Dynamic Pricing Models to Improve Supply Chain Performance: Research and a Case Study Motivated by the Automotive Industry

January 2002
Submission to Council of Logistics Management
Doctoral Dissertation Award

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, manufacturing industries can increase profits and improve supply chain performance.

To illustrate these benefits, we discuss strategies that incorporate pricing, production scheduling, and inventory control under production capacity limits in a multi-period horizon, where unmet sales are lost. We consider a full planning strategy under deterministic demand as well as partial planning strategies under stochastic demand, where one decision (e.g., pricing or production) may be delayed until sales orders are received in previous periods. We also consider the case of simultaneously determining a production strategy and optimal fixed pricing policy under stochastic demand.

Using heuristics we develop and analyze, we perform computational analysis to generate insights for managers about the potential impacts of dynamic pricing. We quantify the profit potential and sales variability due to full planning dynamic pricing, and we suggest that it is possible to achieve significant benefit with few price changes. We compare pricing strategies under stochastic demand and show that delaying the pricing decision is usually more effective than delaying the production decision.

Contents

1 Introduction

1.1 General Introduction

In recent years, a number of industries have used innovative pricing strategies to manage their inventory effectively. For example, techniques such as revenue management have been applied in service industries as varied as the airlines, hotels, and rental car agencies–integrating price, inventory control, and quality of service. In other cases such as the power industry, varying price according to demand or time of day is becoming a tool to utilize capacity more efficiently. The retail industry has also concentrated on coordinating pricing and inventory control, particularly focusing on using price as a market clearing mechanism.

Coordinating pricing with other aspects of the supply chain offers significant opportunity to improve efficiency and profits. Unfortunately, the integration of pricing with production in a manufacturing setting is still in its early stages. There are a number of characteristics that distinguish general manufacturing industries from the industries mentioned previously, including the non-perishability of products and the ability to vary production levels. Furthermore, manufacturing differs from most retail environments in its reordering and capacity characteristics. Manufacturing generally has production in every period that is limited by the capacity of the system, while retailing often involves a single large order at the start of a selling season. One recent example of dynamic pricing in manufacturing is Dell Computers, which uses both dynamic pricing of products and pricing based on market segmentation (Wall Street Journal 2001), but the extent to which pricing is integrated with production and inventory control decisions at Dell is unclear.

Our focus on dynamic pricing as a tool to improve supply chain efficiency in manufacturing was initially motivated by a collaborative effort with a manufacturer of automobiles. Discussions with our Original Equipment Manufacturer (OEM) partner shaped the assumptions of the model and the specific consideration of *integrated* pricing and production policies. 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. In this context, we focus on developing models and analyzing the optimal pricing and production policies.

In addition to the focus of the theoretical research, the collaborative effort with the OEM greatly determined the direction of 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. We present significant material about the industrial perspective and motivation in this abstract.

For the general research problem, we focus on the coordination of pricing, production and distribution decisions for non-perishable products in a multi-period time horizon where production capacity is limited. In particular, we are interested in *full planning* strategies or *partial planning* strategies in a multi-period setting under either deterministic or stochastic demand, where products are non-perishable or have a long selling season. In full planning strategies, all decisions are made in advance before the beginning of the horizon. In contrast, in partial planning strategies, one decision (price or production) is made at the beginning of horizon, and the second decision (production or price) is made at the beginning of each period–after the realization of demand uncertainty in previous periods but before demand is realized in the current period. The strategies that we analyze are driven by general examples such as the following.

1. A manufacturer must make all decisions regarding pricing and production in advance of signing contracts with the retailers of the products. Since demand is non-stationary over time, non-stationary pricing is used to match supply and demand.

2. A supplier has non-stationary demand and uses prices to better match demand and supply in each period. The supplier contracts with a manufacturer over a time horizon, offering the manufacturer fixed prices in advance for planning purposes, but allowing orders to be placed in each period due to the manufacturer's high inventory holding cost and unpredictable demand. The supplier adjusts production in each time period based on previous inventory and expected orders.

3. A manufacturer needs to determine procurement decisions in advance in order to sign contracts with suppliers for part delivery. Thus the manufacturer determines a pro-

duction schedule at the beginning of a time horizon, but makes decisions regarding available inventory and price on a period by period basis.

4. A manufacturer whose primary distribution channel is through catalogs determines a single constant price over the lifetime of the product and wants the profit maximizing price. Production decisions are determined period by period, based on expected demand in present and future periods as well as inventory from previous periods.

These examples demonstrate the strategies that we are interested in, namely ones combining pricing and production decisions within a manufacturing supply chain. We focus our efforts on *dynamic pricing* models, where price changes over time in response to changes in supply or demand, but we also address *fixed pricing* models in some cases, where price is constant over time.

The first example illustrates a Full Planning model, where all decisions are made in advance. Whereas most firms do not make all decisions in advance, the model may be a useful starting point for analyzing the impacts of dynamic pricing.

The remainder of the examples are of Partial Planning strategies. In example two, the decision making firm determines prices for a planning horizon a priori and determines the production decision based on the state of the system and forecast demand. The firm is able to vary the production value based on inventory left over from previous periods or produce to satisfy demand in future periods. We refer to this strategy as *Delayed Production.*

Example three illustrates another partial planning model. The manufacturer plans production at the beginning of the horizon but makes the price decision on a period by period basis. We refer to this strategy as *Delayed Pricing.* In this case price can be used partially as a market clearing mechanism to deal with inventory from previous periods, basically the inventory that results from uncertain demand.

In some cases, a firm may believe that selecting a constant price for a product is the best strategy. However, procurement flexibility in each period may still be desired. This is illustrated in the last example, and the strategy is referred to as *Fixed Pricing.* As before, the price chosen should account for demand uncertainty over the horizon as well as demand seasonality (or non-stationarity) over time. This pricing strategy is often applied when goods are non-perishable or have a long selling season, e.g. furniture products.

In all of the examples described above, it may be profitable to apply *discretionary sales*, that is, to set aside inventory to satisfy future demand, even if the decision means losing sales in the current period. Although choosing to lose sales may seem counter to making profit, the inventory is set aside in situations when it is likely to generate a larger income in the future. This would typically occur if the price in the future is higher or if the future production costs were high.

The purpose of the research in this thesis is thus to develop and analyze multi-period models that integrate pricing, production and inventory control of non-perishable products in a manufacturing environment.

Our specific objective is to address the following research questions.

- What is the impact of dynamic pricing strategies on supply chain performance? For instance, what is the impact of dynamic pricing relative to a fixed pricing policy on profit or production variability?
- How important is it to incorporate stochastic demand models when considering pricing strategies? For example, is it important to take randomness into account even when the firm is interested in a fixed pricing policy?
- Under what situations does one partial planning model, e.g., Delayed Pricing, dominate another partial planning model, e.g., Delayed Production?

In all cases, we are particularly interested in the insights that can be provided to decision makers.

1.2 Automotive Motivation and Perspective

As noted earlier, the research in this thesis is motivated by a collaborative effort with a manufacturer of automobiles. In particular, the driving forces behind our industrial partner's interest in exploring dynamic pricing strategies are the Internet and the DTC model. The DTC model offers significant benefits to the manufacturer, such as better demand information and increased flexibility in matching supply and demand, and the model is a natural fit for innovative pricing strategies.

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 paper publication was contracting with AutobyTel for special Internet services.

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 absurd lengths. A third possibility for dynamic pricing is between OEMs and dealers, although this strategy is not currently in use.

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

1.2.1 Automotive Retail Business Models

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), (3) customer names price and dealers can accept or reject offer (e.g., www.priceline.com), (4) auctions-buyers

Figure 1: The North American Automotive Industry Distribution Model

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.

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

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 thesis, we focus on the last challenge and present a first attempt to address this issue, under the assumption of 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, and these issues are not addressed in the current research. The models presented here are a step towards helping the manufacturer and dealers to balance supply and demand and realize additional profit potential through dynamic pricing.

2 Literature Review

2.1 Examples of Dynamic Pricing in Industry

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. In this section, we provide a few industrial examples of these techniques.

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 (price markdowns). 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. However, many of the retail techniques do not allow for prices to go both up and down nor do they 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. Thus, we focus on situations applicable to a general manufacturing setting rather than focusing exclusively on sales and promotions.

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. As described earlier, evidence also suggests that companies such as Dell Computer have implemented some type of dynamic pricing system based on inventory levels or competition

Many other examples of dynamic pricing 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. The authors 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 weak. 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.

2.2 Research Literature Review

2.2.1 Pricing and Inventory Control

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 period problem, most similar to a "newsvender" 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 a few key 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 allows for lost and discretionary sales; in addition, extensions to our deterministic model include multiple products sharing common production capacity and the addition of production set-up cost. Finally, for the stochastic case, we assume a general demand function without assumptions on the structure, and analyze the partial planning strategies.

In addition to papers combining pricing and inventory, it is important to consider traditional inventory models. Zipkin (2000) provides an excellent review of literature on inventory control. A multi-period inventory model with non-stationary demand, varying production cost and lost sales is most closely related to the stochastic pricing problems considered in this thesis (see Chapter 9 in Zipkin). Of course, in this model, as well as the traditional inventory literature, demand is always satisfied when inventory exists, while in our model it is possible to decide not to satisfy demand even when inventory is available. This difference between our work and the traditional inventory literature leads to the concept of discretionary sales, which is an important strategy in our partial planning models.

Recent work by Scarf (2000) addresses an inventory model when sales are discretionary, as we also assume. Scarf considers a model with production capacity limits and set-up costs where price is given and inventory policy is the decision factor; this is similar to our Delayed Production strategy after a pricing decision has been made. The optimal policy for his model is of the (s,S) form, and the optimal discretionary sales is dependent on realized demand, whereas in our Delayed Production Strategy the optimal policy has an order-up-to and save-up-to target in each period, which are both independent of initial inventory and realized demand.

2.2.2 Specific Pricing Strategies

Many of the pricing and inventory papers in a multi-period setting are focused on finding a price in each period of the horizon, however some research has also been directed towards finding a single optimal price over an multi-period horizon. Kunreuther and Schrage (1973) provide upper and lower bounds on the optimal price, whereas Gilbert (1999) finds the optimal price under a less general model of demand. Gilbert (2000) considers a problem with multiple products sharing common production capacity and demonstrates a procedure for finding the optimal fixed price for each product. In the first two papers, production has a set-up cost, whereas in the multiple product work there is no set-up cost. Little work has been done on the fixed pricing problem with stochastic demand. This research is pertinent to the Fixed Pricing strategy we discuss.

Van Mieghem and Dada (1999) is another important paper relevant to our work, in which the authors explicitly consider price postponement versus production postponement strategies. They focus on a single-period, two-stage process with an initial decision, e.g. production decision, followed by a realization of demand, followed by another decision, e.g.

pricing decision. Thus, price postponement as outlined by Van Mieghem and Dada is different from Delayed Pricing in our model since in their case the postponed decisions are made after demand is realized; similarly, production postponement differs from Delayed Production in the sequence of events. They find that conditions dictate whether price postponement or production postponement is more valuable to a firm. Specifically they show that the former is likely to be more valuable if demand variability, marginal production, and holding costs are low. Their paper also addresses the decision of capacity investment and considers competition.

3 The Pricing Problem: Research Summary

To address the research questions described, we focus on two types of models. We first consider the Full Planning model under the assumption of deterministic demand, then we consider Partial Planning models under stochastic demand. In both cases, we use analytical and empirical analyses to generate insights.

3.1 Full Planning and a Resource Allocation Problem

In the Full Planning model, we determine prices and production jointly for a single product over a multi-period horizon. Production (X) is limited by the available capacity in each period, which is known and may vary over time. We assume that production cost is linear but the results can be extended to a model in which there is an additional production set-up cost. We allow for parameters such as inventory holding cost (h), production cost (k) and capacity (q) to vary over time (denoted by t).

In this model, we assume that demand is non-stationary over time and that the relationship between price (P) and demand is known; additionally, we assume that revenue ($R(D)$) in the Full Planning model is concave in each period with respect to sales or satisfied demand (D). If product is not available to meet customer demand, we assume the sale is lost. We focus on the pricing and production decisions and the impact of dynamic pricing on the supply chain under the assumption of a monopolistic decision-making firm. The goal is to maximize profit over the horizon, which comprises revenue, inventory holding cost and production cost.

A simple version of this problem is represented below, where I represents the inventory decision, and the price decision is determined according to the known demand function:

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

Some dynamic pricing problems may also be viewed as a special case of a resource allocation problem. 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.

The seminal paper by Federgruen and Groenevelt (1986) is directly related to the Full Planning problem. In this paper, the authors analyzed a general resource allocation problem where the constraint set forms a polymatroid and the objective function belongs to a class of functions referred to as *weakly concave* functions. They show that in this case the greedy algorithm generates an optimal solution. Unfortunately, while the constraint set of the Full Planning pricing problem forms a polymatroid, its objective function does not belong to the class of weakly concave functions.

Although the Full Planning Pricing Problem (PP) can be formulated as a resource allocation problem, we show that it may not necessarily have an objective that fits into the class of weakly concave cost functions defined by Federgruen and Groenevelt (1986). Thus we introduce a new class of functions, called *lightly concave*, and we show that the greedy algorithm provides the optimal solution for any problem with a lightly concave objective over a polymatroid feasible region. Although the result regarding the greedy algorithm and lightly concave objectives is a general one, we specifically show that the Full Planning Pricing Problem fits into the class we define. Since the constraint set of the pricing problem characterizes a polymatroid, these results imply that a greedy algorithm generates an optimal policy for the Full Planning model. In this case, the application of the greedy algorithm to problem PP is to increase sales according to the maximal contribution to profit, while solving a network flow to determine the best production schedule.

There are a number of extensions to the Full Planning problem that we consider, many of which also can be solved with the greedy algorithm. One important problem is where firms must price and plan production for multiple products that are sharing common production capacity. In addition, it is possible to extend Full Planning to account for variable lead-time in each period where the lead-time is a deterministic value. Firms may also want to consider a case with multiple classes of customers distinguished by their sensitivity to lead-time, which can be accommodated in our framework. Finally, we describe Full Planning with production set-up costs, which can be solved with a dynamic program that incorporates the greedy algorithm. For all of the extensions described, we assume that revenue curves are concave, and there are no demand diversions among products or customers.

Research on the Full Planning model including insights generated from computational analysis and the automotive perspective appears in Biller et al (2002), and the general resource allocation results are available in Chan et al (2002a).

3.2 Partial Planning

One of the limitations of the Full Planning scenario is that demand is deterministic, i.e., the relationship between price and demand is known exactly. Thus, the focus of Partial Planning Strategies is to add a stochastic element to demand, while incorporating decision-making in each period.

For the Partial Planning Strategies, we make many of the same assumptions as for the Full Planning Strategy, however, our demand function is more general. We assume that demand in each period is a non-stationary, general stochastic function, $d_t(P_t, \epsilon_t(P_t))$, that depends on the price in period t, P_t, and a random term with a known distribution that may also depend on price, $\epsilon_t(P_t)$. Furthermore, we assume demand in each period is independent and that expected demand in period t for a price P_t, $\overline{d}_t(P_t)$, is decreasing in P_t. We do not assume a particular distribution on demand, and in the most general case we allow demand to depend on the prices in all periods.

As described previously, in Partial Planning models we make an initial decision at the beginning of the horizon and delay another decision until the beginning of each period. Specifically, in Delayed Production, pricing decisions are made at the beginning of the horizon while production decisions are made period by period. For a given price vector P, the problem of determining production and sales is defined as a Markov Decision problem with the initial inventory (I) in a period as the state of the system.

As in classic inventory literature, we define Y_t to be the inventory level after production and before demand is realized, and let the facility decide on the amount of inventory Z_t to save for future sales. Let $J_t(I)$ represent the expected profit from the beginning of period t until the end of the time horizon. We use the phrase "profit-to-go" to refer to profit in the current and future periods. Let $G_t(Y)$ represent the expected profit-to-go with Y units of product available after production. Let the cumulative demand distribution for a given price P_t in period t be $\Psi_t^{P_t}$, and let the corresponding demand distribution be $d\Psi_t^{P_t}$ or $\psi_t^{P_t}$.

We formulate a dynamic program with the initial inventory as the state of the system. Given a price vector, the optimal expected profit in period t is as follows:

$$J_t(I_t) = \max_{Y_t: I_t \leq Y_t \leq I_t + q_t} -k_t(Y_t - I_t) + G_t(Y_t), \tag{1}$$

where

$$\begin{aligned} G_t(Y_t) \;\; = \max_{Z_t: 0 \leq Z_t \leq Y_t} \;\; & \{ \textstyle\int P_t(min(D_t, Y_t - Z_t))) d\Psi_t^{P_t}(D_t) \\ & -h_{t+1} Z_t - \textstyle\int h_{t+1}(max(0, Y_t - Z_t - D_t)) d\Psi_t^{P_t}(D_t) \\ & + \textstyle\int J_{t+1}(Z_t + max(0, Y_t - Z_t - D_t)) d\Psi_t^{P_t}(D_t) \}. \end{aligned} \tag{2}$$

The first term in (1) is production cost, and the first term in (2 is expected revenue. The remaining terms in (2 include inventory holding cost for units that are set aside, the expected inventory holding cost of the remaining units, and the expected profit-to-go for future periods, J_{t+1}. In the case that t is the last period, (2) salvage value replaces holding cost and it includes no profit-to-go.

We show that given an initial price vector, the expected profit-to-go functions are concave, and thus the optimal production and inventory policy has a special structure that allows us to efficiently find the optimal policy. Specifically, we show that the optimal policy is characterized by two parameters, both of which are time dependent but are independent of the inventory level at the beginning of each period. The first parameter is an order-up-to level or modified base stock policy (S, where $S = Y_t^*$), where the decision is to produce as much as possible to bring inventory up to that value ($max[q_t, Y_t^* - I_t^*]$). The second parameter is a save-up-to level or modified base discretionary sales policy (Z, where $Z = Z_t^*$), where as much as possible of that value is set aside for future sales (limited by capacity and inventory), thus resulting in a policy we call (S, Z). Furthermore, these decisions are independent of the realization of demand, thus the decision to accept or reject customers may be made as the customers arrive into the system.

Using a deterministic approximation of the problem with mean demand curves, we suggest an heuristic to generate an effective pricing vector for the general Delayed Production Strategy with dynamic prices. We show that the profit from the deterministic problem PP is an upper bound on the expected profit for Delayed Production, and we further characterize the worst-case performance with a data dependent bound. Finally, since the Fixed Price strategy with stochastic demand is a special case of Delayed Production, our results imply that a simple search on all possible prices finds the optimal Fixed Price policy.

In a similar manner, in Delayed Pricing, production decisions are made at the beginning of the horizon while pricing decisions are delayed to the beginning of each period. Again, we formulate the problem as a dynamic program with initial inventory as the state of the system. However, unlike Delayed Production, we show that the expected profit is not concave, and thus the optimal policies are more complex. Specifically, the optimal price and discretionary sales decisions depend on the initial inventory level, and we show that knowledge of customer demand can improve the discretionary sales decision. We further analyze the relationship between optimal price and inventory level at the beginning of each period, and we show that the optimal price in a period is not a decreasing function of inventory. Therefore, given a production schedule, we provide a dynamic program that generates the optimal pricing vector, and we suggest an heuristic for determining a production schedule based on a deterministic approximation of the problem.

Full details of the research on partial planning models is available in Swann (2001) and Chan et al (2002b).

3.3 Selections from a Case Study: Managerial Insights

Lastly in the thesis, we review industrial implementations of dynamic pricing and revenue management scenarios and describe a case study performed with our industrial partner. Using their data, we perform computational analysis on all of the pricing strategies we analyze, in order to generate insights for companies. The following describes some goals of the analysis: (i) identify situations where dynamic pricing provides significant increases in profit compared to fixed pricing. (ii) examine the impacts of dynamic pricing on supply chain performance (iii) determine if one partial planning strategy is preferred over the other.

A selection of graphs supporting the computational analysis and managerial insights from the final chapter of the thesis is presented below with a brief description.

3.3.1 Demand Seasonalities

Using initial data from our partner, we generate demand representing different types of products such as computers, cars, seasonal apparel, etc., to examine the effect of dynamic pricing under different product characteristics. The graph below exhibits the demand scenarios used in graphs following. The first seasonality case (SEAS1) is based on seasonal variability in demand experienced in the automotive industry (low demand in winter, high in spring, etc.). The second seasonality case (SEAS2) is similar, except that high demand occurs at the beginning of the horizon; for example, some retail clothing industries such as sporting goods.

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.

The demand variability over time for each of the scenarios is depicted in Figure 2.

3.3.2 Impact on Profit: Full and Partial Planning

For all of the strategies, we are interested in addressing the effect of dynamic pricing on profit compared to fixed pricing. Figure 3 shows the percentage increase in profit due to Full Planning dynamic pricing over fixed pricing (called *profit potential*). The figure illustrates that dynamic pricing may significantly increase profit, particularly when demand variability is high.

Figure 4 shows profit potential due to the Delayed Production heuristic and Delayed Pricing heuristic over the Fixed Pricing strategy, all under stochastic demand. For the Partial Planning computational analysis, we are particularly interested in the comparison between the Delayed Production and Delayed Pricing heuristics; the analysis shows that the

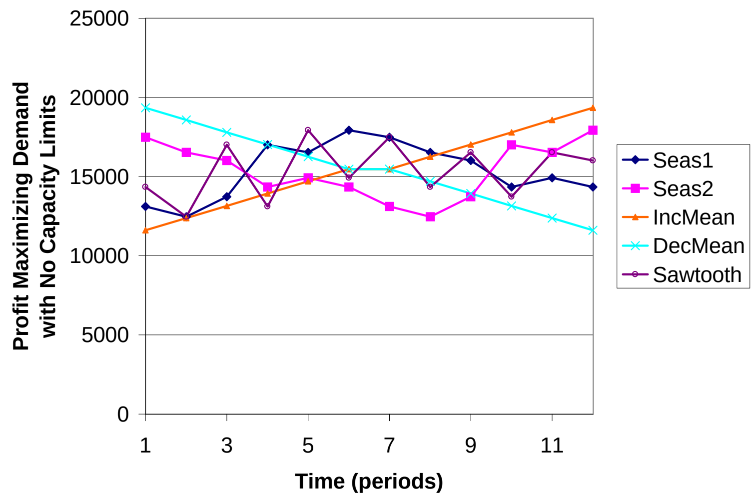

Figure 2: Demand Scenarios: Variabilities represent different product types.

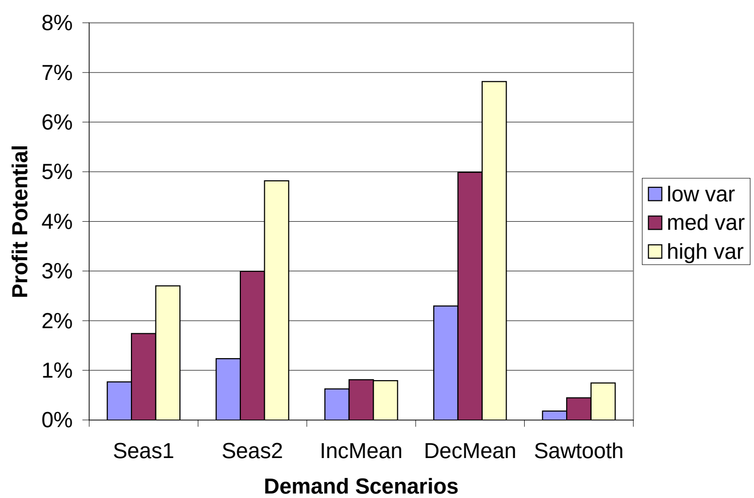

Figure 3: Profit Potential by Scenario Type and Variability Level

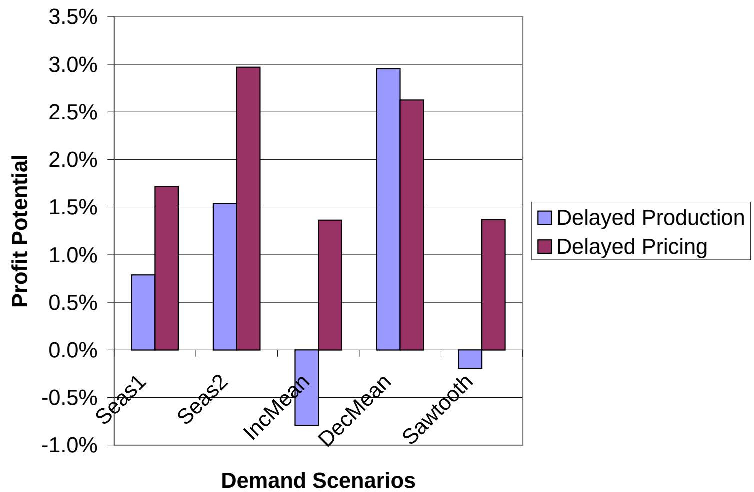

Figure 4: Impact of Type of Seasonal Variability under Partial Planning Models

Delayed Pricing heuristic is usually more profitable. Exceptions to this are under certain types of demand seasonalities or when production cost is high and the available capacity is high (not shown). The computational analysis for Partial Planning is particularly useful in estimating the magnitude of profit increase from dynamic pricing; our results indicate the profit increase may be as much as 2 - 7 % compared to fixed pricing.

An interesting result from both figures is that product characteristics affect the relative performances of all of the dynamic pricing strategies. In particular, dynamic pricing is generally more effective when demand is higher at the beginning of the season (DecMean or Seas2), similar to a markdown scenario. In these case, there also may not be an opportunity to build up inventory to meet demand, so dynamic pricing may be used to better match supply and demand.

3.3.3 Impact on Supply Chain: Variability of Sales

Dynamic pricing may also have an impact on the performance of the supply chain, as indicated by analysis of the Full Planning problem. Figure 5 displays the sales over time under fixed pricing, while the sales over time under dynamic pricing are shown in Figure 6.

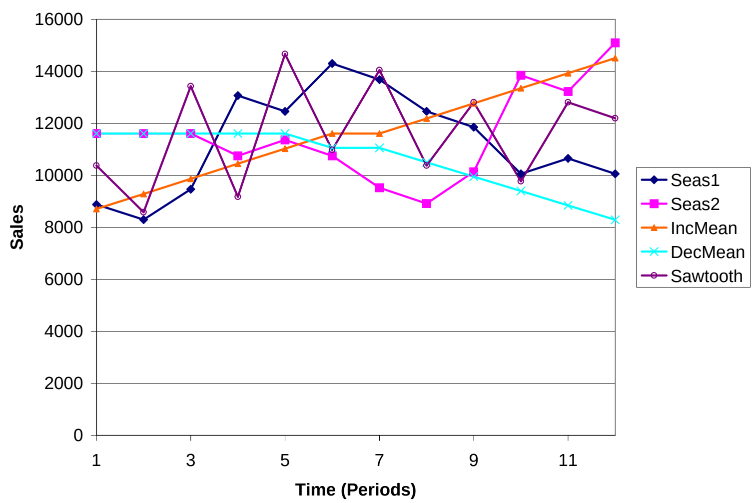

Figure 5: Variability of Sales under Fixed Pricing Policy

Other potential impacts on the supply chain are reduced variability in production schedule, increases in average sales, and decreases in average price. The latter two effects seem to be due to the more effective match between supply and demand under dynamic pricing. The decrease in average price is a particularly important result, since we do not consider the impact of competition in our model.

3.3.4 Prices Resulting from Dynamic Pricing

An important question for industry practitioners is the amount of variability in price that results from dynamic pricing; an example of this is shown in Figure 7. The graph depicts

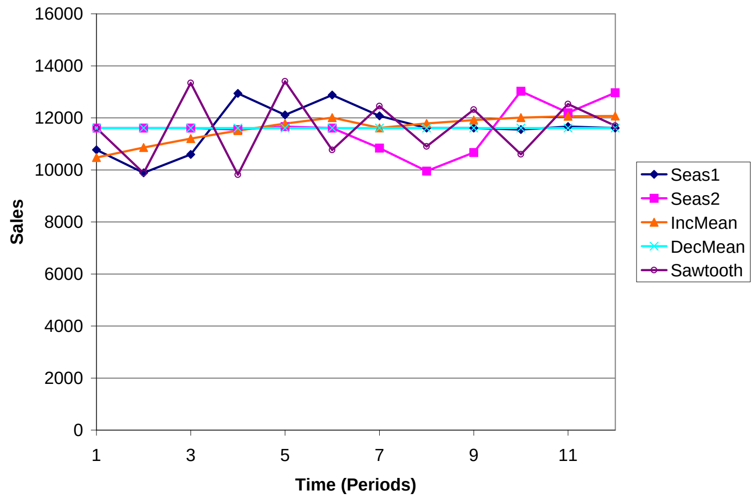

Figure 6: Variability of Sales under Dynamic Pricing Policy

the optimal dynamic price over time, as a percentage of the optimal fixed price; for this example of a typical mid-sized vehicle, prices vary by as much as 11% of the fixed price over the horizon, similar in value to some rebate promotions used in the automotive industry.

Analysis also indicated that effective dynamic pricing techniques may not need to implement a price change in every time period. Figure 8 shows the percentage of the total profit potential achieved by increasing frequency of price changes. In most cases, one price every third period resulted in significant profit potential.

Finally, other analysis is performed such as considering the impact of Full Planning dynamic pricing when there are multiple products in the portfolio. The full research results are available in Biller et al (2002) and Chan et al (2002b), and Swann (2002) provides a case study for classroom use.

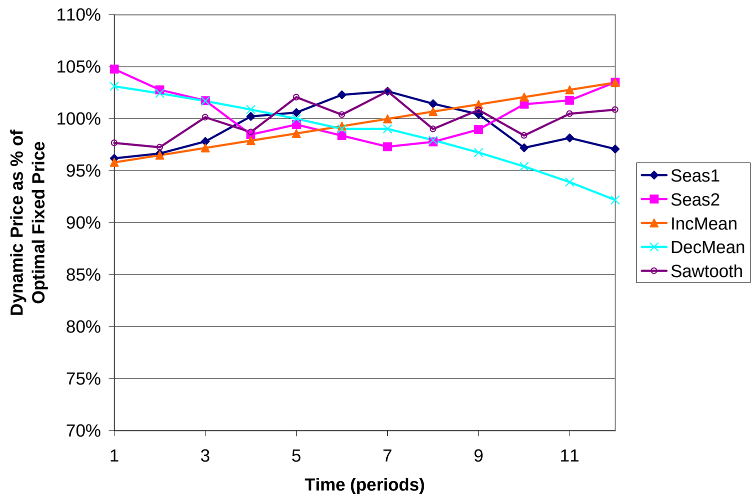

Figure 7: Variability of Optimal Prices under Dynamic Pricing Policy

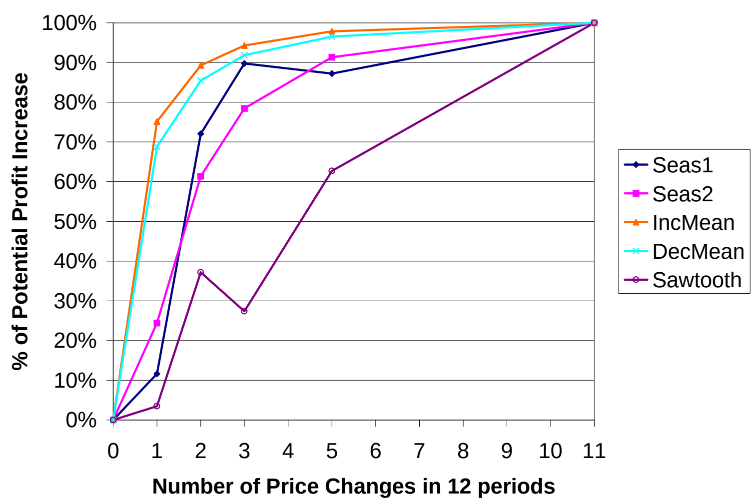

Figure 8: Percentage of Potential Profit Increase due to Number of Price Changes

4 Contribution and Conclusion

In this thesis, we consider coordinated pricing and production models motivated by our work with a manufacturer of automobiles. We present a Full Planning model coordinating price and production under time-varying deterministic demand, lost sales, and production capacity limits, and we demonstrate that the greedy algorithm provides the optimal pricing and production schedule for the problem we describe. Furthermore, we extend our results to a general class of functions with a lightly concave objective over a polymatroid feasible region for which the greedy algorithm gives the optimal solution. We describe a number of extensions to the model, including a multi-product problem.

We also describe a number of Partial Planning strategies, motivated by the variety of problems that firms must address, and under the assumption of stochastic demand. For all of the strategies, Delayed Production, Delayed Pricing, and Fixed Pricing, we provide methods to make pricing and production decisions and generate insights based on structural properties of the problems. We introduce the general concept of discretionary sales, where inventory is set aside to satisfy future demand even if sales are lost in the current period.

Finally, we demonstrate some potential benefits from implementing dynamic pricing strategies in a manufacturing environment. Indeed, our computational results suggest that companies that experience variability in demand curves or production capacity may significantly benefit from dynamic pricing, and that dynamic pricing may have significant impact on the supply chain. The case study we developed based on data from the automotive industry 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) significant profit potential may be attained with a few price changes (5) delaying the price decision is usually more effective than delaying the production decision.

There are a number of future directions suggested by the results in this research and by industry motivated situations. Of significant interest to us is the problem with multiple products and multiple parts, where there is limited supply of parts and limited common production capacity. This problem is of significant concern to industry as well, since its application arises in many situations. We are also interested in the problem where a firm has full flexibility on price and production decisions and makes both of these decisions in each period before customer demand is realized.

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 which could include service differentiation, and (3) exploring the consequences of dynamic pricing in a hybrid make-to-order make-to-stock environment.

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 integration of OR and economic techniques in a manufacturing environment.

5 References

Baker, W., M. Marn, and C. Zawada (2001). Price Smarter on the Net. *Harvard Business Review*, **79**, 2, pp. 122-126.

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. Forthcoming in the Dynamic Pricing issue of the Electronic Commerce Journal.

Chan, L. M. A., D. Simchi-Levi, and J. L. Swann (2002a). Applications of the Greedy Algorithm to Resource Allocation Problems and Three Types of Concavity. Working Paper, Georgia Tech, available at http://www.isye.gatech.edu/~jswann/publications.

Chan, L. M. A., D. Simchi-Levi, and J. L. Swann (2002b). Dynamic Pricing for Manufacturing with Stochastic Demand and Discretionary Sales. Working Paper, Georgia Tech, available at http://www.isye.gatech.edu/~jswann/publications.

Cook, T. (1998). SABRE Soars. *ORMS Today*, June, pp. 26-31.

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.

Federgruen, A., and H. Groenevelt (1986). The Greedy Procedure for Resource Allocation Problems: Necessary and Sufficient Conditions for Optimality. *Operations Research*, **34**, 6, pp. 909-918.

Federgruen, A. and A. Heching (1999). Combined Pricing and Inventory Control under Uncertainty. *Operations Research*, **47**, 3, pp. 454-475.

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.

Gilbert, S. M. (1999). Coordination of pricing and multi-period production for constant priced goods. *European Journal of Operational Research*, **114**, 2, pp. 330-337.

Gilbert, S. M. (2000). Coordination of pricing and multiple-period production across multiple constant priced goods. *Management Science*, **46**, 12, pp. 1602-1616.

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.

Ibaraki, T. and N. Katoh (1988). Resource Allocation Problems, MIT Press, Cambridge, Massachusetts.

J.D. Power Automotive and Associates (2000). 1999 New Autoshopper.com Study, in *The Power Report*, January, Augora Hills, CA.

Kay, E. 1998, "Flexed pricing," Datamation, Vol. 44, No. 2 (February), pp. 58-62.

Kunreuther and Schrage (1973). Joint Pricing and Inventory Decisions for Constant Priced Items. *Management Science*, **19**, 7, pp. 732-738.

Scarf, H.E. (2000). Optimal Inventory Policies When Sales are Discretionary. Cowles foundation Discussion Paper No. 1270. August 2000.

Swann, J. (2001). Dynamic Pricing Models to Improve Supply Chain Performance. Northwestern University, Dec. 2001.

Swann, J. (2002). The Right-Price Consultants: A Case Study on Dynamic Pricing, chapter to appear in Cases in Revenue Management, eds. I. Yeoman and U. McMahon-Beattie.

Thowsen, G. T. (1975). A Dynamic, Nonstationary Inventory Problem for a Price/Quantity Setting Firm. *Naval Research Logistics*, **22**, 3, pp. 461-476.

Van Mieghem and Dada (1999). Price versus Production Postponement: Capacity and Competition. *Management Science*, **45**, 12, pp. 1631-1649.

Wall Street Journal. Lean Machine: How Dell Fine-Tunes Its PC Pricing. June 8, 2001.

Washington Post. On the Web, Price Tags Blur; What you Pay Could Depend on Who You Are. Sep 27, 2000.

Whitin, T. M. (1955). Inventory Control and Price Theory. *Management Science*, **2**, 1, pp. 61-68.

Zabel, E. (1972). Multi-Period Monopoly Under Uncertainty. *J. Economic Theory*, **5**, pp. 524-536.

Zipkin, P. (2000). *Foundations of Inventory Management.* McGraw-Hill, Burr Ridge, IL.