

Dynamic Pricing and the Direct-to-Customer Model in the Automotive Industry*

Stephan Biller[†]
Lap Mui Ann Chan[‡]
David Simchi-Levi[§]
Julie Swann[¶]

forthcoming April 2005
Electronic Commerce Journal, Issue on Dynamic Pricing

Abstract

The Internet is changing the automotive industry as the traditional manufacturer and dealer structure faces increased threats from third party e-tailers. Dynamic pricing together with the Direct-to-Customer business model can be used by manufacturers to respond to these challenges. Indeed, by coordinating production and inventory decisions with dynamic pricing, the automotive industry can increase profits and improve supply chain performance. To illustrate these benefits, we discuss a strategy that incorporates pricing, production scheduling, and inventory control under production capacity limits in a multi-period horizon. We show that under concave revenue curves, a greedy algorithm provides the optimal solution, and we describe extensions to the model such as multiple products sharing production capacity. Using computational analysis, we quantify the profit potential and sales variability due to dynamic pricing, and we suggest that it is possible to achieve significant benefit with few price changes.

1 Introduction

The influence of the Internet and e-commerce on the economy in general, and supply chain management in particular, has been tremendous. The ability to dynamically change pricing of products is an important revolution in the retail and manufacturing industries, driven in large part by the Internet and the Direct-to-Customer (DTC) model. This business model, used by industry giants such as Dell Computers and Amazon.com, allows companies to quickly and easily change prices based on parameters such as demand variation, inventory levels, or production schedules. Further, the model enables manufacturers to collect demand data more easily and accurately (Washington Post 2000).

Many examples of dynamic pricing can be found on the Internet; for example, online auctions allow buyers to bid on everything from spare parts to final goods. However, the *integration* of pricing, production and distribution decisions in retail or manufacturing environments is still in its early stages. One recent example is Dell Computers, which uses both dynamic pricing of products and pricing based on market segmentation (Wall Street Journal 2001); the extent to which pricing is integrated with production and inventory control decisions at Dell is unclear. The implementation of dynamic pricing strategies in the manufacturing and

*Research supported in part by ONR Contracts N00014-90-J-1649 and N00014-95-1-0232, NSF Contracts DDM-9322828, DMI-9732795 and DMI-0085683, a research grant from the Natural Sciences and Research Council of Canada (NSERC) a grant from General Motors, and a grant from the eBusiness Center at MIT.

[†]General Motors, Research and Development Center, Warren, MI

[‡]School of Management, University of Toronto, Toronto, Ontario

[§]Dept. of Civil and Environmental Engineering and the Engineering Systems Division, MIT

[¶]School of Industrial and Systems Engineering, Georgia Institute of Technology

retail industries has the potential to radically improve supply chain efficiencies in much the same way as revenue management has changed airline, hotel and car rental companies.

Our focus on dynamic pricing as a tool to improve supply chain efficiency in manufacturing is motivated by a collaborative effort with a manufacturer of automobiles. We focus on the coordination of pricing, production and distribution decisions for non-perishable products in a multi-period time horizon, where planning in advance is a key element. We allow for periodically varying parameters such as demand curves, capacity limits, holding costs and production costs. The individual revenue curves are concave, and the objective is to maximize profit; general properties of the problem allow it to be solved efficiently with the greedy algorithm. Extensions to the model include the addition of production set-up cost and allowing multiple products to share common production capacity.

An extensive computational study with demand curves received from our industrial partner demonstrates that the integration of supply chain functions such as pricing and production may result in large benefits. The most obvious is due to an increase in profit based on a better match between supply and demand. However, other, perhaps less obvious, benefits can be significant as well. These benefits include reduction in demand variability seen by the manufacturer, which in turn results in smoother production schedules, more stable ordering policies, etc. We consider performance measures such as profit, variability in sales, and price variability. Our objective is to provide insight into the impact of dynamic pricing on supply chain performance using data from the automotive industry.

1.1 Industrial Motivation

The driving forces behind the interest in our industrial partner to explore dynamic pricing strategies are the Internet and the DTC model. The DTC model offers significant benefits to the manufacturer, including better demand information and increased flexibility in matching supply and demand.

The use of the DTC model in the automotive retail industry has been growing. In 1999, 40% of all new vehicle buyers used the Internet during their shopping; this number is estimated to grow to 55% in 2000 (J.D Power Automotive and Associates 2000). While currently most customers use the Internet to inquire invoice prices but purchase the car from franchise dealers, there are many third parties (www.AutobyTel.com or Microsoft's CarPoint at www.carpoint.msn.com, etc.) trying to convince customers to use their buying services. Indeed, in the last few years we have seen an evolution in the Internet pricing arena from information sites where invoice prices can be found, to referral sites where price quotes can be requested, and finally to sites that post transaction prices and sometimes allow purchasing on the web.

The growth of these developments on the Internet is of great concern to the Original Equipment Manufacturers (OEMs) and the traditional dealers. First, disintermediation might result in third parties rather than dealers "owning" customer relations and consumer data. Second, if a specific third party becomes too powerful (e.g., Microsoft CarPoint), it could dictate prices and wipe out manufacturers' margins. Recognizing the opportunity and, at the same time, fighting the threat of third parties, OEMs have launched their own DTC e-commerce initiatives; General Motors even founded a new division called e-GM, and at the time of the paper was contracting with AutobyTel for special Internet services. Today, e-commerce is a way of doing business at GM and the capabilities, which were rapidly developed in the former e-GM, have been successfully re-integrated into the relevant GM organizations (Purchasing, Order Fulfillment, and Sales).

These e-initiatives are aimed at building a system that will eventually allow customers to order custom-built vehicles, and at the same time will enable manufacturers, in cooperation with their dealers, to coordinate production and pricing. Thus, OEMs and their dealers can simultaneously optimize system performance while balancing supply and demand.

Typically, manufacturers' prices are driven by a number of factors including competition, capacity utilization, market share targets, and profits. Dealers' prices are primarily driven by automobile cost and market forces. Currently, in the automotive industry, dynamic pricing occurs at two levels: (1) at the dealer, who negotiates with each customer, (2) at the manufacturer, who makes extensive use of promotion and rebates, but very rarely of price increases. In fact, competitive price pressures led to a decline of invoice prices (the price a dealer pays to the manufacturer) by 0.7% in 1999 (J.D. Power and Associates). A rare

exception is Chrysler's PT Cruiser whose price was raised a couple of times by the manufacturer even before the vehicle was released for production. This, however, is a case where the manufacturer priced the vehicle too low initially and realized that they would not have enough capacity to satisfy demand or that customer wait-times would increase to extraordinary lengths. A third possibility for dynamic pricing is between OEMs and dealers, although this strategy is not currently in use.

Initial discussions with our OEM partner motivated the consideration of integrated pricing and production policies. Subsequent discussions shaped the assumptions of the model as well as the focus of the research. The research, however, is intentionally applicable in a very general manufacturing setting where the manufacturer has price control, disregarding elements specific to automotive retail such as the dealer system. The models presented would be applicable in an Internet-based automotive retail process where price may go up or down and customers expect to find a non-negotiable price and delivery time on the web.

In addition to the focus of the theoretical research, the collaborative effort with the OEM greatly shaped the computational analysis presented in this paper. The analysis is based on demand curves and production costs which were obtained from the OEM and represent typical vehicles from its product profile. The discussions with the OEM have been of significant benefit in determining useful questions and reasonable approaches to addressing them.

The remainder of this paper is organized as follows. Section 2 surveys dynamic pricing and implementations in various industries. In Section 3 we review dynamic pricing research, outline our mathematical model and discuss efficient problem solving methods. We perform computational analysis in Section 4 and focus on insights for practitioners. We conclude with a short summary and suggestions for future research directions.

2 Survey of Industry Practices

2.1 General Overview

Dynamic pricing techniques have received much attention in recent years from companies trying to improve profitability. These methods, which integrate pricing and inventory strategies to influence market demand, provide controls for companies to improve the bottom line. The focus of the review in this section is on implementation of these techniques in practice. (See Elmaghraby and Keskinocak (2003) or Chan et al (2004) for more complete reviews of dynamic pricing, including academic research.)

Dynamic pricing, which we define as changing prices over time without necessarily distinguishing between different types of customers, has been employed for ages but has traditionally been used only for sales or promotions. For example, fashion clothing retailers may offer discounts later in the season to reduce inventory, and this discount is the same to all customers at a given time. An exception to this is the pricing strategy that Amazon.com used briefly in 2001, differentiating among different types of customers and providing different prices to those customers for the same book or music item (Washington Post 2000).

However, many of the retail techniques do not account for the capacity limitations that exist in a manufacturing environment. In addition, many of the techniques limit the initial supply of product or allow only limited restocking during the time horizon. In this paper, we focus on dynamic pricing techniques based on characteristics such as inventory levels and variability in customer demand taking into account production capacity, rather than focusing exclusively on sales and promotions. Our model also incorporates pricing and production decisions for all periods.

2.2 Dynamic Pricing in Manufacturing and E-tailing

Dynamic pricing techniques have been implemented along with new business models developed for e-commerce. E-tailing, or retailing over the Internet, lends itself well to dynamic pricing since changing prices frequently is much more cost-effective electronically than physically. Smith et al. (1999) review empirical evidence of products online and through conventional outlets that suggests Internet markets are more

efficient than conventional markets with respect to price levels and menu costs, i.e., the cost of making changes to catalogs.

Pricing techniques employed in e-tailing are not limited to those based on inventory control or product differentiation like revenue management but include auctions, increasing discounts based on total volume accumulated by individual consumers, etc. In addition, companies can use the Internet distribution channel as a means to collect accurate demand data, which can then be used to determine effective dynamic pricing strategies.

Kay (1998) describes dynamic pricing at Boise Cascade Office Products, where many products are sold on-line. Boise Cascade states that prices for the 12,000 items ordered most frequently on-line might even change as often as daily. Evidence also suggests that companies such as Amazon.com or Dell Computer have implemented some type of dynamic pricing system based on inventory levels or competition. Agrawal and Kambil (2000) document price changes on products offered at each of the companies over a long time horizon and show that the prices are anything but fixed. Another article on Dell Computer Corp. (Wall Street Journal 2001) describes how the company uses various pricing strategies including dynamic pricing, differential pricing for different markets, and dynamic pricing of supplier components. The PC maker is able to use dynamic pricing because of its direct-to-customer model, but the method of determining prices is not described.

Many other examples of dynamic pricing in e-tailing are described by Baker et al. (2001). In one example, an electronics supplier changed prices more quickly than its competitors and realized an additional \$25 million in profit. In another case, the price of concerts and events was adjusted to match supply and demand (Tickets.com), resulting in as much as 45% more revenue per event in some cases. In yet another, differential pricing was used on customers who needed an electronic component immediately rather than with a more flexible lead-time. Baker et al. suggest ways to improve pricing strategies that incorporate the information and flexibility that is available through Internet channels.

More pertinent to the automotive industry are examples of dynamic pricing in manufacturing. For example, Campbell Foods installed a system to control prices based on such factors as inventory level (Kay 1998). However, overall, documentation of manufacturers using dynamic pricing is rare. One of the few manufacturing applications is described by Harris and Pinder (1995). In their illustration, a repair facility for industrial transformers could differentiate prices based on whether the transformers are on scheduled maintenance or are on “emergency order status” requiring immediate attention.

In the next section, we provide background on the automotive industry and examine the effect of the Internet on business models in the industry.

2.3 Introduction to Automotive Retail

The current franchise system of automotive retail has its roots at the beginning of the century. In fact, the first formal franchise agreement was signed in 1898 between a Pennsylvania bicycle dealer and the Winston Motor company (A.T. Kearny 1998). Franchise agreements assigned car dealers an exclusive territory in exchange for providing “adequate service and suitable facilities” (Sloan 1963). The manufacturer promised to sell vehicles at a discount to the dealer who would then sell the vehicle to the customer at a price predetermined by the manufacturer. In addition, dealers would advance large sums of money to the manufacturer before the car was even produced.

These franchise agreements allowed manufacturers to grow quickly since they received “venture capital” from their dealers, did not have to carry any finished goods inventory, and could invest all capital in their manufacturing systems. In addition, manufacturers could operate their factories without paying attention to demand fluctuation since dealers absorbed the demand variability in their inventory.

Dealers, on the other hand, also benefited from the exclusivity agreement for new products and aftersales services. Furthermore, dealers began to accept trade-ins when customers wanted to buy a new model and seized the additional opportunity to sell used vehicles. This point in time marks the end of manufacturer-determined fixed pricing since the two transactions (new vehicle sell and trade-in buy) could be booked separately by the dealer. Therefore, lowering the new vehicle price could be achieved through artificially




Figure 1: The North American Automotive Industry Distribution Model

increasing the value of the trade-in. Effectively, the dealer could now set the price for every customer individually. As the industry grew, manufacturers became more powerful and state legislatures passed laws to protect dealers from unfair competition by the manufacturers. In particular, these laws prevent manufacturers from setting fixed prices (the MSRP is a **suggested** retail price) and from selling new vehicles directly to customers.

Currently, the following business model is predominant in the automotive industry in North America (see Figure 1). During production planning, OEMs determine vehicle allocation to dealers based on a “turn-and-earn” system, i.e., production is sequentially allocated to the dealer who has the smallest number of days of supply of a particular vehicle. This system is used to ensure fair treatment of dealers, regardless of their size. When the vehicle is released for transportation at the plant, the dealer’s account is charged. To offset outbound transit time and time on the dealer’s lot, the OEM subsidizes the interest rate for the charged amount for a predetermined amount of time. The amount charged to the dealer is fixed and independent of the volume of vehicles bought. When the dealer receives the vehicle, he sets the price, and sells it to the customer.

Until recently, the above-described distribution model (manufacturer produces, dealer sets prices and sells) has prevailed. However, automotive distribution and retailing is now undergoing rapid change with the advent of third parties that are trying to sell cars through the Internet. Currently, there are 5 basic automotive Internet retail models, also displayed in Table 1: (1) third parties generate leads (e.g. www.autobytel.com), (2) manufacturers generate leads and search dealer inventory (e.g., www.gmbuypower.com), (3) customer names price and dealers can accept or reject offer (e.g., www.priceline.com), (4) auctions-buyers offer price bids for a given vehicle (e.g., www.ebay.com), (5) third parties sell vehicles on the web at posted prices (e.g., www.carsdirect.com). Of course, in a number of the cases, particularly models such as 1, 3, or 5, the entry of the third party may lower dealer margins by reducing the average price paid by consumers.

In response to the competition from third parties, manufacturers are searching for ways to sell vehicles over the Internet and move from the traditional “push” system towards a “pull” system. The challenge is to determine who sets prices for vehicles (dealers or manufacturers), how to move towards a customer pull system within the current dealer structure, and how to balance supply and demand. In the following, we focus on the last challenge and present a first attempt to address this issue. *We assume a make-to-order environment in which customers order their vehicles at a price set by the manufacturer.* The vehicles are delivered through a dealer who gets a commission for servicing the sale. Clearly, legal issues such as antitrust laws have to be addressed before such a system could be implemented in the automotive industry. These issues are not addressed in the current research. The models presented here will help the manufacturer and dealers to balance supply and demand and realize additional profit potential through dynamic pricing. However, we do not address a number of practical issues, such as availability and pricing of options (e.g., sunroof) but discuss these issues in the extension section.

3 The Pricing Problem

3.1 Pricing Research Review

As far as we know, Whitin (1955) is the first to suggest the need to consider joint pricing and inventory control strategies in a non-perishable environment such as retailing. In this paper, Whitin examined a single

Model	Customer Action	Response	E-tail Example
1	Requests quote	Third party generates leads from dealers	www.autobytel.com
2	Requests quote	Manufacturer generates leads from dealers	www.gmbuypower.com
3	Offers price bid for a desired vehicle	Dealers accept or reject offer	www.priceline.com
4	Offers price bid for a specific vehicle	Dealers accept highest bid if meets minimum	www.ebay.com
5	Requests vehicle at posted prices	Third party buys from dealers and resells	www.carsdirect.com

Table 1: Basic Automotive Internet Retail Models and Corresponding Examples

period problem, most similar to a “newsboy” problem, and determines a single price and supply quantity. Numerous other researchers have considered price determination and restocking in a multi-period setting. For example, Thowsen (1975), and Zabel (1972) both consider multi-period models with a convex ordering cost structure. The retail industry, particularly fashion items with seasonality, has also seen application of price differentiation policies and coordination of inventory control, in some cases under the name “yield management”. For instance, Gallego and van Ryzin (1994) analyzed the dynamic adjustment of price as a function of inventory and length of remaining sales; the demand was stochastic but restocking was not allowed. A thorough review of both single and multi-period models combining pricing and inventory strategies can be found in Eliashberg and Steinberg (1991).

In a manufacturing environment, a dynamic pricing model must determine prices as well as inventory levels. However, unlike retail dynamic pricing models, a manufacturing model must also schedule production and account for limitations in production capacity. In the case of multiple products, the strategy must also reflect shared production capacity among products. To date, few models have encompassed all of the necessary requirements.

The most notable exception, however, is the work by Federgruen and Heching (1999), who address the problem of determining optimal pricing and inventory control strategies under demand uncertainty; their model can also be extended to cover capacity limits on production. Indeed, their model is similar to our model, except for two important differences: Federgruen and Heching allow for stochastic demand but require fully backlogged stockouts, i.e., a customer purchasing an item which is out of inventory would receive the product as soon as it becomes available. In contrast, our model is deterministic but allows for lost sales; in addition, extensions to our model include multiple products sharing common production capacity and the addition of production set-up cost.

Some dynamic pricing problems may also be viewed as a special case of resource allocation problems. In these problems, we are given a fixed amount of resources, e.g., production and distribution capacity. Our objective is to allocate the resources to activities, e.g., production and distribution, so as to maximize a certain objective function, e.g., profit. For a comprehensive examination of resource allocation problems and algorithms, see Ibaraki and Katoh (1988).

An important algorithm for resource allocation problems is the greedy algorithm, which is known to be optimal under certain conditions. The greedy approach assigns one unit of resource at each iteration to the activity which contributes most favorably to the objective until the constraint set is tight or no activity is found. This algorithm is also known as a marginal allocation or incremental algorithm.

Chan, Simchi-Levi and Swann (CSS) (2002) showed that a certain class of resource allocation problem can be solved with the greedy algorithm. In particular, they defined a class of functions called *lightly concave*, and considered problems with a polymatroid feasible region, showing that the greedy algorithm provides the optimal solution for problems in this class. In the following sections we will review their result and show that

the pricing problem we consider falls in the class of problems that can be solved by the greedy algorithm.

3.1.1 Preliminary Notation and Previous Results

In order to present the pertinent results from CSS (2002), we present necessary notation below.

Consider a finite index set $\{1, 2, \dots, E\}$, also referred to as set E , and let V be a non-negative real function defined on 2^E , i.e. $V : 2^E \rightarrow \mathcal{R}$, where $2^E = \{S | S \subseteq E\}$.

Björner and Ziegler (1992) define a polymatroid as follows: A pair (E, V) , consisting of a finite ground set E and a function $V : 2^E \rightarrow \mathcal{R}$, is called a *polymatroid* if for all $S, T \subseteq E$:

1. $f(\emptyset) = 0$;
2. $S \subseteq T$ implies $V(S) \leq V(T)$;
3. $V(S \cap T) + V(S \cup T) \leq V(S) + V(T)$.

Let $F = (E, V)$ be a polymatroid and $f(x)$ be a cost function defined on $x \in \mathcal{N}^E$. We focus on the following general resource allocation problem, referred to as Problem $P(f, F)$:

$$P(f, F) : \max \left\{ f(x) | x \in F \right\}. \quad (1)$$

We say that $x \in F$ is a *global optimum* of $P(f, F)$ if and only if $f(x) \geq f(y)$ for all $y \in F$.

The main result of CSS (2002) is that the *greedy algorithm* solves $P(f, F)$ under certain assumptions on the cost function $f(\cdot)$ and the constraint set F . This algorithm can be described as follows:

Greedy Algorithm:

Step 0: $x = 0$;

Step 1: Find $i \in E$ with $x + e^i \in F$, $f(x + e^i) \geq f(x)$ and $f(x + e^i) \geq f(x + e^j)$ where $j \in E$ and $x + e^j \in F$

Step 2: If no such $i \in E$ exists, stop.

Step 3: $x = x + e^i$ and go to step 1.

CSS (2002) defined a class of functions called *lightly concave*. A function $f(\cdot)$ is lightly concave with respect to a polymatroid if it satisfies:

- (L1) if $y \geq x$, $f(x) \geq f(x + e^i)$, then $f(y) \geq f(y + e^i)$, $i \in E$
- (L2) if $y, x + e^i \in F$, $y > x$, $y_i = x_i$ with $f(x + e^i) \geq f(x + e^j)$ for all $x + e^j \in F$, then there exists $y + e^i - e^l \in F$ such that $f(y + e^i - e^l) \geq f(y)$ and $y_l > x_l$.

The main result from CSS (2002) is shown in the following theorem:

Theorem 3.1 (CSS (2002) Main Result) *If a function $f(\cdot)$ is lightly concave with respect to a polymatroid feasible region F , then the greedy algorithm generates an optimal solution for $P(f, F)$, defined in (1).*

We will show in following sections that the pricing problem we consider falls within the class defined by CSS (2002), and thus the greedy algorithm provides an efficient method for solving these problems.

3.2 Mathematical Formulation of the Pricing Problem

Consider a single facility that must determine prices and production scheduling for a single product over a finite horizon. For each period $t = 1, \dots, T$, let X_t , D_t , and I_t be the amount of product produced, the demand satisfied, and the end of period inventory, respectively.

The facility may produce a maximum of Q_t products in period t , and production cost incurred in period t is k_t per unit produced. Production costs are initially assumed to be linear. Inventory holding cost at a rate of h_t dollars per unit is charged for any inventory carried from period $t - 1$ to t . All cost and capacity parameters may vary from period to period.

We assume that demand is a non-increasing function of the price of the product, and these demand curves may vary from period to period. Thus, by determining the satisfied demand or sales in each period, we simultaneously determine the price of the product in each period. In the model, there is no time lag between a price change and the corresponding change in demand. We also assume that demand occurs in discrete units, and thus so does production. Since in our model we allow for limits on production, it may not be possible to satisfy all observed demand that occurs in a period; the amount of demand that occurred but was not satisfied is lost sales.

We allow for upper and lower limits on price or demand as well. Lower limits on demand ensure a minimum market share of a product, and price bounds may be used to stay within reasonable ranges compared to competitors. In addition, methods for estimating demand curves may provide parameters that are valid only within certain ranges.

The revenue occurring in each period, $R_t(D_t)$, is the selling price times the amount sold, or $P_t * D_t$. It is assumed that the revenue function is a concave function of the sales in each period. Linear demand curves are one example of demand-price functions that satisfy this property. The revenue function also allows for bounds on price or demand in any period t .

The objective of the pricing problem is to maximize total profit over the T periods, considering revenue, holding costs, and production costs in each period. Beginning inventory is zero, and there are production capacity limits in each period. The pricing problem, referred to as Problem PP, can be formulated as:

$$\begin{aligned}
 \text{(PP)} \quad & \max && \sum_{t=1}^T (R_t(D_t) - h_t I_t - k_t X_t) \\
 & \text{subject to} && I_1 = 0 \\
 & && I_{t+1} = I_t + X_t - D_t, \quad t = 1, 2, \dots, T \\
 & && X_t \leq Q_t, \quad t = 1, 2, \dots, T \\
 & && I_t, X_t, D_t \text{ integer} \geq 0, \quad t = 1, 2, \dots, T.
 \end{aligned}$$

Observe that in this model, the decision variables are the inventory level at the beginning of the period, I_t , production level X_t , and satisfied demand D_t . Since demand is a non-increasing function of price, demand satisfied in period t , D_t , will uniquely determine the product price, P_t . Our objective is to maximize total revenue minus holding and production costs. The first constraint indicates that there is no inventory at the beginning of period 1. The second set of constraints balances the inventory at each period. The third set of constraints ensures that production capacities are not exceeded.

3.3 Theoretical Results

Problem PP can be described as a min-cost network flow problem with (negative) convex cost and thus standard network flow algorithms for convex cost flows can be applied, see Ahuja, Magnanti and Orlin (1993). However, there are also extensions of the pricing problem which cannot be solved by network flow algorithms, for example, certain multi-product models sharing common components (see Swann 2001).

In this section, we outline a proof of the following theorem. The complete proof of all theorems and properties in this section are available in Biller et al. (2002).

Theorem 3.2 *If the revenue functions are concave in Problem PP, then the objective function $f(\cdot)$ is lightly concave with respect to a polymatroid feasible region F .*

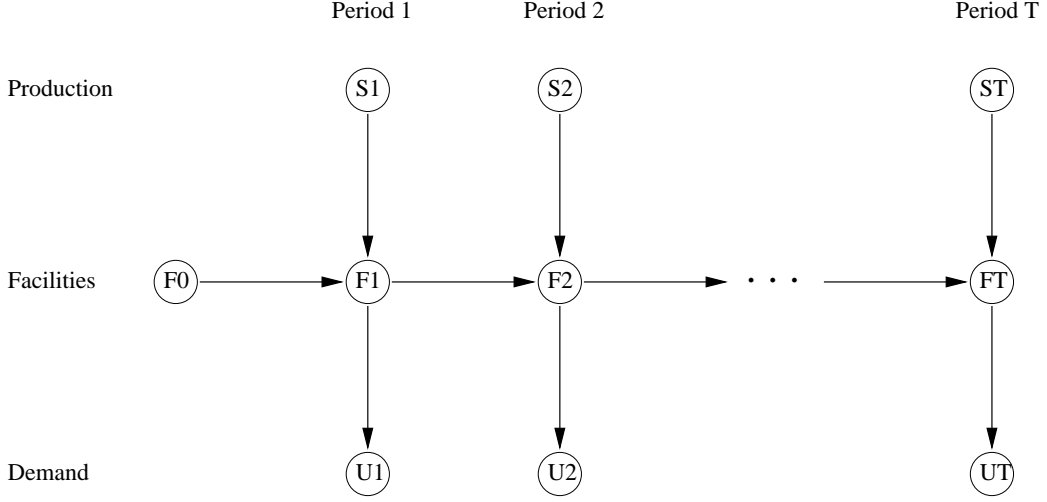


Figure 2: Network Formulation of Problem PP

We begin by showing that Problem PP has a polymatroid feasible region, then we show that the objective function is lightly concave as defined in CSS (2002). Of course, the implication is that the greedy algorithm solves pricing problem PP. Furthermore, following the same method, it is possible to show that the multi-product/multi-component problems mentioned above can also be solved by the greedy algorithm. The specific application of the greedy algorithm to Pricing Problem PP can be found in the Appendix of this article.

To show that the feasible region of Problem PP is a polymatroid, we first show it as a network-based model, see Figure 2 for a graphical depiction.

For each period $t = 1, 2, \dots, T$, nodes F_t, S_t and U_t represent the facility, production and demand satisfied, respectively. Node F_0 represents the beginning of the planning horizon, i.e., the state of the facility at the beginning of period 1. The directed arcs (F_{t-1}, F_t) , (S_t, F_t) and (F_t, U_t) indicate the inventory flow from the previous period, the production flow to the facility and the flow of demand satisfied at period t , respectively. Consider the directed network G with node set $N = \{F_0, F_t, S_t, U_t | t = 1, 2, \dots, T\}$ and arc set $A = \{(F_{t-1}, F_t), (S_t, F_t), (F_t, U_t) | t = 1, 2, \dots, T\}$, and define the following variables:

x_{ij} = flow on arc $(i, j) \in A$;

D_t = demand satisfied at period $t, t = 1, 2, \dots, T$.

To describe the feasible region of Problem PP as a network flow model, we impose the following constraints on the network G :

$$\begin{aligned}
 0 &\leq x_{S_t F_t} \leq Q_t, t = 1, 2, \dots, T, \\
 x_{F_t U_t} &= D_t, t = 1, 2, \dots, T, \\
 \sum_{l|(j,l) \in A} x_{jl} - \sum_{l|(l,j) \in A} x_{lj} &= 0, \quad j \in \{F_t | t = 1, 2, \dots, T\}, \quad \forall \\
 x_{F_0 F_1} &= 0, \\
 x_{ij} &\geq 0, \text{ and integer } \forall (i, j) \in A, \\
 D_t &\geq 0, t = 1, 2, \dots, T.
 \end{aligned} \tag{2}$$

The first set of constraints ensures that the production capacities are not exceeded, while the second set of constraints guarantees that D_t units of demand are satisfied at each period t . The third set of constraints is a flow conservation constraint, and zero inventory at the beginning of period 1 is guaranteed by the fourth constraint.

Let $D = (D_1, D_2, \dots, D_T)$ and $x = (x_{ij})_{(i,j) \in A}$. Megiddo (1974) showed that the set

$$\{D \in \mathcal{N}^T \mid (D, x) \text{ satisfies the constraints in (2) for a given vector } x\}$$

is a polymatroid.

To formulate Problem PP as Problem $P(f, F)$, we need to define a function $f(\cdot)$. Thus, for any given $D \in F \subseteq \mathcal{N}^T$, let $f(D)$ be the optimal objective value of

$$\begin{aligned} \text{(PP(D)) } \quad & \max_x \quad \sum_{t=1}^T (R_t(D_t) - h_t I_t - k_t X_t) \\ & \text{subject to } \quad I_1 = 0 \\ & \quad I_{t+1} = I_t + X_t - D_t, \quad t = 1, 2, \dots, T \\ & \quad X_t \leq Q_t, \quad t = 1, 2, \dots, T \\ & \quad I_t, X_t, D_t \text{ integer } \geq 0, \quad t = 1, 2, \dots, T. \end{aligned}$$

Otherwise, let $f(D) = -1$. Hence, Problem PP is equivalent to Problem $P(f, F)$ with $f(\cdot)$ being the function induced by the real value function f .

To show that the greedy algorithm generates an optimal solution for the pricing problem with concave revenue functions, it remains to prove that the objective function of Problem PP is lightly concave. To do this, we condition on the periods in which to increase demand and production, and further subdivide cases according to the inventory levels between the periods (see Biller et al. (2002) for details).

3.4 Model Extensions and Limitations

There are a number of additional applications that the model covers. Below we list some that have lightly concave cost functions over a polymatroid feasible region, and thus can be solved with the greedy algorithm.

1. *Multi-product systems* The model presented in this paper may be extended to cover supply chains with m , $m \geq 1$, products each of which is assembled from a set of parts and shares common production capacity. We assume that revenue curves exist for each product at each time period and that there are no demand diversions among different products.
2. *Variable Leadtimes:* Consider models in which a customer places an order in period j , $j = 1, 2, \dots, T$, and the product is delivered within L_j time periods. In this case, D_t represents demand satisfied at time t , $t = 1, 2, \dots, T$. In Problem PP, the revenue function of demand would need to be modified to account for all periods in which an order could have been placed to be satisfied by a delivery in time period t .
3. *Different Classes of Customers:* In practice, customers are sometimes distinguished by their responsiveness to different leadtimes. For instance, consider a model with two types of customers; customers who insist on receiving the product immediately and those who are willing to wait up to L time periods. This problem fits in the class of problems described in this paper if the firm has information on the demand curves by customer type and wait time, and the revenue functions are concave.
4. *Production Set-up Cost:* An important extension of the model is the addition of production set-up cost. In this case, it is possible to show that an optimal policy is *consecutive*. That is, in an optimal policy, production in a specific period, say t , will be used to satisfy demand in consecutive periods, say periods $r, r+1, r+2, \dots, s$. However, due to production capacity limits, it is not true that production in period t will necessarily satisfy all of the demand in periods r and s . Indeed, this observation leads to a dynamic program where the states are the initial inventory levels in the time periods, and the greedy algorithm is used within the dynamic program to determine the optimal demands satisfied between periods r and s given the initial inventory levels of those two periods. Of course, this algorithm is exponential and may not be efficient for large problem instances.

There are a number of limitations to the model that we consider. For example, in this paper we consider dynamic pricing in a monopolistic scenario. Clearly a more realistic scenario accounts for price competition. To date, few models have examined dynamic pricing under systemic constraints such as production capacity while also considering competition; this represents an ongoing area of challenge and interest. In addition, our research in this paper does not consider strategic buying practices, where customers time their purchases according to the price. Indeed, reality suggests that some customers will purchase strategically while others purchase according to immediate need. Incorporating this behavior in a multi-period setting is also a significant challenge but one of interest to researchers and industry alike. Finally, the model we present assumes that customer demand behaves according to deterministic demand curves. Adding a stochastic component to demand is an ongoing research interest of ours; more discussion of our stochastic pricing models is available in Swann (2001).

4 Analysis for Managerial Insights

In this section, we try to provide insights for practitioners regarding dynamic pricing policies in manufacturing. We quantify the effect of dynamic pricing on profit as well as identify contributors to the change in profit. Other impacts of dynamic pricing are considered, in particular, the effect of pricing policies on sales. We also examine the amount of variability in price observed under an optimal dynamic pricing policy and the frequency of price changes needed to attain the profit potential. We consider the effects that product characteristics have on the performance of dynamic pricing, and we consider the application of dynamic pricing to multiple products.

4.1 Computational Details

To generate insights about dynamic pricing, we compare the solution from the dynamic pricing problem PP, with the solution obtained from a fixed price problem where only a single price is allowed for all of the periods.

The fixed price problem, which requires a single price for all periods and integer demands, is referred to as Problem *DFP*. Problem *DFP* has a non-linear objective and integer variables, so we also find it useful to introduce Problem *LPFP*, which is the linear relaxation of Problem *DFP*.

The profit potential due to dynamic pricing is defined as $Z_{PP}^*/Z_{DFP}^* - 1$, where Z^* indicates the optimal objective value of the problem being solved. This indicates the additional profit that may be obtained through using a different price in every time period. However, since Problem *LDFP* is much easier to solve than Problem *DFP*, and since we know that $Z_{FLP}^* \geq Z_{FP}^*$, we know that the profit potential due to dynamic pricing is at least as large as $Z_{PP}^*/Z_{LDFP}^* - 1$. We use this ratio in all computation results unless otherwise indicated. Indeed, we solved a number of cases of Problem *DFP* using a non-linear integer solver available at NEOS and found the profit potential ratios to be almost identical.

In the analysis, we examine five demand scenarios, displayed in Figure 3: two types of seasonality, increasing mean, decreasing mean, and sawtooth. The first seasonality case (SEAS1) is based on seasonal variability in demand experienced in the automotive industry. Generally, there is low demand for cars in the winter, high in the late spring, etc. This scenario incorporates quarterly seasonality factors, with some variability of demand within each of the quarters. The second seasonality case (SEAS2) is similar, except that high demand occurs at the beginning of the horizon. This example is applicable to some retail clothing industries such as sporting goods; sales of ski equipment are generally high in the winter and low in the summer.

The increasing mean scenario (INCMEAN) occurs if sales undergo a learning, or word-of-mouth, effect, for example, a musical CD that builds in sales as satisfied customers influence friends to buy. Demand in high technology industries-like computer manufacturing-leads to the decreasing mean scenario (DECMEAN), where sales decline as newer products cannibalize sales of older products. The sawtooth scenario (SAW), which contains some randomness in the pattern, is not motivated by a particular example but was chosen as a contrast to the other demand scenarios.




Figure 3: Demand Scenarios: Variabilities represent different product types.

The demand curve used in the analysis represents the demand of a typical mid-size car in a monthly period. This demand curve, which is a linear function of price, is referred to as the “standard demand curve”. Variability in customer demand is effected through changes in the demand curve parameters and thus in the demand curve itself. For each of the demand scenarios, the demand curve varied in each period, but the average curve over the entire 12 period horizon was the standard curve. There are no bounds on demand for the scenarios analyzed, and therefore there were no lost sales in all cases discussed in the computational study.

Three levels of demand variability were examined for each of the scenarios namely, low, medium, and high; the medium variability case is displayed in Figure 3. The variability levels correspond to the method described in the Appendix. Production cost is that of a typical mid-size car. Capacity was set to be 75% of the demand level that maximizes profits with no capacity constraints, or the *optimal uncapacitated demand*. Holding cost in each period was 1.5% of the optimal price determined by the optimal uncapacitated demand.

4.2 Insights about Dynamic Pricing

In the following, we describe the impact of dynamic pricing on profit, including the sources of the additional profit potential, as well as the impact on variability in price and variability in sales. We also examine how much of the total additional profit potential can be achieved through a few price changes.

4.2.1 Impact of Dynamic Pricing on Profit

The profit potential due to dynamic pricing depends on the type of demand scenario as well as the amount of demand variability. Overall, our analysis indicates that additional profit potential due to dynamic pricing ranges from 1-7% over and above the optimal fixed pricing profit.

Clearly, if demand is more variable, we should expect a greater potential benefit from dynamic pricing (Figure 4). For the first seasonality scenario, low demand variability leads to 0.8% profit potential, compared to 2.7% for high demand variability. Intuitively this makes sense, as dynamic pricing has increased flexibility in each period, allowing for a better match between supply and demand. Put differently, with no variability in the supply chain the optimal policy does not vary over time, and hence the profit potential due to dynamic pricing must be insignificant.

Profit potential also depends on the type of demand variability. We find that profit potential is largest in SEAS2 and DECMEAN. Only in those scenarios is the demand significantly higher at the beginning of the

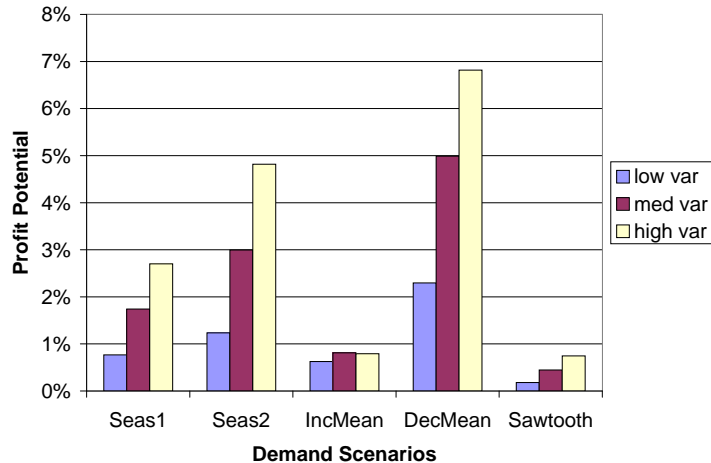


Figure 4: Profit Potential by Scenario Type and Variability Level

horizon than the available production capacity. Therefore, we conclude that dynamic pricing has its largest potential if demand is high relative to capacity at the beginning of the horizon. This is also intuitive since the manufacturer has the largest potential to increase profits through dynamic pricing if demand cannot be filled out of inventory or production. On the other hand, a fixed pricing strategy performs quite well in the SEAS1 or INCMEAN scenario because the manufacturer can build inventory in the slow season and sell it in peak season. In addition to larger profit potential under the DECMEAN scenario, dynamic pricing is easier to implement in this scenario, since customers are much more accepting of price drops over time than price increases. Overall, we conclude that the manufacturer should determine the type of demand variability before implementing a costly dynamic pricing system.

Earlier results also indicate that the profit potential depends on the amount of available capacity; specifically, the profit potential often increases as available capacity decreases. It is important to point out, that in general, the automotive industry suffers from over-capacity. However, there are particular cases where this is not true. For example, production capacity for large trucks in general, and specific models in particular is quite limited. As described before, Chrysler’s PT Cruiser provides an example of this, where production capacity was much less than demand for the vehicle.

4.2.2 Sources of Profit Potential

The question remains as to where the increase in profits (Figure 4) is obtained. For the case of medium variability, Figure 5 depicts the contribution to the profit potential due to production cost, inventory cost, change in revenue due to change in sales, and change in revenue due to change in price. The change in revenue from fixed pricing to dynamic pricing is calculated as follows:

$$\text{Revenue Change due to Sales} = (\text{Dynamic Sales} - \text{Fixed Sales}) * \text{Dynamic Price}$$

$$\text{Revenue Change due to Price} = \text{Fixed Sales} * (\text{Dynamic Price} - \text{Fixed Price})$$

where the Dynamic Price is calculated as the weighted average price over the horizon. Indeed, the entire change in revenue from fixed to dynamic pricing is simply the sum of the two values. Note in the figure that the total contribution to profit potential is normalized to 100%. Contributions above the x-axis indicate an increase in profit potential due to dynamic pricing; contributions below the x-axis show a decrease in profit potential. For example, if the contribution of inventory cost appears above 0, it implies that inventory costs for dynamic pricing were lower than for fixed pricing.

The graph yields several insights about dynamic pricing policies. First, profit potential is due to a number of factors, and the source of the potential is not always the same for every scenario. For example, in SEAS1

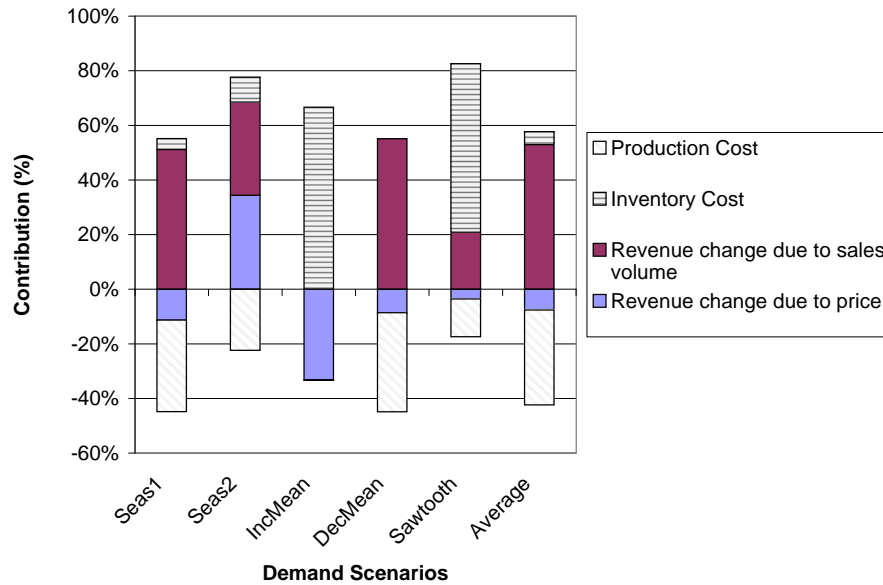


Figure 5: Sources of Profit Potential by Scenario Type

the potential is largely due to increased revenue from increased sales volume, but in the INCMEAN scenario, the potential is attributed to a decrease in inventory costs. This suggests that rules of thumb and heuristics applied independently of the demand structure are likely to work badly.

While the source of profit potential is not the same in every scenario, in most cases revenue due to sales volume shows a positive profit contribution. This implies that sales volume, and hence market share, is generally higher for dynamic pricing than for fixed pricing. Of course, an increase in sales volume is accompanied by an increase in production cost and a decrease in price (due to the structure of the demand curve).

4.2.3 Variability of Sales

As mentioned above, the automotive industry is moving to a build-to-order environment and dynamic pricing is a potential tool to match supply and demand. However, to evaluate dynamic pricing in a build-to-order environment, it is also important to consider the variability of sales because sales variability may lead to increased costs throughout the supply chain and increased variability in cash flow.

Figure 6 shows sales over time under a fixed pricing policy for each of the demand scenarios with medium variability. In contrast, Figure 7 depicts sales over time for the optimal dynamic pricing solution. Clearly, dynamic pricing reduces the sales variability significantly in all demand scenarios and can, therefore, be an effective tool when inventory buffers are not available to absorb demand variability. Similarly, dynamic pricing reduces production variability and therefore the Bullwhip Effect.

4.2.4 Variability in Price

Customers may not readily accept large changes in price for the same products. Thus, another important performance measure for dynamic pricing is the variability in price that results from the optimal policy.

Figure 8 depicts the optimal dynamic price over time, as a percentage of the optimal fixed price for the medium variability case. Prices vary by as much as 11% over time. Since it is very difficult to get customers to accept price increases, it would be most difficult to implement dynamic pricing in the INCMEAN scenario, and the easiest in the DECMEAN scenario. Fortunately, the benefits of dynamic pricing are the largest in




Figure 6: Variability of Sales under Fixed Pricing Policy

the DECMEAN and the smallest in the INCMEAN scenario. For the scenario relevant to the automotive industry (SEAS1), price varies by as much as 8% with the peak in the spring. This reflects a variability of up to \$1200 on a typical midsize car—an amount similar to typical incentives. We conclude therefore that dynamic pricing is an implementable alternative in the automotive industry.

The figure depicting the optimal dynamic price over time also suggests guidelines for determining prices according to the type of scenario. For instance, in the DECMEAN scenario demand over time decreases, and so does the optimal price over time. This combination of demand and price behavior are seen for products like fashion clothing or laptop computers. In the SEAS1 scenario, where demand begins low, increases, and decreases again over the horizon, the optimal price follows a similar pattern (low, high, low). Similar structures of optimal prices can be determined for the other scenarios.

4.2.5 Frequency of Price Changes

In addition to being sensitive to price changes, customers may also be sensitive to the frequency of price changes, in particular for expensive items such as cars. However, it is possible to gain a significant profit potential from dynamic pricing without changing the price frequently. To illustrate this point, observe that in the previous analyses, we allowed price changes in any of the 12 periods (11 changes). In this section, we compare this scenario to scenarios with 1, 2, 3, and 5 price changes for the medium demand variability case. Note that 0 price changes constitutes fixed pricing and 11 price changes refers to our previous scenario, which we will refer to as “total dynamic pricing”.

Our model requires that the price changes be evenly spaced. That is, if we allow 1 price change during the 12 periods of a year, this price change has to occur after six months. This clearly limits the flexibility of our model. Thus, our results in this section are lower bounds on the profit potential that can be achieved.

For each demand scenario, we determine the percentage of the profit potential due to dynamic pricing relative to the profit potential that can be achieved through total dynamic pricing. The results are shown in Figure 9. The graph shows that a significant additional profit can be obtained with a few price changes. Most of the time, a price change every 3 periods results in 80% or more of the profit potential.

Unlike the other scenarios, in Sawtooth significant profit potential was not obtained until 5 or even 11 price changes over the horizon. One likely reason for this is depicted in Figure 10, which shows the actual profit potential achieved by each of the price changes. In the Sawtooth scenario there was simply less potential profit to attain (0.4% for the case with medium variability). Additionally, the nature of that scenario (demand alternating up and down) leads to the intuition that price changes are needed in every




Figure 7: Variability of Sales under Dynamic Pricing Policy

period to match supply and demand most efficient.

The graph also indicates that, when evenly spaced, more frequent price changes do not always result in higher profit potential. For example, for SEAS1, 3 price changes achieves 90% of the profit potential, but 5 price changes only obtains 87% of the potential. This is because of the aforementioned characteristic of our model, where pricing is required to be evenly spaced over the horizon. Price changes that are evenly spaced are not necessarily able to match the variability in demand when it occurs. If the 5 price changes in this example occurred during any periods, then additional profit potential would have been obtained. However, even considering this limitation, we conclude that in our scenarios, most benefits of dynamic pricing can be achieved through a few price changes.

4.2.6 Multiple Products

Many firms manage portfolios with numerous products; thus the performance of dynamic pricing as the number of products increases is important to assess. In the analysis of multiple products, we assume that all products experience the same seasonality effects and that production capacity is shared among products. We also assume that there are no diversions among products. That is, a change in the price for one product does not affect the demand for another product. In this scenario, demand curves represent vehicles that appeal to various consumer market segments, such as luxury car, SUV, small pickup, etc.

The formulation of this problem is a simple extension of the single product pricing problem. Each of the decisions: demand or price, production, and inventory, must now be made for each product. We assume that the products share common production capacity in each period, so the production capacity constraint applies to the sum of all products in a period. The demand curves for each product are independent of each other. A full outline of a general multiple product problem is provided in Swann (2001).

As described in Section 3.4, the dynamic pricing multiple product problem without demand diversions may be solved using a greedy algorithm. Alternatively, if the demands are large, a linear model will be quite close to the optimal integer solution. For the test cases that follow, we solve the problems using AMPL with the MINOS solver. The multiple product problem with a fixed price over time for each product was solved in a similar way.

The results of dynamic pricing on a multiple product portfolio are shown in Figure 11. In general, the profit potential from dynamic pricing tends to decrease with the number of products. Observe that in the DECMEAN1 scenario, the profit potential is highest when there is one product and lower when there are 18 products. This suggests that balancing the mix of products in the portfolio may be as important or




Figure 8: Variability of Optimal Prices under Dynamic Pricing Policy

more important than tweaking the price of individual products. In the case with many products competing for capacity, a fixed pricing strategy can use the mix of products to achieve much of the same profit as the dynamic pricing case. The trends relating profit potential to number of products hold for cases with increased variability as well; to see this, compare SEAS1 with the same demand scenario except with higher demand variability (SEAS1, high var).

Of course, this trend does not hold true in all cases. Most notably, dynamic pricing on two or three products sometimes has a better performance than dynamic pricing on only one product. One possible explanation for this is that a little flexibility (here, the ability to choose the ratio of products in a portfolio), may have a significant effect. It is interesting to note that this effect may depend on the type of product characteristics. For this example, SEAS1 has highest profit potential from dynamic pricing with two products, whereas for DECMEAN, the highest profit potential is with one product.

The scenario called DECMEAN2 is another exception to the trend of decreased profit potential. In particular, dynamic pricing performs better on the 3-product case than in the 2-product case. This scenario is similar to DECMEAN1, except that the order of the products that are added (from 1 to 18) is reversed. The reversed portfolio shows that the relative profitability of the products that are added to the portfolio also affects the profit potential due to dynamic pricing.

Although the profit potential from dynamic pricing tends to decrease with the number of products, the graph also indicates that the profit potential for a particular scenario may stabilize. As described before, intuition suggests this result is due to the nature of seasonality of different products and the level of variability of the customer demand.

5 Summary and Future Directions

In this paper, we describe the current competitive threats to OEMs and the traditional dealer franchise system. We present a dynamic pricing model that can be applied in a manufacturing environment, incorporating capacity limitations and inventory holding costs. The model led to the following insights: (1) profit potential from dynamic pricing can be significant, (2) dynamic pricing is a useful lever to absorb demand variability, (3) the potential benefits of dynamic pricing depend on the type of demand variability, (4) in our scenarios, price changes may be as high as about 10% of fixed price, and (5) a significant profit potential may be attained with a few price changes. Although these conclusions are based on the linear demand curve representing a typical mid-size car, computational experiments with other linear demand curves yield similar

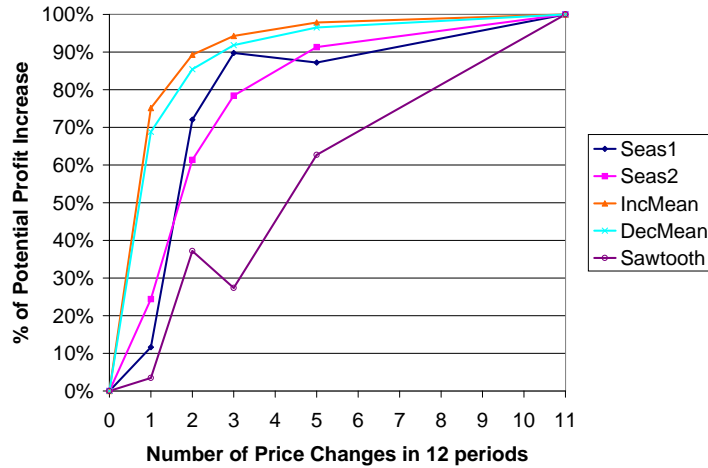


Figure 9: Percentage of Potential Profit Increase due to Number of Price Changes

insights (see Swann 2001 for instance).

The results obtained so far suggest a number of directions for future studies. For example, in the multiple product case, adding demand diversions among products may lead to further insights. In addition, we would like to extend our research to general multi-product/multi-component problems, where multiple products share production capacity and a limited supply of common parts.

The problem of strategic behavior is also one of considerable interest. Although our model does not account for it, customers are likely to exhibit strategic behavior in dynamic pricing environments, where they plan their buying purchases based on their expectation of price changes over time. The easiest way to incorporate this is to measure demand functions in situations where customers are strategizing their purchases; the result will be higher sensitivity to price changes (in the form of higher demand elasticities for instance). However, a more explicit way to model this is to allow demand in each period to depend on the prices in all of the periods. The impact of this behavior is to reduce profit relative to what is indicated by our results. This is true since in a model that captures strategic behavior, customers will plan their purchases so as to buy product at times of low prices.

We also would like to take the opportunity to discuss future directions from the automotive industry's standpoint—however, these extensions could be applied to other manufacturing industries as well. We see three major areas of interest to the automotive industry: (1) extending dynamic pricing to options (e.g., sunroof), (2) assessing benefits of revenue management in the automotive industry, and (3) exploring the consequences of dynamic pricing in a hybrid make-to-order make-to-stock environment. In the following, we describe these extensions.

We presented models that integrated production, inventory, and pricing decisions for a single product. While these models demonstrate the benefits of dynamic pricing on the vehicle level, it would be desirable to extend those models to dynamic pricing of options. We suspect that pricing vehicles and options simultaneously while using dynamic pricing as a lever to smooth demand for vehicles and options would increase profits significantly.

In addition, we would like to encourage work to explore the impact of service differentiation in the manufacturing industry. For more than two decades, revenue management, combining dynamic pricing with inventory control and service differentiation, has been a major contributor to the airlines' profits. We believe that the manufacturing industry could benefit tremendously from applying revenue management techniques to increase revenues through service differentiation. Specifically, applying revenue management to the car industry suggests that customers who are willing to wait longer to get a vehicle will pay a lower price than customers who want to get their vehicle immediately. Another possibility for revenue management in the




Figure 10: Profit Potential due to Number of Price Changes

automotive industry is to differentiate based on demographics such as car ownership history, where the amount and timeframe of the discount are the decision variables.

Finally, we are interested in the effects of dynamic pricing on a hybrid make-to-order make-to-stock environment. It is crucial for a manufacturer to know how much production should be allocated to replenishing dealer inventory and how much should be allocated to orders to maximize profits. In addition, we would be interested in models that consider two-agent (manufacturer and dealer) pricing models for vehicles and options. The interactions in a hybrid distribution system become much more complex and, at this point, it is even unclear whether a vehicle in stock should be more expensive (since the customer gets faster service) or less expensive (since the dealer tries to keep high inventory turns).

While it is obvious that the models presented here are too simplistic to be used in the actual pricing of automobiles, we hope that the analysis demonstrates to both practitioners and researchers of the value of such models to gain insight into actual pricing decisions. In addition, we hope that we have stimulated interest in the arena of dynamic pricing and look forward to the continued applications of OR techniques in an e-tailing environment.

6 References

References

- [1] Agrawal, V. and A. Kambil 2000, "Dynamic Pricing Strategies in Electronic Commerce," working paper, Stern Business School, New York University.
- [2] Ahuja, R. K., T. L. Magnanti and J. B. Orlin (1993). Network Flows: Theory, Algorithms and Applications. Prentice Hall, Englewood Cliffs, New Jersey.
- [3] Chan, L.M.A., Z. J. M. Shen, D. Simchi-Levi and J. Swann (2004). "Coordination of Pricing and Inventory Decisions: A Survey and Classification". Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era. D. Simchi-Levi, S. D. Wu and Z. J. M. Shen, Kluwer Academic Publishers: pg. 335-392.

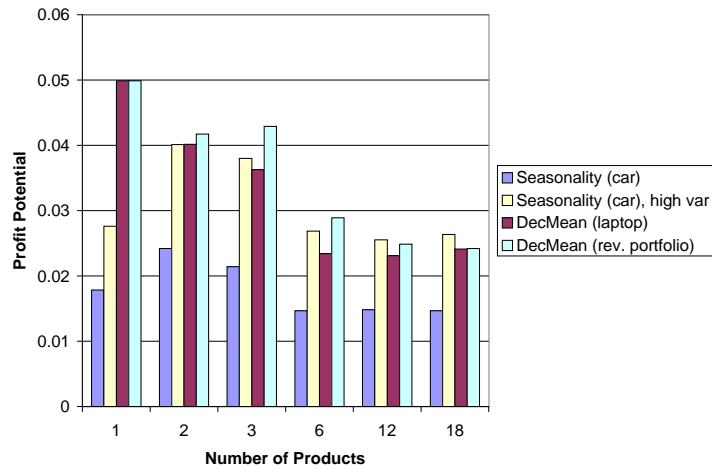


Figure 11: Profit Potential for Multiple Products Sharing Capacity

- [4] Elmaghraby, W. and P. Keskinocak (2003), “Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices and Future Directions,” *Management Science*, **49**, No.10, 1287-1309.
- [5] A.T. Kearny (1998), “The Future of Automotive Distribution”, *FT Automotive*, pg. 4.
- [6] Baker, W., M. Marn, and C. Zawada (2001). Price Smarter on the Net. *Harvard Business Review*, **79**, 2, pp. 122-126.
- [7] Biller, S., L. M. A. Chan, D. Simchi-Levi, and J. L. Swann (2002). Dynamic Pricing and the Direct-to-Customer Model in the Automotive Industry: Technical Results and Proofs. Technical Report, Georgia Institute of Technology, available at <http://www.isye.gatech.edu/~jswann/publications>.
- [8] Björner, A. and G. M. Ziegler (1992). Introduction to Greedoids, in *Encyclopedia of Mathematics and Its Applications: Matroid Applications*, N. White (ed.), Cambridge University Press.
- [9] Chan, L. M. A., D. Simchi-Levi, and J. L. Swann (2002). Applications of the Greedy Algorithm to Resource Allocation Problems and Three Types of Concavity. Working Paper.
- [10] Eliashberg, J. and R. Steinberg (1991). Marketing-Production Joint Decision Making, in *Management Science in Marketing, Handbooks in Operations Research and Management Science*, J. Eliashberg and J.D. Lilien (eds.), North Holland.
- [11] Federgruen, A., and A. Heching 1999, “Combined pricing and inventory control under uncertainty,” *Operations Research*, Vol. 47, No. 3, pp. 454-475. 21
- [12] Gallego, G. and G. van Ryzin (1994). Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons. *Management Science*, **40**, 8, pp. 999-1020.
- [13] Harris, F., and J. Pinder 1995, “A revenue management approach to demand management and order booking in assemble-to-order manufacturing,” *Journal of Operations Management*, Vol. 13, pp. 299-309.
- [14] Ibaraki, T. and N. Katoh (1988). *Resource Allocation Problems*, MIT Press, Cambridge, Massachusetts.

- [15] J.D Power Automotive and Associates 2000, “1999 New Autosshopper.com Study,” in The Power Report, January, Augora Hills, CA.
- [16] Johnson, S. M. (1957). Sequential Production Planning over Time at Minimum Cost. *Management Science*, **3**, 4, pp. 435-437.
- [17] Kay, E. 1998, “Flexed pricing,” *Datamation*, Vol. 44, No. 2 (February), pp. 58-62.
- [18] Megiddo, N. (1974). Optimal Flows in Networks with Multiple Sources and Sinks. *Mathematical Programming*, **7**, pp. 97-107.
- [19] Sloan, A. P. 1963, *My Years with General Motors*, Currency and Doubleday, New York, NY.
- [20] Smith, M. D., J. Bailey, and E. Brynjolfsson 1999, “Understanding Digital Markets,” forthcoming in *Understanding the Digital Economy*, eds. E. Brynjolfsson and B. Kahin, MIT Press, available at ecommerce.mit.edu/papers/ude.
- [21] Swann, J. L. 2001, “Dynamic Pricing Models to Improve Supply Chain Performance”, Doctoral Dissertation, Northwestern University, available at www.isye.gatech.edu/~jswann.
- [22] Thowsen, G. T. (1975). A Dynamic, Nonstationary Inventory Problem for a Price/Quantity Setting Firm. *Naval Research Logistics*, **22**, 3, pp. 461-476.
- [23] Wall Street Journal. Lean Machine: How Dell Fine-Tunes Its PC Pricing. June 8, 2001.
- [24] Washington Post. On the Web, Price Tags Blur; What you Pay Could Depend on Who You Are. Sep 27, 2000.
- [25] Whitin, T. M. (1955). Inventory Control and Price Theory. *Management Science*, **2**, 1, pp. 61-68.
- [26] Zabel, E. (1972). Multi-Period Monopoly Under Uncertainty. *J. Economic Theory*, **5**, pp. 524-536.

7 Appendix A: Application of the Greedy Algorithm to the Pricing Problem

In the following we consider the specific application of the greedy algorithm to the pricing problem, Problem PP.

- Let $D = (0, 0, \dots, 0)$, $X = (0, 0, \dots, 0)$ and $f(D) = 0$.
- Begin loop: For all i , $i = 1, 2, \dots, T$:
 - Given i , consider $D + e^i$ and solve Problem PP($D + e^i$) (e.g., using a network flow model) with $f(D + e^i)$ as its optimal solution value. Let $X + e^j$ be the new production plan, corresponding to demand $D + e^i$.
 - Calculate the marginal profit contribution of one unit increase of demand: $MP_i^D = f(D + e^i) - f(D)$.
- Choose the period k which maximizes the profit contribution, that is, choose k such that MP_k^D is positive and is the maximum over all periods. If no such period exists, stop the algorithm. $X + e^j$ is the production plan associated with $D + e^k$.
- Determine the maximum increase in demand possible in period k .
 - Begin loop: Let $D = D + e^k$ and $X = X + e^j$.

- Calculate MP_k^D , the marginal profit of the additional unit of demand.
- Stop loop when one of following occurs:
 - * Production is greater than maximum capacity, $X_j > Q_j$,
 - * The marginal profit in a period other than k is higher than the marginal profit at this level of demand: $MP_k^D < MP_i^D$ for some period i ,
 - * The marginal profit is negative, $MP_k^D < 0$.
- After loop ends, let $D = D - e^k$ and $X = X - e^j$, the last values before this loop.
- Return to the loop above, over all periods.

The greedy algorithm may also be modified to take into account lower and upper bounds on demand and price.

8 Appendix B: Details of Demand Curves

The demand curve used in the analysis is linear, and it is a function of the following three parameters: the base price p^{base} , the base demand or volume v^{base} , and the demand elasticity, E . The demand elasticity is the percentage change in quantity / percentage change in price. The curves are determined by the following equation, where V^{new} and P^{new} are the actual demand and price, respectively:

$$V^{new} = v^{base} + E * (v^{base} / p^{base}) * (P^{new} - p^{base}),$$

As described before, the data representing a typical mid-sized car establishes the standard demand curve. Let the base demand and demand elasticity that determine the standard demand curve be referred to as the standard parameters. To obtain the variable demand curves, the base demand and demand elasticity in each period vary as a percentage of the standard values. In certain scenarios, namely INCMEAN and DECMEAN, only the base demand varies over the horizon.

For the analysis, three levels of demand variability are considered: low, medium, and high. The corresponding percent variation levels of base demand as a function of the standard base demand are 90% - 110%, 75% - 125%, and 60% - 140%. The specific percentages used for each scenario are available in Swann (2001). To ensure that demand elasticity values were realistic, the variation was half as much as the base demand. Specifically, the demand elasticity varied as a percentage of the standard elasticity within the following ranges: 95% - 105%, 87.5% - 112.5%, and 80% - 120%. The periodic variation was determined according to the five demand scenarios described in the paper. In all scenarios, the average base demand elasticity over the horizon were equal to the standard parameters.