OPPORTUNITIES AND CHALLENGES IN USING ONLINE PREFERENCE DATA FOR VEHICLE PRICING: A CASE STUDY AT GENERAL MOTORS

Paat Rusmevichientong
School of Operations Research and Industrial Engineering
Cornell University
paatrus@cornell.edu

Joyce A. Salisbury
GM NA Market Research
joyce.a.salisbury@gm.com

Lynn T. Truss
Manufacturing Systems Research Lab
GM R&D Center
lynn.t.truss@gm.com

Benjamin Van Roy
Department of Management Science and Engineering
Stanford University
bvr@stanford.edu

Peter W. Glynn
Department of Management Science and Engineering
Stanford University
glynn@stanford.edu

24 May 2005

ABSTRACT

Developed by General Motors (GM), the Auto Choice Advisor web site (http://www. autochoiceadvisor.com) recommends vehicles to consumers based on their requirements and budget constraints. Through the web site, GM has access to large quantities of data that reflect consumer preferences. Motivated by the availability of such data, we formulate a non-parametric approach to pricing GM vehicles, highlight opportunities and challenges in using online data, and contrast our approach with existing methodologies and traditional data sources. Our analysis provides insights into the current pricing practice and suggests enhancements that may lead to a more effective pricing strategy.

1 Introduction

Availability of customer data and increasingly sophisticated analytical tools enable significant increases in profit through optimization of prices. In addition to airlines and hotels who have been traditional users of revenue management, many companies in retail and manufacturing have started to employ advanced pricing strategies to increase their bottom lines.

In this paper, we develop a new approach to price optimization that leverages data that can be collected through operation of a web site that hosts a "product recommendation engine." Our work is motivated by availability of data from the Auto Choice Advisor web site (www.autochoiceadvisor.com) and an aim towards optimization of GM vehicle prices. Developed by General Motors (GM), the Auto Choice Advisor (ACA) web site recommends vehicles to consumers based on their requirements and budget constraints. Through the web site, GM has access to large quantities of data that reflect consumer preferences. In this paper, we formulate a non-parametric approach to pricing GM vehicles, highlight opportunities and challenges in using data from the ACA web site, and contrast our approach with existing methodologies and traditional data sources.

Table 1 shows a typical record that is generated by a visit to the ACA web site. Each consumer has a consumer identifier (ID), a budget, and a list of recommended vehicles. The recommended list is generated by the web site based on the requirements specified by the consumer. Each vehicle in the list is ranked based on how well it matches the consumer's requirements.

ID	Budget	Recommended Vehicles	Rank
1001	$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

Table 1: An example of a record that is generated during a visit to the ACA web site.

Given a large number of records of this type, one can infer valuable information about demand for and substitutability among products. Availability of such data offers opportunities for optimizing prices in a way that leverages the associated information to increase profits. In this paper, we discuss an approach to conducting this kind of data-driven price optimization.

1.1 Contributions and Literature Review

In this section, we outline the contributions of our paper, and provide a review of related literature. Our research contributes to the literature on multi-product pricing by identifying a new non-parametric formulation that is motivated by the consumer preference data collected from the ACA web site. To put our results in perspective, let us briefly review current research in this area.

The literature on multi-product pricing focuses on understanding interactions and substitutions among products, and finding prices that maximize profit in light of these interdependencies and consumer preferences. Yano and Dobson [Yano and Dobson, 1998] and Chen et al. [Chen et al., 1998] provide reviews of research in this area.

Smith [Smith, 1986] and Oren, Smith, and Wilson [Oren et al., 1984] considered a class of demand functions, which depend on product prices and attributes. They then studied how prices and attributes influence sales. Most results revolve around the case of a single differentiating attribute, such as a metric of quality. Recently, Chen and Hausman [Chen and Hausman, 2000] showed that in the case of a single customer type with a logit choice function, the multi-product pricing problem possesses special structure that admits efficient solution techniques.

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 the theory 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 admit efficient computation -- remains a challenge.

Our non-parametric formulation relates closely to another segment of the literature, which describes the demand process by directly modeling the selection made by each consumer. 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] [Dobson and Kalish, 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, since the integer programming problems are NP-complete, this approach requires applications of various heuristics to determine optimal prices.

Our research revolves around a new kind of assumption on user preferences, motivated by data from the ACA web site. Our formulation assumes that a consumer's preferences among vehicles in her recommended list is given by the ranking inferred by the web site. The consumer is then assumed to select from the recommended list the highest ranked vehicle that she can afford.

We identify a class of relevant problems that possess a certain type of special structure. These problems involve maximization of super-modular functions and are obtained through imposing a "price ladder" -- that is, a constraint on the ordering of prices. Though it is not known whether or not maximization of super-modular functions is NP-complete, the presence of this structure facilitates development of an effective heuristic solution method. (See [Rusmevichientong et al., 2002] for more details.)

In addition to providing a new formulation of multi-product pricing, our work is the first, to the best of our knowledge, that focuses on consumer preference data collected from e-commerce web sites. Approaches to multi-product pricing generally assume that data on consumer preferences are gathered through market research or conjoint analysis [Dobson and Kalish, 1988] [Dobson and Kalish, 1993] [Green, Krieger, and Wind, 2001] [Jedidi and Zhang, 2002] [McBride and Zufryden, 1988]. These methods of acquiring data are often expensive and time-consuming, limiting the size of the dataset.

Our paper also relates to an emerging stream of research that focuses on understanding customer behavior in online markets and e-commerce web sites. Brynjolfsson, Hu, and Smith [Brynjofsson et al., 2003] studied the impact on consumer surplus of increased product variety at online booksellers (see also [Chevalier and Goolsbee, 2002]). A related paper by Smith and Brynjolfsson [Smith and Brynjolfsson, 2001] analyzed the consumer decision-making process at internet shopbot, along with the impacts of search cost on product differentiation [Brynjolfsson et al., 2004].

Through e-commerce web sites, 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 web sites that may facilitate a more effective analysis of the data.

1.2 Organization

In the next section, we provide a brief overview of the current pricing methodology used at GM, motivating the use of online data. Section 3 provides additional details on the ACA web site: its purposes, the web site's organization, a detailed description of the data generated from the site, and a discussion of potential opportunities of such data.

Section 4 outlines a problem formulation based on this dataset, and provides motivations behind it. Section 5 discusses the concept of a price ladder – an ordering of prices. This ordering will be a constraint in our optimization process. Section 6 presents results and new insights, and discusses limitations of our approach and means to remedy them. Finally, Section 7 offers conclusions and directions for future work.

2 Current Pricing Methodology at GM

When GM wants to determine the MSRP (manufacturer's suggested retail price) 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 in the vehicle. Including a sister GM vehicle in the comparison set enables GM to understand the interactions between GM vehicles as the price of each vehicle changes.

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.

Since GM has opportunities to interact closely with the participants, the data collected from the clinic are quite rich, capturing detailed consumer preferences. However, the collection of such data is quite expensive and time-consuming, limiting the size of the dataset, 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 market. Such a small comparison set may not accurately capture the interdependencies among all relevant vehicles.

Since 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.

3 Auto Choice Advisor Website

Developed by GM, the Auto Choice Advisor (ACA) web site provides a tool for new vehicle shoppers looking for unbiased help in determining what to consider during their shopping process. This free web-based service attracts users by simplifying their vehicle selection with smart, impartial advice. Consumers are provided an unbiased vehicle consideration set based on their needs and preferences such as personal usage habits, price, brand and body style preference, as well as the importance of certain features like gas mileage and frequency of repairs.

At the center of the ACA web site is a sophisticated decision engine that recommends ten vehicles that best meet the consumer's needs from over 1,000 vehicles, representing over 250 makes and models available from all major manufacturers offering products in North America.

In return, GM connects with online automotive consumers and receives quick feedback on their preferences, providing real-time consumer-based guidance to GM's product development and marketing teams. GM aggregates the information so that personal information is anonymous and consumers are guaranteed that their personal information will never be sold or used without permission. Since the site's debut in January 2002, over 450,000 sessions have been completed.

3.1 Dataset

Table 1 (shown in Section 1) provides an example of a record that is generated during a visit to the web site. Each consumer has a consumer ID, a budget, and a list of recommended vehicles that is generated based on the requirements specified by the consumer. Each vehicle in the recommended list is ranked based on how well it matches the consumer's requirements, with the Honda Accord Value Sedan being the best match in this example. The web site typically presents ten recommendations. In this example, we list the top five. We should note that the web site may recommend vehicles whose prices exceed the consumer’s budget, provided that the vehicles meet other requirements.

In addition to the information shown in Table 1, the ACA web site keeps track of all information that each consumer specified while she visited the web site, including her responses to questions presented to her. For the purpose of our analysis, we will restrict ourselves to records of the kind shown in Table 1.

3.2 Opportunities

The ACA web site offers an alternative source of data for vehicle pricing. Through the web site, GM has access to large quantities of data that reflect consumer preferences and budget constraints. The list of recommended vehicles provides a proxy for the set of vehicles that the consumer would consider purchasing. Such information, when combined with the consumer's budget, offers an opportunity to determine the appropriate price for each vehicle -- one that maximizes the total profit while meeting the consumers' budget constraints.

We can compare the price of each vehicle with the willingness-to-pay of the consumers in the dataset. For instance, consider those consumers for which the 2002 Saturn L100, a compact GM passenger car, appears as the highest ranked vehicle in their recommended lists. Suppose that 90% of such consumers have a budget over $18,000. Since the current MSRP (manufacturer's suggested retail price) for the vehicle is $16,370, this discrepancy provides us with an opportunity to increase the price of the Saturn, with minimal losses in demand[1].

We can also leverage the size of this dataset. The web site has 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 dataset 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.

3.3 Uses of ACA Data as Surrogate Utility Function

We assume that data from the ACA web site represents *the only available information* we have about consumers' preferences, and we will use the recommended list of vehicles as a surrogate utility function for each consumer (see Assumption 3 in Section 4 for more details). Our optimization model and our choice of the objective function depend on the structure of the available data. Clearly, we can develop more complex and realistic models of each consumer's utility function if we have more information about the consumer's preferences such as the actual detailed inputs specified by each consumer or results from other market research studies.

Despite the limited structure of the available data and the choice of the optimization model considered, we will show how our formulation and analyses still provide useful insights into the current pricing policy at GM and suggest enhancements that may lead to a more effective pricing strategy.

4 Problem Formulation and Underlying Assumptions

Section 3.2 highlights opportunities in using data from the ACA web site for pricing vehicles. This section introduces a problem formulation based on this dataset and identifies assumptions underlying our problem formulation, including a discussion of merits and shortcomings of each assumption. Let $A = \{1,2, \ldots, n\}$ denote the set of GM vehicles and $\bar{A} = \{n+1, \ldots, N\}$ the set of

---

[1] We make the simplifying assumption that all of these web consumers are in the market and ready to purchase.

competitors' vehicles. We have a set of M consumers *1, …, M*. Each consumer j has a budget $b_j > 0$ and a list $Z_j = (z^j_1, \ldots, z^j_{nj})$ of n_j vehicles. The consumer specifies her budget b_j when she visits the web site. The list Z_j represents the list of recommended vehicles generated by the web site based on the criteria specified by the consumer. The first assumption deals with the behavior of the decision engine that generates recommendations for the consumers.

Assumption 1 (Property of the Recommended List): The list Z_j is sorted by the degree to which each vehicle matches the consumer's requirements, with z_1^j representing the best match.

GM has conducted a market research study and verified that, among users of the ACA web site, a significant proportion agree with the recommended list returned by the decision engine. Thus, we feel this is a reasonable assumption.

Given the consumers' budgets and the list of recommended vehicles, we want to determine the prices of GM vehicles that maximize revenue generated from the sample of consumers in our dataset. Denote these prices by a vector $p = (p_1, p_2, \ldots, p_n) \in R_+^n$. The next assumption describes competitors' reactions to changes in the prices of GM vehicles.

Assumption 2 (Competitors' Responses): Regardless of changes in the prices of GM vehicles, the prices of competitors' vehicles remain constant.

Assumption 2 is quite stringent, but we believe that if the prices of GM vehicles do not change significantly, this assumption provides a reasonable approximation. Incorporating competitors' responses to our price changes is beyond the scope of this paper.

Unfortunately, our dataset does not contain purchase records. To make use of this data, we need to make an assumption on how each consumer will select a vehicle to purchase from her recommended list. Throughout this paper, we will make the following assumption. We will say that a consumer can *afford* a vehicle if its price is no greater than her budget.

Assumption 3 (Choice Function): The consumer will purchase from the recommended list the highest ranked vehicle that she can afford.

Under this assumption, the consumer whose record is shown in Table 1 would purchase a Saturn L100 if and only if she can afford it but not the Honda Accord Value Sedan. Given the structure of the data from the ACA web site, we believe that Assumption 3 serves as a reasonable approximation for consumers' behaviors. For additional analyses on other choice functions, the reader is referred to the first author's dissertation [Rusmevichientong, 2003].

Our objective is to choose a price vector that maximizes revenue generated from the sample of consumers represented in our records:

$$\max_{p \in R_+^n} \sum_{j=1}^{M} R_j(p),$$

where $R_j(p)$ denotes the revenue generated from the j[th] consumer under a pricing strategy p.

Under Assumption 3, $R_j(p)$ corresponds to the price of the highest ranked vehicle under the consumer's budget if the vehicle is a GM vehicle, and zero otherwise. We will focus on the problem of revenue maximization. However, our approach can be extended to accommodate profit maximization for certain forms of profit functions. For a treatment of this extension, see [Rusmevichientong, 2003].

4.1 Properties of the Choice Function

In the previous section, motivated by the data structure, we introduced the choice function (Assumption 3) and its associated revenue function for our optimization problem. Through a series of examples, we will show that our choice function exhibits many desirable properties such as the substitution effect and overall demand sensitivity to price changes. All examples assume that we only consider GM vehicles.

Example 1 (Substitution Effect): In this example, we will show how, under our choice function, consumers will substitute one vehicle with another, if the price of her original selection suddenly changes. Consider a consumer with the following profile: budget b=$10 and recommended list Z= (1, 2). Suppose we initially set the prices of vehicle 1 at $10 and vehicle 2 at $5. Then, according to Assumption 3, the consumer will purchase vehicle 1, which corresponds to the highest ranked vehicle that she can afford. However, if the price of vehicle 1 increases beyond $10 while the price of vehicle 2 remains constant, the consumer will now substitute vehicle 2 for vehicle 1 since she can no longer afford her original selection.

Example 2 (Overall Demand Sensitivity to Changes in Price): Although the previous example showed how the selection of each consumer changes when prices change, it is not clear if such changes will result in corresponding changes in the overall demand. This example shows how our choice function allows for such changes.

Suppose we have two consumers with the following profiles: b_1 = \$8, Z_1 = (1, 2, 3), b_2 = \$5, and Z_2 = (2, 3). Assume that we initially set the prices of all three vehicles as follows: p_1=\$10, p_2= \$8, p_3 = \$5. Then, according to Assumption 3, the overall demand for vehicle 2 is one unit, since only consumer 1 who will purchase vehicle 2. However, if the price of vehicle 2 decreases to less than \$5, the overall demand will increase to two units because the second consumer can now afford to purchase it. Similarly, if the price of vehicle 2 increases above \$8, the overall demand for this vehicle will decrease to zero.

Changes in the price of a vehicle affect not only its overall demand; they also impact overall demands for other vehicles. Suppose the price of vehicle 2 increases beyond \$8 while the prices of vehicle 1 and 3 remain constant, i.e. p_1=\$10 and p_3=\$5. Then, we not only see a decrease in the overall demand for vehicle 2 from one unit to zero, the demand for vehicle 3 will also increase to two units, since the first consumer will now substitute vehicle 3 for vehicle 2.

5 The Price Ladder

We can show that our optimization problem is NP-complete in the strong sense. To obtain a tractable solution, we use the concept of a price ladder -- an ordering of prices among vehicles, leading to the following assumption.

Assumption 4 (Price Ladder): We have *a priori* an ordering of prices, say $p_1 \le p_2 \le \cdots \le p_n$.

The constraint on the ordering of prices arises naturally in the automotive industry. 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. For our pricing strategy to make sense, the L300 should have a higher price than the L200, whose price should exceed that of the L100.

A similar ordering of prices exists for vehicles from different models. For instance, the price of a Prizm -- a compact passenger car -- should always be less than the price of a Corvette – a high-end sports car.

Our approach assumes that we have a complete ordering of prices among vehicles. In practice, many such orderings exist, and we can search these price ladders, applying our optimization techniques to each. Moreover, we can use an ordering implied by the current price of the vehicles as a starting point.

Given the price ladder, we then optimize the same objective function subject to a price ladder constraint. If we let $D_n = \{ p \in R_+^n : p_1 \le \cdots \le p_n \}$, it turns out that our objective function

$$R(p) = \sum_{j=1}^{M} R_j(p),$$

when restricted to the domain D_n, possesses a special property that enables us to obtain an algorithm for approximating the revenue-maximizing price. For more details on the algorithm, the reader is referred to [Rusmevichientong et al., 2002].

6 Experimental Results

In this section, we describe the dataset and the methodologies used in our analysis.

6.1 Dataset

The dataset used in our analysis consists of 83,813 consumers, representing visitors to the ACA web site from 1/1/02 until 9/30/02. We only consider those visitors who spent at least 3 minutes at the web site and specified a budget of $60,000 or less. We impose these constraints in order to rule out consumers who may not be truly interested in purchasing vehicles.

We divide the dataset 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 dataset can be divided into three categories[2]:

[2] Throughout this paper, each consumer refers to a record that is generated by the ACA web site during a visit. We assume that the record reflects an individual's preference for a vehicle she is currently shopping for. However, it is possible that the individual owns multiple vehicles, and the record only reflects her preference for a vehicle of a particular type (e.g. the one that she uses for commuting from home to work). The Consumer Category, therefore, really represents what's in her current recommendation list.

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 customers as "non-GM Only".

Recall that the competitors' prices are assumed to be fixed. Hence, we can identify affordable non-GM vehicles. Moreover, under Assumption 3, 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 2 shows the number of consumers in each category and their average budgets in the validation sample[3].

Consumer Category	Number of Consumers	% of Total	Avg. Budget	Std 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

Table 2: The number of consumers in each category in the validation sample.

Although each consumer’s recommended list typically has ten recommendations, we use only the top five vehicles because we feel that these vehicles most accurately reflect the consumer preference and budget constraint. 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.

6.2 Methodologies

The experiment compares three pricing policies: a "rounded MSRP", a greedy policy, and the coordinated pricing policy generated by the algorithm discussed in Section 5. The "rounded MSRP" of each vehicle corresponds to the existing MSRP (manufacturer's suggested retail price)

[3] These average budgets are based on these web consumers’ inputs and do not necessarily reflect the budgets of actual consumers in the market.

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 web site 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.

Although we use rounded MSRPs to determine total sales, 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 which the i^{th} vehicle appears in the recommended list. The greedy price of the i^{th} vehicle is the smallest price that maximizes revenue from these consumers, assuming that they will purchase the i^{th} vehicle if its price meets their budgets.

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.

6.3 Results and Discussion

The results from Table 3 and Table 4 show the average price of GM vehicles, the sales volume[4], and the total revenue generated under the three pricing policies, when applied to the training and the validation samples, respectively. We see from the table that our coordinated pricing policy yields a 6% improvement in revenue over the rounded MSRP and about 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 sales for each pricing policy, we compute the revenue generated from each consumer in the validation sample under the pricing policy. By definition, consumers in the "non-GM Only" category will not generate any revenue for GM, regardless of the pricing policy. Since we want to determine if our coordinated pricing policy yields a significant increase in sales, we focus only on those consumers in the validation sample who are either in the "GM Only" or the "Both GM and non-GM" category.

Once we determine the revenue generated from the consumers in these two categories, we then estimate the sample standard deviation of these revenues and multiply the result by $\sqrt{M}$, where M denotes the total number of consumers in both categories.

[4] 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.

Pricing Policy	Avg. Price	Vol (units)	Sales (millions)	Std (millions)
Rounded MSRP	$28,432	17,094	$400	$1.73
Greedy	$30,668	18,280	$416	$1.58
Coordinated	$28,807	18,386	$422	$1.73

Table 3: Performance of various pricing policies on the training sample.

Pricing Policy	Avg. Price	Vol (units)	Sales (millions)	Std (millions)
Rounded MSRP	$28,432	17,016	$399	$1.73
Greedy	$30,668	18,151	$414	$1.58
Coordinated	$28,807	18,243	$421	$1.74

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

6.3.1 <u>Deviation from Optimality</u>

In this section, we wish to assess the difference between the revenue generated under the coordinated pricing policy and the maximum revenue R^* that we can generate from all consumers in the training sample, measuring an effectiveness of our algorithm. Recall from the beginning of Section 6 that, for our optimization, we only need to consider those consumers who are either "GM Only" or "Both GM and non-GM". Thus, R^* is bounded above by the sum of the budgets of these consumers, which turns out to be $594,634,000. Thus, we have

$$R^* \leq \$594{,}634{,}000 .$$

Since our coordinated pricing policy yields a total sales of $422,351,000 (see Table 3) and

$$\frac{\$422{,}351{,}000}{\$594{,}634{,}000} \geq 0.7103 ,$$

this implies that the revenue from our coordinated pricing differs from the optimal by at most 29%.

By using a more detailed analysis, we can significantly improve the above performance bound. It turns out that, during the course of computing the coordinated pricing policy, our algorithm also produces two price vectors $p^L \in D_n$ and $p^U \in D_n$ that *provably* form lower and upper bounds for the optimal price $p^* \in D_n$, i.e. $p_i^L \leq p_i^* \leq p_i^U$ for all $i = 1, 2, \ldots n$ [see Rusmevichientong et al., 2002]. We will use this information to improve our estimate of the maximum revenue that can be extracted from all consumers.

For any price vector $p \in D_n$ and consumer j, let $R_j(p)$ denote the revenue from consumer j under price p. Then, we have the following inequalities

$$R^* = \max_{p \in D_n:\, p^L \le p \le p^U} \sum_{j=1}^{M} R_j(p) \le \sum_{g=1}^{G} \max_{p \in D_n:\, p^L \le p \le p^U} \sum_{j \in P_g} R_j(p)$$

where $P_1, P_2, \ldots, P_G$ corresponds to *any* partition of the consumers in mutually exclusive and totally exhaustive groups.

We obtain an improved bound on the optimal revenue by judiciously selecting a partition of the consumers. Among consumers who have at most two vehicles in their (revised) recommended lists, we partition them based on the first and second ranked vehicles. Those with three or more vehicles in their lists, we group them based on the first and second ranked vehicles, along with their budgets. This procedure results in a total of 2,933 distinct groups of consumers, with an average of 6.9 consumers per group. Note that consumers who do not purchase any GM vehicle under p^L were removed from consideration since they will never purchase any GM vehicle under the optimal price p^*.

Since the number of consumers in each group is relatively small, we can apply exhaustive search to determine the maximum revenue that we can obtain from each group. It turns out that, under this partition,

$$\sum_{g=1}^{G} \max_{p \in D_n:\, p^L \le p \le p^U} \sum_{j \in P_g} R_j(p) = \$477{,}816{,}000$$

which implies that $R^* \le \$477{,}816{,}000$. Since our coordinated pricing policy yields a total revenue of \$422,351,000 and

$$\frac{\$422{,}351{,}000}{\$477{,}816{,}000} \ge 0.8839$$

it follows that that the revenue from our coordinated pricing actually differs from the optimal by at most 12%, a significant improvement over the previous error bound.

The error bound of 12% provides a conservative and worst-case estimate of the difference between the optimal revenue and the revenue under our coordinated pricing policy. From our experience, however, the difference between these two quantities in practice will be much smaller. Since we cannot compute the optimal revenue exactly, this error bound serves as a rough approximation on the maximum obtainable increase in revenue over the coordinated pricing policy.

6.3.2 Coordinated Pricing and Substitution Effects

From Table 4, we see that the average price assigned by the coordinated pricing policy is about \$400 higher than the average MSRP. Despite this increase in the average price, the coordinated pricing policy yields higher unit sales and revenues than the rounded MSRP. This seems surprising since conventional wisdom suggests that an increase in price should lead to a decrease in demand.

Figure 1 provides a partial explanation. The figure compares the rounded MSRP and the price generated by our coordinated pricing policy for all 301 GM vehicles. For each vehicle, we plot its rounded MSRP and the relative difference between its price under the coordinated pricing policy and its rounded MSRP. Despite the increase in the average price, the coordinated pricing policy assigns a lower price to some GM vehicles, notably those vehicles with MSRP less than $18,000. The increase in the average price comes primarily from expensive vehicles, especially those vehicles with MSRP higher than $38,000.

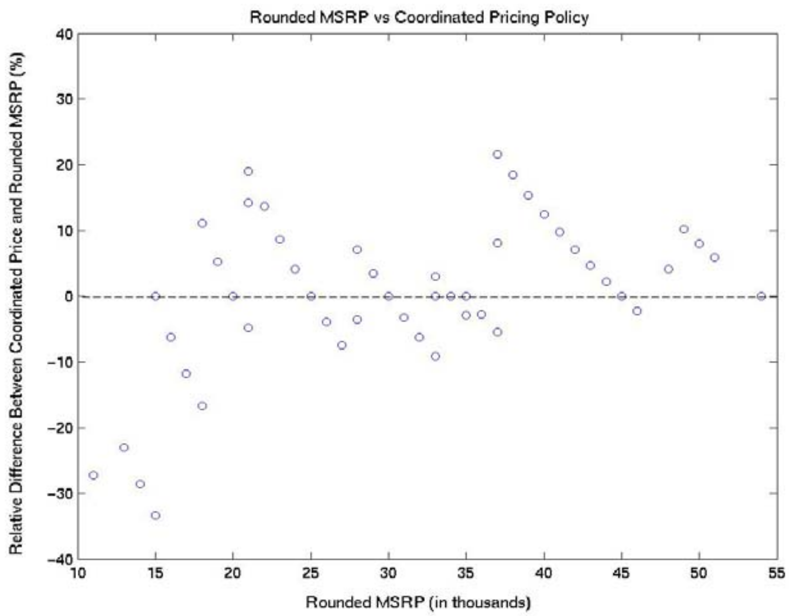

Figure 1: For each GM vehicle, we plot its rounded MSRP and the relative difference between its price under the coordinated pricing policy and its rounded MSRP. Each circle corresponds to a GM vehicle. However, since we consider rounded MSRPs, many vehicles end up having the same price.

Figure 1 suggests that by strategically increasing the price of certain vehicles, we can increase the sales volumes even when the average price increases. Substitution between GM vehicles provides another explanation for this phenomenon, as demonstrated by the following example.

Consider Table 5 which shows a hypothetical consumer with a budget of $18,000 and a list of recommended vehicles that consists of the Chevy Prizm LSi, the Saturn L100, the Dodge Stratus SXT Sedan, the Chevy Malibu Sedan, and the Mitsubishi Lancer LS. The first, second, and fourth vehicles in the list represent GM vehicles. According to our assumption, the consumer will purchase the highest ranked vehicle that she can afford. So, if we increase the price of the Chevy Prizm beyond $18,000, the consumer will substitute the Saturn for the Prizm, provided that the price of the Saturn remains below the budget.

ID	Budget	Recommended Vehicles	Rank

		Chevy Prizm LSi (GM)	1st
		Saturn L100 (GM)	2nd
1002	$18,000	Dodge Stratus SXT Sedan	3rd
		Chevrolet Malibu Sedan (GM)	4th
		Mitsubishi Lancer LS	5th

Table 5: An example that illustrates how substitutions between vehicles enable us to obtain a more effective pricing strategy. See Section 6.3.2 for more details.

Of course, in practice we need to consider the preferences and budget constraints of all consumers simultaneously. Nonetheless, this simple example illustrates how substitutions between vehicles enable us to strategically increase the price of certain vehicles without sacrificing sales.

6.3.3 Market Share

Table 6 shows the average price, the sales volume, and the total revenue for both GM and non-GM vehicles under the two pricing policies. The table shows that the increase in revenue under the coordinated pricing policy comes from both the competitors' reduction in sales volumes and the increase in new sales.

Vehicle Type	Rounded MSRP			Coordinated		
	Avg. Price	Vol (units)	Sales (millions)	Avg. Price	Vol (units)	Sales (millions)
GM	$28,432	17,016	$399	$28,807	18,243	$421
non-GM	$26,363	19,157	$406	$26,363	18,897	$403

Table 6: Sales of GM and non-GM vehicles under the rounded MSRP and the coordinated pricing policy in the validation sample.

Interestingly, the increase in sales volume and revenue only happens to selected vehicles. Table 7 shows the sales of different GM brands. Some brands, notably Buick, GMC, and Oldsmobile, actually lose sales volumes under the coordinated pricing policy. In fact, GMC and Oldsmobile even show a decrease in total revenue.

Brand	Rounded MSRP			Coordinated		
	Avg. Price	Vol (units)	Sales (millions)	Avg. Price	Vol (units)	Sales (millions)
Buick	$28,778	2,676	$63	$29,889	2,566	$66
Cadillac	$46,667	159	$7	$48,667	153	$7
Chevy	$27,752	7,685	$192	$28,051	8,097	$204
GMC	$29,541	1,469	$45	$29,703	1,408	$44
Oldsmobile	$26,611	290	$7	$27,222	267	$6

Pontiac	$24,724	989	$22	$24,759	1,066	$23
Saab	$36,222	82	$3	$39,000	85	$3
Saturn	$17,615	3,666	$62	$16,769	4,601	$68

Table 7: Sales among GM brands under the rounded MSRP and the coordinated pricing policy in the validation sample.

Even for brands that show a significant increase in total revenue, only specific models produce increased revenue. For instance, under the coordinated pricing policy, the Saturn increases its sales volume by 935 units and its revenue by $6 million. As demonstrated in Table 8, these increases occur primarily within the L-Series Sedan model and the S-Series model. The 3-D Coupe actually experiences a decline in total revenue, while the L-Series Wagon and the Vue see a decrease in volume[5].

Model	Rounded MSRP			Coordinated		
	Avg. Price	Vol (units)	Sales (millions)	Avg. Price	Vol (units)	Sales (millions)
3-D Coupe	$15,500	238	$4	$12,500	295	$3
L-Series Sedan	$18,667	1,787	$31	$20,000	2,390	$37
L-Series Wagon	$21,500	96	$2	$22,500	93	$2
S-Series	$12,667	576	$6	$9,333	864	$7
Vue	$20,333	969	$19	$20,000	959	$18

Table 8: Sales among different models of the Saturn family under the rounded MSRP and the coordinated pricing policy in the validation sample.

6.3.4 Sales by Budgets

Table 9 relates the sales volumes under the two pricing policies with consumers' budgets. The table shows that the coordinated pricing policy increases sales by targeting those consumers with low budgets. Under the coordinated pricing policy, the sales of GM vehicles to consumers with a budget of $20,000 or less increase by 1,251 units. This increase comes at the expense of a decrease in margin, as indicated by the reduction in the average price paid by those consumers with a budget between $10,000 and $20,000 (see Table 10).

Consumer Budget (in thousands)	Total Consumers	Vehicles Purchased Under Rounded MSRP (units)			Vehicles Purchased Under Coordinated (units)		
		GM	non-GM	None	GM	non-GM	None
≤$10	3,499	-	466	3,033	485	412	2,602
$11-$20	10,759	3,063	5,528	2,168	3,829	5,279	1,651
$21-$30	15,223	7,161	7,590	472	7,147	7,604	472

[5] These examples only illustrate the methodology and are not meant to represent actionable results.

$31-$40	7,382	3,739	3,617	26	3,734	3,641	7
$41-$50	3,971	2,378	1,592	1	2,373	1,597	1
$51-$60	1,039	675	364	-	675	364	-
Total	41,873	17,016	19,157	5,700	18,243	18,897	4,733

Table 9: Sales volumes among consumers with different budgets in the validation sample.

Interestingly, the coordinated pricing policy and the rounded MSRP have roughly the same sales volumes among consumers with a budget greater than $20,000. However, as indicated by Table 10, the consumers in this budget range, on average, pay a higher price under the coordinated pricing policy.

The results in Table 9 and Table 10 show that our coordinated pricing policy increases sales by increasing market share among consumers with low budgets through a decrease in price. At the same time, to maintain profit margin, the coordinated pricing policy increases the average price paid by consumers with higher budgets. This result is consistent with the literature on price discrimination, since 3rd degree price discrimination allows the seller to include more customers by charging lower prices for low-budget customers and higher prices for those with higher willingness-to-pay [Luenberger, 1995 and Wilson 1993].

Consumer Budget (in thousands)	Total Consumers	Avg. Price Paid by Consumers Under Rounded MSRP			Avg. Price Paid by Consumers Under Coordinated		
		GM	non-GM	none	GM	non-GM	none
≤$10	3,499	N/A	$9,813	N/A	$8,441	$9,806	N/A
$11-$20	10,759	$16,014	$14,734	N/A	$14,768	$14,762	N/A
$21-$30	15,223	$21,256	$21,297	N/A	$21,792	$21,271	N/A
$31-$40	7,382	$27,117	$27,480	N/A	$27,596	$27,505	N/A
$41-$50	3,971	$30,937	$29,162	N/A	$32,678	$29,178	N/A
$51-$60	1,039	$33,701	$33,203	N/A	$36,027	$33,203	N/A
Total	41,873	$23,447	$21,171	N/A	$23,093	$21,302	N/A

Table 10: Average price paid by consumers with different budgets in the validation sample.

6.3.5 <u>Conquest Sales</u>

Recall from Table 2 that our dataset consists of three different categories of consumers: "GM Only", "Both GM and non-GM", and "non-GM Only". Table 11 shows the sales volumes and revenue of GM and non-GM vehicles among the three categories, under the two pricing policies.

Consumer Category	Rounded MSRP		Coordinated	
	GM	non-GM	GM	non-GM

	Vol (units)	Sales (mil)	Vol (units)	Sales (mil)	Vol (units)	Sales (mil)	Vol (units)	Sales (mil)
GM Only (13,116 consumers)	12,252	$293	N/A		12,592	$305	N/A	
Both GM and non-GM (8,847consumers)	4,764	$106	920	$16	5,651	$116	660	$13
Non-GM Only (19,910 consumers)	N/A		18,237	$390	N/A		18,237	$390

Table 11: Sales among different groups of consumers in the validation sample.

Since we cannot influence consumers who are in the "non-GM Only" category, an important measure that determines the effectiveness of our coordinated pricing policy is the total conquest sales -- the additional increase in sales of GM vehicles to consumers who are in the "Both GM and non-GM" category. The results from Table 11 suggest that the coordinated pricing policy yields an increase in such sales. This conclusion is further supported by Table 12 which shows that the coordinated pricing policy generates sales from those consumers who would have bought non-GM vehicles under the rounded MSRP. In Table 12, we divide those consumers who consider both GM and non-GM vehicles into nine categories, depending on their purchasing tendencies under the two pricing policies. For each policy, the consumer either purchases a GM vehicle, a non-GM vehicle, or neither.

We see from Table 12 that, among 920 consumers who would buy non-GM vehicles under the rounded MSRP, 372 of these consumers (approximately 40% conquest rate) would switch to GM vehicles under the coordinated pricing policy. Moreover, our coordinated pricing policy generates 666 new sales from a total of 3,163 consumers who did not purchase any vehicles under the rounded MSRP, approximately 21% of the total. The increase in new sales under the coordinated pricing policy comes at the expense of 151 lost sales, 112 of which go to our competitors, while the remaining 39 consumers can no longer afford any vehicle.

Vehicles Purchased Under Rounded MSRP	Vehicles Purchased Under Coordinated			Total
	GM	non-GM	None	
GM	4,613	112	39	4,764
non-GM	372	548	-	920
None	666	-	2,497	3,163
Total	5,651	660	2,536	8,847

Table 12: Sales among those consumers who are in the "Both GM and non-GM" category in the validation sample.

6.4 Limitations and Possible Improvements

The results in Section 6.3 demonstrate the performance of our coordinated pricing policy, provide insights into the current pricing practice, and indicate opportunities for improvements. Despite these promising results, our analysis has its limitations, which need to be addressed before this research can be applied in industrial settings. This section outlines the shortcomings and identifies means to remedy them.

6.4.1 The ACA Web Site's Design

The ACA web site tracks each visitor using a "session ID" -- a text file generated by the web site that uniquely identifies each visitor during her session. These session IDs are temporary, and they expire after 25 minutes. Thus, each record in our dataset represents a unique consumer within a twenty-five-minute time window. If the same consumer returns to the site after 25 minutes, she will be issued a new session ID and treated as a new consumer. As a result, the dataset may contain multiple records for the same consumer, creating distortions in the dataset.

Using a permanent cookie may alleviate this problem. As long as the user does not delete it, the permanent cookie will reside on her computer and will never expire, providing a unique identifier when she returns to the web site. The ACA web site does not use permanent cookies in order to maintain a trusted environment. (The privacy policy on the web site provides more details.)

Another approach to this problem is to entice each consumer to create a login name, enabling us to authenticate each user. The web site already allows each consumer to create a login name and save her recommended list for future visits. The majority of consumers, however, do not use this feature.

Creating login names and requiring consumers to provide their personal information will also help us track their purchase records. Currently, all personal automobile purchases in the United States must be done through the dealers. Thus, in the absence of each visitor's personal information, we cannot determine if she has purchased any vehicle from her recommended list. The absence of any purchase record forces us to make assumptions on consumer purchasing tendencies, as in Assumption 3. Such an assumption may not truly reflect the consumer behavior.

Instead of asking consumers to provide their personal information, with a site redesign, we may be able to obtain information on consumer purchasing tendencies through less intrusive means. After showing each consumer her list of recommended vehicles, we may ask her to select a vehicle that she would most likely purchase, if any. The information generated from this interaction can potentially provide us with a proxy for consumer purchasing behavior.

6.4.2 Bias in Web Data

Section 6.4.1 identifies different measures that we can take to ensure that our dataset truly reflects the information provided by each visitor to the web site. However, we still need to ensure that these visitors are representative of the entire population of potential automotive buyers.

Research has indicated that consumers who use the Internet to research their vehicle purchases exhibit different demographic and personal characteristics than the general population, creating an inherent bias in our dataset. For a meaningful analysis, we must correct for this bias, mitigating its impact. We do not currently have a systematic approach to such a correction, and this represents an ongoing research topic.

Despite these difficulties, the report by J.D. Power [Power, 2000] -- an automotive consultancy -- provides hope. The report estimated that, in 1999, over 40% of all new vehicle buyers used the Internet during their shopping, and this number was projected to grow to 55% in 2000. We believe that this number will continue to increase in the future.

6.4.3 Models of Consumers' Selections

Our formulation assumes that each consumer will purchase the highest ranked vehicle that she can afford (Assumption 3). We feel that this assumption provides an intuitive approximation, forming the basis for more complex models. Even with this assumption, our analysis should provide useful insights to the current pricing practice. If we have purchase records, we can incorporate such information and consider more complex models of consumer behavior.

One can also consider other formulations based on this data set. For instance, we may assume that the consumer will select from the recommended list the cheapest, or alternatively the most expensive, product that she can afford. For treatments of such formulations, see [Rusmevichientong et al., 2002].

6.4.4 Production Constraints

Our formulation assumes that we can produce all of the vehicles needed, imposing no constraint on production capacities. Our analysis only focuses on understanding how we can optimize for the vehicle prices given the consumers' preferences and budget constraints. It should be interesting to explore how we can integrate capacity constraints into our formulation.

6.4.5 Use of MSRP

Our analysis compares the coordinated pricing policy with the rounded MSRP. This comparison might be misleading since the MSRP does not represent the actual transaction price of each vehicle, since it depends on many factors such as financing options and promotional incentives. Unfortunately, we do not have access to the transaction price of each vehicle. However, if this information is available, we can incorporate it into our analysis, by replacing the MSRP with the corresponding transaction price.

In addition to the vehicle's price, incentives and financing options can potentially influence the consumer's purchasing decision. These factors are not included in our current formulation. Exploring how we can incorporate such information into our formulation is a topic for future research.

6.4.6 Competition

Our analysis assumes that competitors will not respond to our pricing decisions. Consequently, the increase in sales under the coordinated pricing policy may be overstated. Competitors will invariably respond to our price changes, lessening the effectiveness of our pricing strategy, and ultimately reducing the potential gains in new sales for GM.

Our future research will focus on understanding the effect of competition on our pricing strategy, and how we should price the vehicles in such an environment.

6.4.7 Incorporating Data from Other Market Research Studies

Our approach to pricing relies exclusively on consumers' preference data generated from the ACA web site. However, GM also has other sources of preference data, such as the results from the physical property clinics discussed in Section 2. Incorporating such data into our pricing models represents an interesting research direction.

7 Conclusions

Using the consumer preference data collected from the ACA web site, we developed a non-parametric approach to pricing GM vehicles. 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.

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 customize our product prices to meet consumer needs. Our analysis uses the data to determine a pricing strategy for GM. It should be interesting to explore other avenues for using this type of data.

ACKNOWLEDGMENTS

This research was supported by General Motors through the GM-Stanford Collaborative Research Lab on Work Systems, by a Stanford Graduate Fellowship, by the National Science Foundation, under NSF CAREER Grant ECS-9985229, and by the Office of Naval Research, under grant MURI N00014-00-1-0637. We also thank our GM colleagues Steve Holland, Jay Ostahowski, Charlie Rosa, and Jeff Tew for helpful comments and suggestions.

REFERENCES

[Brynjolfsson et al., 2003] E. Brynjolfsson, Y. (Jeffrey) Hu, M. D. Smith, “Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers,” *Management Science*, vol. 49, no. 11, pp. 1580-1596, November 2003.

[Brynjolfsson et al., 2004] E. Brynjolfsson, A. A. Dick, M. D. Smith, “Search and Product Differentiation at an Internet Shopbot,” working paper, 2004.

[Chen et al., 1998] F. Chen, J. Eliashberg, P. Zipkin, “Customer Perferences, Supply-Chain Costs, and Product-Line Design,” *Product Variety Management: Research and Advances* (T.-H. Ho and C.S. Tang, eds.), Kluwer Academic Publisher, November 1998, pp. 143-144.

[Chen and Hausman, 2000] K. D. Chen and W. H. Hausman, “Technical Note: Mathematical Properties of the Optimal Product Line Selection Problem Using Choice-Based Conjoint Analysis,” *Management Science*, vol. 46, no. 2, pp. 327-332, 2000.

[Chevalier and Goolsbee, 2003] J. Chevalier and A. Goolsbee, “Measuring Prices and Price Competition Online: Amazon and Barnes and Noble,” *Quantitative Marketing and Economics*, 1:203-222.

[Dobson and Kalish, 1988] G. Dobson and S. Kalish, "Positioning and Pricing A Product Line," *Marketing Science*, vol. 7, no. 2, pp.107-125, 1988.

[Dobson and Kalish, 1993] G. Dobson and S. Kalish, "Heuristics for Pricing and Positioning a Product-Line Using Conjoint and Cost Data," *Management Science*, vol. 39, no. 2, pp. 160-175, February 1993.

[Green and Krieger, 1985] P. E. Green and A. M. Krieger, "Models and Heuristics for Product Line Selection," *Marketing Science*, vol. 4, no. 1, pp. 1-19, 1985.

[Green, Krieger, and Wind, 2001] P. E. Green, A. M. Krieger, and Y. Wind, "Thirty Years of Conjoint Analysis: Reflections and Prospects," *Interfaces*, 31 (3), S56-S73.

[Guadagni and Little, 1983] P. M. Guadagni and J. D. C. Little, "A Logit Model of Brand Choice Calibrated on Scanner Data," *Marketing Science*, vol. 2, no. 3, pp. 203-238, Summer 1983.

[Hanson and Martin, 1990] W. Hanson and K. Martin, "Optimal Bundle Pricing," *Management Science*, vol. 36, pp. 155-174, 1990.

[Jedidi and Zhang, 2002] "Augmenting Conjoint Analysis to Estimate Consumer Reservation Price," *Management Science*, vol. 48, no. 10, pp. 1350-1368, October 2002.

[Luenberger, 1995] D. G. Luenberger, *Microeconomic Theory*, McGraw-Hill, 1995.

[McBride and Zufryden, 1988] R. D. McBride and F. S. Zufryden, "An Integer Programming Approach to the Optimal Product Line Selection Problem," *Marketing Science*, vol. 7, no. 2, pp. 126-140, 1988.

[Mirman and Sibley, 1980] L. J. Mirman and D. Sibley, "Optimal Nonlinear Prices for Multiproduct Monopolies," *Bell Journal of Economics*, vol. 11, pp. 659-670, 1980.

[Oren et al., 1984] S. Oren, S. Smith, and R. Wilson, "Product Line Pricing," *Journal of Business*, vol. 57, no. 1, pt. 2, pp. S73-S99, 1984.

[Oren et al., 1987] S. Oren, S. Smith, and R. Wilson, "Multi-product Pricing for Electric Power," *Energy Economics*, pp. 104-114, April 1987.

[Power, 2000] J. D. Power Automotive and Associates, "1999 New Autoshopper.com Study," *The Power Report*, Agoura Hills, CA, January 2000.

[Rusmevichientong, 2003] P. Rusmevichientong, "A Non-Parametric Approach to Multi-Product Pricing: Theory and Application," Ph.D. dissertation, Stanford University, 2003.

[Rusmevichientong et al., 2002] P. Rusmevichientong, B. Van Roy, and P. W. Glynn, "A Non-Parametric Approach to Multi-Product Pricing," to appear in *Operations Research*, 2002.

[Smith, 1986] S. Smith, "New Product Pricing in Quality Sensitive Markets," *Marketing Science*, vol. 5, no. 1, pp. 70-87, 1986.

[Smith and Brynjolfsson, 2001] M. D. Smith and E. Brynjolfsson, “Consumer Decision-making at an Internet Shopbot: Brand Still Matters,” *Journal of Industrial Economics*, vol. 49, no. 4, pp. 541-558.

[Spence, 1980] A. M. Spence, "Multi-Product Quantity-Dependent Prices and Profitability Constraints," *Review of Economic Studies*, vol. 47, pp. 821-841, 1980.

[Wilson, 1993] R. R. Wilson, *Nonlinear Pricing*, Oxford University Press, Oxford, 1993.

[Yano and Dobson, 1998] C. A. Yano and G. Dobson, "Profit Optimizing Product Line Design, Selection and Pricing with Manufacturing Cost Considerations: A Survey," *Product Variety Management: Research and Advances* (T.-H. Ho and C.S. Tang, eds.), Kluwer Academic Publisher, November 1998, pp. 145-176.