

A Dynamic Pricing Model for Coordinated Sales and Operations

Moritz Fleischmann¹, Joseph M. Hall², and David F. Pyke²

¹RSM Erasmus University, PO Box 1738, 3000 DR Rotterdam, The Netherlands

²Tuck School of Business, Dartmouth College, Hanover, NH 03755

MFleischmann@rsm.nl, joseph.m.hall@Dartmouth.edu, david.f.pyke@Dartmouth.edu

January 3, 2006

ABSTRACT

Recent years have seen advances in research and management practice in the area of dynamic pricing and revenue management. At the same time, researchers and managers have made dramatic improvements in operations and supply chain management. The interactions between pricing and operations/supply chain performance, however, are not as well understood. In this paper, we examine this linkage by developing a deterministic dynamic programming model that captures a price/demand effect as well as a stockpiling/consumption effect – price and market stockpile influence demand, demand influences consumption, and consumption influences the market stockpile. The decision variable is the unit sales price in each period. Through the market stockpile, pricing decisions in a given period explicitly depend on decisions in prior periods. We develop analytical insights into the nature of optimal pricing policies and establish conditions under which the optimal prices converge to a constant (i.e., situations in which price promotion is suboptimal). We also analytically and numerically illustrate cases where the optimal prices vary over time. In particular, we show that price dynamics may be driven by both (a) revenue effects, due to nonlinear market responses to prices and/or inventory, and (b) cost effects, due to economies of scale in operations.

1. INTRODUCTION

Recent advances in information technology have generated remarkable opportunities for both marketing and supply chain management. In the marketing area, firms are increasingly able to understand individual customer preferences and therefore to dynamically adjust prices. This greatly enhances their ability to optimize revenue generation. For instance, Northern Group Retail Ltd. analyzes sales data by item, by store, to dynamically adjust prices at the store level (Green (2003)). Likewise, Sears is pursuing promotion optimization (the timing and depth of discounts) as well as markdown optimization (the timing and level of markdowns). Accenture estimates that major retailers as a group could improve gross margins by up to \$20 billion by using such technology ("Sears Bolsters Merchandise Performance" (2004)), while Gartner/G2 estimates that individual retailers can expect to see 1% to 9% increases in gross margins (Bernstein (2004)). Software firms, such as 4R Systems, Sabre, and PROS Revenue Management have developed pricing optimization solutions that have been successfully employed in many companies.

In the supply chain area, firms can use technology to improve their visibility to, and management of, costs and lead times. Examples of dramatic savings abound. GM, for instance, captures enormous amounts of data about vehicles coming off lease, and in 2002 used this data and an electronic auction, to sell over 300,000 cars, yielding increased returns totaling \$180 million (Welch (2003)). Firms such as Honeywell, IBM and Intel have saved millions of dollars by using vendor managed inventory and its cousin JIT II (Andel (1997)). Software from Manugistics, SAP, *i2* and others has helped many companies dramatically improve service and reduce cost.

Unfortunately, in spite of these advances, our experience suggests that companies and software firms have not made linkages between pricing and operations in a particularly rich way. Decisions are often made in silos, with little communication across functions; and decision support tools are not advanced enough to truly help. Furthermore, although there has been considerable progress in the academic literature, many questions remain unanswered. Yet we have observed an exploding interest in applying dynamic pricing strategies beyond the well-established cases of the airline and hospitality

industries. Companies as diverse as John Deere, Ford, Radio Shack, Peapod, and J.C. Penney are actively pursuing such initiatives. Some firms use home-grown software, while others employ packaged solutions. Yet in every case, an explicit, well-informed linkage between pricing and operations/supply chain is missing. Even when firms use the term enterprise profit optimization (EPO), the reality is that the application is fundamentally traditional revenue or yield management.

We believe that a major opportunity for competitive advantage is for firms to link innovations in marketing with those in supply chain management, allowing them to refine pricing, capacity, production and inventory decisions. Indeed, Forrester Research recently argued that “manufacturers need composite processes that link sell-side functions to supply chain activities” (Radjou *et al.* (2003)).¹ Such coordination could give marketing managers visibility to true costs and responsiveness as they make pricing and promotion decisions, and it could provide supply chain managers with a better understanding of pricing structures as they make capacity and inventory decisions. The result would be joint profit optimization rather than myopic optimization of revenues or costs.

Any treatment of dynamic pricing must recognize that there may be a strong dependency of demand across periods. A promotion today that generates a spike in demand will typically be followed by a demand trough (Neslin, 2002). Ideally, a pricing decision made today should optimally account for all future effects of the decision. For some product categories in both business-to-business and business-to-consumer contexts, this includes consumption effects (Ailawadi & Neslin (1998)), i.e., more product on the shelf or in the pantry leads to higher levels of consumption.

In this paper, we contribute to the literature on dynamic pricing by developing a deterministic, finite-horizon dynamic programming model that captures a price/demand effect as well as a stockpiling/consumption effect. The decision variable is the unit sales price in each period. We model full dynamic pricing in which the pricing decision in a given period affects demand, which in turn affects consumption. Hence, decisions in a given period explicitly depend on decisions in prior periods.

¹The sales and operations planning (S&OP) process is one valuable approach to facilitate this linkage (Pyke & Johnson (2004)). However, S&OP processes rarely, if ever, address joint optimization.

Traditional operations models typically assume exogenous demand. More recent inventory-pricing models assume endogenous yet independent demand. Herein, we model history-dependent endogenous demand, and we develop analytical insights into the nature of optimal prices and promotions. We also qualitatively explore insights about inventory and operations issues. For tractability, our model is somewhat stylized – we assume deterministic demand and consumption effects, for instance, and we assume that it is possible to observe inventory in the channel with some certainty. We believe, however, that even in this stylized environment, the analysis is insightful and yields valuable results.

The rest of this paper is organized as follows. In Section 2 we review the relevant literature. In Section 3, we present our dynamic model of demand response to pricing and we derive sufficient conditions for the optimality of convergent prices. A study of this model in different operations contexts is presented in Section 4. Section 5 contains discussion and extensions to other environments, and concludes with directions for future research.

2. LITERATURE REVIEW

The last two decades have seen significant research progress on the interaction of pricing and operations. Some of the research has focused on cost-driven dynamics, while the remainder considers revenue-driven dynamics. In the first category, the models divide into pricing/inventory models on the one hand, and pricing/queueing models on the other, while in the second category, researchers have addressed revenue management, retail markdowns and promotions.

2.1 Cost-Driven Dynamics

For an extensive review of the literature on pricing/inventory models, see Chan *et al.* (2004) who cite some 161 papers and books, and Elmaghraby & Keskinocak (2003). In general, these models assume infinite capacity. They set prices based on inventory levels and determine optimal inventory (cycle stock or safety stock) policies accordingly. One insight is that the financial benefit of dynamic pricing when capacity is unconstrained is often less than it is in a capacitated environment. Recent papers in this stream include Federgruen & Heching (1999), Chen & Simchi-Levi (2002a, 2002b), Zhu & Thonemann

(2003), Ray *et al.* (2005), and You (2005). Blattberg *et al.* (1981) find evidence from simple inventory models for consumers and retailers, as well as from empirical data, that promotions serve to transfer inventory holding cost to consumers when both parties act to minimize their own costs. In spite of many excellent contributions in this literature stream, our experience suggests that software firms and managers have not made much progress in capturing the rich linkages between inventory and pricing. The fundamental reason for this gap is that most models assume independence of demand across periods and do not capture promotions or fully dynamic pricing.

The literature on pricing/queueing models is somewhat disparate. Some models jointly determine prices and capacity, and perhaps lead times (So & Song (1998), Van Mieghem & Dada (1999), Boyaci & Ray (2003)), while others jointly determine prices and production quantities (Chan *et al.* (2002), Araman & Ozer (2005)). One insight from these papers is that dynamic pricing can yield significant benefits over static pricing, especially when capacity is tight or demand is non-stationary. Other models use prices to control the length of a queue at a computer resource or production facility (Mendelson & Whang (1990), Hall *et al.* (2005), Maglaras & Zeevi (2005)). The underlying problem is similar to revenue management in that production slots are perishable resources. Examples of management practice that explicitly ties together pricing with queueing, lead time or capacity decisions are few. Our conversations with managers suggest that many are pursuing more rigorous pricing decisions, and more sophisticated operations decisions, but they are doing so on parallel tracks.

2.2 Revenue-Driven Dynamics

Traditional revenue or yield management models are used extensively in the airline and hospitality industries, which face perishable products with uncertain demand, and possibly multiple customer segments (see Talluri & Van Ryzin (2004) and the extensive references therein). Managers and software solutions have incorporated the best of the research with extremely beneficial results (Fleischmann *et al.* (2004)).

The retail markdown problem addresses when, and by how much, to reduce prices at the end of a selling season. This problem is very similar to revenue management, but it generally includes a declining

demand pattern as well as an initial procurement decision. For a detailed discussion of this vast stream of literature see Talluri & Van Ryzin (2004). Retail markdown research has seen less extensive application in industry and software than revenue management, but every indication is that this is changing rapidly.

Research on promotions has primarily been pursued in the marketing area. Some papers address the reasons for running promotions, while others estimate their effects on demand. Firms run promotions for new product introductions, to respond to competition, to drive store traffic, and to capture price discrimination (Neslin, 2002).

Several authors refer to reference-price effects for explaining demand response to price promotions. Demand relies not only on current price, but also on a reference price that reflects past prices. Greenleaf (1995) uses dynamic programming to optimize retailer profits using market-level estimates of demand in response to promotions and shifting reference prices. His numerical experiments suggest that, with constant unit costs, retailers may “cycle” their prices – alternating high and low, for instance, with the high price serving to raise the market’s reference price. Kopalle *et al.* (1996) show that the profitability of cyclical pricing relies on an asymmetric market response, namely customers valuing gains more than losses.

See Neslin (2002) for an extensive review of the marketing literature on promotions. Additional recent references include Pauwels *et al.* (2002), Srinivasan *et al.* (2004), Macé & Neslin (2004), Popescu & Wu (2005), and Natter *et al.* (2005). The exact effect of price on demand remains uncertain, although there is evidence that net sales may increase as a result of a promotion – even with the pre- and post-promotion dip in sales – but the profit effect for retailers is often negative. What is clear from the promotions literature is that the available models are primarily descriptive, with few optimization approaches. Furthermore, the operations and supply chain dynamics are essentially ignored.

There are several exceptions to this latter point, including Huchzermeier *et al.* (2002), Sogomonian & Tang (1993), and Iyer & Ye (2000). Sogomonian & Tang (1993) develop a model that shows that integrating production decisions with promotion decisions can increase profit and decrease inventory. Iyer & Ye (2000) show that if the uncertainty of the demand effect of promotions is high, it

may be more profitable for the retailer not to promote at all. They also show that as customers' inventory holding cost decreases, their stockpiling increases, which in turn suggests that retailers will promote less frequently. Software and management practice are beginning to get more sophisticated, but the research suggests that most managers use simple heuristics when making promotion decisions.

One further dimension of the promotions literature is the effect of product consumption on optimal dynamic pricing. For example, Ailawadi & Neslin (1998) and the references therein describe how increased household inventory can increase consumption because of fewer stockouts and higher usage rates (see also Assuncao and Meyer, 1993). The authors empirically test consumption as a continuous function of inventory and show that it fits data for yogurt and ketchup better than the commonly used constant usage rate model. See also Gourville & Soman (2002) who point out that consumption is dependent on a number of factors and varies over time.

Our observation is that the cost-driven dynamics literature is primarily carried out by researchers in operations management and relies on simplistic demand models. On the other hand, the revenue-driven dynamics research has a simplistic, if any, picture of costs and the supply chain. Our goal is to increase coordination between cost and revenues by analyzing a richer demand model in different cost environments. We employ a concave, nondecreasing consumption function, similar to that in Ailawadi & Neslin (1998), and we model demand as a function of price and market inventory, using a deterministic function. We examine both multi-period and infinite horizon cases, and we assume a single monopolistic decision maker. The decision variable is the price charged in each period, and we allow for fully dynamic pricing.

3. THE DEMAND MODEL

3.1 Model Introduction

The kernel of our model is a method of capturing how consumers behave in aggregate when they face (possibly) varying prices for an item over time. We allow past decisions to impact current demand via tracking of the aggregate stockpile of an item amongst consumers:

M_t = external stockpile in the market at the beginning of period t .

This stockpile forms the state variable for our problem. The decision for the firm is the price to offer in period t :

p_t = unit sales price.

We allow market demand to respond, in a possibly non-stationary way, to current price and current market stockpile: $D_t(p_t, M_t)$. We assume that demand is non-increasing in price and market stockpile.

In effect, the market stockpile serves as a partial memory of all past pricing decisions made by the firm.

We assume that the market stockpile is consumed according to a possibly non-stationary consumption function $C_t(M_t + D_t)$, which allows for current period consumption of current period demand. This

consumption function is assumed non-decreasing in market stockpile. For convenience, we also define

$\tilde{C}_t(M) \equiv M - C_t(M)$ as the remaining stockpile at the end of the period (post-consumption). The state

transitions for this problem are described in a natural way by:

$$M_{t+1} = M_t + D_t - C_t = \tilde{C}_t(M_t + D_t).$$

The events in each period occur in the following order: (1) the firm sets price based on knowledge of current market stockpile; (2) consumer demand occurs, based on price and market stockpile; (3) consumer consumption occurs, based on new (post-demand) market stockpile; (4) the firm accurately estimates or otherwise measures post-consumption stockpile. Both demand and consumption are assumed to be deterministic in our model. These assumptions are made for tractability and parsimony; relaxing them presents possibly rich opportunities for further work.

The cycle of purchasing, stockpiling and consumption reflected in these modeling assumptions follows normative and empirical models developed in the marketing literature (see, e.g., Ailawadi & Neslin (1998) and the references therein). As with Greenleaf (1995) we model consumer behavior in aggregate, not at the individual level. However, we note that under the strong assumption of homogeneous individual consumers with equal beginning individual stockpiles, our model would capture

behavior at both the individual and aggregate levels. Application of our model to business-to-business settings is also feasible. In such settings, sell-into and sell-through data, combined with a smaller number of customers, would likely make our informational assumptions less restrictive.

The internal operations of the firm are described in a rather compact way by a total cost function that may capture either production or procurement costs or both: $K_t(\cdot)$. This is a rather elementary view of the operations of the firm, but it does allow study of various types of operations, albeit in a simplified way. In particular, we ignore possible inventory at the firm in order to allow for a single dimensional state-space. More details on the cost function are provided below in Section 4.

The immediate net return (profit) of the firm is given by:

$$\Pi_t = \Pi_t(p_t, M_t) = p_t D_t(p_t, M_t) - K_t(D_t(p_t, M_t)).$$

We seek an optimal pricing policy under both a finite and infinite horizon, assuming a discount factor of α . Following standard dynamic programming notation, let

$$V_N(M) = \text{terminal value,}$$

$$W_t(p, M) = \Pi_t(p, M) + \alpha V_{t+1}(M + D_t - C_t),$$

$$V_t(M) = \sup_p W_t(p, M), \text{ and}$$

$$V^\alpha(M) = \lim_{t \rightarrow \infty} V_t(M).$$

Our goal is to study solutions to this problem under various forms of the functions D_t , C_t , and K_t in order to gain insights into when different types of pricing policies (e.g., stationary or cyclical) are optimal. Further, we wish to gain some insights into how these policies interact with the production or procurement activities of the firm.

Throughout, we assume all functions are twice differentiable. State space and decision variables are assumed to be continuous. We denote partial derivatives with subscripts wherever the meaning is

unambiguous: e.g., $D_{pM} = \frac{\partial^2 D}{\partial p \partial M}$.

3.2 Convergence of the Optimal Pricing Policy

In this section we explore conditions that lead the optimal dynamic pricing policy to converge to a constant price. Our analysis of the optimal pricing policy relies on basic structural properties of dynamic programs (see e.g. Section 4.7 of Puterman (1994)). In particular, we rely on the result that a dynamic program possesses an optimal policy that is monotone in the state space if the value function is submodular. We specify conditions that preserve submodularity under the dynamic programming recursion and we use these conditions as the basis for sufficient conditions for the optimality of a constant pricing policy in our model. In subsequent sections, we explore conditions which lead to non-constant optimal pricing policies.

We begin our analysis by specifying conditions under which the value function is non-increasing in the market stockpile. In other words, under these conditions a decrease in the market stockpile always benefits the firm.

Lemma 1. *Assume that the following conditions hold for $t = 1, \dots, N - 1$:*

1. $\Pi_t(p, M)$ is non-increasing in M for all p ,
2. $0 \leq D_t(p, M) - D_t(p, M + \Delta) \leq \Delta$, for all p, M and $\Delta > 0$,
3. $0 \leq C_t(M + \Delta) - C_t(M) \leq \Delta$, for all p, M and $\Delta > 0$, and
4. $V_N(M)$ is non-increasing in M .

Then $V_t(M)$ is non-increasing in M for $t = 1, \dots, N$.

Proof. Assumption 3 implies that $\tilde{C}_t(M)$ is non-decreasing in M . Therefore, Assumptions 2 and 3 imply that M_{t+1} is non-decreasing in M_t for fixed p . The result then follows from induction on t . ■

Condition 2 of Lemma 1 requires that an additional unit of market stockpile reduces demand by at most one unit; this rules out “over-reaction” to market stockpile additions. Similarly, Condition 3 requires that an additional unit of market stockpile does not increase consumption by more than one unit.

Corollary 1 below allows comparison between the solutions of the myopic and full problems: if there is a unique optimal price for the myopic problem, any optimal price for the full problem will equal or exceed it.

Corollary 1. *Let $\bar{p}_t(M)$ denote the smallest maximizer of $\Pi_t(p, M)$. Under the conditions of Lemma 1, any optimal pricing policy $p_t^*(\cdot)$ satisfies $p_t^*(M) \geq \bar{p}_t(M)$ for all M .*

Proof. The result follows immediately from Lemma 1 since a price below $\bar{p}_t(M)$ yields smaller immediate profits and smaller future profits than $\bar{p}_t(M)$. ■

Proposition 1 below serves as the key ingredient for the structural results of this section by establishing that there exists an optimal decision rule that is monotone in the state variable. Later, we rely on this result to characterize monotonicity conditions for the sequence of market stockpile levels under an optimal pricing strategy, and thus claim that the optimal price will converge to a constant.

Proposition 1. *Assume that*

$$\frac{D}{(D_p)^2} \left(D_{pp} \frac{1+D_M}{D_p} - D_{pM} \right) - \frac{D_M + 2}{D_p} + K'' \geq 0$$

for all p and M with $D(p, M) > 0$. Then there exists an optimal decision rule $p^(M)$ for which the post-consumption market stockpile is non-decreasing in M .*

Proof. See attached companion document. ■

The assumption of Proposition 1 can be easily checked by direct analysis of the functional forms of the demand and cost functions. The proposition provides a monotonicity result regarding the market stockpile. As an aside we note that the optimal pricing policy itself is monotone in the special case that aggregate market demand always replenishes the market stockpile to a certain price-dependent target level:

Corollary 2. *Assume that the conditions of Proposition 1 hold and that $D_M = -1$ for all p and M with $D(p, M) > 0$. Then there exists an optimal decision rule $p^*(M)$ which is non-decreasing in M .*

Proof. According to Proposition 1, the optimal post-consumption stockpile

$$X^*(M) = \tilde{C}(M + D(p^*(M), M))$$

is non-decreasing in M . Differentiation yields $X_M^* = \tilde{C}'(1 + D_p p_M^* + D_M) \geq 0$ and therefore

$$p_M^* \leq -(1 + D_M) / D_p . \blacksquare$$

In general, the optimal prices may not be monotone in M . However, Proposition 1 implies that they converge:

Corollary 3. *Assume that the conditions of Lemma 1 and Proposition 1 hold for all t and that*

$$\lim_{M \rightarrow \infty} D(\bar{p}(M), M) = 0 .$$

Then there is an optimal infinite horizon pricing policy for which the sequence of market stockpile values M_t converges monotonically to a constant. In particular, the prices effected under this policy converge to a constant price.

Proof. According to Proposition 1, the sequence (M_t) is monotone. The sequence is naturally bounded from below by 0. The condition on \bar{p} implies that for M large enough consumption exceeds demand. Therefore, the sequence (M_t) is also bounded from above and we can conclude that it converges. The corresponding prices converge due to the continuity and monotonicity of D and C . ■

We summarize a slightly weaker form of the above results in the following theorem.

Theorem 1. *Assume that the following conditions hold for all t , p , and M*

1. D_t is concave in p
2. D_t is submodular
3. $D_M \geq -1$
4. $C_M \leq 1$
5. K_t is convex
6. $\lim_{M \rightarrow \infty} D_t(0, M) = 0$.

Then there exists an optimal infinite horizon decision rule $p_i^(M)$ for which the effected prices converge to a constant.*

Proof. The result follows immediately from Proposition 1 and Corollary 3. ■

In spite of allowing arbitrary dynamic pricing policies, Theorem 1 provides sufficient conditions that a constant pricing policy is optimal in the long run. In other words, price promotions under these conditions would be suboptimal. This raises the question of whether the conditions are somehow unusual or counterintuitive. As it happens, the conditions of Theorem 1 appeal to our intuition: (1) as we lower price, demand does not grow at an increasing rate; (2) as the market stockpile gets smaller, customers become less price sensitive; (3&4) a unit increase in the market stockpile reduces demand by at most one unit and increases consumption by at most one unit; (5) the cost of goods exhibits no economies of scale; and (6) demand vanishes as the stockpile grows large. In the next section we expand on this exploration by investigating scenarios in which these conditions hold and other scenarios where they do not.

4. PRICING POLICIES IN DIFFERENT OPERATIONS ENVIRONMENTS

In this section, we analyze the potential benefits of dynamic pricing in a number of ‘prototype’ environments. We first explore a “base case” where all key effects are linear and then expand our scope in two different directions: non-linear market effects and non-linear operations effects. Our goal is to demonstrate that in some contexts, a richer, non-stationary endogenous demand model can reveal behavior that cannot be captured under the assumption of stationary demand. While our models of operations effects are rather simple, they do mark a first attempt to understand these effects. We expect future work to elaborate on these models and allow for a greater level of operational fidelity.

4.1 Base Case – Fully Linear Model

4.1.1 Optimality of Constant Pricing

We start with the simplest interesting case of our model, in which costs, demand, and consumption are all stationary and linear. In addition, we truncate demand at zero. That is, let

$$K_t(d) = kd, \text{ with } k > 0$$

$$D_t(p, M) = [a - bp - gM]^+, \text{ with } a, b \geq 0 \text{ and } 0 < g \leq 1; \text{ and}$$

$$C_t(M) = cM, \text{ with } 0 < c \leq 1.$$

Note that the latter condition implies that $\tilde{C}_t = (1 - c)M \geq 0$. Moreover, we assume the terminal reward to be zero. We show that under these circumstances it is optimal in the long-run to charge a constant price.

Proposition 2. *The full linear model defined above fulfills the conditions of Theorem 1. In particular, an optimal pricing policy exists that converges to a constant.*

Proof. Verification of the conditions of Theorem 1 is immediate. ■

Proposition 2 reveals that this purely linear environment does not offer any incentives for dynamic price changes. Note therefore that this is also an environment in which the typical modeling assumption of stationary demand can be justified. Intuitively, in a linear, monopolistic world, spreading demand equally across periods is the best one can do. Thus, one needs some kind of nonlinear effects to justify price dynamics.

4.1.2 Explicit Solution for Linear Model

The simple form of the above linear model allows for an explicit solution, which we explore in this subsection. For the undiscounted, average-profit case (i.e. $\alpha = 1$), one can obtain the optimal equilibrium price p^* and corresponding market stockpile M^* by equating M_t and M_{t+1} and then differentiating the profit per period. Straightforward algebra yields:

$$M^* = \frac{1}{2} \left(\frac{a}{b} - k \right) \frac{b(1-c)}{1-(1-c)(1-g)},$$

$$p^* = \frac{1}{2} \left(\frac{a}{b} + k \right),$$

$$\Pi(p^*, M^*) = \frac{1}{4} \left(\frac{a}{b} - k \right)^2 \frac{bc}{1-(1-c)(1-g)}.$$

These values hold for $k < a/b$. Otherwise, the solution degenerates and it is optimal not to sell any items at all.

For the discounted version of the linear model, we can explicitly calculate the value function and optimal price under a slight relaxation of our assumptions. Since this is fairly exceptional for a dynamic program, we explore this structure here.

If we do not rule out negative pricing and negative demand, the dynamic program outlined in Section 4.1.1 possesses some interesting structural characteristics: namely, an optimal pricing policy that is linear in the state and a value function that is quadratic in the state. This model forms a “demand side” analog of the linear-quadratic production models of Holt *et al.* (1960).

Theorem 2. *If the terminal value function has the following quadratic form:*

$$V_N(M) = r_N + s_N M + u_N M^2,$$

for scalars r_N , s_N , and $u_N \leq \frac{g}{2b}$, then in each period t the optimal price p_t^* will be a linear, decreasing function of M :

$$p_t^* = \frac{a + kb - \alpha b(1-c)s_{t+1} - 2\alpha ab(1-c)^2 u_{t+1}}{2b(1 - \alpha b(1-c)^2 u_{t+1})} - \frac{g + 2\alpha b(1-c)^2(1-g)u_{t+1}}{2b(1 - \alpha b(1-c)^2 u_{t+1})} M,$$

and the resulting value function will be quadratic in M with the following recursions for the scalars:

$$r_t = \frac{(a - bk)^2 + 2\alpha b(2r_{t+1} + (1-c)(a - bk)s_{t+1}) + \alpha^2 b^2(1-c)^2(s_{t+1}^2 - 4r_{t+1}u_{t+1})}{4b(1 - \alpha b(1-c)^2 u_{t+1})},$$

$$s_t = -\frac{g(a - bk) + \alpha b(1-c)[(g - 2)s_{t+1} - 2(1-c)(a - bk)u_{t+1}]}{2b(1 - \alpha b(1-c)^2 u_{t+1})}, \text{ and}$$

$$u_t = \frac{g^2 + 4\alpha b(1-c)^2(1-g)u_{t+1}}{4b(1 - \alpha b(1-c)^2 u_{t+1})} \leq \frac{g}{2b}.$$

Proof. See attached companion document ■

The condition $u_N \leq \frac{g}{2b}$ is a sufficient condition for the problem to be concave in every period. This condition also ensures that $u_t \geq 0$ and $s_t \leq 0$ for all t , which together guarantee that $V_t(M)$ is convex, first decreasing and later increasing in M . As with the models of Holt et al., the quadratic value function presents a quandary: the value function eventually increases in M due to negative demand and negative pricing, a situation which is difficult to imagine in practice. However, the rationalization provided by Holt et al. also applies here: in most applications of this model, the problematic (increasing) portion of the value function is unlikely to be visited and becomes practically irrelevant. An illustrative numerical example is available in the attached companion document.

4.2 Revenue Driven Price Dynamics

4.2.1 Nonlinear Market Response

We now address non-linear effects and investigate how they may justify price variations. We start by considering the demand side. To this end, we consider an exponential demand function, which is common in the marketing literature (see for example Jeuland & Shugan (1988) and Stavins (1997)), and incorporate a dependence on the market stockpile. The cost side remains unchanged, i.e., we still assume constant unit costs. Consumption is general. Specifically, let

$$D_t(p, M) = ae^{-bp-gM} \text{ with } a, b, g \geq 0 \text{ and } K_t(d) = kd, \text{ with } k > 0.$$

This model does not satisfy the conditions of Theorem 1. Specifically, the demand function violates Conditions 1 and 2 and possibly Condition 3, depending on parameter values. The more general condition in Proposition 1 is that

$$\frac{D}{(D_p)^2} \left(D_{pp} \frac{1+D_M}{D_p} - D_{pM} \right) - \frac{D_M+2}{D_p} = \frac{1-gD}{bD} \geq 0$$

for all p and M . This condition is satisfied for the above functional forms if and only if $ga \leq 1$. Thus, for $ga > 1$ dynamic price changes may potentially outperform constant pricing in this setting.

We illustrate numerically a case where we do find dynamically varying prices to be superior to constant pricing. To this end, we have implemented a dynamic programming algorithm for our model in a spreadsheet. The spreadsheet essentially consists of two worksheets. The rows of the first sheet contain the values of $V_t(M)$ for $t = n, \dots, N$ for a discrete grid of M -values. The second sheet uses the values of $V_n(M)$ to calculate $W_{n-1}(p, M)$ on a two-dimensional grid. Taking the maximum of each column yields $V_{n-1}(M)$. A Visual Basic macro copies these values to a new row of the first worksheet and reduces the period index by 1 until $n = 0$.

A few comments are in order regarding our approach to creating a discrete problem. For the decision space (i.e., the values of p) we use equal step sizes. For the state space (i.e. the values of M) exponential step sizes have proven beneficial as they allow for higher precision in the relevant range of states. In addition, we used linear interpolation to account for the fact that the state grid is not closed under the transition rule. In other words, choosing a price p in a given state M_n on the grid may lead to a new stock pile level M_{n+1} that is not on the grid. In this case, we interpolate $V_{n+1}(M_{n+1})$ in the calculation of $W_n(p, M_n)$, based on the two nearest grid points. This interpolation proves crucial for the accuracy of the solution.

We would like to point out that our spreadsheet model is not intended to showcase a particularly efficient solution approach. Rather, we designed it to support our analysis of the qualitative properties of the proposed dynamic pricing model. We found it to serve this purpose well. Its chief advantage is the easy accessibility of all results and calculations. For example, we could directly inspect the value function $W(p, M)$ to understand problem structure. Similarly, we built numerous graphs that helped us understand the properties of the underlying model. We found the computational effort very acceptable for this purpose. Solving a 100-period model, on a $p \times M$ grid of 50×30 required a few seconds on a Pentium IV PC, while a 100×250 grid required about one minute.

Below we present the numerical solution to the above exponential demand model for the parameter values as listed in Table 1. Note that $ga > 1$, thus the convergence conditions do not hold. Figure 1 depicts the optimal pricing policy and corresponding system behavior for an initial market stockpile of 10 units. The bars in the upper part of the figure indicate the optimal prices in consecutive periods. The line graphs in the lower part of the figure refer to different product volumes: the dotted line shows the market stockpile, the bold line the demand volume, and the dashed line the consumption volume. Clearly, the optimal pricing policy does not converge to a constant price in this example. Instead, we observe a 7-period cyclical pattern. Table A1 in the attached companion document lists the corresponding numerical values for one such cycle. Each cycle contains one dominant sales period, in which a low price in combination with a low market stockpile initiate a large demand spike, and thus a sharp increase of the market stockpile in the next period. Virtually all profits are earned in the dominant sales period. Subsequently, the optimal prices increase, resulting in very low demand, if any. As a result, consumption gradually depletes the market stockpile at a rate dictated by $c = 0.5$. Once the stockpile falls below a certain threshold, the optimal price drops sharply, and the cycle repeats itself.

To a certain extent, this pricing policy resembles a promotion strategy. Promotional price cuts concentrate demand in certain periods, rather than distributing it evenly over time. In order for such a strategy to be profitable, the demand response to the price cut needs to be sufficiently strong. In this example, this effect is driven, in particular, by the fact that $D_M > 1$ for small values of M . In other words, a “starved” market overreacts, in the sense that for a given price, a drop of the market stockpile by one unit triggers a demand increase of more than one unit. It is worth relating this effect to the impact of loss aversion in reference-price models (Greenleaf (1995), Kopalle et al. (1996)). In both cases, the profitability of promotion pricing relies essentially on an asymmetric market response. The fact that, in our model, demand between promotions vanishes almost entirely is due to the fact that we model only a single market segment. In Section 5 we discuss how to extend our model to multiple market segments.

Plotting the optimal price p^* as a function of M supports the above observations (see Figure A1 in the companion document). In contrast with the linear model analyzed in the previous section, p^* is not monotone in M . For large stockpile levels the optimal price increases as the stockpile decreases, just as in the linear model. However, when the stockpile decreases far enough, the optimal price suddenly drops sharply. This behavior follows from the shape of the value function $W(p, M)$. Figure 2 shows that for a given M the value function has two local maxima in p . As M increases, the value of the left maximum (filled circle) decreases relative to the right maximum (open circle). At the point where the right maximum becomes the global maximum the optimal price jumps discontinuously.

The above example highlights a case in which a truly dynamic pricing policy is superior to constant pricing. We note, however, that we found that a constant price was optimal in many exponential demand scenarios. In the next section we provide a more systematic sensitivity analysis by reducing the complexity of the underlying problem.

4.2.2 Reduced Action Space: On-Off Pricing

The optimal pricing policy illustrated in Figure 1 has a rather complex, general form. However, the resulting pattern of demand suggests a simpler heuristic: as soon as the market stockpile drops to a certain critical level M_{low} , charge a price of p . Thereafter, do not sell anything until the market stockpile has fallen to M_{low} again, at which time the cycle repeats. This heuristic policy results in a cyclical pattern in which a single period with positive demand is followed by a number of periods (possibly zero) of no demand. In other words, we eliminate “off-promotion” sales entirely. (We refer again to Section 5 for an extension to multiple market segments.) Within this class of policies, the pricing problem is reduced to finding the optimal values of M_{low} and p .

In order to solve this problem, it is helpful to express p in terms of the cycle length n . One can then determine the optimal value of M_{low} for each n and then find the optimal cycle length through enumeration. Note that an optimal cycle length of unity corresponds to a constant pricing policy. We

show that under the conditions of Lemma 1 a cycle length of unity is always optimal, provided that unit costs are constant.

Theorem 3. *Assume that Conditions 2 and 3 of Lemma 1 hold and that unit costs k are constant.*

Then the discounted profit of a cyclical policy, which sets a price of $p^ \geq k$ every n periods, where $n \geq 2$, and does not sell anything in the intermediate periods, is no larger than the discounted profit achieved by charging p^* in each period.*

Proof. See attached companion document. ■

We now apply the cycle analysis to the example of Section 4.2.1. To this end, we first write the returns of the cyclical policy over an infinite horizon as $W(p, M_{low}, n) = (p-k) D(p, M_{low}) / (1-\alpha^n)$. Using the stability condition $M_{low} = (1-c)^n (M_{low} + D(p, M_{low}))$, and filling in the specific demand function, we can express p in terms of n and M_{low} . We get $p = - (1/b)[\ln(\gamma M_{low} / a) + g M_{low}]$, with $\gamma = (1-c)^{-n} - 1$. Inserting this value in the expression for W above, differentiating by M_{low} , and equating the resulting expression to zero yields $1 + kb + 2g M_{low} + \ln(\gamma M_{low} / a) = 0$.

We solve this equation numerically for M_{low} in a simple spreadsheet. Note that the left-hand side is increasing in M_{low} and therefore the solution is unique. We can now plot the maximum returns over an infinite horizon for different cycle lengths as shown in Figure 3. For our example, we obtain an optimal cycle length of $n = 7$, with $p = 5.02$ and $M_{low} = 2.17$ units. Note that this is very much in line with the more general optimal policy discussed in Section 4.2.1. The performance of both policies is very similar, with the heuristic trailing in profit by just 0.7%. In contrast, constant pricing is significantly inferior. We obtain the best constant pricing policy from the above analysis by setting $n = 1$. We get $p = 7.39$ and $M_{low} = 16.31$ units. Table 2 summarizes these results.

The on-off heuristic pricing policy also provides a basis for a more efficient sensitivity analysis. Figures A2 – A4 in the attached companion document illustrate the impact of different parameters on the

optimal cycle length. In summary, we find that the benefit of dynamic pricing relative to constant pricing increases with increasing market size, increasing sensitivity of demand to both price and market stockpile, decreasing unit costs, and increasing consumption ratio.

In conclusion, the example presented in this section shows how market factors can justify periodic price variations. The profit increase relative to a constant pricing policy relies on a strong demand surge in response to a price cut. This effect is dependent on a high sensitivity of demand to both price and market stockpile. It also requires that an increasing market stockpile sufficiently increases consumption, thereby limiting the effect of forward buying. All of these factors are well known in the marketing literature (see e.g. Ailawadi & Neslin (1998)). The contribution of our model is that it captures these effects analytically.

4.3 Cost Driven Price Dynamics

4.3.1 Non-linear Costs

We now turn to the operations aspects of our model and investigate the impact of a non-linear cost function. To this end, we consider a parametric cost function of the form $K(D) = kD^l$. For $l < 1$ we have increasing returns to scale, whereas $l > 1$ implies decreasing returns to scale. For $l = 1$ we obtain the original linear cost function.

In order to separate cost effects from revenue effects we return to the linear demand and consumption functions introduced in Section 4.1. Taking the linear model as a starting point, one might expect that increasing returns to scale on the supply side might favor batching of demand and therefore possibly lead to optimal price variations. In contrast, decreasing returns to scale seem to call for spreading demand uniformly and thus for setting constant prices. Theorem 1 confirms this intuition. Since a linear demand function satisfies Conditions 1-3 and 5 of Theorem 1, application of the theorem depends solely on the curvature of the cost function. We illustrate the implications in two examples below.

4.3.2 Increasing Returns to Scale

We start with the case of increasing returns to scale, i.e., $l < 1$. Specifically, we let $k = 48$ and $l = 0.5$ while the demand and consumption functions are as in the example of Section 4.1. The resulting optimal policy turns out to have the same structure as in the revenue-driven dynamic pricing example in Figure 1 (see Figure A5 in the companion document for the exact plot of the new example). Of note is that the cyclical on-off heuristic is strictly optimal in this case. All demand occurs in the first of every four periods, with a price promotion serving to induce that demand. However, the underlying motivation for this policy is now very different. Instead of a demand effect, it is supply-side economies of scale that drive the concentration of sales in cyclical intervals. Recall that our model does not include inventories at the firm. Therefore, batch production also implies batch sales. We discuss how to relax this admittedly strong assumption in Section 5.

In all of our experiments with the above functional forms we found an on-off pricing policy to be optimal (possibly with a cycle length of $n = 1$). We apply the heuristic procedure introduced in Section 4.2 to investigate the optimal cycle length for an on-off policy. We obtain $p = [a - M_{low}(g+\gamma)] / b$, where again $\gamma = (1-c)^n - 1$. Moreover, for $l \neq 0$ the new optimality condition for M_{low} is

$$2M_{low}(g + \gamma(1 - g - \gamma)) + a\gamma - bkl\gamma^l M_{low}^{l-1} = 0.$$

For $l = 0$ the last term on the left-hand side disappears. As in Section 4.2.2, this equation can be solved numerically and used to evaluate the returns for different cycle-lengths (see Figure A6 in the companion document for details). It is worth noting that choosing $n = 1$, i.e. constant pricing, is infeasible in this example in the sense that it does not allow for any profits to be made. To see this, note that the consumption ratio $c = 0.5$ implies that $D = M_{low}$ for constant pricing. Thus, the equilibrium condition yields the price-demand curve

$$p = a/b - D(1 + g)/b = 10 - 0.09D,$$

which is below the unit cost curve $K(D)/D = 48D^{-0.5}$ for all values of D .

Before addressing the impact of the parameter l , we first consider the alternative case of decreasing returns to scale.

4.3.3 Decreasing Returns to Scale

As discussed, decreasing returns to scale (reflected by $l > 1$) act as a driver for smoothing the demand pattern. In the case of linear demand, the optimal pricing policy thus remains constant. This result also follows from Theorem 1. Further, decreasing returns to scale on the supply side may reduce or even eliminate revenue-driven pricing cycles. We illustrate this effect in the example discussed in Section 4.2. To this end, we combine exponential demand and polynomial costs and analyze again the performance of different cycle lengths for the on-off pricing heuristic. In the solution procedure, the expression for p remains the same as in Section 4.2. The optimality condition for M_{low} now reads $1 + 2gM_{low} + \ln(\gamma M_{low}/a) + kbl(\gamma M_{low})^{l-1} = 0$. Figure 4 shows the optimized returns as a function of the cycle length for different values of l from 0.6 to 1.4 in steps of 0.2. The case $l = 1$ coincides with the original example of Section 4.2. For $l < 1$ we find that supply-side economies of scale add even further to the benefit of dynamic price variations, in line with the results of Section 4.3.2. However, for $l > 1$ this benefit decreases and eventually disappears. In that case, increasing costs, e.g., due to congestion, outweigh the additional revenues of deviating from constant pricing.

The above examples highlight the interplay of demand and supply-side effects on dynamic pricing policies. Factors exist on both the revenue and operations sides of the issue that may justify or rule out dynamic price variations. The final choice between constant and dynamic pricing can only be made by *jointly* considering supply and demand, i.e., costs and revenues. This result argues strongly for coordination between operations and marketing in any pricing policy decisions.

5. DISCUSSION & EXTENSIONS

In this paper we have examined the linkage between pricing and operations by employing a finite horizon dynamic programming model. We develop sufficient conditions under which the optimal prices

converge to a constant. As it happens, these conditions support our intuition and apply in a broad range of environments. For instance, they rule out excessive reactions of demand to price changes and of demand and consumption to changes in the market stockpile. Perhaps the strongest condition is that the cost of goods exhibits no economies of scale, which is true, for instance, when congestion effects are significant. However, when setup costs are significant, we will observe economies of scale, which work against constant pricing. Hence, we extend the analysis to different operating environments. First we prove the optimality of constant pricing when costs, demand and consumption are all stationary and linear. Then we examine the effects of nonlinear market response and illustrate the pattern of dynamic optimal prices. Finally, we examine the effects of increasing and decreasing returns to scale in unit costs. Table 3 captures, in a simplified way, these effects, along with price and market inventory effects on demand. When operations exhibits increasing returns to scale, and demand responds strongly to decreasing price and market inventory, there is clearly a strong incentive for dynamic prices. However, when these forces are decreasing and weak, optimal prices are constant. When the operations and market forces pull in opposite directions, the need for coordination is magnified because it is not clear, ex ante, whether constant or dynamic prices are optimal. Furthermore, if dynamic prices are warranted, it is not clear how sharply they should vary over time.

Our model has several limitations that suggest fertile areas for future research. We comment on four of these below.

Although our model captures the effects of the market stockpile, it does not allow for stock-keeping at the firm. Hence, production must match demand in each period. One extension would be to introduce an additional state variable for inventory, I_t , and a decision variable for production quantity, x_t , as in traditional inventory models. These adjustments would make the model “symmetric” in the sense that it would have an internal stock, I , which is replenished by production and depleted by demand, and an external stock, M , which is replenished by demand and depleted by consumption. The two-dimensional state space of this model would increase complexity significantly – i.e. the optimal policy would be more difficult to characterize, and numerical solutions would require conventional

implementation of dynamic programming algorithms rather than a spreadsheet. Inventory at the firm provides the option of decoupling demand from procurement and thereby revenues from costs. Therefore, the impact of operations on the optimal pricing policy should decrease. In particular, stock-keeping provides an alternative option for achieving economies of scale (or avoiding diseconomies of scale). Dynamic pricing should become less (more) attractive in the case of increasing (decreasing) returns to scale. In the case of constant returns to scale, we expect little change since there is no need for storing inventory. Note that while inventory loosens the ties between the supply and demand effects, it does not eliminate them since excessive inventory is still costly.

Because our current model is purely deterministic, a second extension would be to consider a stochastic market response where both demand and consumption would include a stochastic term. The literature is rich with both additive and multiplicative stochastic demand models (Talluri & Van Ryzin (2004)). The same approach, in theory, could be used for stochastic consumption, with the additional restriction that consumption cannot exceed the available stockpile. Although model formulation would be straightforward, it is not obvious whether the analysis would produce reasonable analogues to our deterministic theoretical results, because a stochastic model will certainly not converge to a constant. Spreadsheet solution of a dynamic program should, in principle, remain applicable, although the calculations would become significantly more complex. One may be able to analytically determine criteria that guarantee a monotone pricing policy wherein price changes serve only as a response to stochastic disturbances, rather than as an active inducement to dynamics. But this remains for further research. In any case, stochastic market behavior should prompt a response that increases price fluctuations. On the other hand, market behavior would be more difficult to fine tune with price adjustments, and therefore, it would be more difficult to achieve a specific cycle length. Constant pricing, due to its smoothing effect, may thus become relatively more attractive.

A third extension would be to expand the analysis to additional market segments. For instance, one could model a second segment that does not stockpile inventory by replacing D with $D + D_1$ in the profit function, where D_1 is independent of M , the market stockpile. The rest of the model would remain

unchanged. One would expect that market segmentation would provide an additional driver of price dynamics. The optimal price for each segment should typically be different, and cyclical price changes may provide a means for addressing both segments. In order for this to work, the stockpiling segment must be more price sensitive than the non-stockpiling one. In terms of our current results, we would not observe sales between “promotion” instances dropping to zero, since we would still be selling to the non-stockpiling segment. This segmentation effect is the main driver of the price dynamics in the models of Iyer & Ye (2000) and Huchzermeier et al. (2002).

Finally, any discussion of promotions raises the issue of competition. Because our model is purely monopolistic, a fourth extension would be to explicitly capture competitive dynamics. Given the complexity of the price/demand, market stockpile/demand and market stockpile/consumption effects, a game theoretic approach seems intractable. It is likely, however, that the results would generate a coordinated equilibrium that would serve to induce cyclic dynamic prices. However, many unknowns remain. For instance, what would be the effect of competitor sales on the market stockpile, and how should one model this effect? Further, how should one model the competitive response to each firm’s strategy? And finally, if the model contains stochastic responses, how should that be captured?

References

- Ailawadi, K., & Neslin, S. A. (1998). The Effect of Promotion on Consumption: Buying More and Using It Faster. *Journal of Marketing Research*, 35, 390-398.
- Andel, T. (1997). Warehousing in Manufacturing: From Inventory Holding to Product Molding. *Transportation & Distribution*, 38(9), 80-84.
- Araman, V. F., & Ozer, O. (2005). *Capacity and Inventory Management in the Presence of a Long-Term Channel and a Spot Market*. Unpublished Working Paper, Stanford University, Stanford, CA.
- Assuncao, J. L., & Meyer, R. J. (1993). The Rational Effect of Price Promotions on Sales and Consumption. *Management Science*, 39(5), 517-535.

- Bernstein, C. (2004). The Top 10 Technologies to Watch -- Promotion and Merchandise Optimization: Taking the Next Step. *Executive Technology*, <http://www.executivetechnology.com/ViewCA.cfm?ID=311>.
- Blattberg, R. C., Eppen, G. D., & Lieberman, J. (1981). A Theoretical and Empirical Evaluation of Price Deals for Consumer Nondurables. *Journal of Marketing*, 45(1), 116-129.
- Boyaci, T., & Ray, S. (2003). Product Differentiation and Capacity Cost Interaction in Time and Price Sensitive Markets. *Manufacturing & Service Operations Management*, 5(1), 18-36.
- Chan, L. M. A., Shen, Z. J. M., Simchi-Levi, D., & Swann, J. L. (2004). Coordination of Pricing and Inventory Decisions: A Survey and Classification. In D. Simchi-Levi, S. D. Wu & Z.-J. M. Shen (Eds.), *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era* (pp. 335-392). Boston, MA: Kluwer Academic Publishers.
- Chan, L. M. A., Simchi-Levi, D., & Swann, J. (2002). *Dynamic Pricing Strategies for Manufacturing with Stochastic Demand and Discretionary Sales*. Unpublished Working Paper, Georgia Institute of Technology, Atlanta.
- Chen, X., & Simchi-Levi, D. (2002a). *Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Finite Horizon Case*. Unpublished Working Paper, MIT, Cambridge, MA.
- Chen, X., & Simchi-Levi, D. (2002b). Coordinating Inventory Control and Pricing Strategies with Random Demand and Fixed Ordering Cost: The Infinite Horizon Case. *Mathematics of Operations Research*, 29(3), 698-723.
- Elmaghraby, W., & Keskinocak, P. (2003). Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions. *Management Science*, 49(10), 1287-1309.
- Federgruen, A., & Heching, A. (1999). Combined Pricing and Inventory Control under Uncertainty. *Operations Research*, 47(3), 454-475.

- Fleischmann, M., Hall, J. M., & Pyke, D. F. (2004). Smart Pricing: A Review of Recent, and Some Seminal, Work Linking Pricing Decisions with Operational Insights. *Sloan Management Review*, 45(2), 9-13.
- Gourville, J., & Soman, D. (2002). Pricing and the Psychology of Consumption. *Harvard Business Review*(September), 91-96.
- Green, H. (2003, November 24). The Web Smart 50: Northern Retail Group. *Business Week*, 94-96.
- Greenleaf, E. A. (1995). The Impact of Reference Price Effects on the Profitability of Price Promotions. *Marketing Science*, 14(1), 82-104.
- Gupta, D., Hill, A. V., Bouzdine-Chameeva, T., & Ahn, H. (2003). *A Pricing Model for Clearing End of Season Retail Inventory*. Unpublished Working paper, University of Minnesota, Minneapolis, MN.
- Hall, J. M., Kopalle, P. K., & Pyke, D. F. (2005). *Static and Dynamic Pricing of Excess Capacity in a Make-to-Order Environment*. Unpublished Working Paper, Dartmouth College, Hanover, NH.
- Holt, C. C., Modigliani, F., Muth, J. F., & Simon, H. A. (1960). *Planning Production, Inventories and Work Force*. Englewood Cliffs, New Jersey: Prentice-Hall.
- Huchzermeier, A., Iyer, A., & Freiheit, J. (2002). The Supply Chain Impact of Smart Customers in a Promotional Environment. *Manufacturing & Service Operations Management*, 4(3), 228-240.
- Iyer, A. V., & Ye, J. (2000). Assessing the Value of Information Sharing in a Promotional Retail Environment. *Manufacturing & Service Operations Management*, 2(2), 128-143.
- Jeuland, A. P., & Shugan, S. M. (1988). Channel of Distribution Profits When Channel Members Form Conjectures. *Marketing Science*, 7(2), 202-210.
- Kopalle, P., Rao, A.G., & Assunção, J.L. (1996). Asymmetric Reference Price Effects and Dynamic Pricing Policies. *Marketing Science*, 15(1), 60-85.
- Macé, S., & Neslin, S. A. (2004). The Determinants of Pre- and Postpromotion Dips in Sales of Frequently Purchased Goods. *Journal of Marketing Research*, 41(3), 339-350.

- Maglaras, C., & Zeevi, A. (2005). Pricing and Design of Differentiated Services: Approximate Analysis and Structural Insights. *Operations Research*, 53(2), 242-262.
- Mendelson, H., & Whang, S. (1990). Optimal Incentive-Compatible Priority Pricing for the M/M/1 Queue. *Operations Research*, 38(5), 870-883.
- Natter, M., Reutterer, T., Mild, A., & Taudes, A. (2005). *An Assortment-Wide Decision-Support System for Dynamic Pricing and Promotion Planning in Diy Retailing*. Unpublished Working paper, Vienna University of Economics and Business Administration, Vienna, Austria.
- Neslin, S. A. (2002). *Sales Promotion* (Vol. 1): Marketing Science Institute.
- Pauwels, K., Hanssens, D. M., & Siddarth, S. (2002). The Long-Term Effects of Price Promotions on Category Incidence, Brand Choice, and Purchase Quantity. *Journal of Marketing Research*, 39(4), 421.
- Popescu, I., & Wu, Y. (2005). *Dynamic Pricing Strategies under Repeated Interactions*. Unpublished Working paper, INSEAD, Fontainebleau, France.
- Puterman, M. L. (1994). *Markov Decision Processes*. New York: John Wiley & Sons.
- Pyke, D. F., & Johnson, M. E. (2004). Real-Time Profit Optimization: Coordinating Demand and Supply Chain Management. *ASCET*, 6, 2-6.
- Radjou, N., Orlov, L. M., & Herbert, L. (2003). Helping Supply Chain Cope with Demand. *Forrester Report*(June).
- Ray, S., Gerchak, Y., & Jewkes, E. M. (2005). Joint Pricing and Inventory Policies for Make-to-Stock Products with Deterministic Price-Sensitive Demand. *International Journal of Production Economics*, 97(2), 143-158.
- Sears Bolsters Merchandise Performance. (2004, 07/26/04). *Top of the Net*, RetailSystems.com, <http://www.retailsystems.com/index.cfm?PageName=PublicationsTONArticle&ArticleId=3921>.
- So, K. C., & Song, J.-S. (1998). Price, Delivery Time Guarantees and Capacity Selection. *European Journal of Operational Research*, 111(1), 28-49.

- Sogomonian, A. G., & Tang, C. S. (1993). A Modeling Framework for Coordinating Promotion and Production Decisions within a Firm. *Management Science*, 39(2), 191-203.
- Srinivasan, S., Pauwels, K., Hanssens, D. M., & Dekimpe, M. G. (2004). Do Promotions Benefit Manufacturers, Retailers, or Both? *Management Science*, 50(5), 617-629.
- Stavins, J. (1997). Estimating Demand Elasticities in a Differentiated Product Industry: The Personal Computer Market. *Journal of Economics and Business*, 49(4), 347.
- Talluri, K., & Van Ryzin, G. (2004). *Revenue Management*. Boston: Kluwer Academic Publishers.
- Van Mieghem, J. A., & Dada, M. (1999). Price Versus Production Postponement: Capacity and Competition. *Management Science*, 45(12), 1631-1649.
- Welch, D. (2003, November 24). The Web Smart 50: Gm. *Business Week*, 88.
- You, P.-S. (2005). Inventory Policy for Products with Price and Time-Dependent Demands. *Journal of the Operational Research Society*, 56, 870–873.
- Zhu, K., & Thonemann, U. W. (2003). *Coordination of Pricing and Inventory Control across Products*. Unpublished Working paper, Stanford University, Stanford, CA 94305-4026.

APPENDIX: TABLES AND FIGURES

α	a	B	g	c	k
0.95	7000	0.6	0.1	0.5	3

Table 1: Parameter Settings for Example 4.2.1

Policy	Profit in perpetuity	Percent Deviation
Optimal full dynamic pricing	1867.6	0
Best on-off pricing	1854.2	0.7%
Constant pricing	1430.3	23.4%

Table 2: Performance Comparison of Heuristics

	Strong Response of Demand to Price and Market Inventory Reductions	Weak Response of Demand to Price and Market Inventory Reductions
Increasing Returns to Scale in Operations	Strong	Weak
Decreasing Returns to Scale in Operations	Weak	None

Table 3: Incentives for Dynamic Pricing




Figure 1: Optimal Pricing Policy in Example 4.2.1




Figure 2: Value Functions for Different Stockpile Values




Figure 3: Performance of On-Off Pricing Heuristic for Different Cycle Lengths




Figure 4: Impact of Supply-Side Returns to Scale on Benefit of Dynamic Pricing