MARKETING SCIENCE INSTITUTE

MSI Reports

Developing Optimal Store-level Pricing Strategies for an Automotive Aftermarket Retailer (06-105)

Murali K. Mantrala, P. B. (Seethu) Seetharaman, Rajeeve Kaul, Srinath Gopalakrishna, and Antonie Stam

2006

WORKING PAPER SERIES

ISSUE ONE

NO. 06-001

MSI Reports

2006

WORKING PAPER SERIES

ISSUE ONE

NO. 06-001

The Marketing Science Institute supports academic research for the development—and practical translation—of leading-edge marketing knowledge on issues of importance to business performance. Topics are identified by the Board of Trustees, which represents MSI member corporations and the academic community. MSI supports academic studies on these issues and disseminates findings through conferences and workshops, as well as through its publications series.

Marketing Science Institute
1000 Massachusetts Avenue
Cambridge, MA
02138-5396

Phone: 617.491.2060
Fax: 617.491.2065
www.msi.org

Working Paper Series
The articles that appear in *MSI Reports* have not undergone a formal academic review. They are released as part of the MSI Working Paper Series, and are distributed for the benefit of MSI corporate and academic members and the general public.

Subscriptions
Annual subscriptions to *MSI Reports* can be placed online at www.msi.org. Questions regarding subscriptions may be directed to pubs@msi.org.

Single reports
Articles in *MSI Reports* are available in downloadable (PDF) format at www.msi.org.

Past reports
MSI working papers published before 2003 are available as individual hard-copy reports; many are also available in downloadable (PDF) format. To order, go to www.msi.org.

Corporate members
MSI member company personnel receive all MSI reports (PDF and print versions) free of charge.

Academic members
Academics may qualify for free access to PDF (downloadable) versions of MSI reports and for special rates on other MSI print publications. For more information and to apply, go to "Academic and Quantity Discounts" in the Publications section of www.msi.org.

Classroom use
Upon written request, MSI working papers may be copied for one-time classroom use free of charge. Please contact MSI to obtain permission.

Search for publications
See the searchable publications database at www.msi.org.

Submissions
MSI will consider a paper for inclusion in *MSI Reports*, even if the research was not originally supported by MSI, if the paper deals with a priority subject, represents a significant advance over existing literature, and has not been widely disseminated elsewhere. Only submissions from faculty members or doctoral students working with faculty advisors will be considered. "MSI Working Paper Guidelines" and "MSI 2004-2006 Research Priorities" are available in the Research section of www.msi.org.

Publication announcements
To sign up to receive MSI's electronic newsletter, go to www.msi.org.

Change of address
Send old and new address to pubs@msi.org.

Developing Optimal Store-level Pricing Strategies for an Automotive Aftermarket Retailer

Murali K. Mantrala, P. B. (Seethu) Seetharaman, Rajeeve Kaul, Srinath Gopalakrishna, and Antonie Stam

Pricing research has largely focused on consumer packaged goods and grocery retailing, overlooking retail sectors with different demand environments. Using a national database, this study develops a micromarketing pricing model for the aftermarket automotive industry.

Murali K. Mantrala is Sam M. Walton Distinguished Professor of Marketing at the College of Business, University of Missouri-Columbia.
P.B. Seetharaman is Associate Professor of Management at the Jesse H. Jones Graduate School of Management, Rice University.
Rajeeve Kaul is Director of Pricing Optimization at AutoZone.
Srinath Gopalakrishna is Associate Professor of Marketing and
Antonie Stam is Leggett & Platt Distinguished Professor of Information Systems, both at the College of Business, University of Missouri-Columbia.

Report Summary

Much of extant retail pricing research has focused on grocery retailers, neglecting other major retailing sectors that face very different demand environments. One such sector is automotive aftermarket retailing, in which highly specialized “hard parts” are stocked, the vast majority of which sell five or fewer units a year.

In this paper, authors Mantrala, Seetharaman, Kaul, Gopalakrishna, and Stam create a model for optimal store-level pricing of specialty hard parts for a large U.S. specialty automotive aftermarket chain. They use two years’ worth of weekly store-level sales data from 800 stores located across the U.S. to formulate and estimate a heterogeneous multinomial logit model of store-level demand for three quality variants (good, better, and best) for 27 subclasses of hard parts.

The model offers insights into store-level SKU price elasticities and generates profit-maximizing prices for the three quality variants in the multiple subclasses. This price optimization is done separately for each store in the sample, meeting the collaborating company’s stated goal of micromarketing store-level SKU prices.

The team finds that in many stores, the company would benefit from lowering its prices, while in a smaller fraction of stores, the retailer would benefit from raising the prices. The team also finds that the retailer would benefit from reducing the size of the gap between the prices of its “good” and its “best” variants.

Contrary to findings in the packaged-goods industry, they find that changing the price of the highest-quality variant affects the demand for a lower-quality brand significantly less than the amount by which changing the price of the lower-quality variant affects the demand for the highest-quality variant, a difference that may be due to the durable nature of automotive hard parts. ■

Introduction

Many retailers' category management efforts center on optimal retail category pricing (i.e., pricing of brands or variants in different product categories) and, accordingly, retail category pricing has received considerable research attention in recent years (e.g., Zenor 1994; Kim, Blattberg, and Rossi 1995; Montgomery and Bradlow 1999; Anderson and Vilcassim 2001; Basuroy, Mantrala, and Walters 2001; Chintagunta 2002). A convergence of significant advances in both the science (e.g., statistical modeling and estimation of consumer price response functions) and technology (computing power, data availability) of marketing management has made smarter pricing a critical success factor in today's retailing environment (e.g., Montgomery 2004; Shankar and Bolton 2004). The grocery retail sector in particular has benefited from improvements in pricing decision support systems that have been fueled by the widespread availability of consumer transaction data collected via optical bar code scanners. Two important trends in this context are micromarketing and demand-based management, which emphasize tailoring supermarket retail chain prices to each individual store's clientele, environment, and performance conditions (e.g., Hoch et al. 1995; Shankar and Krishnamurthi 1996; Montgomery 1997; Chintagunta, Dube, and Singh 2003).[1] Micromarketing represents an interest on the part of managers in combining the advantages of large operations with the flexibility of independent, "customized" neighborhood stores.

So far, however, few retail micromarketing model applications have been reported in the literature. Further, extant research has focused on grocery retailers, largely overlooking other major retailing sectors with different demand environments. One such sector is automotive aftermarket retailing that caters to the DIY ("do-it-yourself") segment—consumers who maintain and repair their vehicles themselves. The 2002 *Aftermarket Factbook* (published by the Automotive Aftermarket Industry Association, or AAIA) reported that the DIY category witnessed a 5% compounded annual growth from approximately $22 billion in 1991 to $35 billion in 2001.

The objective of the present research is to develop (and apply) an optimal category-pricing solution for a specialty automotive aftermarket retailer with more than 3,000 stores in the United States. The solution will include an approach for customizing prices across these stores. This research both extends extant retail pricing research to an important retailing sector and expands current knowledge regarding the determinants and pricing implications of store-level price elasticities of multiple variants in product categories.

Automotive Aftermarket Industry

The automotive aftermarket industry supplies replacement parts (excluding tires), accessories, maintenance items, batteries, and automotive fluids for cars and light trucks (pickup trucks, vans, minivans, and sport utility vehicles). Replacement parts include radiators, brake pads, fan belts, starters, alternators, batteries, shock absorbers, spark plugs, transmission parts, clutches, engines, and transmissions.

Aftermarket segments

The automotive aftermarket is composed of two consumer segments: DIY and DIFM ("do-it-for-me"). Our research focuses on retailers serving the DIY segment. In *The Aftermarket Consumer: Do-It-Yourself or Do-It-For-Me* (2005), the AAIA reports that 70% of U.S. households have someone who has personally performed vehicle maintenance or repair in 2004, and 62% of DIYers are doing as much or more work now than they were five years ago. The primary motivation for DIYers is saving money. This motivation has been a major driver of the automotive aftermarket's growth in recent times, which have been characterized by (1) an economic downturn, which leads to longer vehicle retention and a greater number of cars coming off warranty and needing more repairs; (2) the threat of terrorism and other factors increasing the number of miles driven annually; (3) an increase in the cost of replacement parts as a result of technological changes in recent models of vehicles; and (4) an increasing number of light trucks and sport utility vehicles, which require more expensive parts, resulting in higher average sales per customer (Automotive Aftermarket Industry Association 2002).

Automotive market structure

Traditionally, the retailing of automotive parts in the DIY category has been highly fragmented, with competition occurring between multiple retailing formats, including national and regional automotive parts specialty retailers, wholesalers or jobbers, independent operators, automobile dealers that supply parts, and discount stores and mass merchandisers that carry automotive products. However, as automotive replacement parts have proliferated, discount stores and mass merchandisers have found it increasingly difficult (because of their broad scope) to maintain the wide and deep selection of high-quality brand name automotive parts and accessories that DIY customers demand—branded products such as Monroe, Bendix, Purolator, and ACDelco. This has given a strong competitive advantage to specialty automotive parts retailers, who have the distribution capacity, sophisticated information systems, and knowledgeable sales staff necessary to offer a broad selection and value to DIY customers. The AAIA (2005) report indicates that from January 2004 to January 2005, 68% of DIYers purchased parts and supplies from a specialty auto parts store, with quality being the most important criterion driving store selection.

Strategies of specialty automotive aftermarket retailers

Based on interviews with our collaborative research partner, it appears that specialty automotive aftermarket retailer strategies focus on enhancing operating margins by (1) improving purchasing efficiencies with vendors; (2) streamlining supply chain and distribution network management to optimize the inventory mix and maximize distribution capacity; (3) leveraging overall scale to reduce other operating expenses as a percentage of sales; (4) implementing category management processes to customize store-level merchandise assortments, in-stock availability, merchandising, and marketing initiatives; (5) using well-trained sales associates who provide expert assistance in product selection and installation to achieve a level of customer service that outstrips that of mass merchandisers; (6) investing in fully integrated point-of-sale store-level "look-up" systems that can be used for refining cross-selling, pricing, marketing, and merchandising strategies; (7) replenishing inventory accurately and rapidly.

Specialty automotive aftermarket retailers typically carry between 10,000 and 70,000 in-store stock-keeping units (SKUs). Sales of replacement parts tend to account for a majority of an automotive aftermarket retailer's net sales and usually generate higher gross margins than maintenance items or general accessories.

How specialty automotive retailer demand compares with grocery retailer demand

From the viewpoint of demand-based management, there are three important differences between specialty automotive aftermarket retailing and grocery retailing: the product mix, the category management process, and the availability of demand data.

First, unlike grocery retailers' product assortments, those of the auto part retailer can be classified as either "front-of-the-counter" parts such as chemicals, floor mats, and mechanical security components or "behind-the-counter or "back-of-the-store" application-specific ("hard") parts, such as brakes, mufflers, alternators and so forth—items whose purchase is typically failure related. Back-of-the-store hard parts are often unique to a store. A major complication in automotive aftermarket demand assessment is that many hard parts (about 50% in the case of our collaborative research partner) move less than once a year and 90% of them move less than five times a year. Such slow-moving items are stocked because of the perceived high opportunity cost from losing a customer who has to wait for an item not in stock. Further, due to heterogeneity in SKU demand across stores, 50% of items penetrate only half the stores, requiring demand assessment to be done at the store level rather than at a store cluster level.

Second, the traditional category management process that originated in the grocery industry (e.g., Basuroy, Mantrala, and Walters 2001) does not take into account the application parts side of the aftermarket business, where the bulk of the auto parts store's inventory and capital assets are tied up. While front-of-the-counter items are relatively amenable to grocery-style retail category management, behind-the-counter or application-specific hard parts are less so because, as mentioned above, they are typically failure-related items that have limited turnover. A major objective of AAIA's Category Management Committee is to adapt the grocery model to its version of category management, called enhanced line review (Automotive Aftermarket Industry Association 2003).

Third, unlike in grocery retailing, syndicated point-of-sale data are not yet available from major auto parts retailers. Currently available mass merchandiser data are not sufficient to guide auto part retailer pricing decisions because mass merchants do not carry application products such as brakes and engine parts. Over the past few years, AAIA's Category Management Committee has worked with the NPD Group (a global sales and marketing information provider) and other companies to establish product category definitions for collecting and compiling point-of-sales data from specialty retailers, discount stores, mass merchandisers and independent operators.

AAR's background and strategy

Our collaborative research partner, whom we refer to here as AAR is a leading U.S. automotive aftermarket retailer whose annual sales are several billion dollars. AAR sells thousands of auto and light truck parts, chemicals, and accessories through more than 3,000 stores in 42 U.S. states. More specifically, AAR sells over 1,000 product categories with 0–200 items in each category per store. The bulk of AAR's business comes from DIY consumers. AAR has completed a merchandise category management study, defining product categories and their overall roles they play in the firm's total business (e.g., the engine parts category plays the role of core or destination category) and defining category hierarchies and major segments. AAR has the ability to vary prices across stores and seeks to design and implement models for the optimal pricing of its product assortments, accounting for demand heterogeneity across stores. Currently, AAR's stores are grouped into five zones, and prices vary by zone. The prices are reviewed globally once a year, but the process allows for occasional local price adjustments during the year in response to changing demand conditions or competitive pricing (e.g., the price of an item in a store may be adjusted to within ±10% of its price—or the price of a similar item—in, say, a Wal-Mart outlet in the same area).

Research Objectives and Model Development

For the present research, we sought to find the optimal micromarketing pricing for one category of AAR's hard parts, which we shall call "HPX," as AAR does not wish to disclose the identity of this focal test category for competitive reasons. A unique feature of auto part retailing is that hard-part categories have many product subcategories defined by automobiles' make, model, and year of manufacture. That is, only alternatives within a specific subset or subcategory (called subclass henceforth) of the HPX category can fit a particular type of vehicle defined by its make, model, and year. Thus, each HPX subclass pertains to a different automobile, and therefore, a different "automobile market." In all, AAR stocks about 30 HPX subclasses. Further, each HPX subclass typically comprises three alternatives or SKUs of varying quality and price grade: Good, or entry grade; Better, or mid-level grade; and Best, or premium grade. Therefore, to determine optimal prices, we develop demand models at the HPX subclass level using store-level data on the history of sales and prices of each of the three variants in the various HPX subclasses. Such single-category analysis is justified because consumers at an automotive aftermarket retailer typically buy one unit of a particular hard part on any shopping occasion, whereas in a supermarket context consumers typically buy a basket of goods (therefore warranting the application of multicategory demand models, see, e.g., Seetharaman et al. 2005).

Summarizing, our collaborative research goals are, first, to derive and understand the determinants of product subclass demand elasticities at the store level, considering substitution between good, better, and best variants in a subclass, as well as a "no purchase" option and, second, to derive and validate optimal prices for the three product variants within each subclass of HPX sold at AAR stores around the country. We focus on everyday, or regular, price changes (and not price promotions), as these are of great importance to the company. As pointed out by Montgomery (1997), most of a retailer's profits come from products sold at their everyday prices, which are amenable to store-level customization.

Substantively, our research makes the following contributions to marketing practice: First, as already mentioned, the bulk of extant work on micromarketing-related price-demand modeling and optimization in the marketing literature focuses on consumer packaged goods and grocery retailing (e.g., Montgomery 2004). The automotive aftermarket industry is a very different retailing sector, worth studying and understanding from a marketing and pricing standpoint. Second, understanding the determinants of store-level price elasticities is important for micromarketing. Retail pricing research in the supermarket context has largely failed to uncover an effect of observable characteristics—such as consumer demographics—on estimated price elasticities of demand.[2] The nature of automotive aftermarket retailing and our national database, which spans a much wider range of demographics and store characteristics than, for example, those considered by Hoch et al. (1995), allows us to uncover significant effects of store characteristics on stores' estimated demand functions. In this regard, our research contributes importantly to emerging studies of retail price discrimination and price customization (e.g., Iyer and Seetharaman 2003).

We develop our subclass-level model in four steps. First, we present the multinomial logit (MNL) model that is used to model consumer demand for stockkeeping units (SKUs), or quality variants, within each HPX subclass. Second, using a latent-class specification, we accommodate the effects of heterogeneity across AAR stores in the parameters of the MNL model, where the heterogeneity is driven by both observed store characteristics and unobserved factors. Third, we discuss how to estimate the parameters of the proposed heterogeneous MNL model. Fourth, we show how to employ an empirical Bayes procedure, which uses the estimated heterogeneous MNL model as a "prior" for a store and the store's data as the likelihood, to estimate a store-specific demand function separately for each AAR store.

Multinomial logit (MNL) model of SKU demand within a subclass

The MNL model was first applied to scanner panel data by Guadagni and Little (1983), and since then it has been extensively used to model consumer demand for packaged goods in numerous applications. In our setting, we believe the MNL model is a more appropriate demand model than, say, the widely used log-log demand model (Reibstein and Gatignon 1984) for two reasons: first, the log-log demand model would involve nine price parameters as opposed to just one price parameter in the case of the MNL model, and second, the log-log demand model would not yield a globally concave category profit function for the retailer, making the calculation of optimal prices difficult (see Anderson and Vilcassim 2001). The MNL has been shown to be a well-behaved model in terms of yielding sensible own- and cross-price elasticities of demand as well as optimal prices in a category profit maximization problem (see, e.g., Chintagunta, Dube, and Singh 2003). In fact, the MNL is a more appropriate demand model for our setting than it is for the packaged-goods context in which it has been most often used. In the packaged goods context, consumers often buy multiple SKUs within the same product category, e.g., multiple flavors of yogurt or ice cream, or multiple units of a given brand (e.g., two 2-liter bottles of Coke). Both these behaviors violate the MNL model's basic underlying assumption of discrete choice, that is, that the consumer will buy one object out of many. In contrast, in the automotive aftermarket setting, a consumer shopping for a hard part visits a store to buy just one SKU within a product category and just one unit of the SKU.

The MNL model develops a consumer's probability Pr_{ist} of choosing alternative i within a subclass at store s at time t, from a collection of alternatives N = {Good, Better, Best, None}, where Good, Better, and Best refer to the three quality variants within the subclass, and None refers to the no-purchase option. Accordingly, Pr_{ist} is given by:

$$Pr_{ist} = \frac{e^{V_{ist}}}{\sum_{j \in N} e^{V_{jst}}}, \; i \in N; s = 1, \ldots, S; t = 1, \ldots, T, \tag{1}$$

where S is the total number of stores and the attractiveness of alternative i, with $i = 0$ denoting the no-purchase option, at store s to the shopper at time t, V_{ist}, is given by:

$$\begin{aligned} V_{ist} &= \alpha_i + \beta P_{ist} \quad \forall\, i \in \{Good, Better, Best\}, \\ V_{0st} &= 0, \end{aligned} \quad s = 1, \ldots, S; t = 1, \ldots, T. \tag{2}$$

In Equation 2, α_i denotes the consumer's intrinsic preference for SKU i, β represents the consumer's price sensitivity, and P_{ist} is the price of SKU i in store s at time t.

Assuming a sufficiently large number of shoppers for the hard part at store s in each time period t have choice probabilities for alternatives identical to those expressed in equation (1), the aggregate market share for alternative i in store s at time t is given by[3]:

$$S_{ist} = \frac{e^{V_{ist}}}{\sum_{j \in N} e^{V_{jst}}}, \; i \in N; s = 1, \ldots, S; t = 1, \ldots, T. \tag{3}$$

The parsimonious MNL specification of market shares in Equation 3 involves only four (unknown) parameters (α_1, α_2, α_3, and β) to fully characterize the market shares of the three SKUs of the product subclass as well as the no-purchase option. By explicitly allowing for the no-purchase option as an alternative in the consumer's choice process, the model in Equation 3 allows the three SKUs' prices to affect not only their relative market shares (say, $S_{Good,s,t} / S_{Better,s,t}$), but also their share relative to the no-purchase option (say, $S_{Good,s,t} / S_{0,s,t}$). From a normative standpoint, including the influence of the no-purchase option in the MNL is important for the retailer since an MNL that ignores the no-purchase option would predict infinite optimal prices for all SKUs for any assumed values of model parameters (we will expand on this in the discussion of price optimization, below).

The market share model in Equation 3 restricts its parameters $\{\alpha_1, \alpha_2, \alpha_3, \beta\}$ to be identical across stores, which may not be a realistic assumption. For example, it is possible that market shares of different HPX SKUs are more responsive to price changes within stores in the Midwest than in stores in the West. More generally, some stores may be characterized by systematically higher values of price sensitivity than others because of variation in readily observable as well as unobservable characteristics of their environments. Below, we extend the proposed MNL model to allow for such heterogeneity in the model parameters arising from observed and/or unobserved characteristics of stores.

Incorporating heterogeneity in MNL parameters across AAR stores

Using the latent class formulation of Kamakura and Russell (1989), we assume that the MNL parameter vector $\boldsymbol{\delta} = (\alpha_1, \alpha_2, \alpha_3, \beta)'$ follows a multivariate discrete distribution across stores, which is semiparametric in the sense that it has K supports (or "store segments"), whose locations are given by $\boldsymbol{\delta}_1, \boldsymbol{\delta}_2, \ldots, \boldsymbol{\delta}_K$ and masses are given by $\boldsymbol{\pi}_1, \boldsymbol{\pi}_2, \ldots, \boldsymbol{\pi}_K$ (that sum to 1). The locations and masses of this distribution are not specified according to any known parametric discrete distribution and are allowed to be unrestricted. Usually, based on managerial judgment and/or statistical tests, the value of K is chosen to be much smaller than S. For example, as shown in our application below, we use S = 800 AAR stores in the estimation, but allow each store to belong to, with some probability, one of K = 3 segments. That is, we allow for heterogeneity in MNL parameters across stores that may be due to stores being in different demographic neighborhoods with different characteristics (e.g., store size) etc. Thus, we specify the a priori masses of the heterogeneity distribution in the following manner (see, e.g., Gupta and Chintagunta 1994):

$$\pi_{sk} = \frac{e^{W_s \gamma_k}}{\sum_{l=1}^{K} e^{W_s \gamma_l}}, s = 1, \ldots, S; k = 1, \ldots, K. \tag{4}$$

where $\boldsymbol{\pi}_{sk}$ represents the (prior) probability of store s belonging to support k of the heterogeneity distribution, $\boldsymbol{W}_s = (1, w_{s1}, \ldots, w_{sL})$ is a row-vector of L variables characterizing store s, and $\boldsymbol{\gamma}_k$ the corresponding L x 1 column-vector of parameters capturing the effects of store characteristics on the probability of the store belonging to support k. For identification purposes, $\boldsymbol{\gamma}_k$ is restricted to be a zero-vector for one of the K segments (say, segment 1), and therefore the estimated $\boldsymbol{\gamma}_k$'s for the remaining segments should be interpreted as the effects of store characteristics on the relative probability of the store belonging to segment k rather than segment 1.

Under this latent-class specification of heterogeneity, the aggregate market share S_{iskt} for alternative i in store s at time t, given that store s belongs to segment k, can be written as follows:

$$S_{iskt} = \frac{e^{V_{iskt}}}{\sum_{j \in N} e^{V_{jskt}}}, I \in N; s = 1, \ldots, S; k = 1, \ldots, K; t = 1, \ldots, T. \tag{5}$$

where

$$\begin{array}{l} V_{iskt} = \alpha_{ik} + \beta_k P_{ist} \quad \forall\ i \in \{Good, Better, Best\}, \\ V_{0skt} = 0, \end{array} \quad s = 1, \ldots, S; k = 1, \ldots, K; t = 1, \ldots, T. \tag{6}$$

Of course, one does not know, a priori, the identity of the segment k that store s belongs to. We address this issue when we discuss estimation.

We use the latent-class methodology—as opposed to the hierarchical Bayes methodology—to incorporate heterogeneity across stores for four reasons. First, the methodology has been widely investigated in numerous choice-modeling applications since Kamakura and Russell (1989) and is almost as successful as the hierarchical Bayes methodology in explaining most of the variance that arises on account of heterogeneity (Andrews, Ainslie, and Currim 2002). Second, the latent-class methodology is easier to automate for implementation purposes than the computationally more demanding hierarchical Bayes methodology. Third, with the latent-class methodology, segmentwide price changes are easier to communicate to managers and easier to fine-tune and use on an ongoing basis. Last, the proposed methodology facilitates the transfer of knowledge to the marketing client, which is an important goal of this collaborative research.

Estimation of the parameters of the heterogeneous MNL model

The parameters of the heterogeneous MNL introduced above can be estimated by maximizing the likelihood function L:

$$L = \prod_{s=1}^{S}\prod_{k=1}^{K}\pi_{sk}[\prod_{t=1}^{T}\prod_{i=0}^{J}S_{iskt}{}^{N_{ist}}] \qquad (7)$$

where J is the total number of alternatives (including the no-purchase option), S_{iskt} is given by Equation 6, and N_{ist} is the number of units of SKU i sold in store s during week t. For the no-purchase alternative, N_{0st} is computed as $\mathrm{Max}_t(\Sigma_j N_{jst}) - \Sigma_j N_{jst}$. In other words, we treat the maximum weekly sales of the product (over all weeks) observed in store s during the study period as the weekly market potential at the store (i.e., the total number of consumers who walk in to store s during any given week, looking to buy in the product subclass). Therefore, this weekly market potential minus the observed product sales during a given week is taken to be the number of no purchases during that week.

In the estimation of Equation 7, we do not know, a priori, the identity of the segment k to which store s belongs. Therefore, we make the assumption that store s belongs to segment 1 with probability π_{s1}, segment 2 with probability π_{s2}, and so on, and then compute the weighted average of the conditional likelihood functions implied by each segment, where the weights are the segment membership probabilities π_{sk}, given by Equation 4.

We use the Schwartz Bayesian Criterion (SBC) to determine the optimal value of K. Specifically, starting with $K = 1$, we estimate heterogeneous MNL models for increasing values of K until the SBC fit criterion stops improving from adding more supports. Sometimes one may choose to stop adding additional supports when the size of a newly uncovered segment is not big enough to attract any managerial interest. Thus, managerial judgment also plays a role in determining the optimal value of also K.

Estimation of store-specific demand parameters

The estimated parameters of the heterogeneous MNL, $K, \{\alpha_{1k}, \alpha_{2k}, \alpha_{3k}, \beta_k, \pi_k\}_{k=1}^{K}$, are used in an empirical Bayes procedure to estimate store-specific demand parameters for store s, as shown below (e.g., Kamakura and Russell 1989; Rossi and Allenby 1993).

Step 1: Compute the store's conditional likelihood function $L_s(\boldsymbol{\delta}_k)$ given that the store belongs to segment k:

$$L_s(\delta_k) = \prod_{t=1}^{T}\prod_{i=0}^{J}S_{iskt}{}^{N_{ist}} \quad s = 1, \ldots, S; k = 1, \ldots, K. \qquad (8)$$

Repeat this calculation for each segment, to get $L_s(\boldsymbol{\delta}_1), \ldots, L_s(\boldsymbol{\delta}_k)$.

Step 2: Use Bayes rule to get the store's posterior segment membership probability for segment k:

$$\pi_{sk}^{post} \propto \pi_{sk} * L_s(\delta_k), s = 1, \ldots, S; k = 1, \ldots, K. \qquad (9)$$

In other words, the priors are equal to the probability masses, π_{sk}, estimated in the maximum-likelihood routine (above). Repeat this calculation for each segment, to get $\pi_{s1}^{post}, \ldots, \pi_{sk}^{post}$ (appropriately renormalizing them to add up to 1). The store-specific demand parameters are now given by $\{\alpha_{1k}, \alpha_{2k}, \alpha_{3k}, \beta_k, \pi_{sk}^{post}\}_{k=1}^{K}$. The store-specific demand function is given by:

$$S_{ist}^{post} = \sum_{k=1}^{K}\left[\pi_{sk}^{post}\frac{e^{V_{iskt}}}{\sum_{j\in N} e^{V_{jskt}}}\right], s = 1,\ldots,S; k = 1,\ldots,K; t = 1,\ldots,T \quad (10)$$

where V_{iskt} is given by Equation 6.

The above two-step procedure is repeated for all AAR stores in the sample to get store-specific demand functions for all subclasses of HPX.

Price optimization

Assuming that the demand function is constant over time, we drop the subscript t in Equation 10. Then, the retailer's unconstrained profit optimization problem (see, e.g., Chintagunta 2002), is given by:

$$\text{Max}_{\{\text{P}_{\text{js}}\}_{j=1}^{J}} (\text{Profit}_s) = \sum_{j=1}^{J}(P_{js} - C_j)S_{js}^{post}M, \quad (11)$$

where C_j stands for the marginal cost (known to the retailer) associated with SKU j, M is the market size, and S_{js}^{post} is given by Equation 10. The solution to Equation 11 is a triple of prices, $\{P_{Good,s}, P_{Better,s}, P_{Best,s}\}$ that maximizes the retailer's profits from the subclass.

The three SKUs' market shares, $S_{Good,s}^{Post}$, $S_{Better,s}^{Post}$, $S_{Best,s}^{Post}$,do not add up to one. Instead, they add up to $1 - S_{0,s}^{Post}$. This is what makes the retailer's category profit maximization yield a sensible interior solution for prices. Including the no-purchase option in the MNL model is crucial from a normative standpoint (as opposed to a descriptive standpoint, see, e.g., Anderson and Vilcassim 2001 for an insightful discussion of this point).

Data and Empirical Analysis

Store-level sales data

The data made available for this research represent over two years' historical sales figures (i.e., T = 104 weeks) for 27 different subclasses of HPX (numbered 1, 2, .., 30 by the company, with subclasses 3, 14, 24 excluded) pulled from AAR's transactions database. Specifically, the data set comprises weekly data on SKU-level sales and prices (within each HPX subclass) over a two-year period across a national sample of 800 AAR stores (i.e., S = 800) randomly drawn from the total population of 3,400 AAR stores. AAR provided data on several store demographic variables such as store size (square feet of space); location (longitude, latitude), type of store (satellite, feeder, hub) etc. and product costs (wholesale prices) for price optimization purposes. Table 1 reports average sales and prices across the 800 stores and 104 weeks for the 27 subclasses. It is clear from this table that the price of the Good variant is always less than the price of the Better variant (except in subclass 27), which is, in turn, always less than the price of the Best variant. This clearly indicates the existence of price/quality tiers in each product subclass.

In terms of average weekly sales per store, the Better variant in subclass 1 and the Good variant in subclass 13 seem to be the highest selling HPX SKU, with average weekly sales of .53 units and .51 units respectively. (A vast majority of the SKUs across all subclasses sell during less than 10% of the weeks.) In general, the "Best" variant is the least-selling variant among the three quality variants in each subclass, except in subclasses 18 and 22. In terms of average prices, the Good variant ranges from \$26.28 to \$129.72 while the Best variant has a price range of \$64.80 to \$224.59 across all subclasses. Also, in comparing the average prices for the "Good" with the "Best" variant, we see that the price gap varies from \$30.47 to \$128.60. These descriptive statistics suggest that our data capture a broad range of prices across subclasses and stores.

Nature of empirical analysis

An HPX subclass involves three alternatives, each representing a different price-quality level. As explained earlier, the heterogeneous MNL

Table 1
Average Weekly Sales and Prices for 27 Subclasses

	Average Weekly Sales (units sold)			Average Weekly Prices		
	Good	Better	Best	Good	Better	Best
Subclass						
1	.37	.53	.03	$26.28	$38.23	$64.80
2	.10	.21	.02	$65.93	$81.12	$126.35
4	.02	.05	.00	$114.09	$151.06	$188.50
5	.07	.10	.00	$125.93	$162.51	$197.94
6	.16	.18	.02	$92.59	$128.06	$166.21
7	.02	.03	.00	$129.72	$154.10	$199.99
8	.02	.04	.00	$112.60	$155.59	$199.99
9	.00	.05	.01	$86.99	$133.01	$161.47
10	.04	.04	.01	$92.46	$115.93	$148.75
11	.05	.09	.00	$95.99	$134.09	$224.59
12	.46	.13	.12	$66.32	$83.74	$133.29
13	.51	.06	.06	$64.99	$93.92	$134.11
15	.22	.06	.05	$64.38	$89.17	$132.19
16	.31	.06	.04	$64.09	$84.11	$134.15
17	.17	.04	.09	$66.39	$89.90	$134.82
18	.25	.07	.25	$73.36	$100.04	$134.05
19	.13	.14	.05	$100.37	$119.04	$140.69
20	.23	.05	.02	$66.32	$95.29	$133.41
21	.27	.09	.07	$76.34	$88.76	$134.19
22	.07	.02	.10	$100.91	$129.82	$141.54
23	.08	.08	.04	$95.90	$114.06	$140.72
25	.37	.00	.01	$63.84	$100.94	$134.21
26	.00	.21	.00	$87.99	$95.07	$118.46
27	.06	.01	.03	$86.36	$81.46	$131.93
28	.08	.02	.01	$90.52	$115.17	$143.08
29	.02	.05	.02	$76.81	$92.36	$134.10
30	.12	.12	.04	$69.56	$87.15	$131.86

model is estimated separately for each subclass using the maximum-likelihood technique, and then the estimated MNL demand model for a given subclass is used to derive a store-specific demand model. The estimated store-specific demand model serves as the input for a store-level price optimization problem for the retailer, which is solved to obtain store-specific optimal SKU prices for that subclass. This price optimization exercise is done store by store to obtain optimal SKU prices for the given subclass for each of the 800 stores in the sample. The same procedure is repeated across the 27 subclasses to obtain store-specific optimal SKU prices for all subclasses in our data set. Subsequently, we examine optimized profits and their sensitivity to price changes.

Empirical Results

For a majority of the 27 HPX subclasses, a three-segment solution for unobserved heterogeneity across stores (i.e., $K = 3$) was found to be most appropriate, on account of either model fit (i.e., SBC) or substantive sizes of the estimated

Table 2

Brand Preferences for 27 Subclasses

(Subclasses subsequently excluded from the price optimization are shown in bold

	Segment 1			Segment 2			Segment 3		
Subclass	$\alpha_{Good,1}$	$\alpha_{Better,1}$	$\alpha_{Best,1}$	$\alpha_{Good,2}$	$\alpha_{Better,2}$	$\alpha_{Best,2}$	$\alpha_{Good,3}$	$\alpha_{Better,3}$	$\alpha_{Best,3}$
1	**–2.8140**	**–2.3915**	**–5.0329**	**–2.7324**	**–2.7100**	**–6.2768**	**–2.2272**	**–3.8674**	**–6.9835**
2	–1.3513	.0241	–.8285	–.8493	.1262	–1.0773	–1.9945	–.6030	–1.2453
4	**–3.3385**	**–2.1345**	**–7.9242**	**–2.2458**	**–1.1855**	**–6.3036**	**–3.9008**	**–2.2000**	**–7.1146**
5	**–4.1378**	**–4.0095**	**–10.028**	**–3.5746**	**–3.3397**	**–9.0859**	**–3.2553**	**–2.7651**	**–7.1596**
6	–1.5074	–.7991	–2.5447	–1.4926	–.8487	–2.4501	–1.9731	–1.2231	–2.2525
7	.0245	.9671	–11.100	–1.6790	–.5537	–10.000	–.7436	–.0806	–8.0011
8	–2.6334	–1.4391	–11.003	–3.4510	–1.7912	–12.001	–4.0912	–.4407	–7.0248
9	–8.7265	1.5110	–1.0637	–8.1044	1.4374	–.8424	–7.0376	1.6789	–.3992
10	–.9010	–.2108	–2.1458	–.9748	.1485	–2.4088	–.9422	–.3522	–2.4837
11	–2.8256	–1.9025	–7.1195	–1.7818	–.4571	–8.0278	–2.2169	–1.0318	–6.0339
12	–.6843	–1.5951	–.1787	–.2779	–1.3641	–.1510	–.7779	–1.1289	–.9412
13	–1.2678	–2.7364	–2.0285	–.8715	–2.8926	–2.5094	–1.0406	–2.7154	–2.1147
15	–.0955	–.8186	.4876	–.7201	–1.3308	.0068	–.6494	–.7887	–.1885
16	.5858	–.4949	.8252	.3861	–.2907	1.6195	.4714	.5850	.1839
17	–1.8920	–3.5466	–2.6699	–2.2853	–3.3366	–1.7565	–2.1263	–2.6286	–1.4965
18	–1.5423	–2.7930	–.7778	.2208	.8026	3.3282	1.0809	.3222	3.6311
19	6.1147	7.8335	8.6564	5.5385	7.3170	8.1709	3.1336	4.4075	4.9707
20	–1.8083	–3.1982	–3.7781	–2.0728	–3.2541	–3.4749	–.9276	–2.2639	–2.6071
21	.2382	–.6745	.9513	–.1031	–.5825	.9772	.1334	–.2644	.5031
22	**–1.8247**	**–2.1553**	**–.4356**	**–3.2425**	**–4.3497**	**–2.6770**	**–2.3921**	**–3.2859**	**–1.8677**
23	3.1843	4.5622	5.4749	3.3233	4.3742	5.1412	–.0203	.7601	.9866
25	.1107	–4.2366	–2.4549	–.0720	–3.6325	–1.9084	.4040	–3.2236	–1.6199
26	–3.6933	–.4854	–6.4227	–6.9962	1.0846	–5.2173	–3.2980	–.1150	–7.2705
27	–.4715	–2.9918	–.2155	–.1367	–2.8559	.5028	–.6246	–2.8577	.0578
28	–.1676	–.7453	–.1078	.1004	–.4052	–.2401	–.0407	–1.3786	–1.5720
29	–3.7624	–3.1927	–4.0914	–3.2846	–1.9371	–2.9935	–2.5213	–2.1857	–2.3001
30	–1.7034	–2.2685	–1.9241	–1.8302	–2.8090	–1.5270	–.9598	–2.1294	–1.1908

segment.[4] For ease of exposition and for the sake of consistency across subclasses, we use the three-segment solution for all HPX subclasses, both in the reported estimation results in this section as well as in the price optimization exercise at the store level.[5]

Estimation results for the heterogeneous MNL

Tables 2 and 3 report the estimation results for the three-segment heterogeneous MNL for the 27 subclasses. Specifically, Table 2 reports the estimated brand preferences for the three segments (i.e., $\{\alpha_{Good,1}, \alpha_{Better,1}, \alpha_{Best,1}\}$, $\{\alpha_{Good,2}, \alpha_{Better,2}, \alpha_{Best,2}\}$, $\{\alpha_{Good,3}, \alpha_{Better,3}, \alpha_{Best,3}\}$), while Table 3 reports the estimated price sensitivities for the three segments (β_1, β_2, β_3,), along with the estimated sizes of the three segments (π_1, π_2, π_3).

From Table 2, we can see that the estimated brand preferences are, for the most part, negative. The reason for this is that the no-purchase option is, by far, the dominant option in almost all the subclasses under study. Since the brand

Table 3

Price Sensitivities and Segment Sizes for 27 Subclasses

(Subclasses subsequently excluded from the price optimization are shown in bold)

	Segment 1		Segment 2		Segment 3	
Subclass	β_1	π_1	β_2	π_2	β_3	π_3
1	**.0063**	**.43**	**.0255**	**.36**	**–.0046**	**.21**
2	–.0293	.41	–.0249	.35	–.0283	.24
4	**–.0093**	**.62**	**–.0111**	**.26**	**.0002**	**.12**
5	**.0045**	**.39**	**.0046**	**.34**	**–.0013**	**.28**
6	–.0116	.37	–.0166	.38	–.0525	.25
7	–.0289	.24	–.0214	.24	–.0272	.52
8	–.0120	.31	–.0138	.33	–.0222	.36
9	–.0402	.46	–.0277	.20	–.0354	.34
10	–.0320	.48	–.0247	.30	–.0242	.21
11	–.0116	.40	–.0142	.28	–.0131	.32
12	–.0251	.32	–.0215	.32	–.0204	.35
13	–.0188	.37	–.0082	.36	–.0131	.27
15	–.0295	.29	–.0337	.32	–.0285	.39
16	–.0359	.26	–.0431	.46	–.0540	.28
17	–.0032	.27	–.0140	.31	–.0115	.42
18	–.0170	.22	–.0424	.24	–.0436	.54
19	–.0869	.23	–.0869	.42	–.0682	.35
20	–.0051	.29	–.0094	.33	–.0114	.38
21	–.0346	.42	–.378	.31	–.296	.26
22	**–.0203**	**.48**	**.0008**	**.33**	**–.0084**	**.18**
23	–.0673	.38	–.0617	.16	–.0402	.46
25	–.0248	.16	–.0342	.54	–.0353	.30
26	–.0251	.26	–.0346	.29	–.0159	.45
27	–.0241	.19	–.0350	.57	–.0361	.24
28	–.0388	.22	–.0352	.51	–.0250	.26
29	–.0067	.24	–.0114	.52	–.0107	.24
30	–.0152	.55	–.0201	.16	–.0182	.29

preferences can be understood as capturing baseline shares of the SKUs in the subclass, estimates of these preferences must be interpreted as being relative to the preference for the no-purchase option, set at 0 for identification purposes. Not surprisingly, these estimates turn out be smaller than the estimated preference for the no-purchase option.

In Table 3, we can see that the estimated price sensitivities are, as expected, for the most part negative. The estimated price sensitivities have the wrong sign (that is, they are positive) for four out of the 27 subclasses, namely, 1, 4, 5, and 22. Therefore, these four subclasses are subsequently excluded from the price optimization analysis since positive price coefficients would imply infinite optimal prices—an unrealistic option from a practical standpoint. (These excluded subclasses clearly need to be separately investigated in more depth using an alternative approach.) The estimated sizes of store segments across the remaining subclasses vary from 16% (e.g., segment 1 in subclass 25) to 57% (e.g., segment 2 for subclass 27).

Using the estimated parameters in tables 2 and 3, we compute the own- and cross-price elasticities of the three SKUs within each subclass (i.e., a total of nine price elasticities, composed of three own-elasticities and six cross-elasticities per subclass) by averaging over the three supports of the estimated heterogeneity distribution. These elasticities are reported in Table 4.

For 25 out of the 27 subclasses (that is, for all but subclasses 1 and 5), the elasticities are correctly signed; that is, all own-elasticities are negative, while all cross-elasticities are positive.[6] From Table 4 one can see that the highest estimated own-price elasticities across the three quality variants are in HPX subclass 19. In fact, HPX subclass 19 appears to be the most price-responsive among all subclasses under study.

The average own-price elasticities of Good, Better, and Best variants, across all 27 subclasses in the study, are found to be –1.10, –1.60, and –3.11 respectively. The average cross-price elasticities of demand for Good, Better, and Best variants are found to be .77, .72, and 1.02 respectively. This shows that the demand for the highest-quality variant in a subclass generally responds more to changes in prices of the other two variants than do the demands for the Good and Better variants.

The demand for the Good variant is least responsive to price, and also least vulnerable on price, in subclass 25.[7] This indicates that there is

Table 4

Estimated Price Elasticities for 27 Subclasses

(Subclasses with wrongly signed elasticities are shown in bold)

Subclass	$E_{G \to G}$	$E_{b \to b}$	$E_{B \to B}$	$E_{b \to G}$	$E_{B \to G}$	$E_{G \to b}$	$E_{B \to b}$	$E_{G \to B}$	$E_{b \to B}$
1	**.1849**	**.1645**	**.6808**	**–.2530**	**–.0269**	**–.1022**	**–.0269**	**–.1022**	**–.2530**
2	–1.3033	–.7965	–3.2159	1.4359	.2612	.5111	.2612	.5111	1.4359
4	–.6768	–.4096	–1.6233	.8938	.0030	.3075	.0030	.3075	.8938
5	**.2169**	**.2009**	**.5842**	**–.2793**	**–.0007**	**–.1552**	**–.0007**	**–.1552**	**–.2793**
6	–.8645	–1.9375	–3.8198	1.1007	.1235	1.3321	.1235	1.3321	1.1007
7	–1.8512	–1.8408	–5.2427	2.1990	.0002	1.5495	.0002	1.5495	2.1990
8	–1.5024	–.4554	–3.2522	2.0754	.0008	.3291	.0008	.3291	2.0754
9	–3.1362	–.1629	–5.6285	4.6345	.1953	.0013	.1953	.0013	4.6345
10	–1.3884	–1.5644	–4.0448	1.6645	.0981	1.1867	.0981	1.1867	1.6645
11	–.8052	–.5941	–2.8739	1.1233	.0026	.4242	.0026	.4242	1.1233
12	–.5419	–1.4736	–2.4436	.3730	.4956	.9205	.4956	.9205	.3730
13	–.1758	–1.1360	–1.6211	.1267	.1819	.6979	.1819	.6979	.1267
15	–.6759	–2.1818	–3.4291	.5338	.5965	1.2847	.5965	1.2847	.5338
16	–.7531	–3.0173	–5.4914	.7071	.4487	2.0847	.4487	2.0847	.7071
17	–.3391	–.7550	–.8845	.1470	.4683	.3270	.4683	.3270	.1470
18	–1.5392	–3.3658	–2.7202	.3816	2.3013	1.2088	2.3013	1.2088	.3816
19	–4.8441	–5.3638	–9.4809	4.2016	1.8242	3.2211	1.8242	3.2211	4.2016
20	–.1405	–.7145	–1.0951	.1348	.0939	.4506	.0939	.4506	.1348
21	–1.0085	–2.3753	–3.7464	.6377	.8087	1.5829	.8088	1.5829	.6377
22	–.7457	–1.2456	–.7075	.1813	.8483	.3634	.8483	.3634	.1813
23	–3.1316	–3.5284	–6.2319	2.6238	1.3581	2.0410	1.3582	2.0410	2.6238
25	–.0425	–3.3095	–4.3751	.0241	.0572	2.0658	.0572	2.0659	.0241
26	–2.0320	–.0611	–2.8067	2.1935	.0024	.0547	.0024	.0547	2.1935
27	–.9333	–2.5349	–3.2269	.1689	1.15218	1.9332	1.1522	1.9332	.1689
28	–.7345	–3.1706	–4.3398	.6286	.3801	2.2515	.3801	2.2515	.6285
29	–.5745	–.3995	–1.1267	.5336	.2282	.2016	.2282	.2016	.5336
30	–.4559	–1.1725	–1.8065	–.2530	–.0269	.7165	.4159	.7165	.2963

(G = Good, b = Better, B = Best)
(Key: $E_{b \to G}$ refers to the elasticity of demand for Good with respect to the price of Better)

substantial room to increase the price of the lowest-quality variant in subclass 25. This variant has its highest clout in subclass 19.[8] The Better variant is least responsive to price, and also least vulnerable on price, in subclass 26, while it has the highest clout in subclass 9. The Best variant is least responsive to price in subclass 22, least vulnerable on price in subclass 17, and has the most clout in subclass 18.

The average cross-price elasticities of SKUs with respect to the prices of Good, Better, and Best variants are found to be .99, 1.07, and .46, respectively. This indicates that changing the price of the highest-quality variant affects the demand for a lower-quality brand significantly less than the amount by which changing the price of the lower-quality variant affects the demand for the highest-quality variant. This runs counter to the asymmetric price-competi-

Table 5
Optimal Prices for Each Store Segment

Subclass	Segment 1 $P_{Good,1}$	Segment 1 $P_{Better,1}$	Segment 1 $P_{Best,1}$	Segment 2 $P_{Good,2}$	Segment 2 $P_{Better,2}$	Segment 2 $P_{Best,2}$	Segment 3 $P_{Good,3}$	Segment 3 $P_{Better,3}$	Segment 3 $P_{Best,3}$
2	87.72	97.29	110.22	96.08	105.64	118.57	87.74	97.30	110.23
6	151.95	153.64	190.88	121.38	123.07	160.31	74.79	76.48	113.72
7	115.15	125.16	130.00	125.74	135.74	140.00	115.47	125.47	130.00
8	143.87	149.13	150.00	130.46	135.80	150.00	103.74	109.10	142.25
9	93.06	94.56	121.96	102.54	110.69	138.09	90.55	100.03	127.46
10	84.95	90.88	128.07	97.03	102.96	140.16	96.72	102.65	139.84
11	148.83	175.98	183.58	136.52	163.67	171.27	140.42	167.57	175.17
12	76.83	85.24	105.33	87.39	95.80	115.89	87.72	96.13	116.22
13	88.18	96.59	116.68	168.34	176.75	196.84	115.16	123.57	143.66
15	74.26	80.64	100.73	66.00	72.38	92.47	74.10	80.48	100.57
16	66.52	74.93	95.02	59.98	68.39	88.48	52.54	60.95	81.04
17	366.65	375.06	395.15	106.24	114.65	134.74	124.64	133.05	153.14
18	96.93	104.17	124.26	69.73	76.97	97.06	70.85	78.09	98.18
19	77.12	92.46	108.41	73.34	88.68	104.62	68.55	83.89	99.84
20	257.29	266.85	279.78	160.48	170.04	182.97	147.15	156.71	169.64
21	68.18	75.42	95.51	63.82	71.06	91.15	74.82	82.06	102.15
23	70.60	85.94	101.89	74.40	89.74	105.68	73.17	88.51	104.46
25	78.02	82.40	106.52	63.36	67.73	91.86	63.89	68.26	92.39
26	80.73	90.31	121.32	73.60	81.67	112.23	109.87	119.45	150.45
27	79.01	83.38	107.51	63.99	68.36	92.49	66.09	70.46	94.59
28	66.22	77.05	110.95	70.33	81.16	115.06	83.94	94.77	128.67
29	183.64	192.05	212.14	122.65	131.06	151.15	130.49	138.90	158.99
30	101.47	109.88	129.96	83.52	91.93	112.02	92.56	100.97	121.06

tion effects that have been uncovered in scanner data (see, e.g., Blattberg and Wisniewski 1989), where researchers have found lower-quality brands to be more responsive to price cuts of higher-quality brands. Whether this difference is due to the durable nature of the automotive hard part, as opposed to the nondurable nature of packaged goods, is a question that deserves further study.

Store segment-specific optimal prices

Using the price optimization procedure outlined earlier, we develop optimal SKU prices (reported in Table 5) for each of the three estimated store segments within each subclass (excluding subclasses 1, 4, 5, and 22). Comparing these prices to the average observed prices in Table 1, we find that, on average, the optimal prices are roughly equal to the observed prices, with the average value of optimal prices minus observed prices being \$1.14. However, the standard deviation of optimal prices minus the observed prices is \$52.62, which indicates a fair amount of variability in the difference between the observed and optimal prices. As a percentage of observed average prices, the average deviation of optimal prices from the observed prices is 7.65%, while its standard deviation is 58.4%. This shows that the modeling procedure adopted in this study does indeed prescribe substantively different prices from those AAR currently applies at its stores. Table

5 also shows that the three-segment estimation was effective in terms of designing optimal pricing strategies that yield quite different prices across segments within a given subclass. For instance, the optimal price for the Good variant in subclass 13 varies from $88.18 in segment 1 to $168.34 in segment 2 and $115.16 in segment 3. Similar differentiation is found for other subclasses and variants. Thus the model recommends a range of prices, depending on the store and subclass characteristics.

Effects of store characteristics

To shed light on the estimated effects of eight store characteristics on store segment membership probabilities for segments 1 and 2, relative to segment 3, we report the number of times each estimated effect is significant across the 27 subclasses below.[9]

1. Store size: 21 out of 54 estimates are significant,
2. Full year since opening (i.e., the store has been open at least a year): 27 out of 54 estimates are significant,
3. Longitude: 19 out of 54 estimates are significant,
4. Latitude: 9 out of 54 estimates are significant,
5. Small (satellite) Store: 20 out of 54 estimates are significant,
6. Medium (feeder) Store: 0 out of 54 estimates are significant,
7. Large (hub) Store: 11 out of 54 estimates are significant,
8. Sales Channel: 25 out of 54 estimates are significant.

Overall, 132 out of the 432 estimates (i.e., 30%) are significant, which indicates that these store characteristics have good explanatory power. In other words, store characteristics are at least partly responsible (with unobserved store characteristics being primarily responsible) for the observed variation in both estimated preferences for different quality variants of the same brand as well as for the estimated price sensitivities across AAR stores in our sample. To the extent that store location, as captured by the longitude and latitude variables, explains some of the variation in the estimated demand functions across stores, AAR can take this new information into account as it decides where to put its new stores. Locations associated with lower price sensitivity (i.e., those locations containing stores belonging to the segments with low estimated price sensitivities across subclasses) would likely be good sites for new stores.

Store-specific optimal prices

Using the empirical Bayes procedure, we make store-specific estimates of model parameters. This yields 800 estimated sets of the parameter vector, i.e., $\{\alpha_{1k}, \alpha_{2k}, \alpha_{3k}, \beta_k, \pi_{sk}^{post}\}_{k=1}^{K}$, ($K$ = 800), within each of 23 subclasses (i.e., the original 27 subclasses minus the excluded subclasses of 1, 4, 5, and 22).[10] We then utilize each store's specific parameter estimates in the price optimization procedure, thus developing optimal SKU prices at the store level.

Optimal vs. current store-specific prices

We are now in a position to examine two questions of interest to AAR management : (1) Are the company's stores currently underpricing or overpricing HPX SKUs relative to the optimal prices in various subclasses? (2) Are the current store-level price gaps (i.e., differences between the prices of the Good and Best variants in each subclass) too large or too small? Since the number of store-specific current and recommended SKU prices are too numerous to discuss in their entirety, we examine a sample of four subclasses involving 464 "cases" (subclass-store combinations) (selected out of the 800 stores by management). The results of the comparison between observed and optimal store-specific SKU prices for these four subclasses at the selected stores are summarized in Table 6.

As indicated in Table 6, we find that the optimal prices recommended by our model fall into three price-gap scenarios:

Scenario 1. In this scenario, the optimal prices of the Best and the Good variants fall between the currently observed prices of these variants—the optimal price of the Good variant is higher

Table 6
Comparison of Optimal Store Prices with Observed Prices

Subclass	Number of Cases (subclass-store combinations)	Scenario 1	Scenario 2	Scenario 3
6	35	-	21	14
16	146	-	-	146
21	147	-	-	147
30	136	39	97	-
Total	**464**	**39 (8.4%)**	**118 (25.4%)**	**307 (66.2%)**
Bandwidth (Optimal)		\$28.47	\$30.36	\$28.40
Bandwidth (Observed)		\$48.97	\$46.87	\$43.18

than the observed price of the Good variant, but the optimal price of the Best variant is lower than the observed price of the Best variant. That is, $P_{1G}^* > P_{1G}$ and $P_{1B}^* < P_{1B}$ where P_{1G} and P_{1B} are the actual (observed) prices, and $P_{1G}^*$and $P_{1B}^*$are the optimal prices for the Good and the Best variants respectively. Table 6 indicates that was true in about 8% of the 464 cases included in this analysis.

Scenario 2. Here, the optimal prices are slightly higher than the actual prices in the store. In other words, $P_{1G} < P_{1G}^* < P_{1B} < P_{1B}^*$. We found this scenario to be true in nearly 25% of the sample cases.

Scenario 3. In this situation, the optimal prices are slightly lower than the actual prices in the store. In other words, $P_{1G}^* < P_{1G} < P_{1B}^* < P_{1B}$. We found this scenario to be true in nearly 66% of the cases in our sample.

This preliminary analysis of four subclasses suggests that the retailer would benefit from a slight reduction in the prices of all three variants in each subclass in a majority of the stores. In a smaller fraction of stores, the retailer benefits from raising the prices. (Note that our model can provide specific prices for variants of each subclass at a store.)

It is also interesting to compare the differences between optimal and observed prices of the Best and the Good variants, which we refer to here as the subclass price "bandwidths." As shown in Table 6, across the three scenarios described above, the average observed price bandwidths are much greater, 50 to 70% greater, than the optimum bandwidth. That is, it appears that AAR would benefit not only by adjusting the prices above or below current levels at specific stores, but also by narrowing the price gap between the Good and Best variants at these stores.

More generally, AAR's management found our (unconstrained) recommended prices for 200 out of the 800 stores involved in our analysis to be too high relative to current prices to implement from a practical standpoint (not an unusual occurrence in retail pricing studies, see, e.g., Montgomery 2004). As for the remaining 600 stores in the sample, AAR is planning to implement our optimal pricing recommendations at a random sample of 200 stores and study the actual sales movements in the 24 subclasses (24 instead of 27 because, as mentioned earlier, subclasses 1, 4, 5, and 22 were excluded from the price optimization exercise) at these test stores to see if the recommended prices indeed lead to an improvement in the category profits associated with those subclasses. If improvements are obtained, AAR plans to do a national rollout of the recommended optimal prices across all their stores. Note that the demand functions for all stores not included in our study, can be predicted by coupling our parameter estimates with these stores' characteristics and sales data.

Segment-level profit optimization snalysis

We next use the actual average weekly prices in Table 1 and the optimal prices of variants in Table 5 to conduct a sensitivity analysis, comparing the total profit yielded by the price optimization model with estimated predicted profits at the current prices. Optimal and current profit levels (summed across the three variants) are summarized in Table 7.

Table 7
Actual and Optimal Profits

Subclass	Optimal Profit Segment 1	Optimal Profit Segment 2	Optimal Profit Segment 3	Optimal Profit Total	Actual Profit Segment 1	Actual Profit Segment 2	Actual Profit Segment 3	Actual Profit Total	% Decrease Actual vs Optimal
2	3.29	5.58	2.08	**10.95**	2.78	4.23	1.78	**8.80**	**19.65**
4	2.81	5.39	-[1]	**8.20**	2.43	4.90	-[1]	**7.33**	**10.59**
6	10.44	5.52	.16	**16.13**	9.07	5.25	.07	**14.39**	**10.76**
7	3.70	2.07	1.86	**7.63**	3.14	1.98	1.62	**6.74**	**11.56**
8	4.38	2.22	2.66	**9.26**	4.28	2.15	1.96	**8.39**	**9.38**
9	2.58	7.45	4.63	**14.67**	1.43	6.59	3.19	**11.21**	**23.53**
10	2.29	5.19	4.07	**11.54**	2.01	5.04	3.96	**11.01**	**4.64**
11	2.61	6.13	4.36	**13.11**	2.10	5.31	3.70	**11.11**	**15.22**
12	6.24	10.18	7.83	**24.25**	5.81	9.40	7.29	**22.50**	**7.22**
13	4.20	16.21	8.30	**28.71**	3.84	9.10	6.51	**19.45**	**32.25**
15	7.62	3.56	6.24	**17.42**	6.89	3.02	5.63	**15.54**	**10.84**
16	7.88	6.05	3.26	**17.19**	7.12	4.56	2.44	**14.12**	**17.88**
17	23.40	4.03	7.14	**34.56**	7.44	3.64	6.07	**17.16**	**50.36**
18	6.28	14.24	16.02	**36.53**	6.01	9.58	11.36	**26.95**	**26.23**
19	21.09	17.31	9.38	**47.78**	9.05	5.67	3.24	**17.97**	**62.39**
20	11.68	4.33	8.98	**24.98**	3.00	1.98	4.20	**9.18**	**63.24**
21	7.33	5.41	9.17	**21.91**	6.21	4.11	8.58	**18.91**	**13.68**
23	11.22	13.67	3.79	**28.69**	5.05	8.07	2.71	**15.84**	**44.80**
25	6.87	3.38	4.78	**15.03**	6.42	3.28	4.67	**14.38**	**4.34**
26	2.69	5.06	8.81	**16.56**	2.67	4.70	8.10	**15.47**	**6.56**
27	6.69	4.66	2.87	**14.22**	6.37	3.49	2.03	**11.89**	**16.34**
28	2.60	4.11	6.02	**12.73**	1.84	3.27	5.88	**10.99**	**13.66**
29	3.32	4.41	5.99	**13.72**	2.23	3.90	5.14	**11.27**	**17.87**
30	5.13	3.09	6.81	**15.03**	4.67	2.95	6.36	**13.98**	**6.99**
Total				**460.79**				**334.57**	**27.39**

[1] The profit figures for subclass 4, segment 3 are omitted due to numerical problems; the posterior probability for this category equals zero, hence this does not affect our analysis.

From Table 7, we see that for most subclasses (17 out of 24), the model-based total profits at current prices are within 20% of the optimal total profits. For the other subclasses, however, the actual profits are 23.5% to 63.2% lower than the optimized profits. Subclasses 17, 19, 20 and 23 show the largest differences, suggesting that these subclasses have substantial potential for improvement if the optimal pricing strategy is implemented. Summing across all subclasses, the model-based predicted profits at current prices are 27.4% lower than the predicted profits at the optimal prices.

Subsequently, we explored the sensitivity of predicted profits to uniform deviations (of –10%, –20%, –30%, +10%, +20% and +30%) from the variants' optimal price levels in each subclass and segment. In general, the impact of a 10% deviation (increase or decrease) from the optimal prices of variants in each subclass/store segment leads to a 1 to 4% reduction from the optimized profit, while a ±30% deviation from optimal prices results in, for most subclasses, at least a 20% reduction from optimum profits. Further, consistent with the findings described in the previous paragraph, the predicted profit levels for subclasses 19 and 23 are most sensitive to these deviations from the optimal prices, while the profits achieved in subclasses 13, 17, 20, and 30 appear to be the least sensitive (the profit reduction not exceeding 10% even with deviations from optimal prices of ±30%). Importantly, our sensitivity analyses reveal that the profit reduction effects of deviations from optimal price levels are asymmetric. Specifically, *profits are typically much more sensitive to underpricing than overpricing*. This suggests that enhanced sales at lower prices do not adequately compensate for reduced margins, implying that store managers should proceed with much more care when considering a price reduction than when considering an increase.

Overall, the sensitivity analysis of the optimal model shows that the model behaves reasonably and that such model-based explorations of pricing strategies can be of great value to management.

Conclusions and Directions for Future Research

This study integrates both descriptive and normative aspects of demand modeling in marketing research. While the statistical methodologies used in our study have received great attention in the academic literature, they have not been embraced with any degree of familiarity in marketing practice, far less in automotive aftermarket retailing. We believe that the publication of this work will encourage managers in a range of industries to employ modeling-based pricing strategies for their brands or products. It may also encourage managers to customize prices across stores, based on our proposed procedures. One caveat is in order. In this study, we ignored the effect of prices at competing retailers on demand at the focal retailer; we did so because prices at competing retailers were not available to us. Future researchers could further the work of this study by supplementing our demand data with prices at competing retailers, obtained using mystery shoppers, and recalibrating the price elasticities both across quality grades and across retailers. ■

Notes

1. Iyer and Seetharaman (2003) provide empirical evidence that such micromarketing strategies govern prices and service decisions in gasoline retailing.

2. Important exceptions include Hoch et al. (1995) and Chintagunta, Dube, and Singh (2003), which show that demographic and competitive variables describing the trading areas of stores explain a lot of the variation in estimated price elasticities across stores.

3. Allenby and Rossi (1991) show that even if consumers are heterogeneous in their choice probabilities, as long as all consumers are exposed to the same marketing variables on a given shopping trip (as in our case), the aggregate MNL model is a good approximation to the mixed logit that would capture the suitably summed brand choice probabilities across heterogeneous consumers. If the heterogeneity distribution across different consumers entering the product market in each week does not change from one week to another (an assumption supported by our collaborator) then, given marketing mix homogeneity, Equation (3) is an appropriate representation of the aggregate market shares even if consumers are heterogeneous.

4. In other words, going from the three-segment to the four-segment solution yielded a segment that was substantively too small, in terms of its estimated size, to be of much practical value to the retailer.

5. For all subclasses, we found the correlations between the estimated residuals of the demand functions and the prices of the three brands to be quite small (of magnitude less than or equal to .1), which precludes concerns about correcting for possible price endogeneity in the estimation.

6. For subclasses 4 and 22, although the estimated price coefficient is wrongly signed (is positive) for one of the three estimated segments (see Table 3), these wrongly signed segments are sufficiently small in size that we do not obtain wrongly signed aggregate price elasticities.

7. Vulnerability of an SKU is defined as the sum of the elasticities associated with the SKU's sales with respect to its competitors' prices (Kamakura and Russell 1989).

8. Clout of an SKU in a subclass is defined as the sum of the elasticities associated with the sales of the subclass's other SKUs with respect to the focal SKU's price (Kamakura and Russell 1989).

9. Since there are 27 subclasses, and two estimates per subclass (i.e., one each for segments 1 and 2), we have 54 estimates for each store characteristic. Interested readers can obtain the actual estimates from the authors.

10. For obvious reasons, we do not report these 800*23 sets of four-dimensional parameter estimates in the paper. They are available from the authors upon request.

References

Allenby, G. M., and P. E. Rossi (1991) "There Is No Aggregation Bias: Why Macro Logit Models Work." *Journal of Business and Economic Statistics* 9, 1–14.

Anderson, Eric, and Naufel J. Vilcassim (2001), "Structural Demand Models for Retailer Category Pricing." Chicago, Ill.: University of Chicago, Working Paper.

Andrews, Rick L., Andrew Ainslie, and Imran S. Currim (2002), "An Empirical Comparison of Logit Choice Models with Discrete Versus Continuous Representations of Heterogeneity." *Journal of Marketing Research* 39 (4), 479–87.

__________, Asim Ansari, and Imran S. Currim (2002), "Hierarchical Bayes versus Finite Mixture Conjoint Analysis Models: A Comparison of Fit, Prediction and Partworth Recovery." *Journal of Marketing Research* 39 (1), 87–98.

Automobile Aftermarket Industry Association (2005), *The Aftermarket Consumer: Do-It-Yourself or Do-It-For-Me?* Bethesda, Md.: Automobile Aftermarket Industry Association.

__________ (2002), *Aftermarket Factbook*. Bethesda, Md.: Automobile Aftermarket Industry Association.

__________ (2003), *Enhanced Line Review for Application Parts: A Guide to Managing Hard Parts in the Aftermarket*. Bethesda, Md.: Automobile Aftermarket Industry Association.

Basuroy, Suman, Murali K. Mantrala, and Rockney G. Walters (2001), "The Impact of Category Management on Retailer Prices and Performance: Theory and Evidence." Journal of Marketing 65 (4), 16-32.

Blattberg, Robert, and Kenneth Wisniewski (1989), "Price-Induced Patterns of Competition." Marketing Science 8 (4), 291-309.

Chintagunta, Pradeep K. (2002), "Investigating Category Pricing Behavior at a Retail Chain." *Journal of Marketing Research* 39 (2), 141–54.

__________, J. P. Dube, and Vishal Singh (2003), "Balancing Profitability and Customer Welfare in a Supermarket Chain." *Quantitative Marketing and Economics* 1 (1), 111–47.

Guadagni, Peter M., and John D. C. Little (1983), "A Logit Model of Brand Choice Calibrated on Scanner Data." *Marketing Science* 2 (2), 203–38.

Gupta, Sachin, and Pradeep K. Chintagunta (1994), "On Using Demographic Variables to Determine Segment Membership in Logit Mixture Models." *Journal of Marketing Research* 31 (1), 128–36.

Hoch, Stephen J., Byung-Do Kim, Alan L. Montgomery, and Peter E. Rossi (1995), "Determinants of Store-Level Price Elasticity." *Journal of Marketing Research* 32 (1), 17–29.

Iyer, Ganesh, and P. B. Seetharaman (2003), "To Price Discriminate or Not: Product Choice and the Selection Bias Problem." *Quantitative Marketing and Economics* 1 (2), 155–78.

Kamakura, Wagner A., and Gary J. Russell (1989), "A Probabilistic Choice Model for Market Segmentation and Elasticity Structure." *Journal of Marketing Research* 26 (4), 379–90.

Kim, Byung-Do, Robert C. Blattberg, and Peter E. Rossi (1995), "Modeling the Distribution of Price Sensitivity and Implications for Optimal Retail Pricing." *Journal of Business and Economic Statistics* 13 (2), 291–304.

Montgomery, Alan L. (1997), "Creating Micro-Marketing Pricing Strategies Using Supermarket Scanner Data." *Marketing Science* 16 (4), 315–37.

__________ (2004), "The Implementation Challenge of Pricing Decision Support Systems for Retail Managers." Pittsburgh, Penn.: Carnegie Mellon University, Working Paper.

__________, and Eric T. Bradlow (1999), "Why Analyst Overconfidence about the Functional Form of Demand Models Can Lead to Overpricing." *Marketing Science* 18 (4), 569–83.

Reibstein, David J., and Hubert Gatignon (1984), "Optimal Product Line Pricing: The Influence of Elasticities and Cross-Elasticities." *Journal of Marketing Research* 21 (3), 259–67.

Rossi, Peter E., and Greg M. Allenby (1993), "A Bayesian Approach to Estimating Household Parameters." *Journal of Marketing Research* 30 (2), 171–82.

Seetharaman, P. B., Siddhartha Chib, Andrew Ainslie, Peter Boatwright, Tat Chan, Sachin Gupta, Nitin Mehta, Vithala R. Rao, and Andrei Strijnev (2005), "Models of Multi-Category Choice Behavior." *Marketing Letters* 8 (3), 297–305.

Shankar, Venkatesh, and Ruth N. Bolton (2004), "An Empirical Analysis of Determinants of Retailer Pricing Strategy." *Marketing Science* 23 (1), 28–49.

__________, and L. Krishnamurthi (1996), "Relating Price Sensitivity to Retailer Promotional Variables and Pricing Policy." *Journal of Retailing* 72 (3), 249–73.

Vilcassim, Naufel J., and Pradeep K. Chintagunta (1995), "Investigating Retailer Product Category Pricing From Household Scanner Panel Data." *Journal of Retailing* 71 (2), 103–28.

Zenor, Michael J. (1994), "The Profit Benefits of Category Management." *Journal of Marketing Research* 31 (2), 202–13.