and $\omega = (b, Z) \in \Omega$, $C(p, \omega) \in \{0\} \cup Z$. Moreover, $p_{C(p, \omega)} \leq b$ if $C(p, \omega) \in \{0\} \cup A$ and $d_{C(p, \omega)} \leq b$ if $C(p, \omega) \in \bar{A}$.*

It follows from Assumptions 1.1 and 1.2 that to find the optimal prices for our products, it suffices to consider only price vectors in $[0, B]^n$.

We make the following assumption on the choice function to ensure that $R(\omega, \cdot)$ is upper semicontinuous for each $\omega \in \Omega$, and therefore, that the maxima in (1.1) and (1.2) are attained in $[0, B]^n$.

ASSUMPTION 1.3. *For all $\omega \in \Omega$, the set $\bigcup_{i \in A} \{p \in \mathfrak{R}_+^n: C(p, \omega) = i, p_i \geq \alpha\}$ is closed for all $\alpha \in \mathfrak{R}$.*

In words, the assumption is that the set of price vectors resulting in a selection from among our products forms a closed subset. Note that $\bigcup_{i \in A} \{p \in \mathfrak{R}_+^n: C(p, \omega) = i, p_i \geq \alpha\}$ is the set of prices that will lead to a selection from our products and yield a revenue of at least α . In general, each individual set $\{p \in \mathfrak{R}_+^n: C(p, \omega) = i, p_i \geq \alpha\}$ will not be closed. However, we only require that the union of these sets be closed.

For any two vectors $x, y \in \mathfrak{R}^n$, we write $x \sim y$ if the elements of the two vectors share a common ordering; i.e., if for all i, j , $x_i \geq x_j$ implies $y_i \geq y_j$. The following assumption allows for increasing prices of our unchosen products and, within limits, the chosen product, without changing the consumer's choice. Note that $\mathbf{1}(\cdot)$ denotes an indicator function that takes the value 1 if the condition in the parentheses is true, and 0 otherwise.

ASSUMPTION 1.4. For all $\omega = (b, Z) \in \Omega$ and $p, p' \in \mathfrak{R}_+^n$ such that $p \sim p'$ and $p \leq p'$, if $\mathbf{1}(C(p, \omega) \in A)p'_{C(p, \omega)} \leq b$, then $C(p', \omega) = C(p, \omega)$.

To interpret this assumption, first consider increases in prices of unchosen products. These increases should not make them any more desirable, and therefore they should remain unchosen. For the chosen product, if $R(\omega, p) = b$, then the assumption does not allow for a price increase. However, when $R(\omega, p) < b$, there is some slack and the price of the chosen product can be increased without changing its chosen status.

The following examples demonstrate that choice functions of practical interest satisfy Assumptions 1.2–1.4.

EXAMPLE 1.1 (THE RANK PRICING MODEL). For any $p \in \mathfrak{R}_+^n$, $\omega = (b, Z) \in \Omega$, let $C^{\text{RANK}}(p, \omega)$ denote the highest-ranked product in Z that meets the consumer's budget. Recall that the list $Z = (z_1, z_2, \dots, z_\ell)$ is sorted by the degree to which each product matches the consumer's requirements, with z_1 representing the highest-ranked product.

The RANK PRICING Model leverages the special ranked-list structure of our data set, and it will be the primary model that we will use to test the performance of our algorithms in §5.

Clearly, C^{RANK} satisfies Assumption 1.2. Moreover, as product prices increase, the set of affordable products gets smaller. So, if the original selection by the consumer remains affordable, the same product will continue to have the highest rank because the rank is independent of price. Thus, this choice function also satisfies Assumption 1.4.

It remains to show that C^{RANK} satisfies Assumption 1.3. Fix $\omega = (b, Z) \in \Omega$ and $\alpha \in \mathfrak{R}$. We wish to show that $\bigcup_{i \in A} \{p \in \mathfrak{R}_+^n : C^{\text{RANK}}(p, \omega) = i, p_i \geq \alpha\}$ is a closed subset of \mathfrak{R}_+^n . By relabeling if necessary, we may assume that $Z = (i_1, i_2, \dots, i_k)$. In addition, we can assume that $\alpha \leq b$; otherwise, the result is trivially true.

We will now show that it suffices to consider the case when Z contains only our products, i.e., $i_m \in A$ for all $m = 1, 2, \dots, k$. Because competitors' prices remain fixed, we can first eliminate all competitors' products with prices higher than b . Then, let l denote the smallest integer such that i_l represents a competitors' product with price under b . If no such l exists, we are done. Otherwise, we can remove all products i_{l+1}, \dots, i_k because the consumer will never purchase them. We can also effectively remove i_l because it will not impact the set of prices leading to a selection of our products. Thus, in every case, we can eliminate all competitors' products from our consideration. Note that when $l = 1$, the set Z is effectively an empty set.

Under the RANK PRICING Model, the consumer will choose $i_m \in Z$ if and only if its price is under the budget b and the prices of the first $m - 1$ products exceed b , i.e., for all $m = 1, \dots, k$,

$$\begin{aligned} \{p \in \mathfrak{R}_+^n : C^{\text{RANK}}(p, \omega) = i_m\} \\ = \{p \in \mathfrak{R}_+^n : p_{i_m} \leq b < \min\{p_{i_1}, \dots, p_{i_{m-1}}\}\}. \end{aligned}$$

Let $D_1 = \{p \in \mathfrak{R}_+^n : \alpha \leq p_{i_1} \leq b\}$, and for $m = 2, \dots, k$, let D_m be defined by

$$D_m = \{p \in \mathfrak{R}_+^n : \alpha \leq p_{i_m} \leq b < \min\{p_{i_1}, \dots, p_{i_{m-1}}\}\}.$$

Note that the closure of D_m , denoted by \bar{D}_m , is given by

$$\bar{D}_m = \{p \in \mathfrak{R}_+^n : \alpha \leq p_{i_m} \leq b \leq \min\{p_{i_1}, \dots, p_{i_{m-1}}\}\}.$$

It follows from the definition of D_m that

$$\bigcup_{m=1}^k D_m = \bigcup_{i \in A} \{p \in \mathfrak{R}_+^n : C^{\text{RANK}}(p, \omega) = i, p_i \geq \alpha\}.$$

Thus, it suffices to show that $\bigcup_{m=1}^k D_m$ is a closed set. To prove this result, we will show by induction that

$$\bigcup_{m=1}^k D_m = \bigcup_{m=1}^k \bar{D}_m.$$

Clearly, this is true for $k = 1$ because D_1 is closed by definition. Suppose that for $\ell < k$,

$$\bigcup_{m=1}^{\ell} D_m = \bigcup_{m=1}^{\ell} \bar{D}_m.$$

By definition, we have

$$\begin{aligned} \bar{D}_{\ell+1} \setminus D_{\ell+1} \\ = \{p \in \mathfrak{R}_+^n : \alpha \leq p_{i_{\ell+1}} \leq b = \min\{p_{i_1}, \dots, p_{i_\ell}\}\} \\ \subseteq \{p \in \mathfrak{R}_+^n : b = \min\{p_{i_1}, \dots, p_{i_\ell}\}\} \\ = \bigcup_{m=1}^{\ell} \{p \in \mathfrak{R}_+^n : p_{i_m} = b \leq \min\{p_{i_s} : s \neq m, 1 \leq s \leq \ell\}\} \\ \subseteq \bigcup_{m=1}^{\ell} \{p \in \mathfrak{R}_+^n : p_{i_m} = b \leq \min\{p_{i_1}, p_{i_2}, \dots, p_{i_{m-1}}\}\} \\ \subseteq \bigcup_{m=1}^{\ell} \bar{D}_m, \end{aligned}$$

where the last inequality follows from the definition of \bar{D}_m . Thus, it follows that

$$\bar{D}_{\ell+1} \cup \left(\bigcup_{m=1}^{\ell} \bar{D}_m \right) = D_{\ell+1} \cup \left(\bigcup_{m=1}^{\ell} \bar{D}_m \right) = \bigcup_{m=1}^{\ell+1} D_m,$$

where the last equality follows from the inductive assumption. This completes the proof. \square

EXAMPLE 1.2 (THE MIN PRICING MODEL). For any $p \in \mathfrak{R}_+^n$, $\omega = (b, Z) \in \Omega$, let $C(p, \omega)$ denote the cheapest product in Z that meets the budget constraint b . The MIN PRICING Model clearly satisfies Assumption 1.2.

Unfortunately, the MIN PRICING Model does not satisfy Assumption 1.4, because the consumer's selection may

change from our products to the competitors' as the prices of our products increase.

Nonetheless, by removing the competitors' products, we can convert the MIN PRICING Model to an equivalent one that satisfies Assumption 1.4. Consider any consumer with a profile $\omega = (b, Z)$. Let $u(\omega)$ denote the price of the cheapest competitors' product in Z . If Z contains no competitors' product, we set $u(\omega)$ to b , i.e.,

$$u(\omega) = \begin{cases} \min\{d_i: i \in Z \cap \bar{A}\} & \text{if } Z \cap \bar{A} \neq \emptyset, \\ b & \text{otherwise.} \end{cases}$$

Because the competitors' prices remained fixed, the consumer will purchase our products if and only if our price beats $u(\omega)$. Thus, for the purpose of optimizing the total revenue from each consumer, we can remove all competitors' products from the profile ω , provided that we replace the budget b by $\min\{b, u(\omega)\}$.

Once the competitors' products are removed, the MIN PRICING Model satisfies Assumption 1.4 because for any two price vectors that have the same ordering, they both have the same cheapest product.

Finally, using the same argument as in Example 1.1, we can also show that the MIN PRICING Model satisfies Assumption 1.3.

1.2. Literature Review

Our research contributes to the literature on multiproduct pricing by identifying a new nonparametric formulation that is motivated by consumer preference data collected from the ACA website. To put our results in perspective, let us briefly review related research in this area.

The literature on multiproduct pricing focuses on understanding interactions and substitutions among products, and finding the prices that maximize profit in light of these interdependencies and consumer preferences. Shocker and Srinivasan (1979) and Sen (1982) provide reviews of research in this area.

A dividing line separates research on modeling of demand processes into two categories. One approach involves postulating a demand function, which is taken to be a function of price and product attributes. Such formulations facilitate the study of how prices and attributes influence sales. Most results revolve around the case of a single differentiating attribute, such as a metric of quality. Examples include the work of Smith (1986) and Oren et al. (1984).

In each of the cases we have mentioned, a specific parametric form is imposed to model a demand process. There is also research aimed at developing a general theory (Oren et al. 1987), but to apply a theory of multiproduct pricing in practice, a specific functional form is again required. Selecting the appropriate parametric form—one that models salient features of the true demand function, yet is simple enough to allow efficient computation—remains a challenge.

Our nonparametric formulation relates closely to a segment of the literature that focuses on modeling consumer purchasing behavior and on finding the optimal prices give such models. In this line of research, each consumer associates a utility with each product, and she will choose the product that maximizes her utility subject to her budget constraint (Mirman and Sibley 1980, Spence 1980). Given the utility values, pricing problems are then formulated as mixed-integer programming problems (Dobson and Kalish 1988, 1993; Green and Krieger 1985; Hanson and Martin 1990; McBride and Zufryden 1988). The integer programming formulation is quite flexible, allowing for a wide range of constraints. However, the integer programming problems are NP-complete.

Our research revolves around a new kind of assumptions on consumer preferences, motivated by data from the ACA website. We identify a class of relevant problems by imposing a constraint on the ordering of prices. For each model of consumer behavior, we then develop an efficient algorithm that computes approximately revenue-maximizing prices based on the data.

In addition to providing a new formulation of multiproduct pricing, our work is the first, to the best of our knowledge, that focuses on consumer preference data collected from e-commerce websites. Approaches to multiproduct pricing generally assume that data on consumer preferences are gathered through market research or conjoint analysis (Dobson and Kalish 1988, 1993; McBride and Zufryden 1988). These methods of acquiring data are often expensive and time consuming, limiting the size of the data set.

Through e-commerce websites, we can obtain preference information on a large number of consumers. Thus, online consumer preference data provides us with new opportunities to understand consumer behavior and to customize our product prices to meet consumer needs. Our research identifies opportunities and challenges in using such data, and suggests changes in the design of e-commerce websites that may facilitate a more effective analysis of the data.

1.3. Contributions and Organization

Motivated by the availability of data on consumer preferences from e-commerce websites, this paper presents a nonparametric formulation of multiproduct pricing. We offer an original algorithm for solving a large class of problems. We prove new theoretical results about the algorithm and formulation. Finally, we present a case study using a real data set generated from an e-commerce website of a large automobile manufacturer. In this section, we discuss the organization of the paper, and in doing so, we lay out the algorithms, theory, and experimental results that make up our contributions.

Our approach optimizes product prices based on data. However, the data represent only a sample of all consumers. To ensure that we have enough data for meaningful optimization, we study sample size requirements in §2. It turns out that the required number of samples exhibits

a dependence of $O(n \log n)$ on the number of products n . This suggests that our approach is scalable to problems involving large numbers of products.

Section 3 assesses the computational complexity of our optimization problem, showing that it is NP-complete in the strong sense. Motivated by this result, we identify a restricted class of problems by imposing a price ladder—an ordering of prices—as a constraint in our optimization process. The relevance of such a constrained model is motivated by examples from the automotive industry. When this constraint is imposed, our optimization problem becomes a special case of supermodular function maximization. It is not known whether this class of problems is NP-hard. Section 4 presents a heuristic algorithm for our price-ladder-constrained optimization problem.

In §5, we apply our algorithm to a real data set from the ACA website, focusing on the RANK PRICING Model. We provide additional background on the ACA website, and contrast our approach with the current pricing methodology employed at GM. Our analysis provides insights into the current pricing practice, and suggests improvements that may lead to a more effective pricing strategy. Finally, §6 presents our conclusions and discusses potential future research.

2. Sample Complexity

Our formulation computes prices based on data. The data represent only a sample of all consumers. For the results to be meaningful, we need a sufficiently large sample size. However, for our approach to be practical, the required sample size must scale well with the number of products. We will show that the number of required samples is $O(n \log n)$, where n denotes the number of products. Moreover, this sample complexity bound is independent of the underlying distribution from which the data is generated.

To prove this result, we will first establish an upper bound on the covering number of the set of functions mapping consumer profiles to revenue, each element of the set being identified by a price vector.

2.1. Covering Numbers of Revenue Functions

Consider a normed space $(X, \|\cdot\|)$. An ϵ -cover of a set $S \subseteq X$ is a set $T \subseteq X$ such that for each $x \in S$ there exists a $y \in T$ with $\|x - y\| < \epsilon$. The ϵ -covering number $N(\epsilon, S, \|\cdot\|)$ is the cardinality of the smallest ϵ -cover of S .

Recall that Ω denotes the set of all possible consumer profiles and \mathbb{F} denotes the Borel σ -algebra generated by this set. For any probability measure μ on the measurable space (Ω, \mathbb{F}) , we define an L_1 norm $\|\cdot\|_{1, \mu}$ on the space of \mathbb{F} -measurable functions as follows:

$$\|f\|_{1, \mu} = \int |f(y)| \mu(dy).$$

Let $\mathcal{F} = \{R(p, \cdot) \mid p \in [0, B]^n\}$ denote a collection of revenue functions indexed by $p \in [0, B]^n$. Recall that for

any $p \in [0, B]^n$ and $\omega \in \Omega$, $R(p, \omega)$ denotes that revenue that is generated from a consumer with a profile ω under a price vector p . The following theorem provides an upper bound on the covering number of \mathcal{F} .

THEOREM 2.1. *Under Assumptions 1.2 and 1.4, for any probability measure μ on (Ω, \mathbb{F}) , and $\epsilon \in (0, B)$,*

$$N(\epsilon, \mathcal{F}, \|\cdot\|_{1, \mu}) \leq \left(\frac{2B(n+2)}{\epsilon} \right)^n.$$

PROOF. We begin by defining a set of $U \subset [0, B]$ for which $T = \{R(p, \cdot) \mid p \in U^n\}$ constitutes an ϵ -cover of \mathcal{F} . Let

$$U_1 = \left\{ k \frac{\epsilon}{2} \mid k = 0, 1, \dots, \left\lfloor \frac{2B}{\epsilon} \right\rfloor \right\} \cup \{B\}.$$

Note that this is simply a uniformly spaced grid. Define a distribution function $F: \mathfrak{R}_+ \rightarrow [0, 1]$ as follows:

$$F(\beta) = \mu(\{\omega = (b, Z) \in \Omega \mid b \geq \beta\}) \quad \forall \beta \geq 0,$$

and let $U_2 \subset [0, B]$ be defined by

$$U_2 = \left\{ F^{-1} \left(k \frac{\epsilon}{2Bn} \right) \mid k = 1, \dots, \left\lfloor \frac{2Bn}{\epsilon} \right\rfloor \right\},$$

where $F^{-1}(\alpha) = \max\{\beta \mid F(\beta) \geq \alpha\}$ represents an inverse of F . Note that F and F^{-1} are both nonincreasing functions. The set U_2 can be thought of as a grid that becomes uniformly spaced when mapped to cumulative probabilities.

Let $U = U_1 \cup U_2$. Note that

$$|U| \leq \left\lfloor \frac{2B}{\epsilon} \right\rfloor + 2 + \left\lfloor \frac{2Bn}{\epsilon} \right\rfloor \leq \frac{2B(n+2)}{\epsilon},$$

where the last inequality follows from the fact that $\epsilon < B$. The resulting set T has a cardinality

$$|T| \leq \left(\frac{2B(n+2)}{\epsilon} \right)^n.$$

We now set out to establish that T constitutes an ϵ -cover of \mathcal{F} . Consider an arbitrary vector $p \in [0, B]^n$. Let $\hat{p}_i = \min\{u \in U \mid u \geq p_i\}$ for $i = 1, \dots, n$. Note that \hat{p}_i is well defined because $B \in U$. We will show that $\|R(p, \cdot) - R(\hat{p}, \cdot)\|_{1, \mu} < \epsilon$, and the fact that T is an ϵ -cover follows from this. Note that by our construction of U , for all $i \in A$,

$$|p_i - \hat{p}_i| < \frac{\epsilon}{2} \quad \text{and} \quad |F(p_i) - F(\hat{p}_i)| \leq \frac{\epsilon}{2Bn}.$$

For any $i = \{0\} \cup A \cup \bar{A}$, let S_i denote the set of customer profiles that will lead to a selection of product i under the price vector p , i.e.,

$$S_i = \{\omega \in \Omega: C(p, \omega) = i\}.$$

Note that the sets S_i form a partition of Ω .

Because $p \leq \hat{p}$ and both p and \hat{p} have the same ordering, it follows from Assumption 1.4 that for any $i \in \{0\} \cup A$ and $\omega \in S_i$,

$$0 = R(p, \omega) = R(\hat{p}, \omega),$$

which implies that

$$\|R(p, \cdot) - R(\hat{p}, \cdot)\|_{1, \mu} = \sum_{i \in A} \int_{S_i} |R(p, \omega) - R(\hat{p}, \omega)| \mu(d\omega).$$

For $i \in A$, let us partition the set S_i into two parts: S_i^1 and S_i^2 , where

$$S_i^1 = \{\omega \in S_i: \hat{p}_{C(p, \omega)} \leq b\} \quad \text{and} \quad S_i^2 = \{\omega \in S_i: \hat{p}_{C(p, \omega)} > b\}.$$

It follows from Assumption 1.4 once again that $C(p, \omega) = C(\hat{p}, \omega)$ for all $\omega \in S_i^1$. Thus, for any $\omega \in S_i^1$,

$$|R(p, \omega) - R(\hat{p}, \omega)| = |p_i - \hat{p}_i| < \frac{\epsilon}{2},$$

which implies that

$$\begin{aligned} \sum_{i \in A} \int_{S_i^1} |R(p, \omega) - R(\hat{p}, \omega)| \mu(d\omega) \\ < \frac{\epsilon}{2} \sum_{i \in A} \mu(S_i^1) = \frac{\epsilon}{2} \mu(\cup_{i \in A} S_i^1) \leq \frac{\epsilon}{2}, \end{aligned}$$

where the equality follows from the fact that S_i is disjoint.

Consider any $\omega \in S_i^2$. By definition, we have $p_i \leq b < \hat{p}_i$, which implies that

$$\mu(S_i^2) \leq \mu(\{\omega: p_i \leq b < \hat{p}_i\}) = F(p_i) - F(\hat{p}_i) \leq \frac{\epsilon}{2Bn}.$$

Thus,

$$\sum_{i \in A} \int_{S_i^2} |R(p, \omega) - R(\hat{p}, \omega)| \mu(d\omega) \leq B \sum_{i \in A} \mu(S_i^2) \leq \frac{\epsilon}{2},$$

where the first inequality follows from the fact that $|R(p, \omega) - R(\hat{p}, \omega)| \leq B$, and the last inequality follows from the fact that $A = \{1, 2, \dots, n\}$. Putting everything together, we have

$$\begin{aligned} \|R(p, \cdot) - R(\hat{p}, \cdot)\|_{1, \mu} \\ = \sum_{i \in A} \int_{S_i} |R(p, \omega) - R(\hat{p}, \omega)| \mu(d\omega) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon. \quad \square \end{aligned}$$

2.2. Sample Complexity for Pricing Problems

Recall that we want to find

$$p^* \in \arg \max_{p \in [0, B]^n} E_{\mathcal{P}}[R(p, \omega)],$$

where the expectation is taken with respect to the distribution \mathcal{P} . Because \mathcal{P} is not known explicitly, we

draw independent samples $\omega_1, \dots, \omega_M$ according to \mathcal{P} , and instead compute

$$\hat{p}^M \in \arg \max_{p \in [0, B]^n} \frac{1}{M} \sum_{j=1}^M R(p, \omega_j).$$

We hope that the expected revenue under \hat{p}^M ,

$$E_{\mathcal{P}}[R(\hat{p}^M, \omega)],$$

would be close to the optimal expected revenue,

$$E_{\mathcal{P}}[R(p^*, \omega)],$$

for a sufficiently large sample size M .

The following theorem relates covering numbers to rates of uniform convergence. This result is adapted from Corollary 1 in Haussler (1992).

THEOREM 2.2. *Consider the probability space $(\Omega, \mathbb{F}, \mathcal{P})$. Let $\omega_1, \dots, \omega_M \in \Omega$ be drawn independently according to \mathcal{P} . Let \mathcal{F} be any set of \mathbb{F} -measurable functions from Ω to $[0, B]$. Then, for any $\epsilon \in (0, B)$,*

$$\begin{aligned} \Pr \left\{ \sup_{f \in \mathcal{F}} \left| E_{\mathcal{P}}(f) - \frac{1}{M} \sum_{j=1}^M f(\omega_j) \right| > \epsilon \right\} \\ \leq 4 \sup_{\mu} \mathcal{N} \left(\frac{\epsilon}{16}, \mathcal{F}, \|\cdot\|_{1, \mu} \right) e^{-\epsilon^2 M / 128 B^2}, \end{aligned}$$

where $E_{\mathcal{P}}(f) = \int f(y) \mathcal{P}(dy)$ denotes the expectation of f under the probability measure \mathcal{P} , and μ denotes an arbitrary probability measure on the measurable space (Ω, \mathbb{F}) .

The following theorem is an immediate corollary of Theorems 2.1 and 2.2. The proof of this result is given in Appendix A.

THEOREM 2.3. *Under Assumptions 1.1–1.4, for any $0 < \delta < 1$, $0 < \epsilon < B$, if*

$$M \geq \frac{128 B^2}{\epsilon^2} \left(\ln \frac{4}{\delta} + n \ln \frac{2B}{\epsilon} + n \ln(n+2) \right),$$

then

$$\Pr \{ |E_{\mathcal{P}}[R(p^*, \omega)] - E_{\mathcal{P}}[R(\hat{p}^M, \omega)]| > 2\epsilon \} \leq \delta$$

and

$$\Pr \left\{ \left| E_{\mathcal{P}}[R(p^*, \omega)] - \frac{1}{M} \sum_{j=1}^M R(\hat{p}^M, \omega_j) \right| > \epsilon \right\} \leq \delta.$$

The above result shows that the number of required samples increases with the number of products at the rate of $O(n \ln n)$. For a sufficiently large sample size M , the above theorem also shows that the expected revenue under \hat{p}^M is close to the optimal expected revenue. This result implies that our pricing policy \hat{p}^M , which is computed from the samples, when applied to the general population would perform well relative to the optimal policy p^* .

In addition to showing that the price vector \hat{p}^M performs well relative to the optimal vector p^* , it follows from Theorem 2.3 that the optimal expected revenue is close to the sample average revenue, $(1/M) \sum_{j=1}^M R(\hat{p}^M, \omega_j)$, for a sufficiently large sample size M . Because we can compute the sample average revenue, this result provides us with information on the magnitude of the optimal expected revenue.

3. Computational Complexity

Let \mathcal{X}_+ denote the set of nonnegative integers. The computational problem we would like to solve is captured formally by:

GENERAL PRICING

Instance:

- Finite sets of products $A = \{1, \dots, n\}$ and $\bar{A} = \{n+1, \dots, N\}$.
- Prices of competitor products d_{n+1}, \dots, d_N .
- A maximum budget $B \in \mathcal{X}_+$.
- A revenue function R that is based on a choice function that satisfies Assumptions 1.2–1.4 and that can be evaluated in polynomial time.
- A set $\mathcal{C} = \{\omega_1, \dots, \omega_M\}$, where $\omega_j = (b_j, Z_j)$, $b_j \in \mathcal{X}_+$, $b_j \leq B$, and Z_j is an ordered list of products in $A \cup \bar{A}$.

Problem: Find a price vector

$$\hat{p}^M \in \arg \max_{p \in [0, B]^n} \sum_{j=1}^M R(p, \omega_j).$$

In this section, we establish that GENERAL PRICING is NP-complete in the strong sense. Therefore, unless $P = NP$, there is no pseudopolynomial-time algorithm—that is, an algorithm with a run time polynomial in N , M , and B —that computes \hat{p}^M .

In an effort to avoid dealing with an NP-complete problem, we consider in §3.2 modifying the formulation by adding a price-ladder constraint. We explain why this modification results in a relevant model and discuss the new optimization problem, which turns out to be a special case of supermodular function maximization with real-valued variables. Although many researchers have worked on it, the question of whether or not supermodular function maximization is NP-hard remains an open problem.

3.1. NP-Completeness

As discussed in Example 1.1, for any set of competitor products and prices, the revenue function R^{RANK} of the RANK PRICING Model is based on a choice function that satisfies Assumptions 1.2–1.4. Further, it is easy to see that R^{RANK} can be evaluated in polynomial time. Hence, the following computational problem is a special case of GENERAL PRICING:

RANK PRICING

Instance:

- Finite sets of products $A = \{1, \dots, n\}$ and $\bar{A} = \{n+1, \dots, N\}$.
- Prices of competitor products d_{n+1}, \dots, d_N .
- A maximum budget $B \in \mathcal{X}_+$.
- A set $\mathcal{C} = \{\omega_1, \dots, \omega_M\}$, where $\omega_j = (b_j, Z_j)$, $b_j \in \mathcal{X}_+$, $b_j \leq B$, and Z_j is an ordered list of products in $A \cup \bar{A}$.

Problem: Find a price vector

$$\hat{p}^M \in \arg \max_{p \in [0, B]^n} \sum_{j=1}^M R^{\text{RANK}}(p, \omega_j).$$

It is easy to see that GENERAL PRICING, and therefore RANK PRICING, is in NP. In this subsection, we will establish that RANK PRICING, and therefore GENERAL PRICING, is NP-complete in the strong sense.

Consider a class of RANK PRICING problems for which:

- there are no competitor products,
- the maximum budget B is an integer, and
- each j th consumer's budget b_j is an integer.

Note that for each pricing problem in this specific class, there is a vector $\hat{p}^M \in \mathcal{X}_+^n \cap [0, B]^n$ that attains the maximum in

$$\max_{p \in [0, B]^n} \sum_{j=1}^M R^{\text{RANK}}(p, \omega_j).$$

Further, it is easy to obtain this integer-valued vector from any optimal solution to the original continuous optimization problem (in polynomial time). Hence, the optimization problem can be replaced by

$$\max_{p \in \mathcal{X}_+^n \cap [0, B]^n} \sum_{j=1}^M R^{\text{RANK}}(p, \omega_j)$$

without loss of generality. If this optimization problem can be solved in polynomial time, so can the following binary decision problem:

RANK PRICING DECISION

Instance:

- A finite set of products $A = \{1, 2, \dots, n\}$.
- A maximum budget $B \in \mathcal{X}_+$.
- A set $\mathcal{C} = \{\omega_1, \dots, \omega_M\}$, where $\omega_j = (b_j, Z_j)$, $b_j \in \mathcal{X}_+$, $b_j \leq B$, and Z_j is an ordered list of products in A .

Question: Is there a price vector $p \in \mathcal{X}_+^n \cap [0, B]^n$ such that

$$L \leq \sum_{j=1}^M R^{\text{RANK}}(p, \omega_j),$$

where there are no competitor products?

Via a reduction from SIMPLE MAX CUT provided in Appendix B, we obtain the following result.

THEOREM 3.1. *RANK PRICING DECISION is NP-complete in the strong sense.*

We have the following corollary.

COROLLARY 3.1. *GENERAL PRICING and RANK PRICING are both NP-complete in the strong sense.*

3.2. Price Ladders and Supermodularity

In an effort to avoid having to solve an NP-complete problem, we consider a version of GENERAL PRICING modified by introducing a *price-ladder* constraint—a constraint on the ordering of prices. In particular, we will restrict the price vector to a set $\mathcal{D}_n \subseteq \mathfrak{R}_+^n$ defined by

$$\mathcal{D}_n = \{p \in \mathfrak{R}_+^n : p_1 \leq p_2 \leq \dots \leq p_n\}.$$

The constraint on the ordering of prices arises naturally in many applications. In practice, marketing managers typically know the relative magnitude of the price for each product. This knowledge comes from strategic considerations, characteristics of the market for each product, and other practical considerations. For example, a company may offer two versions of a product: a basic model marketed to a general population, and a premium version aimed at more affluent markets. The premium model should have a higher price than the basic model.

Price-ladder constraints occur in the automotive industry, where a vehicle often comes with different trim packages. For instance, consider a 2002 Saturn L-series, a compact passenger vehicle offered by GM. The L-series comes in three standard trim packages: L100, L200, and L300. The L300 has more features than the L200, which in turn has more features than the L100. In this case, the L300 should have a higher price than the L200, whose price should exceed that of the L100.

These examples describe a situation common to companies with large product lines, suggesting the assumption on the ordering of prices is reasonable in many settings. With this assumption, marketing managers can influence the optimal pricing policy by restricting the ordering among prices of products. The restriction ensures that the price of each product makes sense and meets other business constraints. In addition, we can consider multiple orderings of the prices, with each ordering corresponding to possibly different business constraints.

In certain situations, a natural ordering of prices may not be available. For instance, consider a company that has multiple product lines, and each product line has its own price-ladder constraint. It is not clear how we can order the prices of products from different product lines. In this case, we can use as a starting point the ordering implied by the current price of each product, and consider permutations of this ordering.

Thus, we wish to solve instead the following constrained optimization problems:

$$\max_{p \in \mathcal{D}_n} \sum_{j=1}^M R(p, \omega_j).$$

Note that \mathcal{D}_n is a sublattice of \mathfrak{R}_+^n . The following lemma identifies a special property of the revenue function when restricted to \mathcal{D}_n .

LEMMA 3.1. *Consider any choice function C that satisfies Assumptions 1.2 and 1.4 and its associated revenue function R . Then, for each $\omega \in \Omega$, $R(\cdot, \omega)$ is supermodular when restricted to \mathcal{D}_n , i.e.,*

$$R(p, \omega) + R(p', \omega) \leq R(p \wedge p', \omega) + R(p \vee p', \omega) \\ \forall p, p' \in \mathcal{D}_n, \omega \in \Omega,$$

where $p \wedge p' = (p_1 \wedge p'_1, \dots, p_n \wedge p'_n)$ and $p \vee p' = (p_1 \vee p'_1, \dots, p_n \vee p'_n)$.

PROOF. Fix $\omega \in \Omega$. To simplify notation, for any $p \in \mathcal{D}_n$, we will denote $C(p, \omega)$ by $C(p)$ and $R(p, \omega)$ by $R(p)$. Note that $p \wedge p' \leq p$, $p' \leq p \vee p'$, and all four vectors have the same ordering. If $C(p \wedge p') \in \{0\} \cup \bar{A}$, it follows from Assumption 1.4 that $C(p \wedge p') = C(p) = C(p') = C(p \vee p')$, and we are finished because the revenue under all four price vectors will be zero.

So, suppose that $C(p \wedge p') = l \in A$ for some l . This implies that either $p_l \leq b$ or $p'_l \leq b$. We will consider the case where $p_l \leq b$ because the proof for the case where $p'_l \leq b$ is exactly the same by symmetry. Because $C(p \wedge p') = l$ and $p_l \leq b$, it follows from Assumption 1.4 that $C(p) = l$.

At this point, there are two cases to consider: $p'_l \leq b$ and $p'_l > b$.

Case 1. $p'_l \leq b$. By Assumption 1.4 once again, we conclude that $C(p') = l$ as well. We also have $p_l \vee p'_l \leq b$, which implies that $C(p \vee p') = l$ by Assumption 1.4, and therefore,

$$l = C(p \wedge p') = C(p) = C(p') = C(p \vee p'),$$

and we are finished because

$$p_l + p'_l = (p_l \wedge p'_l) + (p_l \vee p'_l).$$

Case 2. $p'_l > b$. In this case, we have $p_l \leq b < p'_l$. Because $C(p) = C(p \wedge p') = l$, it follows that

$$R(p) = p_l = p_l \wedge p'_l = R(p \wedge p').$$

To prove our desired result, it suffices to show that $R(p') \leq R(p \vee p')$. We will prove this by showing that $C(p') = C(p \vee p')$.

Because $p'_1 \leq p'_2 \leq \dots \leq p'_n$ and $b < p'_l$, it follows from Assumption 1.2 that

$$C(p') \in \{0\} \cup \bar{A} \cup \{1, 2, \dots, l-1\}.$$

If $C(p') \in \{0\} \cup \bar{A}$, we are finished because we would have $C(p') = C(p \vee p')$ by Assumption 1.4. On the other hand, if $C(p') = s \in \{1, 2, \dots, l-1\}$ for some s , then we have $p'_s \leq b$. Moreover, we also have $p_s \leq p_l \leq b$, which implies that $p_s \vee p'_s \leq b$, and by Assumption 1.4, we conclude that $C(p \vee p') = s$. In all cases, we see that $C(p') = C(p \vee p')$. \square

It follows from Lemma 3.1 that a function

$$\bar{R}(p) = \sum_{j=1}^M R(p, \omega_j) \quad \forall p \in \mathcal{D}_n,$$

is supermodular if the revenue function R is generated by a choice function that satisfies Assumptions 1.2 and 1.4. Hence, GENERAL PRICING becomes a problem of supermodular function maximization once a price-ladder constraint is added. If decision variables are constrained to take on binary values, there are polynomial-time algorithms for

supermodular function maximization (Cunningham 1985; Grötschel et al. 1981, 1988; Iwata et al. 2001; Schrijver 1971). The problem of optimizing supermodular functions with real decision variables—as in our pricing problem—has also been one of interest to the supermodular games literature (Topkis 1998, Veinott 2002). Unfortunately, to the best of our knowledge, it is not known whether this problem is NP-hard or whether it can be solved in polynomial time. In the next section, we develop a heuristic algorithm for this sort of optimization.

Because GENERAL PRICING is NP-complete without the price-ladder constraint, we should expect that the corresponding objective is not supermodular. This is verified by the following example.

EXAMPLE 3.1. Consider the choice function C^{RANK} of the RANK PRICING Model from Example 1.1. Suppose that $A = \{1, 2\}$, $\bar{A} = \emptyset$, and $\omega = (b, Z)$, with $b = 3$ and $Z = (1, 2)$. Let $p = (1, 5) \in \mathcal{D}_n$ and $p' = (4, 2) \notin \mathcal{D}_n$. Then, $p \wedge p' = (1, 2)$, $p \vee p' = (4, 5)$, and

$$\begin{aligned} R^{\text{RANK}}(p, \omega) + R^{\text{RANK}}(p', \omega) \\ = 1 + 2 > 1 + 0 = R^{\text{RANK}}(p \wedge p', \omega) + R^{\text{RANK}}(p \vee p', \omega). \end{aligned}$$

4. An Approximation Algorithm

In this section, we present a heuristic approximation algorithm for price optimization subject to a price-ladder constraint, together with posterior performance bounds. Denote the objective function by

$$\bar{R}(p) = \sum_{j=1}^M R(p, \omega_j) \quad \forall p \in \mathcal{D}_n.$$

Under Assumption 1.1, consumer budgets are bounded by B . To find the maximum of \bar{R} , it suffices to consider only vectors $p \in \mathcal{D}_n \cap [0, B]^n$. Because we do not know how to optimize the prices of all products simultaneously, our heuristic will iteratively optimize the prices of individual products, one at a time, while fixing the prices of others. The heuristic will cycle through the products repeatedly. To make this description more precise, define a function $T: \mathcal{D}_n \cap [0, B]^n \rightarrow \mathcal{D}_n \cap [0, B]^n$ as follows:

$$T(p) = (T_1(p), T_2(p), \dots, T_n(p)) \quad \forall p \in \mathcal{D}_n \cap [0, B]^n,$$

where

$$T_n(p) = \arg \max_{p_{n-1} \leq x_n} \bar{R}(p_1, \dots, p_{n-1}, x_n),$$

and for $i = n-1, \dots, 1$,

$$T_i(p) = \operatorname{argmax}_{p_{i-1} \leq x_i \leq T_{i+1}(p)} \bar{R}(p_1, \dots, p_{i-1}, x_i, T_{i+1}(p), \dots, T_n(p)),$$

and we choose the greatest maximizer (no greater than B) in case of ties. The function T computes the optimal price for each product one at a time. Here, $T_i(p)$ represents the optimal price for the i th product, while the prices of other products remain fixed at $p_1, \dots, p_{i-1}, T_{i+1}(p), \dots, T_n(p)$. The following lemma shows that T is a bounded nondecreasing mapping.

LEMMA 4.1. For any $p \in \mathcal{D}_n \cap [0, B]^n$,

$$0 \leq T(p) \leq Be,$$

where e denotes the vector of all ones. Also, for any $p, p' \in \mathcal{D}_n \cap [0, B]^n$, if $p \leq p'$, then $T(p) \leq T(p')$.

PROOF. The boundedness of T follows from the definition. It remains to show that T is a nondecreasing mapping. Fix $p, p' \in \mathcal{D}_n \cap [0, B]^n$ with $p \leq p'$. Because \bar{R} is supermodular, $p \leq p'$, and

$$T_n(p) = \arg \max_{p_{n-1} \leq x_n} \bar{R}(p_1, \dots, p_{n-1}, x_n)$$

and

$$T_n(p') = \arg \max_{p'_{n-1} \leq x_n} \bar{R}(p'_1, \dots, p'_{n-1}, x_n),$$

it follows from a standard result in theory of supermodular functions (Topkis 1998, Veinott 2002) that $T_n(p) \leq T_n(p')$. By induction, we can show that $T_\ell(p) \leq T_\ell(p')$ for all ℓ . \square

Let $\bar{p} \in \mathcal{D}_n \cap [0, B]^n$ denote the greatest maximizer of \bar{R} , whose existence is ensured by the following lemma.

LEMMA 4.2. Consider any choice function C that satisfies Assumptions 1.2–1.4 and its associated objective function \bar{R} . The set of maximizers of \bar{R} on $\mathcal{D}_n \cap [0, B]^n$ is a nonempty compact sublattice of \mathfrak{R}_+^n and has least and greatest elements.

PROOF. By Assumption 1.3, \bar{R} is upper semicontinuous, and by Lemma 3.1, it is supermodular on $\mathcal{D}_n \cap [0, B]^n$. Therefore, a maximizer of \bar{R} on $\mathcal{D}_n \cap [0, B]^n$ exists, and it follows from Topkis (1998) and Veinott (2002) that the set of maximizers on $\mathcal{D}_n \cap [0, B]^n$ is a nonempty closed sublattice of \mathfrak{R}_+^n . Because $\mathcal{D}_n \cap [0, B]^n$ is nonempty and bounded, the set of maximizers is also a nonempty compact sublattice of \mathfrak{R}_+^n . Therefore, it has least and greatest elements. \square

Define sequences of vectors $\langle x^{(k)} \in \mathcal{D}_n: k \geq 0 \rangle$ and $\langle y^{(k)} \in \mathcal{D}_n: k \geq 0 \rangle$ according to

$$x^{(k+1)} = T(x^{(k)}) \quad \text{and} \quad y^{(k+1)} = T(y^{(k)}) \quad \forall k,$$

with $x^{(0)} = 0$ and $y^{(0)} = Be$. The following lemma establishes that each sequence converges to a fixed point of T . These fixed points also provide bounds on \bar{p} . A similar result appears in the theory of supermodular games (Topkis 1998, Veinott 2002).

LEMMA 4.3. There exist $p^L, p^U \in \mathcal{D}_n \cap [0, B]^n$ such that

$$\lim_{k \rightarrow \infty} x^{(k)} = p^L \leq \bar{p} \leq p^U = \lim_{k \rightarrow \infty} y^{(k)}.$$

Moreover, both p^L and p^U are fixed points of T , i.e., for $p \in \{p^L, p^U\}$,

$$\begin{aligned} \bar{R}(p_1, \dots, p_n) \\ = \max_{p_{i-1} \leq x_i \leq p_{i+1}} \bar{R}(p_1, \dots, p_{i-1}, x_i, p_{i+1}, \dots, p_n) \quad \forall i. \end{aligned}$$

PROOF. Boundedness of T implies that $x^{(0)} = 0 \leq T(x^{(0)}) = x^{(1)}$, and it follows from the monotonicity of T that

$$x^{(k)} \leq x^{(k+1)} \quad \forall k \geq 0.$$

Thus, $\langle x^{(k)} \in \mathcal{D}_n: k \geq 0 \rangle$ is a bounded nondecreasing sequence. Hence, it converges to a limit. A similar argument applies to the sequence $\langle y^{(k)} \in \mathcal{D}_n: k \geq 0 \rangle$. Because $x^{(0)} \leq \bar{p} \leq y^{(0)}$ and $T(\bar{p}) = \bar{p}$, it follows from Lemma 4.1 that $p^L \leq \bar{p} \leq p^U$. \square

The following corollary shows that for the RANK PRICING Model considered in Example 1.1, we can upper bound the number of iterations required for convergence.

COROLLARY 4.1. *For the RANK PRICING Model, the sequences $\langle x^{(k)} \in \mathcal{D}_n: k \geq 0 \rangle$ and $\langle y^{(k)} \in \mathcal{D}_n: k \geq 0 \rangle$ converge to their limits in at most nM iterations.*

PROOF. Because the competitors' prices remain fixed, using the same argument as in Example 1.1, we can assume without loss of generality that each customer profile ω_j contains only our products.

Let $V = \{b_j: j = 1, \dots, M\}$ denote the set of consumers' budgets. Consider any product whose price does not belong to the set V . Under the RANK PRICING Model, we can increase the product price to the next-highest budget in V without sacrificing sales. Because T always chooses the greatest maximizer, it follows that $x_i^{(k)}, y_i^{(k)} \in V$ for all $k \geq 0$ and $i = 1, \dots, n$.

In the proof of Lemma 4.3, we show that the sequence $\langle x^{(k)} \in \mathcal{D}_n: k \geq 0 \rangle$ is nondecreasing. Thus, the value of at least one coordinate must increase at each iteration. Each coordinate changes its value at most M times because $|V| \leq M$. The sequence therefore converges in at most nM iterations. A similar argument applies to the sequence $\langle y^{(k)} \in \mathcal{D}_n: k \geq 0 \rangle$. \square

By Lemma 4.3, p^L and p^U bound \bar{p} . If p^L equals p^U , then we get the optimal solution. However, we may obtain $p^L < \bar{p} < p^U$, as the following example demonstrates.

EXAMPLE 4.1. Consider a choice function C^{RANK} for the RANK PRICING Model in Example 1.1. In this example, we have two products, $A = \{1, 2\}$, $\bar{A} = \emptyset$, and five customer profiles: $\omega_j = (b_j, Z_j)$ for $1 \leq j \leq 5$, where

$$Z_1 = Z_2 = Z_3 = Z_4 = Z_5 = (2, 1)$$

for all j , and

$$b_1 = 11, \quad b_2 = 21, \quad b_3 = 34, \quad b_4 = 44, \quad b_5 = 55.$$

Let $\mathcal{D}_2 = \{p \in \mathfrak{R}_+^2: p_1 \leq p_2\}$. Table 2 shows the value of $\bar{R}^{\text{RANK}}(b_i, b_j)$ for all $(b_i, b_j) \in \mathcal{D}_2$. By definition, we have

$$R^{\text{RANK}}(p, \omega_j) = \begin{cases} p_2 & \text{if } p_2 \leq b_j, \\ p_1 & \text{if } p_1 \leq b_j < p_2 \\ 0 & \text{otherwise.} \end{cases}$$

Because $p_1 \leq p_2$, it follows that

$$R^{\text{RANK}}(p, \omega_j) = \max_{i=1,2} \{p_i 1[p_i \leq b_j]\}.$$

It follows from Table 2 that

$$p^L = (b_1, b_3) < (b_2, b_4) = \bar{p} < (b_3, b_5) = p^U.$$

In our experience, the price vectors p^L and p^U perform well, yielding a relatively high objective value. However, we can increase the objective value by performing a local search with p^L as a starting point. To formalize this idea, let $w^0 = p^L$. Because $w^0 = T(w^0)$, we will perturb w^0 and generate a new sequence by repeatedly applying the mapping T . Our heuristic will consider an ensemble of perturbations, with each one corresponding to a small increase in a coordinate of w^0 . To make this statement more precise, let r_i denote the perturbation of w_i^0 by some small constant $\epsilon > 0$, i.e.,

$$r_i = w_i^0 + \epsilon \quad \forall i = 1, \dots, n,$$

and define $w^i \in \mathcal{D}_n \cap [0, B]^n$ as follows:

$$w_l^i = \begin{cases} w_l^0 & \text{if } l < i, \\ r_i & \text{if } l = i, \\ w_l^0 \vee r_i & \text{if } l > i. \end{cases}$$

The vector w^i defines a perturbation of w^0 in the i th coordinate. For $l > i$, we require that $w_l^i = w_l^0 \vee r_i$ to ensure that the perturbed vector w^i maintains the same ordering of prices as w^0 . For each i , our heuristic generates a sequence $\langle T^k(w^i): k \geq 0 \rangle$ by iteratively applying the mapping T . Let s^i denote the limit of the sequence, i.e., $s^i = \lim_{k \rightarrow \infty} T^k(w^i)$. After we obtain s^1, \dots, s^n , we choose s^{i^*} such that

$$\bar{R}(s^{i^*}) = \max_{i=1, \dots, n} \bar{R}(s^i).$$

If $\bar{R}(s^{i^*}) > \bar{R}(w^0)$, we repeat the process with $w^0 = s^{i^*}$ as a new starting point; otherwise, the algorithm terminates.

In general, ϵ can be chosen arbitrarily. However, in the case of the RANK PRICING Model, we can design our perturbation mechanism that exploits its special property. Let

Table 2. Values of $\bar{R}^{\text{RANK}}(p_1, p_2)$ for Example 4.1.

$\bar{R}^{\text{RANK}}(p_1, p_2)$	b_1	b_2	b_3	b_4	b_5
b_1	55	95	*124 ¹	121	99
b_2		84	123	*130 ¹	118
b_3			102	122	*123 ¹
b_4				88	99
b_5					55

¹ Those pairs of (p_1, p_2) that correspond to either p^L , \bar{p} , or p^U are marked with an *.

$V = \{b_j: j = 1, 2, \dots, M\}$. From the proof of Lemma 4.1, we know that it suffices to only consider product prices that belong to V . Thus, we can redefine r_i as follows:

$$r_i = \min\{v \in V: v > w_i^0\} \quad \forall i = 1, \dots, n.$$

In this case, r_i represents the smallest budget that lies above w_i^0 .

From our practical experience, the heuristic has worked well. The sequence $\langle T^k(w^i): k \geq 0 \rangle$ converges quickly, even though we cannot guarantee its convergence nor establish any theoretical bound on its convergence rate. The initial choice of $w^0 = p^L$ in the heuristic is somewhat arbitrary. Alternative starting points will lead to different performance. However, $w^0 = p^L$ seems to work well from our experience, yielding a final price vector that results in high revenue. Section 5 presents experimental results that validate the performance of this methodology.

In the remainder of this section, we develop posterior performance bounds for our approximation algorithm. Let R^L denote the revenue generated by p^L given in Lemma 4.3. Also, let p^H denote the price vector obtained from our heuristic approximation algorithm with p^L as an initial starting point, and let R^H denote the associated revenue. Further, let R^* denote the optimal revenue, i.e., the revenue generated by prices \bar{p} .

The following lemma establishes an upper bound on the optimal revenue. Recall that for any j , we have $\omega_j = (b_j, Z_j)$, where Z_j denotes an ordered list of recommended products.

LEMMA 4.4. Let $X \subseteq \{1, 2, \dots, M\}$ be defined by

$$X = \left\{ j: \min_{i \in Z_j \cap A} p_i^L \leq b_j \right\}.$$

Then,

$$R^L \leq R^H \leq R^* \leq \sum_{j \in X} b_j.$$

PROOF. It follows from the definition of our heuristic that $R^L \leq R^H$. Consider any $j \notin X$. This implies that $b_j < \min_{i \in Z_j \cap A} p_i^L$. Because $p^L \leq \bar{p}$ by Lemma 4.3, it follows that $b_j < \min_{i \in Z_j \cap A} \bar{p}_i$. It follows from Assumption 1.2 that $R(\bar{p}, \omega_j) = 0$.

Thus, $R(\bar{p}, \omega_j) = 0$ for all $j \notin X$, which implies that

$$R^* = \sum_{j=1}^M R(\bar{p}, \omega_j) = \sum_{j \in X} R(\bar{p}, \omega_j) \leq \sum_{j \in X} b_j,$$

where the last inequality follows from Assumption 1.2. \square

Our next lemma offers a posterior performance guarantee in terms of the ratio between p^L and p^U . The performance of the heuristic increases as the ratio approaches 1.

LEMMA 4.5.

$$\min_{i=1, \dots, n} \left\{ \frac{p_i^L}{p_i^U} \right\} \leq \frac{R^H}{R^*} \leq 1,$$

where we define $0/0 = 1$.

The proof of this lemma makes use of the following result, which relates the ratio between the revenue under two sets of prices.

LEMMA 4.6. Consider any choice function C that satisfies Assumption 1.2 and 1.4 and its associated revenue function R . Then, for any $\omega \in \Omega$ and $p, p' \in \mathcal{D}_n$ such that $p \leq p'$, we have

$$\frac{R(p, \omega)}{R(p', \omega)} \geq \min_{i=1, \dots, n} \left\{ \frac{p_i}{p'_i} \right\}.$$

PROOF. Fix $\omega = (b, Z) \in \Omega$. To simplify notation, for any $p \in \mathcal{D}_n$, we will denote $C(p, \omega)$ by $C(p)$ and $R(p, \omega)$ by $R(p)$. If $R(p) \geq R(p')$, then the result is trivially true because $p \leq p'$. Therefore, it suffices to consider only the case when $R(p) < R(p')$.

Because $0 \leq R(p) < R(p')$, it must be the case that $C(p') \in A$; otherwise, the revenue under p' will be zero.

The fact that $C(p') \in A$ also implies that $C(p) \in A$. To see this, suppose on the contrary that $C(p) \notin A$. Because $p \sim p'$ and $p \leq p'$, it follows from Assumption 1.4 that $C(p) = C(p')$, implying that $C(p') \notin A$. Contradiction! Therefore, $C(p) \in A$.

Thus, under both price vectors p and p' , the consumer chooses our products. This implies that

$$R(p) = p_{C(p)} \quad \text{and} \quad R(p') = p'_{C(p')}.$$

We will now prove that $C(p) \geq C(p')$. Suppose on the contrary that $C(p) < C(p')$. Then, we have

$$p_{C(p)} \leq p'_{C(p)} \leq p'_{C(p')} \leq b,$$

where the first inequality follows from the fact that $p \leq p'$. The second inequality follows from the fact that $p'_1 \leq p'_2 \leq \dots \leq p'_n$ and our assumption that $C(p) < C(p')$. The final inequality follows from Assumption 1.2 and the fact that $C(p') \in A$. The above inequality implies that $p'_{C(p)} \leq b$. Because $C(p) \in A$, $p \sim p'$, and $p \leq p'$, it follows from Assumption 1.4 that $C(p) = C(p')$. Contradiction! Therefore, $C(p) \geq C(p')$.

Putting everything together, we have

$$\frac{R(p)}{R(p')} = \frac{p_{C(p)}}{p'_{C(p')}} \geq \frac{p_{C(p)}}{p'_{C(p)}} \geq \min_{i=1, \dots, n} \left\{ \frac{p_i}{p'_i} \right\},$$

where the first inequality follows from the fact that $p'_1 \leq p'_2 \leq \dots \leq p'_n$ and $C(p) \geq C(p')$. The final inequality follows from the fact that $C(p) \in A$ and $A = \{1, 2, \dots, n\}$. \square

Here is the proof of Lemma 4.5.

PROOF. It is easy to show that for any $x_1, \dots, x_n \in \mathfrak{R}_+$ and $y_1, \dots, y_n \in \mathfrak{R}_+$,

$$\min_{i=1, \dots, n} \frac{x_i}{y_i} \leq \frac{\sum_{i=1}^n x_i}{\sum_{i=1}^n y_i}.$$

Because $R^L = \sum_{j=1}^M R(p^L, \omega_j)$ and $R^* = \sum_{j=1}^M R(\bar{p}, \omega_j)$, it follows that

$$\frac{R^L}{R^*} \geq \min_{j=1, \dots, M} \frac{R(p^L, \omega_j)}{R(\bar{p}, \omega_j)} \geq \min_{i=1, \dots, n} \left\{ \frac{p_i^L}{\bar{p}_i} \right\} \geq \min_{i=1, \dots, n} \left\{ \frac{p_i^L}{p_i^U} \right\},$$

where the second inequality follows Lemma 4.6, and the final inequality follows from the fact that $p^L \leq \bar{p} \leq p^U$ by Lemma 4.3. Because $R^L \leq R^H$ by Lemma 4.4, the desired result follows. \square

5. A Case Study at General Motors

In this section, we discuss some work from Rusmevichientong (2003) that involves the application of methods developed in this paper to data collected from the ACA website, focusing on the RANK PRICING formulation with a price-ladder constraint. This section presents only a summary of experimental results. For more information on the application of our model and new insights that are generated using our algorithm, the reader is referred to Rusmevichientong (2003).

We will discuss opportunities and challenges in pricing vehicles using consumer preference data collected from the ACA website, and contrast our approach with the existing pricing practice at GM and traditional data sources. Our analysis identifies opportunities to enhance the current pricing strategy, and highlights the benefits of coordinating the prices of all GM vehicles simultaneously.

This section is organized as follows. In §5.1, we provide a brief overview of the current pricing methodology at GM, motivating the use of online data. Then, we present experimental results in §5.2.

5.1. Current Pricing Methodology at General Motors

When GM wants to determine the manufacturer's suggested retail price (MSRP) for a vehicle, it typically conducts a "physical property clinic." During the clinic, participants compare the target vehicle with a sister GM vehicle and up to six other competitive vehicles from the same segment. The participants conduct reviews of each vehicle, rank the vehicles according to their preferences, and complete a discrete-choice-based pricing exercise that assesses relative values of each feature of the vehicle.

The data from the clinic provides an initial estimate of the market share at different prices. The marketing and finance groups then extrapolate the results to reflect the whole segment, not just the vehicles used in the clinic.

The extrapolated demand is then compared to production capacities and production targets. Given this information, GM then tries to determine the vehicle price that most effectively balances market demand, market competitiveness, and internal production constraints.

Because GM has opportunities to interact closely with the participants, the data collected from the clinic are rich, capturing detailed consumer preferences. However, the collection of such data is expensive and time consuming, limiting the size of the data set, with about 200–350 consumers in a typical clinic.

Moreover, participants in a clinic only compare a target vehicle with about 6–7 other vehicles from a vehicle segment that can have as many as 55 vehicles, as in the mid-size vehicle market. Such a small comparison set may not accurately capture the interdependencies among all relevant vehicles.

In contrast, the ACA website has received over 450,000 visitors since its launch in January 2002. Moreover, the set of recommended vehicles is generated based on considerations of over 1,000 vehicles from the entire automotive market, representing over 250 makes and models. The number of consumers and vehicles in this data set far exceeds what is typically available from a clinic's data, enabling us to obtain a more accurate estimate of the demand and the substitutability among all GM vehicles.

In addition, because a clinic is organized to determine a price for a particular vehicle, it is often difficult to use such data to coordinate the pricing strategy across the entire GM product portfolio. To allow for coordination across products, we must understand the relationship among prices of all vehicles; specifically, how a change in the price of one vehicle affects the demand for all others.

5.2. Experimental Results

The data set used in our analysis consists of 83,813 consumers, representing visitors to the ACA website from 1/1/02 until 9/30/02. We only consider those visitors who spent at least three minutes at the website and specified a budget of \$60,000 or less. We impose these constraints to rule out consumers who may not be truly interested in purchasing vehicles.

We divide the data set into a training sample with 41,940 consumers and a validation sample with 41,873 consumers. We compute the vehicle prices using data from the training sample, and evaluate the performance of these prices on the validation sample. The consumers in our data set can be divided into three categories:

(1) Those consumers who only consider GM vehicles, and thus only have GM vehicles in their recommended lists. We will refer to these consumers as "GM Only."

(2) Those consumers who are willing to consider both GM and non-GM vehicles, and at least one GM vehicle is "competitive." This means that at least one GM vehicle is ranked higher than all affordable non-GM vehicles. We will refer to these consumers as "Both GM and non-GM."

(3) Those consumers who are willing to consider both GM and non-GM vehicles, but for whom no GM vehicle is “competitive.” Thus, either the list contains only non-GM vehicles, or there exists an affordable non-GM vehicle that is ranked higher than any other GM vehicles. We will refer to these consumers as “non-GM Only.”

Recall that the competitors’ prices are assumed to be fixed. Hence, we can identify affordable non-GM vehicles. Moreover, under the RANK PRICING formulation, each consumer chooses the highest-ranked vehicle that she can afford. For our optimization, we can thus exclude consumers who are “non-GM Only” because they will not purchase any GM vehicle, regardless of its price. We only need to consider those consumers who are either “GM Only” or “Both GM and non-GM.” Table 3 shows the number of consumers in each category and their average budgets in the validation sample.

Although each consumer’s recommended list typically has 10 recommendations, we use only the top five vehicles because we feel that these vehicles most accurately reflect the consumer preference and budget constraints. We have a total of 1,121 vehicles, 301 of which are GM vehicles and the remaining 820 represent competitors’ vehicles. We assume that the prices of competitors’ vehicles remain constant. Thus, the decision variables are the prices of 301 GM vehicles.

The experiment compares three pricing policies: a “rounded MSRP,” a greedy policy, and the coordinated pricing policy generated by the algorithm discussed in §4. The “rounded MSRP” of each vehicle corresponds to the existing MSRP that has been rounded up to the nearest \$1,000.

We consider the rounded MSRP instead of the existing MSRP because each consumer who visits the ACA website can specify her budget only in the increment of \$1,000. Consequently, the prices under our coordinated pricing policy will also have increments of \$1,000. However, the existing MSRPs of the vehicles vary continuously in the increment of \$1. To ensure a fair comparison with our coordinated pricing policy, we thus round up the existing MSRP of each vehicle to the nearest \$1,000. If a vehicle has an existing MSRP of \$12,030, its rounded MSRP will be \$13,000. We round up the price because the consumers who can afford this vehicle will have a budget of \$13,000 or more.

Table 3. The number of consumers in each category and their average budgets in the validation sample.

Consumer category	Number of consumers	% of total (%)	Avg. budget (\$)	Stdev. of budget (\$)
GM Only	13,116	31.32	30,716	11,016
Both GM and non-GM	8,847	21.13	21,343	11,546
Non-GM Only	19,910	47.55	26,032	10,913
Total	41,873	100	26,509	11,580

Although we use rounded MSRPs to determine total revenue, our coordinated pricing policy takes as a price-ladder constraint the ordering implied by the existing MSRP of each vehicle.

To determine the greedy price of the i th vehicle, we identify those consumers in the training sample in which the i th vehicle appears in the recommended list. The greedy price of the i th vehicle is the largest price that maximizes revenue from these consumers, assuming that they will purchase the i th vehicle if its price meets their budgets. More formally, the greedy price of the i th vehicle, denoted by p_i^G , is defined by

$$p_i^G = \arg \max_{x_i \geq 0} x_i \sum_{j: Z_j \ni i} 1[x_i \leq b_j],$$

where we choose the greatest maximizer in case of ties. If the i th vehicle does not appear in the recommended list of any consumer, we set p_i^G to the rounded MSRP of the i th vehicle.

We should note that the greedy prices typically fail to satisfy the a priori price-ladder constraint. Moreover, the computation of the greedy price ignores possible substitutions among GM vehicles, implicitly assuming that those consumers who consider the i th vehicle will not consider any substitute. In the next section, we will show that by taking into account substitutions among vehicles, as is done in our coordinated pricing policy, we can develop a more effective pricing strategy.

5.2.1. Overall Performance. Table 4 shows the average price of GM vehicles, the sales volume, and the total revenue generated under the three pricing policies. To determine the sales volume, for each consumer we identify the highest-ranked vehicle that she can afford under each pricing policy. We then add up the number of GM vehicles sold under each policy.

We see from the table that our coordinated pricing policy yields a 6% improvement in revenue over the rounded MSRP and a 2% improvement in revenue over the greedy policy. These improvements are significant as indicated by the standard deviations. To determine the standard deviation of revenue for each pricing policy, we compute the revenue generated from each consumer in the validation sample. We then estimate the sample standard deviation of these revenues and multiply the result by \sqrt{M} , where M denotes the number of consumers.

Table 4. Performance of various pricing policies on the validation sample.

Policy	Avg. price	Volume (units)	Revenue (millions)	Stdev. (millions)
Rounded MSRP	\$28,432	17,016	\$399	\$1.73
Greedy	\$30,764	18,147	\$414	\$1.58
Coordinated	\$28,807	18,243	\$421	\$1.74

5.2.2. Deviation from Optimality. We wish to assess the performance of our coordinated pricing policy relative to the optimal policy. Because our coordinated pricing policy is computed using data from the training sample, there are two dimensions in this assessment. We must evaluate the effectiveness of our coordinated pricing policy on both the training and the validation sample.

Let p^L and p^U denote the pricing policies defined in Lemma 4.3. The policy p^L is used in our heuristic to determine the coordinated pricing policy (see §4). These pricing policies are computed using data from the training sample.

We will first determine the effectiveness of our coordinated pricing policy on the training sample. It turns out that

$$\min_{i=1,\dots,301} \frac{p_i^L}{p_i^U} = \frac{p_{30}^L}{p_{30}^U} = \frac{\$15,000}{\$60,000} = 0.25,$$

and it follows from Lemma 4.5 that, on the training sample, the revenue generated from our coordinated pricing policy differs from the optimal revenue by at most 75%.

By applying Lemma 4.4, we can significantly improve the performance guarantee for our coordinated pricing policy. Let \mathcal{M} denote those consumers in the training sample who are in either the “GM Only” category or the “Both GM and non-GM” category. Let $X \subseteq \mathcal{M}$ be defined by

$$X = \left\{ j \in \mathcal{M}: \min_{i \in Z_j \cap A} p_i^L \leq b_j \right\}.$$

It turns out that

$$\sum_{j \in X} b_j = \$576,979,000.$$

Moreover, the revenue from our coordinated pricing policy in the training sample is \$422,351,000 and we have

$$\frac{\$422,351,000}{\$576,979,000} \geq 0.7320.$$

It follows from Lemma 4.4 that, on the training sample, our coordinated pricing policy is at least 73% effective; its revenue differs from the optimal by at most 27%.

Unfortunately, because our coordinated pricing policy is computed using data from the training sample, we cannot apply Lemmas 4.4 and 4.5 to evaluate its performance on the validation sample. However, we can establish a simple performance guarantee by computing an upper bound on the optimal revenue that can be generated from the validation sample. By definition, the optimal revenue cannot exceed the sum of the budget of all consumers who are either in the “GM Only” or the “Both GM and non-GM” category. For the validation sample, this sum turns out to be \$591,694,000. Because the revenue from our coordinated pricing policy on the validation sample is \$421,290,000 and

$$\frac{\$421,290,000}{\$591,694,000} \geq 0.7120,$$

it follows that, on the validation sample, the revenue from our coordinated pricing policy differs from the optimal by at most 29%.

6. Conclusion

Motivated by the availability of data on consumer preferences from the ACA website, we developed a nonparametric approach to multiproduct pricing. We consider a class of models of consumer-purchasing behavior, each of which relates observed data on a consumer’s requirements and budget constraint to subsequent purchasing tendencies. When we do not enforce a price ladder, we showed that these problems are NP-complete in the strong sense. We developed algorithms that address these optimization problems, given a price ladder.

We applied our algorithms to a real data set from the ACA website, validating the performance of our algorithms. Our analysis provides insights into the current pricing policy at GM and suggests improvements that may lead to a more effective pricing strategy. We showed that by taking into account substitutions among vehicles and coordinating the prices of all GM vehicles simultaneously, we can find a pricing strategy that generates higher revenue.

Our work can be extended in many directions. One possibility is to consider a more complex model of consumer-purchasing behavior, incorporating additional information that may influence consumer-purchasing decisions. It is also interesting to explore how we can incorporate production constraints into our formulation, and to understand the impact of competition on product prices. These extensions would provide us with a more realistic model that better reflects consumer-purchasing behavior and current business environment.

As more consumers use the Internet, companies will have access to ever larger quantities of data that reflect consumer preferences. As our work has demonstrated, such data provides us with new opportunities to understand consumer behavior and to optimize product prices. It is interesting to explore other avenues for using this type of data.

Appendix A. Proof of Theorem 2.3

The proof of Theorem 2.3 makes use of the following result. Because the proof of this result follows from elementary analysis, we omit the details.

LEMMA A.1. *For any set X , let $f, g: X \rightarrow \mathfrak{R}_+$ be any function such that*

$$\sup_{x \in X} f(x) < \infty \quad \text{and} \quad \sup_{x \in X} g(x) < \infty.$$

Then,

$$\left| \sup_{x \in X} f(x) - \sup_{x \in X} g(x) \right| \leq \sup_{x \in X} |f(x) - g(x)|.$$

The next lemma provides a simple bound on the difference between the expected revenue under \hat{p}^M and the optimal expected revenue.

LEMMA A.2. For any $\epsilon > 0$, let $P_1(\epsilon), P_2(\epsilon) \in [0, 1]$ be defined by

$$P_1(\epsilon) \equiv \Pr\{|E_{\mathcal{P}}[R(p^*, \omega)] - E_{\mathcal{P}}[R(\hat{p}^M, \omega)]| > 2\epsilon\},$$

$$P_2(\epsilon) \equiv \Pr\left\{\left|E_{\mathcal{P}}[R(p^*, \omega)] - \frac{1}{M} \sum_{j=1}^M R(\hat{p}^M, \omega_j)\right| > \epsilon\right\}.$$

Then,

$$\max\{P_1(\epsilon), P_2(\epsilon)\}$$

$$\leq \Pr\left\{\sup_{p \in [0, B]^n} \left|E_{\mathcal{P}}[R(p, \omega)] - \frac{1}{M} \sum_{j=1}^M R(p, \omega_j)\right| > \epsilon\right\}.$$

PROOF. Because

$$E_{\mathcal{P}}[R(p^*, \omega)] = \sup_{p \in [0, B]^n} E_{\mathcal{P}}[R(p, \omega)],$$

it follows that

$$\begin{aligned} & |E_{\mathcal{P}}[R(p^*, \omega)] - E_{\mathcal{P}}[R(\hat{p}^M, \omega)]| \\ &= E_{\mathcal{P}}[R(p^*, \omega)] - E_{\mathcal{P}}[R(\hat{p}^M, \omega)] \\ &= E_{\mathcal{P}}[R(p^*, \omega)] - \frac{1}{M} \sum_{j=1}^M R(p^*, \omega_j) \\ &\quad + \frac{1}{M} \sum_{j=1}^M R(p^*, \omega_j) - \frac{1}{M} \sum_{j=1}^M R(\hat{p}^M, \omega_j) \\ &\quad + \frac{1}{M} \sum_{j=1}^M R(\hat{p}^M, \omega_j) - E_{\mathcal{P}}[R(\hat{p}^M, \omega)] \\ &\leq E_{\mathcal{P}}[R(p^*, \omega)] - \frac{1}{M} \sum_{j=1}^M R(p^*, \omega_j) \\ &\quad + \frac{1}{M} \sum_{j=1}^M R(\hat{p}^M, \omega_j) - E_{\mathcal{P}}[R(\hat{p}^M, \omega)] \\ &\leq 2 \sup_{p \in [0, B]^n} \left|E_{\mathcal{P}}[R(p, \omega)] - \frac{1}{M} \sum_{j=1}^M R(p, \omega_j)\right|, \end{aligned}$$

where the first inequality follows from the fact that

$$\frac{1}{M} \sum_{j=1}^M R(p^*, \omega_j) \leq \sup_{p \in [0, B]^n} \frac{1}{M} \sum_{j=1}^M R(p, \omega_j) = \frac{1}{M} \sum_{j=1}^M R(\hat{p}^M, \omega_j).$$

Thus, it follows that

$$\begin{aligned} P_1(\epsilon) &\leq \Pr\left\{2 \sup_{p \in [0, B]^n} \left|E_{\mathcal{P}}[R(p, \omega)] - \frac{1}{M} \sum_{j=1}^M R(p, \omega_j)\right| > 2\epsilon\right\} \\ &= \Pr\left\{\sup_{p \in [0, B]^n} \left|E_{\mathcal{P}}[R(p, \omega)] - \frac{1}{M} \sum_{j=1}^M R(p, \omega_j)\right| > \epsilon\right\}. \end{aligned}$$

To establish the remaining inequality, note that by definition

$$E_{\mathcal{P}}[R(p^*, \omega)] = \sup_{p \in [0, B]^n} E_{\mathcal{P}}[R(p, \omega)],$$

$$\frac{1}{M} \sum_{j=1}^M R(\hat{p}^M, \omega_j) = \sup_{p \in [0, B]^n} \frac{1}{M} \sum_{j=1}^M R(p, \omega_j),$$

and it follows from Lemma A.1 that

$$\begin{aligned} & \left|E_{\mathcal{P}}[R(p^*, \omega)] - \frac{1}{M} \sum_{j=1}^M R(\hat{p}^M, \omega_j)\right| \\ &\leq \sup_{p \in [0, B]^n} \left|E_{\mathcal{P}}[R(p, \omega)] - \frac{1}{M} \sum_{j=1}^M R(p, \omega_j)\right|. \end{aligned}$$

Therefore,

$$P_2(\epsilon) \leq \Pr\left\{\sup_{p \in [0, B]^n} \left|E_{\mathcal{P}}[R(p, \omega)] - \frac{1}{M} \sum_{j=1}^M R(p, \omega_j)\right| > \epsilon\right\}. \quad \square$$

Here is the proof of Theorem 2.3.

PROOF. For M satisfying the bound in Theorem 2.3, it is easy to verify that

$$4 \left(\frac{2B(n+2)}{\epsilon}\right)^n e^{-\epsilon^2 M / 128B^2} \leq \delta.$$

It follows from Theorem 2.2 that

$$\begin{aligned} & \Pr\left\{\left|E_{\mathcal{P}}[R(p^*, \omega)] - \frac{1}{M} \sum_{j=1}^M R(\hat{p}^M, \omega_j)\right| > \epsilon\right\} \\ &\leq 4 \sup_{\mu} \mathcal{N}\left(\frac{\epsilon}{16}, \mathcal{F}, \|\cdot\|_{1, \mu}\right) e^{-\epsilon^2 M / 128B^2} \\ &= 4 \left(\frac{2B(n+2)}{\epsilon}\right)^n e^{-\epsilon^2 M / 128B^2} \\ &\leq \delta, \end{aligned}$$

where the second inequality follows from Theorem 2.1. The desired result follows from Lemma A.2. \square

Appendix B. Proof of Theorem 3.1

PROOF. The RANK PRICING DECISION Problem is in NP because we simply need to guess a price vector and verify in polynomial time whether or not it satisfies the stated condition of the problem. Moreover, by the definition of the problem, it suffices to consider only price vectors $p = (p_a; a \in A)$ such that $p_a \in \{b_1, b_2, \dots, b_M\}$ for all $a \in A$.

To show NP-completeness, we will present a reduction from SIMPLE MAX CUT, a known NP-complete problem (Garey and Johnson 1979). Here is the definition of the SIMPLE MAX CUT Problem.

SIMPLE MAX CUT

Instance:

- A graph $G = (V, E)$.
- A positive integer $K \leq |E|$.

Question: Is there a partition of V into disjoint sets V_1 and V_2 such that the number of edges in E that have one end point in V_1 and one end point in V_2 is at least K ?

For each instance of the SIMPLE MAX CUT Problem, we construct the following instance of the RANK PRICING DECISION Problem. Let the set of products $A = V$.

Table 5. Values of $R_{\{u,v\}}^{\text{RANK}}(p_u, p_v)$.

(p_u, p_v)	$R_{\{u,v\}}^{\text{RANK}}(p_u, p_v)$
(1, 1)	$4 = 1 + 1 + 1 + 1$
(1, 2)	$5 = 1 + 1 + 1 + 2$
(2, 1)	$5 = 1 + 2 + 1 + 1$
(2, 2)	$4 = 0 + 2 + 0 + 2$

For each edge $e = \{u, v\} \in E$, we construct four customer profiles: $\omega_j^{\{u,v\}}$ for $1 \leq j \leq 4$, where

$$\begin{aligned} \omega_1^{\{u,v\}} &= (1, (u, v)), & \omega_2^{\{u,v\}} &= (2, (u, v)), \\ \omega_3^{\{u,v\}} &= (1, (v, u)), & \omega_4^{\{u,v\}} &= (2, (v, u)). \end{aligned}$$

For any $p \in \mathfrak{R}_+^{|V|}$, define $R_{\{u,v\}}^{\text{RANK}}(p)$ as

$$\begin{aligned} R_{\{u,v\}}^{\text{RANK}}(p) &= R^{\text{RANK}}(p, \omega_1^{\{u,v\}}) + R^{\text{RANK}}(p, \omega_2^{\{u,v\}}) \\ &\quad + R^{\text{RANK}}(p, \omega_3^{\{u,v\}}) + R^{\text{RANK}}(p, \omega_4^{\{u,v\}}). \end{aligned}$$

We design these ordered pairs so that $R_{\{u,v\}}^{\text{RANK}}$ achieves the maximum when $p_u \neq p_v$, i.e., when the corresponding edge $e = \{u, v\}$ has end points in different sets. Table 5 shows the values of $R_{\{u,v\}}^{\text{RANK}}(p_u, p_v)$ for $p_u, p_v \in \{1, 2\}$.

Define the lower bound $L = 4|E| + K$. The corresponding RANK PRICING DECISION Problem is to determine if there exists a price vector $(p_v: v \in V)$ such that

$$4|E| + K \leq \sum_{\{u,v\} \in E} R_{\{u,v\}}^{\text{RANK}}(p).$$

We will show that the graph $G = (V, E)$ has a valid cut if and only if the corresponding RANK PRICING DECISION Problem has a valid solution. Suppose that (V_1, V_2) is a valid cut for the SIMPLE MAX CUT Problem. Let $\mathcal{C}(V_1, V_2)$ denote the set of edges with end points in the different sets, i.e.,

$$\mathcal{C}(V_1, V_2) = \{\{u, v\} \in E: u \in V_1, v \in V_2 \text{ or } u \in V_2, v \in V_1\}.$$

Define a price vector $(p_v: v \in V)$ as follows:

$$p_v = \begin{cases} 1 & \text{if } v \in V_1, \\ 2 & \text{if } v \in V_2. \end{cases}$$

It follows from Table 5 that for any $\{u, v\} \in E$,

$$R_{\{u,v\}}^{\text{RANK}}(p) = \begin{cases} 5 & \text{if } \{u, v\} \in \mathcal{C}(V_1, V_2), \\ 4 & \text{otherwise.} \end{cases}$$

Hence,

$$\begin{aligned} \sum_{\{u,v\} \in E} R_{\{u,v\}}^{\text{RANK}}(p) &= 5|\mathcal{C}(V_1, V_2)| + 4(|E| - |\mathcal{C}(V_1, V_2)|) \\ &= 4|E| + |\mathcal{C}(V_1, V_2)| \\ &\geq 4|E| + K, \end{aligned}$$

where the last inequality follows from the fact that (V_1, V_2) is a valid cut of the graph G . Hence, $(p_v: v \in V)$ represents a valid solution to the RANK PRICING DECISION Problem.

On the other hand, suppose that $(p_v: v \in V)$ is a valid solution to the RANK PRICING DECISION Problem. Without loss of generality, we may assume that $p_v \in \{1, 2\}$ for all v , because we can always construct another price vector from p that has this property and remains a valid solution to the RANK PRICING DECISION Problem. Let $\mathcal{C} = \{\{u, v\} \in E: p_u \neq p_v\}$. Because $(p_v: v \in V)$ is a valid solution to the RANK PRICING DECISION Problem,

$$\begin{aligned} 4|E| + K &\leq \sum_{\{u,v\} \in E} R_{\{u,v\}}^{\text{RANK}}(p) \\ &= 5|\mathcal{C}| + 4(|E| - |\mathcal{C}|) \\ &= 4|E| + |\mathcal{C}|, \end{aligned}$$

where the first equality follows from the definition of $R_{\{u,v\}}^{\text{RANK}}$ and Table 5. The above inequality implies that $K \leq |\mathcal{C}|$. If we define $V_1 = \{v: p_v = 1\}$ and $V_2 = \{v: p_v = 2\}$, then V_1 and V_2 form a valid cut of G .

The above argument shows that the RANK PRICING DECISION Problem is NP-complete. The only choices of b_j s used in the proof are 1s and 2s, RANK PRICING DECISION is also NP-complete in the strong sense. \square

Acknowledgments

This research was supported in part by General Motors through funds provided to the General Motors/Stanford University Collaborative Work Systems Laboratory, by NSF CAREER Grant ECS-9985229, and by the ONR through grant MURI-N0014-00-1-0637. The authors are grateful to Arthur F. Veinott for sharing with them his deep understanding of lattice programming. They thank Lynn Truss for introducing them to various business activities relating to this work and for providing them with valuable feedback as the work progressed. The authors thank Rajeev Motwani for useful discussions on NP-completeness results and Yinyu Ye for pointers to related literature. They also thank Jay Ostahowski, Charlie Rosa, Joyce Salisbury, and Jeff Tew for helpful comments and suggestions. The first author thanks David Choi for useful discussions on the VC-dimension and Curtis Eaves for encouraging him to explore further applications of this work.

References

- Coy, P. 2000. The power of smart pricing. *Bus. Week* (April 10), 160.
- Cunningham, W. H. 1985. On submodular function minimization. *Combinatorica* 5 185–192.
- Dobson, G., S. Kalish. 1988. Positioning and pricing a product line. *Marketing Sci.* 7(2) 107–125.
- Dobson, G., S. Kalish. 1993. Heuristics for pricing and positioning a product-line using conjoint and cost data. *Management Sci.* 39(2) 160–175.
- Garey, M. R., D. S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Bell Telephone Laboratories, Murray Hill, NJ.

- Green, P. E., A. M. Krieger. 1985. Models and heuristics for product line selection. *Marketing Sci.* **4**(1) 1–19.
- Grötschel, M., L. Lovász, A. Schrijver. 1981. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* **1** 169–197.
- Grötschel, M., L. Lovász, A. Schrijver. 1988. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, New York.
- Hanson, W., K. Martin. 1990. Optimal bundle pricing. *Management Sci.* **36** 155–174.
- Haussler, D. 1992. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Inform. Comput.* **100**(September) 78–150.
- Iwata, S., L. Fleischer, S. Fujishigi. 2001. A combinatorial, strongly polynomial-time algorithm for minimizing submodular functions. *J. ACM* **48**(4) 761–777.
- McBride, R. D., F. S. Zufryden. 1988. An integer programming approach to the optimal product line selection problem. *Marketing Sci.* **7**(2) 126–140.
- McGill, J. I., G. Van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Sci.* **33**(2) 233–256.
- McWilliams, G. 2001. Lean machine: How Dell fine tunes its PC pricing to gain edge in a slow market. *Wall Street J.* (June 8), A1.
- Merrick, A. 2001. Priced to move: Retailers try to get leg up on mark-downs with new software. *Wall Street J.* (August 7), A1.
- Mirman, L. J., D. Sibley. 1980. Optimal nonlinear prices for multiproduct monopolies. *Bell J. Econom.* **11** 659–670.
- Oren, S., S. Smith, R. Wilson. 1984. Product line pricing. *J. Bus.* **57**(1, pt. 2) S73–S99.
- Oren, S., S. Smith, R. Wilson. 1987. Multi-product pricing for electric power. *Energy Econom.* **9**(2) 104–114.
- Pollard, D. 1984. *Convergence of Stochastic Processes*. Springer-Verlag, New York.
- Rusmevichientong, P. 2003. A non-parametric approach to multi-product pricing: Theory and application. Ph.D. dissertation, Stanford University, Stanford, CA.
- Schrijver, A. 1971. A combinatorial algorithm for minimizing submodular functions in strongly polynomial time. *J. Combin. Theory* **B80** 346–355.
- Sen, S. 1982. Issues in optimal product design. *Analytic Approaches to Product and Marketing Planning: The Second Conference*. Marketing Science Institute, Cambridge, MA, 265–274.
- Shocker, A. D., V. Srinivasan. 1979. Multiattribute approaches for product concept evaluation and generation: A critical review. *J. Marketing Res.* **16** 159–180.
- Smith, S. 1986. New product pricing in quality sensitive markets. *Marketing Sci.* **5**(1) 70–87.
- Spence, A. M. 1980. Multi-product quantity-dependent prices and profitability constraints. *Rev. Econom. Stud.* **47** 821–841.
- Topkis, D. 1998. *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.
- Veinott Jr., A. F. 2002. Lectures in supply-chain optimization. Course notes for Management Science and Engineering 361, Stanford University, Stanford, CA.
- Weatherford, L. R., S. E. Bodily. 1992. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. *Management Sci.* **40**(5) 831–844.