PRODUCTION AND OPERATIONS MANAGEMENT

POMS

Vol. 20, No. 5, September–October 2011, pp. 668–680
ISSN 1059-1478 | EISSN 1937-5956 | 11 | 2005 | 0668

DOI 10.3401/poms.1080.01195

Competitive Pricing in a Multi-Product Multi-Attribute Environment

Soulaymane Kachani
Department of Industrial Engineering and Operations Research, Columbia University, New York, New York 10027, USA,
kachani@columbia.edu

Kyrylo Shmatov
Department of Applied Physics and Applied Mathematics, Columbia University, New York, New York 10027, USA,
kis2101@columbia.edu

We address the problem of simultaneous pricing of a line of several products, both complementary products and substitutes, with a number of distinct price differentiation classes for each product (e.g., volume discounts, different distribution channels, and customer segments) in both monopolistic and oligopolistic settings. We provide a generic framework to tackle this problem, consider several families of demand models, and focus on a real-world case-study example. We propose an iterative relaxation algorithm, and state sufficient conditions for convergence of the algorithm. Using historical sales and price data from a retailer, we apply our solution algorithm to suggest optimal pricing, and report on numerical results.

Key words: dynamic pricing; revenue management
History: Received: October 2007; Accepted: April 2010 by Costis Maglaras, after 2 revisions.

1. Introduction

Recently, the problem of multi-attribute transactional pricing received a lot of attention. The papers by Elmaghraby and Keskinocak (2003), McGill and van Ryzin (1999), and Chiang et al. (2007) provide thorough overviews of the areas of dynamic pricing and revenue management. Earlier research concerning heuristic algorithms of product-line pricing may be found in Oren et al. (1984) and Dobson and Kalish (1988). Gallego and van Ryzin (1994) and Lin and Li (2004) also considered a multi-product dynamic pricing problem viewed as a dynamic programming problem, assuming that customers arrive sequentially according to a Poisson process. Kachani and Perakis (2006) borrow modeling tools and solution methods from traffic theory, and propose a fluid approach to multi-product pricing. In addition, Kachani et al. (2007) consider the problem of joint pricing and demand learning in a competitive environment.

A broad body of operations research literature investigates the comparison between the centralized and decentralized profits for firms offering a product line of differentiated products engaging in price competition. Farahat and Perakis (2009, 2010) provide upper and lower bounds on the ratio of profits under these two scenarios. Perakis and Roels (2007) consider a similar problem, measuring the efficiency of supply chains that use price-only contracts. Cachon (2004) and Lariviere and Porteus (2001) study the possible profit loss due to decentralization in a setting consisting of a manufacturer and a retailer. Bernstein and Federgruen (2003) focus on an analogous problem in a two-echelon distribution system, in which a supplier distributes a product to competing retailers, the demand rate for which depends on all of the retailers' prices similar to our premise. Adida and DeMiguel (2010) also study the competition in a supply chain with multiple manufacturers competing in quantities to supply a line of products to multiple retailers who compete in quantities to satisfy stochastic consumer demand. Perakis and Zaretsky (2008) approach the problem of competition in a supply chain with exogenous demand using the fluid models' framework.

The marketing literature focused on studying consumer demand characteristics is similarly vast. Dubé and Manchanda (2005) provide a very useful empirical study of the differences in demand traits based on geography. Tellis and Franses (2006) are concerned with the optimal data interval to estimate advertising carryover.

Our interest in the area of multi-product pricing stemmed from a close collaboration during the past several years between the authors and several firms, including a leading office products retailer. This retailer shared detailed historical operational data and business insights with the authors, which motivated the models proposed in the paper and led to the case study and numerical implementation discussed in this paper.

Based on our work in this paper, the approach we proposed was implemented by firms belonging to a wide range of business-to-business (B2B) industries, including chemicals, manufacturing, distribution and high-tech, and was implemented in a leading multi-purpose pricing software.

In multi-product pricing problems, the pricing decisions for different products are interdependent due to the fact that the demand for one product may depend on the prices of other products produced by the same firm or its competitors. Thus, efficient pricing should account for cross-elasticities among products. The complexity of a pricing problem grows significantly as the number of products increases. As a result of this and the inability to make accurate demand predictions, in most practical settings, multi-product dynamic pricing problems are approximately solved by decoupling across products and solving a large number of single-product problems. We propose a more efficient solution for multi-product pricing problems, which suggests simple but efficient pricing policies.

In this paper, we address the problem of optimal multi-product pricing, both in monopolistic and oligopolistic settings. We estimate the parameters of demand models using historical data and consider different demand models studied extensively in the econometrics literature. The most relevant demand models to our setting were introduced in Raz and Porteus (2003), Besanko et al. (1998), and Chan and Seetharaman (2004). The general problem is to maximize profits of a company producing a line of multi-attribute products, acting in a monopolistic or an oligopolistic setting. We consider five different models of consumer demand. In two of them, we assume that the demand observed for each product consists of a "deterministic" factor linear in prices and a factor accounting for the cumulative distribution function of the product's "reservation price."

The probabilistic "reservation price" factor is particularly important in B2B markets, especially in price-negotiation settings (see, e.g., Bichler et al. 2002). In these settings, customers issue quote requests for certain quantities of products, and expect to get lower quotes for higher quantities. The historical data available comes in the form of triples of the type (incoming volume request—outgoing price quote—incoming response on whether the deal is closed or rejected). Considering the "reservation price" factor appears to be crucial in dealing with problems of this nature and leads to important theoretical insights and practical gains.

In each of these cases, we propose iterative algorithms that deliver a centralized solution in a monopolistic setting, or a Nash equilibrium in the case of an oligopolistic setting. We study convergence of these algorithms and prove easy-to-apply sufficient conditions for achieving convergence. In the cases of linear and quadratic demand, the algorithm simplifies considerably, as the corresponding inner-step problem is solvable in closed form.

Also, a time-dependent extension to our models, which incorporates the dependence of demand and resulting optimal pricing policies on pricing history through repeated interactions, is an important contribution to the research on reference prices. This extension further develops previous results by Greenleaf (1995) and Popescu and Wu (2007).

We discuss how to obtain the necessary problem parameters from historical data of the case study, and analyze the developed algorithms and evaluate their convergence properties. We also compare centralized and game-theoretic solutions for a real-world problem with historically realized numbers and evaluate the impact of our solutions on the company's profits and the results of "cooperative" and "competitive" behavior with respect to the company's competitors.

In the following sections, we introduce a general framework for competitive multi-product multi-attribute pricing, where cross-interactions between different products and attributes are explicitly modeled. These attributes include volume discounts, distribution channels, delivery lead-times, customer segments, or time (dynamic pricing strategies under repeated customer interactions). We consider five families of demand models widely used in the econometrics literature, and propose an iterative relaxation algorithm that we solve for the cooperative solution in the centralized case, and for the Nash equilibrium in the competitive case. We state sufficient conditions for the algorithm convergence within the five respective demand models. We discuss our implementation and apply the developed algorithms to real-world problems, obtaining promising results for applications both in B2B and business-to-consumer (B2C) markets.

2. Main Problem Statement

2.1. Monopolistic Problem

The decision maker in our setting faces the decision process workflow shown in Figure 1. The input data, either in the form of monthly/weekly prices and sales volume levels or records of each transaction including unsuccessful ones (request-for-price records), are supplied to the demand forecast unit after appropriate handling. Then, one of the five widely used demand models is chosen, and the input data are regressed against those to obtain the model parameters.

These parameters are subsequently supplied to the price optimization engine. Taking into account the optimization objectives and exogenous constraints, non-linear optimization is used to compute the vector of optimal prices. After the suggested optimal prices are reviewed by the business intelligence unit and market impact is tested, the model may be subject to recalibration and adjustment of parameters given the real-life results obtained. At this stage, both changing demand parameters (possibly, switching to a different demand model) and re-optimizing for prices (adjusting the optimization problem constraints) may be undertaken.

Figure 1 Decision Process Workflow Diagram

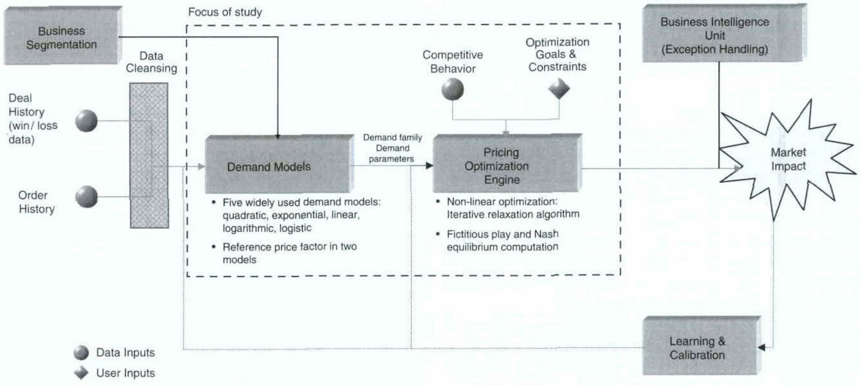

More details regarding the application of this framework are given in section 3.

We consider the optimization problem of simultaneous pricing of a line of N products, both complementary products and substitutes, having M distinct price differentiation classes for each product (e.g., volume discounts, distribution channels, delivery lead-times, customer segments).

Alternatively, our problem may be viewed as the pricing problem of a firm that produces and sells products based on a make-to-order production system so that the firm does not hold inventory of end-products. Then, one may consider average profit per unit-time over a chosen fixed interval $[0, T]$ as the objective function. Furthermore, one may consider M attributes also as the different lead-times quoted for a given product at different prices. At the same time, our framework is readily applicable to dynamic pricing strategies, viewing time as an extra attribute dimension in the presence of history-dependent demand (see repeated interactions extension below for more details).

The overall objective is to maximize the profit function $-\Pi$, defined as

$$-\Pi = \sum_{i,j=1}^{N,M} (p_{ij} - c_{ij}) D_{ij}, \tag{1}$$

subject to

$$c_{ij} \le p_{ij} \le p_{ij}^R, \quad i = 1, \ldots, N, j = 1, \ldots, M, \tag{2}$$

where c_{ij} and D_{ij} are the corresponding cost and demand, p_{ij}^R is the exogenous upper price threshold. We assume that the demand D_{ij} is a function of the price vector $\vec{p} = \{p_{ij}\}$. Note that we can easily extend our results to the case of arbitrary price limits, so that a condition of the type $p_{ij}^{\min} \le p_{ij} \le p_{ij}^{\max}$ will replace (2).

To estimate the dependence of the demand D_{ij} on prices of different products/product segments, we use the historical sales datasets gathered for both B2C (retail and consumer goods) and B2B (distribution) markets. In the first two cases outlined below, we assume that the demand observed for each product consists of an "aggregate" factor (linear in prices in the quadratic model) and a factor that accounts for the probability distribution of the "reservation price"—the highest price at which the consumer is willing to buy the product. In other words, we assume that the effective "market size" depends on a specified form of the vector of prices, and that each consumer in this market has a random reservation price that he/she compares to the posted price to decide whether to buy. The aggregate factor form is chosen among various models widely used in econometrics as the best fit to historical data available (using historical weekly or monthly sales volume and price data). When we applied in practice these demand models to many large datasets from various companies (both B2B and B2C), we found that they exhibited good statistical fit.

2.1.1. Main (Quadratic Demand) Case. We consider the case where the reservation price of a given product is distributed uniformly on the interval $[p_{ij}^R - \Delta_{ij}, p_{ij}^R]$. The product demand takes the form

$$D_{ij}(\vec{p}) = \frac{p_{ij}^R - p_{ij}}{\Delta_{ij}} \left(\alpha_{ij} - \beta_{ij} p_{ij} + \sum_{k \ne j} \Phi_{ij}^k p_{ik} + \sum_{k \ne i} \Theta_{ij}^k p_{kj} \right). \tag{3}$$

Thus, the problem is to minimize Π—the opposite of the profit function—

$$\Pi = \sum_{i,j=1}^{N,M} (p_{ij} - c_{ij}) \frac{p_{ij} - p_{ij}^R}{\Delta_{ij}} \times \left(\alpha_{ij} - \beta_{ij} p_{ij} + \sum_{k \neq j} \Phi_{ij}^k p_{ik} + \sum_{k \neq i} \Theta_{ij}^k p_{kj} \right), \quad (4)$$

so that

$$c_{ij} \leq p_{ij}^R - \Delta_{ij} \leq p_{ij} \leq p_{ij}^R, \quad i = 1, \ldots, N, j = 1, \ldots, M. \quad (5)$$

2.1.2. Exponential Distribution Case. In this case, we assume that the reservation price of a given product is distributed exponentially with parameter γ_{ij}, i.e., with cumulative distribution function

$$\mathbf{P}_{ij}(p) = 1 - e^{-\gamma_{ij} p}. \quad (6)$$

The product demand functional form is now

$$D_{ij}(\vec{p}) = e^{-\gamma_{ij} p_{ij}} \left(\alpha_{ij} - \beta_{ij} p_{ij} + \sum_{k \neq j} \Phi_{ij}^k p_{ik} + \sum_{k \neq i} \Theta_{ij}^k p_{kj} \right). \quad (7)$$

The profit function $-\Pi$ is maximized in the closed region, where $p_{ij}^R - \Delta_{ij} \leq p_{ij} \leq p_{ij}^R$.

2.1.3. Linear Demand Case. Considering the case where there is no probabilistic factor in the demand formula, we obtain

$$D_{ij}(\vec{p}) = \alpha_{ij} - \beta_{ij} p_{ij} + \sum_{k \neq j} \Phi_{ij}^k p_{ik} + \sum_{k \neq i} \Theta_{ij}^k p_{kj}. \quad (8)$$

In this case, the problem is easily tractable and reduces to a classical result.

In the next two cases, we also assume there is no reservation price factor involved, but use a more complicated functional form for the deterministic factor.

2.1.4. Constant Elasticity Demand Case. We consider the classical power-law form of demand function, namely

$$\ln D_{ij} = \alpha_{ij} - \beta_{ij} \ln p_{ij} + \sum_{k \neq j} \Phi_{ij}^k \ln p_{ik} + \sum_{k \neq i} \Theta_{ij}^k \ln p_{kj}. \quad (9)$$

This formula is used extensively in the econometrics literature, when the customer's choice is assumed to be continuous (see Talluri and van Ryzin 2004a for more details).

2.1.5. Multi-Nomial Logit (MNL) Case. We also consider another demand model used in the econometrics literature to account for discrete customer choice, the MNL model. Here, the *demand function* takes the form (see Besanko et al. 1998, Talluri and van Ryzin 2004a, b for more details)

$$D_{ij} = \frac{e^{\alpha_{ij} - \beta_{ij} p_{ij}}}{1 + \sum_{k,l=1}^{N,M} e^{\alpha_{kl} - \beta_{kl} p_{kl}}}. \quad (10)$$

Parameters in these models are estimated from historical data according to the following rules (refer to section 4 for more practical details):

- M ($j = 1, \ldots, M$) is the number of observable product attributes or customer segments. This number is exogenous in nature. In practice it is usually clear from historical sales data.
- Unit costs c_{ij} can be different across the product's attributes (e.g., different costs for different volume segments due to economies of scale or product attributes, and different costs to serve different customer segments).
- Parameters of the reservation price distribution (e.g., p_{ij}^R, Δ_{ij} in the quadratic case) are determined separately for each product i and product attribute j, for example, from constructing cumulative distribution functions of prices using historical data.
- α, β, Θ, and Φ—demand parameters introduced below—are estimated from historical sales data using linear regression for a given product segment. As it is customary to do so, we will assume that $\alpha_{ij} \geq 0$, $\beta_{ij} \geq 0$ for all $i = 1, \ldots, N$, $j = 1, \ldots, M$, while Θ and Φ can take any real values.

In most real-life cases encountered by the authors, quadratic and exponential models appear to have the most practical value (in terms of goodness of fit), with the quadratic case also providing simpler closed-form solutions. The linear demand model may be viewed as a simpler approximation reproducing classical results. Finally, the constant elasticity and MNL models are prevalent in the econometrics literature and are widely used to estimate aggregate customer demand, in continuous and discrete customer choice cases, respectively.

2.2. Competing Products Problem

We extend the previously stated monopolistic problem to the case where $l = 1, \ldots, L$ competing firms in the oligopolistic setting introduce products $i = 1, \ldots, N_l$ each; $j = 1, \ldots, M$, as before, enumerates differentiation classes of the same product/competitor. We assume that the number of products (not attributes) offered by the lth firm may change with l (if the jth attribute does not apply to a given product, we let the corresponding coefficient equal zero and skip it). We use quadratic demand model in this section, although the results are easily extendable to other demand functional forms, as has been done for the monopolistic case.

The utility of the lth firm is given by

$$U_l = \sum_{i,j=1}^{N_l,M} (p_{lij} - c_{lij}) \frac{p_{lij}^R - p_{lij}}{\Delta_{lij}} \times \left(\alpha_{lij} + \sum_{k \neq j} \Phi_{lij}^k p_{lik} + \sum_{m,k=1}^{L,N_m} \beta_{lij}^{mk} p_{mkj} \right). \tag{11}$$

The price vector belongs to the domain

$$K = \bigotimes_{l,i,j=1}^{L,N_l,M} \left[p_{lij}^R - \Delta_{lij}, p_{lij}^R \right]. \tag{12}$$

We are concerned with the existence of Nash equilibrium in this setting.

In the next section, we will discuss the real-world case study this research is based on, followed by the more technical discussion of the solution algorithm in the subsequent section.

3. Case Study and Practical Implementation

3.1. Context

We now consider the problem specified in the previous section within the settings of a real-world example, which acted as a motivation to this study. Through an industry contact, the authors had access to the sales records of a major office product retailer, and the corresponding records of its major competitor.

The market both companies are in is essentially a duopoly, with a less significant ($<25\%$) market share belonging to several smaller competitors. Most of the sales in both companies occur through two distribution channels: (a) dedicated retail stores and (b) major retail chains.

3.2. Data Handling and Demand Estimation Procedure

The handling of price/sales volume data involves two main steps. Raw price/sales data over a chosen interval of time (e.g., 1 month) are preprocessed into a dataset ready for demand estimation. Then, at step 2, demand parameters for a chosen functional form are estimated via regression on the data sample from step 1.

In the first step, the data are extracted in the form of time series of prices and sales volumes for each product/distribution channel (attribute). Evident outliers are removed along with certain points accounting for specific deals known *a priori* to be clearly irrelevant and hinder the dependence.

Data points are plotted on the price vs. sales volume graph. Specific choice of the time series subset (when necessary, for example, to eliminate major sales events, etc.) is explained. Sales volume data (if the time horizon is monthly) are seasonally adjusted.

The preprocessed data are an input to the regression module that estimates the demand parameters. For each product considered, we assume the demand for it to be a deterministic function of the prices of a subset of all products considered. Which products belong to such subset is a model choice. In practice, several alternative assumptions for the subset of relevant products may be considered, and the best fit chosen (such as the highest R^2 coefficient among the alternatives).

In our study, we adopted a demand functional form (see previous section for the functional forms of demand considered), where the demand of the given product offered through a particular channel depends on its own price, the prices of all products in the product line offered through the same channel, and the prices of the same product in other distribution channels (possessing different attributes). The prices of products differing from the one in question in more

Figure 2 Input Parameters for the 10 Products Produced by the Competing Firms ($L = 2, N = 10, M = 2$), Assuming Full Knowledge About Prices and Full Interaction Between Different Stock Keeping Units, Quadratic Model Case

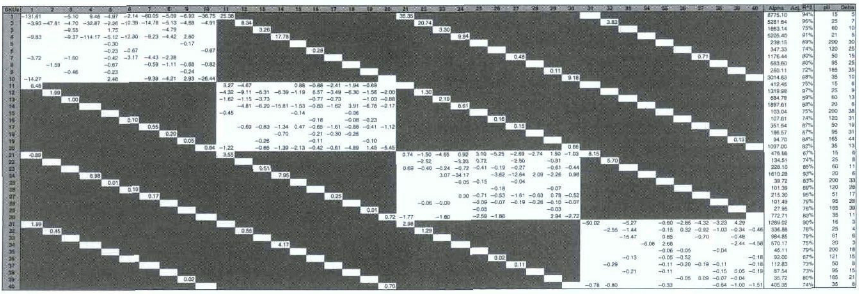

Note: Dark cells indicate terms excluded from the regression by the model adopted than one dimension (different product and different distribution channel, etc.) are not included. Figure 2 demonstrates the cross-dependency matrix between products. Only the white cells need to be filled via regression in step 2. Blue cells have no effect on the demand for a given product (defined by a given row in the matrix).

In the (main) case of quadratic functional form of demand, threshold price p_{ij}^R and price range Δ_{ij} (or respective parameters of the reservation price distribution, *if non-uniform*) are chosen using the data at hand (based on the data's range, etc.). The remaining demand factor (accounting for cross-product influence) is estimated via multi-variate linear regression.

A sales volume series for each product in the product line is linearly regressed against multiple time series of prices of relevant products. Price series with low explanatory power are then removed from the set of regressors. Series of predictor prices that are strongly correlated with other price series present are also removed, and the process is repeated until *satisfactory significance* level is reached.

3.3. Computational Results

We are concerned with how the company's optimal pricing choices and respective profits compare under different scenarios for the company facing a single competitor of comparable size, both having their sales largely concentrated in two distribution channels. First, we assume that one of five different scenarios holds, and compute the profits attained under each of these scenarios. Second, we discuss how the company would fare if it adopted the optimal pricing policy recommended under each scenario while, in fact, facing a "true" competitive environment (with the competitor firm responding to the price change with its own optimal pricing policies).

We consider the following five scenarios:

A. *Centralized optimization solution*: Both firms collude, seeking to maximize the sum of their profits, regardless of the distribution of profits between the two.

B. *Nash equilibrium*: The two firms compete, each assuming that the competitor will respond to any pricing policy changes in an optimal way.

C. *Optimization ignoring cross-product influence*: All cross-product elasticities are assumed to be zero; only the product's own price is assumed by the company to have an effect on the demand for the product.

D. *Optimization ignoring competition*: Cross-product elasticities for products offered by different firms are assumed to be zero; only the prices of products sold by the same firm are thought to have an effect on the demand for the product.

E. *Myopic optimization*: Cross-product influence is present; however, each firm assumes that the competitor holds its prices fixed and does not respond to the change.

These scenarios are compared with the *historically* realized case involving actual prices quoted for the products offered on the same time interval, and were not a result of a comprehensive optimization solution.

The output of the algorithm for the above cases (vs. different assumptions about the actual demand functional form) and the corresponding historical (unoptimized) performance data are presented in Table 1. The actual profit numbers are computed under three different assumptions: quadratic, linear, and exponential functional form of demand, as defined in the previous section. Note that the conclusions are similar under all three assumptions; the sensitivity to demand functional form is discussed in the next subsection.

The numbers in Table 1 are profit values for the firm under study, faced with a competitor of similar size/market power.

In Figure 3, we assume that the different optimization scenarios we consider (A through E) are true (e.g., in case C, that there is indeed no cross-product influence between the products) and compute the firm's profits.

As expected, the profit figures obtained show that the centralized solution A performs much better than the one realized historically on the interval considered (in the quadratic case, the centralized solution exceeds the historically realized profit by 44.4%).

In the competitive case B, the Nash equilibrium solution that utilizes all the interaction data leads, for the quadratic demand model, to a profit improvement of 18.3%. The centralized solution is superior to the Nash equilibrium case as expected.

The solution C that does not take into account cross-elasticities between products appears to be subordinate to the full-interaction centralized solution (exceeds the historical profit only by 14.3%, which is 30.1% less than the collusion case), although it is still advantageous compared with the historical one.

A myopic solution E where one firm optimizes its profits without having full knowledge about the competitor (having access to the competitor's prices but not

Table 1 Profit Values in Different Scenarios vs. Different Assumptions About the Actual Demand Functional Form

Profits	Quadratic (US$)	Linear (US$)	Exponential (US$)
Centralized optimization solution	799,098	777,385	799,076
Nash equilibrium point	654,388	651,894	645,411
Optimization ignoring cross-product influence	632,461	638,270	619,214
Optimization ignoring competition (no cross-firm influence)	666,877	658,811	650,442
Myopic optimization (assuming competitors have fixed prices)	657,952	656,833	643,967
Historically realized		553,220	

Figure 3 **Profits Under Different Optimization Scenarios Assuming These Optimization Scenarios are True (Quadratic Demand Case)**

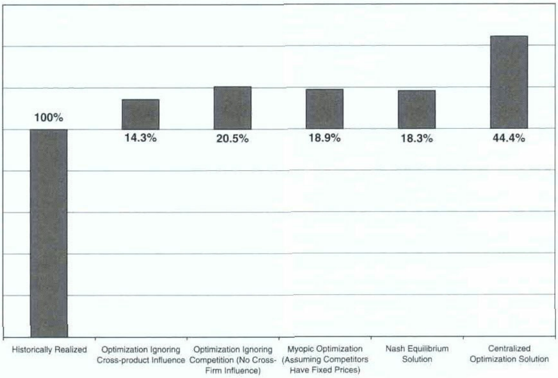

demand parameters) is also (by 23.9% of the historically realized value) worse than the centralized solution.

The experimental runs with real data described here yield in most cases to a vector of optimal prices higher on average than prices realized historically. However, this does not lead to a significant reduction in sales volume because of the cross-elasticity interactions between substitute products.

Differently from the assumption for Figure 3, we consider for the purpose of Figure 4 that firms adopt scenarios A–E, while, in fact, the reality is that both cross-product influence and competition exist. For instance, in case C, we now consider that the firm *mistakenly* ignores cross-product influence, while in case D, we assume that the firm *mistakenly* ignores the influence of competition. Not surprisingly, the profits in cases A and B remain unchanged because these cases properly account for cross-product influence and the existence of competition. However, we would expect cases C, D, and E to perform worse than in Figure 3 where those cases were assumed to reflect reality. In Figure 4, we can observe that neither the fixed-competitors case, nor the case when influence across firms is ignored properly account for the effect of competition. Cross-elasticities for the same products produced by different firms are usually positive (products are substitutes), while self-price elasticity is usually negative. As a result, in the case when a firm

Figure 4 **Profits Under Different Optimization Scenarios Assuming Cross-Product Influence and Competition Exist (Quadratic Demand Case)**

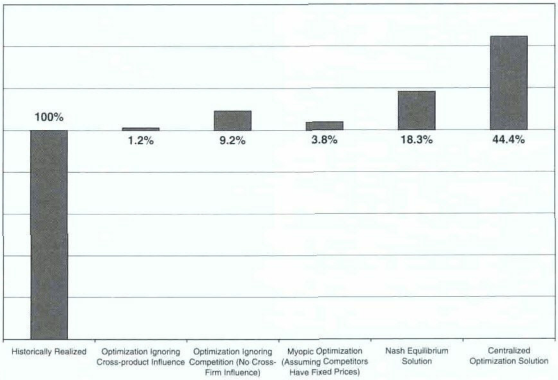

Table 2 Profits Attained when Adopting Optimal Pricing Policies, When Different Functional Demand Models are Relevant

	Optimal for quadratic (US$)	Optimal for linear (US$)	Optimal for exponential (US$)
Quadratic demand	799,098	796,774	770,487
Linear demand	772,408	777,385	740,392
Exponential demand	768,857	760,007	799,076

ignores cross-product interaction, it tends to price conservatively. However, if one accounts for the products' cross-elasticities while keeping competitors' prices fixed, one prices more aggressively, which leads to an even worse outcome due to the reaction of the competition. In other words, in many cases (e.g., when cost firms' cost structures and pricing powers are similar), it is better to ignore competition altogether than to partially (and erroneously) incorporate competitive dynamics (in this case, assuming competitors have fixed prices and performing one-step [myopic] optimization).

Finally, we explored the robustness of the results we obtained to the choice of demand functions. We compared the performance of optimal solutions found assuming a certain demand functional form under scenarios where the "true" demand function can be different. For instance, the vector of optimal prices found in the quadratic model is used to compute the optimal profit attained when profit functions of different models are used. Through this "sensitivity test" (for this and other companies and industries we looked at), we can provide some empirical evidence that the optimal profits under different demand models appear to be reasonably close to each other, each of them exceeding historical profits realized. Thus, the error in choosing a less appropriate demand model among the first three families of demand we used in this paper is not very significant. In Table 2, we provide a numerical illustration of this robustness for this case-study example.

3.4. Practical Application

We applied the methodology described above to over a dozen companies in various industries. Hence, the findings and insights discussed are fairly general and can be applied to a wide range of companies and industries, both in B2B and B2C settings.

Below, we discuss practical insights from the implementation of our models and algorithms across several dimensions:

Data quality and business rules: Data quality varies from one company to another, and even within the same company, from one division to another or from one region to another. Data quality issues include short historical horizon, missing data, erroneous data (e.g., very large or negative volumes and sales), censored data, and availability of data only at the aggregated level (e.g., by day or week or at product family level) instead of a granular level (i.e., transactional level by stock keeping unit). To remedy data quality issues, companies and pricing solution providers use a host of statistical methods to clean the data and estimate missing data.

When companies that suffer from poor data quality implement price optimization solutions, they invest in building robust information and data systems and a host of business rules that override demand-price elasticity coefficients and optimized prices when these numbers do not make business sense (e.g., demand-price elasticities that clearly have the wrong sign or implied price changes that are too high and would impact customer behavior). B2B companies tend to suffer more from poor data quality and are more vulnerable to the effect of that data quality because a small number of customers usually drive a large percentage of volume.

Most of the pricing business rules deal with limiting price changes, volume changes, or market share changes from one period to another, both at the product level and the product family level. In cases of high (respectively, low) inventory levels, constraints on the absolute level of sales are imposed to reduce inventory levels (respectively, avoid stockouts). In some B2B settings, business rules might override optimized prices in more than 50% of the cases until data quality improves. Even with high quality data, it is our experience that business rules still override at least 20% of optimized prices. The answer of some price solution vendors is to add business rules as constraints within the optimization models. However, this tends to be more of a marketing slogan: If added as constraints, many of these business rules would make convex optimization problems non-convex. Solution algorithms developed for convex problems cannot then be proven to converge to a global optimum. In effect, those algorithms then become heuristics, and the optimized prices they output tend anyway to hit the right-hand side of the price change constraints, hence yielding the same prices as if the business rules were implemented *a posteriori*.

Dealing with the impact of censored demand data is attracting more attention in practice. In retail settings, both retailers and consumer good manufacturers are investing in better information systems to capture magnitude of lost sales and time interval of stockouts, as well as implementing statistical tests and methods to correct for censored data.

Segmentation and demand models: Demand estimation and price optimization rely heavily on a good segmentation of the data across price differentiation classes. These price differentiation classes include customer segments, distribution channels, regions, volume discounts, and delivery lead-times. Price optimization vendors and large B2C companies include in their pricing systems automated segmentation models based on statistical and data mining techniques.

As discussed in the previous subsection, our optimization engine works equally fast for all five families of demand functions we introduced in the paper. From our experience with over a dozen companies, the quadratic, exponential, and linear demand models (in this order) outperformed the constant elasticity demand and the MNL models in terms of best statistical fit of the data. The MNL functional form of demand is less efficient for our problem, because it does not capture in the same way the interaction between different products and product attributes, and is harder to calibrate in practice.

Business and technological requirements: Most large firms rely more and more on versatile and industry-independent (as opposed to custom-coded) pricing systems that then offer high configurability and extensibility combined with high performance and scalability. Some systems boast the use of metadata-driven architecture and fully configurable user interfaces (including dashboards) and the ability to handle thousands of pricing requests per second and hundreds of simultaneous users with different roles and pricing (discount) power across many functions, including sales, marketing, finance, and operations.

The main reason behind these business (and implied technological) requirements is that price optimization is only one part of price management. The best price management systems combine best practices in pricing with advanced price optimization and world-class enterprise software. In many settings, the ability of pricing systems to interact with multiple database systems, enterprise resource planning systems, customer relationship management systems, and supply chain management systems, to input or output massive data (millions and sometimes billions of data points), and then aggregate and perform conversions (e.g., currency and unit of measure) and computations/analytics on the data in real time far outweighs in the mind of most executives the ability to obtain slightly more accurate elasticity estimations or optimized prices.

This is what prompted us to develop a generic framework and approach to multi-product multi-attribute pricing, use general demand models, and implement efficient price optimization algorithms that could perform very well under these business constraints. Although stand-alone and customized consulting engagements and software in the area of pricing have great value for strategic pricing decisions and new product launches, or in specific or niche industries, in most settings, tactical and operational dynamic pricing is moving toward the development of plug-and-play demand estimation modules and price optimization engines that can be incorporated without much effort in established price management software and run in real time.

4. Solution Algorithm and Its Convergence

4.1. Algorithm Description

In this section, we provide details on the solution algorithm used and discuss its convergence properties.

The mathematical programming and fixed-point algorithms for non-linear optimization each have their own advantages; however, separately, they either lack the generality or the computational efficiency that is necessary for solving large-scale models (Carey 1977, Harker and Pang 1990). Variational inequalities (VIs) were proposed (Dafermos 1983, Gabay and Moulin 1980, Nagurney 1993) to fill the gap created by mathematical programming and fixed-point approaches. Because of the efficiency required for solving large-scale problems and inherent diagonal dominance of the problem's Hessian, we use the non-linear iterative scheme by Dafermos (1983) within the VI framework, both in monopolistic and oligopolistic case of our problem. Non-linear iterative relaxation algorithm appears to outperform the linearized methods in similar equilibrium problem settings (Pang and Chan 1982).

We consider the centralized problem (1) in this section; later we extend it to the competitive case as well. Where appropriate, we index the price vector $\vec{p} = \{p_{ij}\}_{i,j=1}^{N,M}$ with a single letter $m = 1, \ldots, MN$; we denote by K the domain where the function is minimized, which is a product of corresponding price intervals.

The problem is solved using an iterative algorithm.

The general algorithm is as follows:

- *Initialization step* $m = 0$: Choose $p^R \in K$.
- *Step* m: Compute p^m as the solution of VI

$$\langle g(p^m, p^{m-1}), p - p^m \rangle \geq 0, \quad \forall p \in K. \tag{13}$$

- *Convergence verification*: Step $m = 1, 2, \ldots$ is repeated until

$$|p^m - p^{m-1}| \leq \varepsilon \tag{14}$$

for a predefined tolerance parameter ε.

Here, the generating function $\vec{g}$ is defined as

$$g_m(\vec{p}, \vec{s}) = \frac{\partial \Pi}{\partial p_m}(s_1, \ldots, s_{m-1}, p_m, s_{m+1}, \ldots, s_{MN}), \quad m = 1, \ldots, MN. \tag{15}$$

Convergence of the general algorithm is proven in the subsequent subsections.

On a large set of problems considered by Marcotte (1995), the relaxation algorithm performed better, in terms of computational time and accuracy, than the other considered algorithms such as Newton method and different linearization approaches. Furthermore, as investigated by Pang and Chan (1982), in terms of practical performance, one should expect that the closer the Hessian is to being symmetric, the better the symmetrized Newton method will behave. In the case of the nonlinear relaxation algorithm, the more diagonally dominant the Hessian is (as in the problem at hand), the better the relaxation method is expected to perform.

The solution of the inner subproblem (13) is equivalent to minimizing a function of the form

$$P = \sum_{k=1}^{MN} P_k(s_1, \ldots, s_{k-1}, p_k, s_{k+1}, \ldots, s_{MN}) \qquad (16)$$

with $p = p^m$, $s = p^{m-1}$. This problem splits into MN one-dimensional conditions. In the interior of the feasible set, these conditions are written as

$$\frac{\partial P_k}{\partial p_k}(p_k, \vec{s}) = 0, \quad k = 1, \ldots, MN. \qquad (17)$$

In the case of quadratic demand, these first-order conditions are one-dimensional quadratic polynomials, and the solution is therefore available in closed form. In cases of more involved functional form of demand, a one-dimensional equation is solved numerically.

4.2. Algorithm Convergence

For brevity, we formulate the convergence conditions only in the main case of quadratic demand. Analogous results for the other demand functional forms considered are provided in the supporting information Appendix S1.

The convergence of the relaxation algorithm applied to the general problem of minimizing Π is governed by the following proper diagonal dominance condition on the objective function's Hessian:

THEOREM 1. *A sufficient condition for the relaxation algorithm to converge is*

$$\inf_{p_{ij}} \frac{\partial^2 \Pi}{\partial p_{ij}^2} > \sup_{p_{ij}} \left\{ \sum_{k \neq j} \left| \frac{\partial^2 \Pi}{\partial p_{ij} \partial p_{ik}} \right| + \sum_{k \neq i} \left| \frac{\partial^2 \Pi}{\partial p_{ij} \partial p_{kj}} \right| \right\},$$
$$\forall i = 1, \ldots, N, \; j = 1, \ldots, M. \qquad (18)$$

It is important that the condition (18) automatically provides the problem's convexity, and, as a result, the uniqueness of the minimization problem's solution that the algorithm converges to.

Define the positive semi-axis characteristic function

$$\chi(x) = \frac{\operatorname{sgn} x + 1}{2}. \qquad (19)$$

Now,

THEOREM 2. *A sufficient condition for the relaxation algorithm to converge in the quadratic case is*

$$\frac{2}{\Delta_{ij}} \left(\alpha_{ij} - \beta_{ij}(2p_{ij}^R - c_{ij}) + \sum_{k \neq j} \Phi_{ij}^k \left(p_{ik}^R - \chi(\Phi_{ij}^k)\Delta_{ik} \right) + \sum_{k \neq i} \Theta_{ij}^k \left(p_{kj}^R - \chi(\Theta_{ij}^k)\Delta_{kj} \right) \right) >$$
$$\sum_{k \neq j} \left(\frac{|\Phi_{ij}^k|}{\Delta_{ij}} (p_{ij}^R - c_{ij}) + \frac{|\Phi_{ik}^j|}{\Delta_{ik}} (p_{ik}^R - c_{ik}) \right) + \sum_{k \neq i} \left(\frac{|\Theta_{ij}^k|}{\Delta_{ij}} (p_{ij}^R - c_{ij}) + \frac{|\Theta_{kj}^i|}{\Delta_{kj}} (p_{kj}^R - c_{kj}) \right), \quad \forall i, j. \qquad (20)$$

The condition (20) is not restrictive and holds with real data that we gathered in B2B and B2C settings (see discussion in section 5).

In the general multi-product case, the Hessian matrix has sparse structure (see the proof in the supporting information Appendix S1 for details). One may also consider the situation where only interaction between attributes j and $j \pm 1$ exist; in this case only two subdiagonals are non-zero on each side.

It is important to mention that Theorems 1 and 2 give sufficient conditions for the algorithm convergence. In fact, the algorithm may converge when these conditions are not met. It may be shown (by replacing K in the proof of Theorem 1) that, if the objective function's Hessian is positive definite in a small subdomain of K, the algorithm still converges if the initial point is chosen appropriately within this subdomain. Choosing a smaller domain for K gives weaker conditions that lead to important heuristic uses of the algorithm proposed. For example, as mentioned before, we can formulate a condition analogous to the one in Theorem 2 assuming that the prices change in the region $p_{ij}^{\min} \leq p_{ij} \leq p_{ij}^{\max}$ with arbitrary boundaries independent of Δ_{ij}.

Similar convergence results hold for the other four functional forms of demand. For their exact formulations and proofs, see the supporting information Appendix S1.

It is also worth mentioning that the respective model Hessians are not always diagonally dominant (the diagonal dominance condition holds for every separate row). This is due to the fact that the products considered may differ significantly in price (and therefore parameter scale). However, in practice, in most of the datasets and companies we considered, the algorithm exhibited reliable convergence, to a unique optimum. In cases where the input sales data were not of good quality and multiple local optima existed, we followed practice by enhancing the algorithm with straightforward business rules (heuristics) to ensure that the algorithm converges to prices and expected demand levels that make business sense.

Furthermore, when implemented in Excel and VBA, the algorithm converges quickly (in <5 seconds, including preprocessing, choice of best fit demand function, and estimation of demand parameters) for medium-size problem instances (10 products, 4 attributes, 5 years of historical sales data). Finally, when implemented into software, all the computations of optimal prices can be achieved in real time for large divisions of Fortune 500 corporations (using sharp stopping criteria [e.g., $\varepsilon = 10^{-6}$]).

4.3. Competing Products Problem

Here, we consider the case where $l = 1,\ldots,L$ competing firms in the oligopolistic setting introduce products $i = 1,\ldots,N_l$ each; $j = 1,\ldots,M$, as before, enumerates differentiation classes of the same product/competitor. We assume that the number of products (not attributes) offered by the lth firm may change with l (if the jth attribute does not apply to a given product, we let the corresponding coefficient equal zero and skip it). We use quadratic demand model in this section, although the results are easily extendable to other demand functional forms, as has been done for the monopolistic case.

The utility of the lth firm is given by (11). We are concerned with the existence of Nash equilibrium in this setting.

Gabay and Moulin (1980) states the following:

THEOREM 3. *p^* is a Nash equilibrium if and only if $p^* \in K$ is a solution of the VI*

$$\langle F(p^*), p - p^* \rangle \geq 0, \quad \forall p \in K, \tag{21}$$

where

$$F(p) = \left(-\frac{\partial U_1}{\partial p_{111}}, \ldots, -\frac{\partial U_1}{\partial p_{1N_1M_1}}, -\frac{\partial U_2}{\partial p_{211}}, \ldots, -\frac{\partial U_L}{\partial p_{LN_LM_L}} \right). \tag{22}$$

In our case,

$$\begin{aligned}\frac{\partial U_l}{\partial p_{lij}} = {} & \frac{1}{\Delta_{lij}}(p^R_{lij} + c_{lij} - 2p_{lij}) \\ & \times \left(\alpha_{lij} + \sum_{k \neq j} \Phi^k_{lij} p_{lik} + \sum_{m,k=1}^{L,N_m} \beta^{mk}_{lij} p_{mkj} \right) \\ & + \sum_{k \neq j} \Phi^j_{lik} \frac{1}{\Delta_{lik}} (p_{lik} - c_{lik})(p^R_{lik} - p_{lik}) \\ & + \sum_{k} \beta^{li}_{lkj} \frac{1}{\Delta_{lkj}} (p_{lkj} - c_{lkj})(p^R_{lkj} - p_{lkj}).\end{aligned} \tag{23}$$

The initial methods used to compute Nash equilibria were based on the constructive proof by Lemke and Howson (1964) of the existence of an equilibrium for a bimatrix game, and have come to be known as fixed-point methods. Unfortunately, fixed-point methods have experienced tremendous difficulties in solving large-scale problems. The other traditional approach to solving equilibrium models is mathematical programming. As discussed in Carey (1977), reformulating as a tractable non-linear programming problem required very restrictive assumptions on the model. As a result, solving equilibrium models as optimization problems does not provide a satisfactory alternative to fixed-point methods. As also discussed above for the monopolistic problem, the non-linear iterative relaxation algorithm outperforms such methods in similar equilibrium problem settings.

We solve this problem using the relaxation algorithm described earlier adjusting (15) appropriately by replacing the objective $-\Pi$ with the utilities of each competitor U_l. Theorem 1 then gives sufficient conditions for its convergence when Π is replaced with $-U_l$.

Using the positive semi-axis characteristic function defined above, we formulate the main convergence condition for the competitive case:

THEOREM 4. *A sufficient condition (under the quadratic demand model) for the relaxation algorithm to converge to a Nash equilibrium is*

$$\begin{aligned} & \frac{2}{\Delta_{lij}} \left(\alpha_{lij} + \sum_{k \neq j} \Phi^k_{lij} \left(p^R_{lik} - \chi(\Phi^k_{lij}) \Delta_{lik} \right) \right. \\ & \left. + \sum_{m,k=1}^{L,N_m} \beta^{mk}_{lij} \left(p^R_{mkj} - \chi(\beta^{mk}_{lij}) \Delta_{mkj} \right) \right) \\ & - \frac{2|\beta^{li}_{lij}|}{\Delta_{lij}} (p^R_{lij} - c_{lij}) \\ & > \sum_{k \neq j} \left(\frac{|\Phi^k_{lij}|}{\Delta_{lij}} (p^R_{lij} - c_{lij}) + \frac{|\Phi^j_{lik}|}{\Delta_{lik}} (p^R_{lik} - c_{lik}) \right) \\ & + \sum_{m \neq l} \sum_{k} \frac{|\beta^{mk}_{lij}|}{\Delta_{lij}} (p^R_{lij} - c_{lij}) \\ & + \sum_{k \neq i} \left(\frac{|\beta^{lk}_{lij}|}{\Delta_{lij}} (p^R_{lij} - c_{lij}) + \frac{|\beta^{li}_{lkj}|}{\Delta_{lkj}} (p^R_{lkj} - c_{lkj}) \right). \end{aligned} \tag{24}$$

5. Repeated Interactions Extension

The framework set up above is readily applicable to yet another important extension—dynamic pricing strategies, viewing time as an extra attribute dimension. Our approach provides an efficient method to further elaborate and extend recent research (Greenleaf 1995, Popescu and Wu 2007) on dynamic pricing in the presence of history-dependent demand.

The first step is to consider a system optimal pricing solution for a set of $i = 1, \ldots, N$ products, with the other index $j = 1, \ldots, T$ enumerating a finite set of moments in time. Thus, we consider the main problem analogous to the multi-attribute pricing case

$$\max_{p_i^R - \Delta_i \le p_{ij} \le p_i^R} \sum_{i,j=1}^{N,M} (p_{ij} - c_i) \frac{p_i^R - p_{ij}}{\Delta_i} \times \left(\alpha_i + \sum_k \Phi_{ik} p_{kj} + \sum_{k<j} \Theta_{ik} p_{ik} \right). \tag{25}$$

Technically, we can regress this demand form against the available data as is (in full generality with respect to Θ parameters above), but this may not be the optimal solution, as we may not capture the behavior of demand elasticities with respect to past prices (Θ factors) correctly in this way.

Instead, two more attractive ways are as follows:

A. We assume that the customer memorizes only the previous price he has seen for a given product; thus we only let Θ_{ij-1} be non-zero and regress such model. Here

$$D_{ij} = \alpha_i + \sum_k \Phi_{ik} p_{kj} + \Theta_i p_{ij-1}. \tag{26}$$

B. We introduce the reference price with exponential smoothing:

$$D_{ij} = \alpha_i + \sum_{k \ne i} \Phi_{ik} p_{kj} + \Psi_i (p_{ij} - r_{ij}), \tag{27}$$

$$r_{ij+1} = \kappa r_{ij} + (1 - \kappa) p_{ij} \tag{28}$$

with characteristic length $1/\kappa$.

Excluding r_{ij} using

$$p_{ij} - r_{ij} = p_{ij} + (\kappa - 1) p_{ij-1} + \kappa(\kappa - 1) p_{ij-2} + \cdots + \kappa^{j-2} p_{i1}, \tag{29}$$

we come up with the rule to define Θ parameters in the reference price model.

Using the representation of the reference price in terms of time-dependent prices

$$r_{ij} = \kappa \sum_{k=1}^{j-1} (1 - \kappa)^k p_{ij-k}, \tag{30}$$

we can recast the problem within the basic system–optimal framework used before in (1).

We can derive a convergence condition in terms of the parameters analogous to Theorem 2. The method has the same efficiency in the centralized solution model as dynamic programming. An important improvement of memory-dependent theory may be obtained through competitive extension. In the competitive case, our algorithm becomes more efficient than the classical one. Furthermore, properties of competitive strategies over time may be highly non-trivial.

6. Conclusion

To summarize, in this paper we introduced a generic framework for competitive multi-product multi-attribute pricing with explicitly modeled cross-interactions between different products and attributes. We proposed five different families of demand models widely used in econometrics literature. We described a relaxation algorithm, which we subsequently applied to compute the cooperative solution in the centralized case and the Nash equilibrium solution in the competitive case, and provided sufficient conditions for the algorithm's convergence. We implemented and applied the developed algorithm to a real-world case study, and derived insights.

From a practical viewpoint, our analysis develops an efficient and widely applicable way to address multi-product pricing problems that incorporates cross-product elasticity effects. When applied in practice to a wide range of companies and industries, our approach led to significant revenue improvements over typical heuristic approaches to multi-product pricing.

References

Adida, E., V. DeMiguel. 2010. Supply chain competition with multiple manufacturers and retailers. *Oper. Res.* (forthcoming).

Bernstein, F., A. Federgruen. 2003. Pricing and replenishment strategies in a distribution system with competing retailers. *Oper. Res.* **51**(3): 409–426.

Besanko, D., S. Gupta, D. Jain. 1998. Logit demand estimation under competitive pricing behavior: An equilibrium framework. *Manage. Sci.* **44**(11): 1533–1547.

Bichler, M., J. Kalagnanam, K. Katircioglu, A. J. King, R. D. Lawrence, H. S. Lee, G. Y. Lin, Y. Lu. 2002. Applications of flexible pricing in business-to-business electronic commerce. *IBM Syst. J.* **41**(2): 287–302.

Cachon, G. 2004. The allocation of inventory risk in a supply chain: Push, pull, and advance-purchase discount contracts. *Manage. Sci.* **50**(2): 222–238.

Carey, M. 1977. Integrability and mathematical programming models: A survey and parametric approach. *Econometrica* **45**: 1957–1976.

Chan, T. Y., P. B. Seetharaman. 2004. Estimating dynamic pricing decisions in oligopolistic markets: An empirical approach using micro- and macro-level data. Working paper, Washington University.

Chiang, W.-C., J. Chen, X. Xu. 2007. An overview of research on revenue management: Current issues and future research. *Int. J. Rev. Manage.* **1**(1): 97–128.

Dafermos, S. 1983. An iterative scheme for variational inequalities. *Math. Program.* **26**: 40–47.

Dobson, G., S. Kalish. 1988. Positioning and pricing a product line. *Mark. Sci.* **7**: 107–125.

Dubé, J.-P., P. Manchanda. 2005. Differences in dynamic brand competition across markets: An empirical analysis. *Mark. Sci.* **24**(1): 81–95.

Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices and future directions. *Manage. Sci.* **49**: 1287–1309.

Farahat, A., G. Perakis. 2009. Profit loss in differentiated oligopolies. *Oper. Res. Lett.* **37**(1): 43–46.

Farahat, A., G. Perakis. 2010. Price competition among multiproduct firms. Working paper, Available at http://ssrn.com/abstract=1121904

Gabay, D., H. Moulin. 1980. On the uniqueness and stability of Nash equilibria in noncooperative games. Bensoussan A., P. Kleindorfer, C. S. Tapiero. eds, *Applied Stochastic Control in Econometrics and Management Science*. North-Holland, Amsterdam, 271–292.

Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Manage. Sci.* **40**(8): 999–1020.

Greenleaf, E. A. 1995. The impact of reference price effects on the profitability of price promotions. *Mark. Sci.* **14**: 82–104.

Harker, P. T., J. S. Pang. 1990. Finite-dimensional variational inequalities and nonlinear complementary problems: A survey of theory, algorithms and applications. *Math. Program.* **48**: 161–220.

Kachani, S., G. Perakis. 2006. Fluid dynamics models and their applications in transportation and pricing. *Eur. J. Oper. Res.* **170**: 496–517.

Kachani, S., G. Perakis, C. Simon. 2007. Modeling the transient nature of dynamic pricing with demand learning in a competitive environment. T. Friesz, ed., *Network Science, Nonlinear Science and Game Theory Applied to the Study of Infrastructure Systems*. Springer, New York, 223–268.

Lariviere, M. A., E. L. Porteus. 2001. Selling to the newsvendor: An analysis of price-only contracts. *Manuf. Serv. Oper. Manage.* 3(4): 293–305.

Lemke, C. E., J. T. Howson, Jr. 1964. Equilibrium points of bimatrix games. *J. Soc. Ind. App. Math.* **12**: 413–423.

Lin, K. Y., F. Li. 2004. Optimal dynamic pricing for a line of substitutable products. Working paper, Virginia Polytechnic Institute.

Marcotte, P. 1995. Advantages and drawbacks of variational inequality formulations. Giannessi, F., A. Maugeri eds, *Variational Inequalities and Network Equilibrium Problems*. Plenum Press, New York, 179–194.

McGill, J. I., G. J. van Ryzin. 1999. Revenue management: Research overview and prospects. *Trans. Sci.* **33**: 233–256.

Nagurney, A. 1993. *Network Economics: A Variational Inequality Approach*. Kluwer Academic Publisher, Norwell, MA.

Oren, S., S. Smith, R. Wilson. 1984. Pricing a product line. *J. Bus.* **57**: 73–99.

Pang, J. S., D. Chan. 1982. Iterative methods for variational and complementarity problems. *Math. Program.* **24**: 284–313.

Perakis, G., G. Roels. 2007. The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Manage. Sci.* **53**: 1249–1268.

Perakis, G., M. Zaretsky. 2008. Multiperiod models with capacities in competitive supply chain. *Prod. Oper. Manag.* **17**(4): 439–454.

Popescu, I., Y. Wu. 2007. Dynamic pricing strategies with reference effects. *Oper. Res.* **55**(3): 413–429.

Raz, G., E. L. Porteus. 2003. Empirical estimation of stochastic consumer demand functions. Working paper, Stanford University. Available at http://www2.agsm.edu.au/agsm/web.nsf/AttachmentsByTitle/GalRaz-EmpiricalEstimationofStochasticConsumerDemandFunctions/$FILE/Empirical+Estimation+of+Stochastic+Consumer+Demand+Functions.pdf (accessed date October 27, 2010).

Talluri, K., G. van Ryzin. 2004a. *The Theory and Practice of Revenue Management*. Kluwer, Amsterdam.

Talluri, K., G. van Ryzin. 2004b. Revenue management under a general discrete choice model of consumer behavior. *Manage. Sci.* **50**(1): 15–33.

Tellis, G. J., P. H. Franses. 2006. Optimal data interval for estimating advertising response. *Mark. Sci.* **25**(3): 217–229.

Supporting Information

Additional supporting information may be found in the online version of this article:

Appendix S1. Proofs and Supplementary Cases.

Please note: Wiley-Blackwell is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

PRODUCTION AND OPERATIONS MANAGEMENT
Vol. 20, No. 5, September–October 2011, pp. vii–ix
ISSN 1059-1478 | EISSN 1937-5956 | 11 | 2005 | 0vii
DOI 10.1111/j.1937-5956.2011.01184.x

Management Insights

The Manufacturer's Incentive to Reduce Lead Times
Santiago Kraiselburd, Richard Pibernik, Ananth Raman

The prevailing wisdom in supply chain management says that all else remaining equal, managers should seek to make their supply chains more responsive. In other words, manufacturers should seek to reduce the leadtime to respond to orders from their retailers because this would allow them to reduce stockouts in the supply chain and, hence, would result in higher sales. However, under some circumstances, the opposite is possible. That is, increasing supply chain responsiveness (or reducing leadtimes) can lead to *lower* sales for the manufacturer. Two effects drive sales to be lower with shorter leadtimes: a "safety stock "effect (which comes from the idea that, if lead times are long, retailers must "protect" their service levels by keeping a large inventory, while short lead times decrease the need for such protection), and an "effort effect" (which is driven by the efforts that retailers exert to sell more products). As it turns out, these two effects interact in non intuitive ways, and must be considered before a decision to reduce lead times is implemented.

Real-Time Delay Estimation Based on Delay History in Many-Server Service Systems with Time-Varying Arrivals
Rouba Ibrahim, Ward Whitt

Waiting customers in service systems, such as a hospital emergency department or a call center, are typically unable to estimate their own delay. A long wait, coupled with feelings of uncertainty about the length of that wait, leads to poor service evaluation. For system managers, making delay announcements is a relatively inexpensive way of reducing customer uncertainty about delays, thereby improving customer satisfaction with the service provided. The authors investigate alternative ways to *estimate*, in real time, the delay of an arriving customer in a service system. These delay estimates may be used to make delay announcements. The authors focus especially on delay estimators exploiting recent customer delay history in the system. Delay-history-based estimators are appealing for complicated service systems because they do not exploit information about system parameters and therefore adjust automatically to changes in those parameters. The authors consider the realistic feature of time-varying arrival rates. They show that time-varying arrival rates can introduce significant estimation bias in delay-history-based delay estimators when the system experiences alternating periods of overload and underload. They introduce refined delay-history estimators that effectively cope with time-varying arrival rates together with non-exponential abandonment-time distributions, which are often observed in practice.

Competitive Pricing in a Multi-Product Multi-Attribute Environment
Soulaymane Kachani and Kyrylo Shmatov

In multi-product pricing problems, the pricing decisions for different products are interdependent due to the fact that the demand for one product may depend on the prices of other products produced by the same firm or its competitors. Thus, efficient pricing should account for cross-elasticities among products. The complexity of a pricing problem grows significantly as the number of products increases. As a result of this and the inability to make accurate demand predictions, in most practical settings, multi-product dynamic pricing problems are approximately solved by decoupling across products and solving a large number of single-product problems. In this paper, the authors propose a generic framework and approach to multi-product multi-attribute pricing, use general demand models and develop efficient price optimization algorithms that they have successfully tested and implemented in several industries.

Optimizing Customer Forecasts for Forecast-Commitment Contracts
Elizabeth J. Durango-Cohen, Candace A. Yano

Suppliers of customized products are seeking ways to work with customers to reduce excess production while simultaneously providing greater reliability of supply. Cohen and Yano propose a Forecast-Commitment (FC) contract that resembles what the ASIC manufacturer that motivated this study was already doing vis-à-vis information exchange with customers. But the contract also includes incentives for the parties to act in their mutual