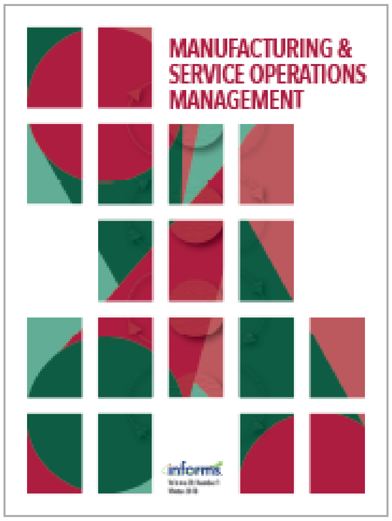

Manufacturing & Service Operations Management

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

The MSOM Society Student Paper Competition: Extended Abstracts of 2003 Winners

Amar Sapra, Jiri Chod, Aydın Alptekinoğlu, Li Chen, A. Gürhan Kök, Victor Martínez-de-Albéniz,

MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 6, No. 1, Winter 2004, pp. 92–112
ISSN 1523-4614 | EISSN 1526-5498 | 04 | 0601 | 0092
DOI 10.1287/msom.1040.0041

The MSOM Society Student Paper Competition: Extended Abstracts of 2003 Winners

As is our tradition at the journal, we are pleased to publish the extended abstracts from the winners of the 2003 MSOM Society Student Paper Competition. We do this to celebrate the achievements of these young scholars and provide you with the opportunity to learn about their work in more detail.

The 2003 prize committee was chaired by Professor Gérard Cachon from the University of Pennsylvania. The other committee members were: Dan Adelman (University of Chicago), Narendra Agrawal (Santa Clara University), Yossi Aviv (Washington University), René Caldentey (New York University), Wedad Elmaghraby (Georgia Institute of Technology), Noah Gans (University of Pennsylvania), Roman Kapuscinski (University of Michigan), Pinar Keskinocak (Georgia Institute of Technology), Constantinos Maglaras (Columbia University), Joe Mazzola (Georgetown University), Serguei Netessine (University of Pennsylvania), Michael Pinedo (New York University), Erica Plambeck (Stanford University), Nils Rudi (University of Rochester), Sergei Savin (Columbia University), Kevin Shang (Duke University), Terry Taylor (Columbia University), Christian Terwiesch (University of Pennsylvania), Brian Tomlin (University of North Carolina), Tunay Tunca (Stanford University).

The 2003 prizewinners are:

First Place
Amar Sapra, Cornell University
Advisor: Peter L. Jackson
"The Martingale Evolution of Price Forecasts in a Market for Supply Chain Capacity"

Second Place
Jiri Chod, University of Rochester
Advisor: Nils Rudi
"Strategic Investments, Trading, and Pricing Under Forecast Updating"

Finalists
Aydın Alptekinoğlu, UCLA
Advisor: Charles J. Corbett
"Mass Customization vs. Mass Production: Variety and Price Competition"

Li Chen, Stanford University
Advisor: Erica L. Plambeck
"Dynamic Bayesian Quantity Control when Some Customers Will Accept a Substitute"

A. Gürhan Kök, Duke University
Advisor: Marshall L. Fisher
"Demand Estimation and Assortment Optimization Under Substitution: Methodology and Application"

Victor Martínez-de-Albéniz, MIT
Advisor: David Simchi-Levi
"A Portfolio Approach to Procurement Contracts"

The Martingale Evolution of Price Forecasts in a Market for Supply Chain Capacity

Amar Sapra
School of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York 14853, as236@cornell.edu

Advisor: Peter L. Jackson
School of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York 14853

1. Introduction

We are seeing the advent of markets in which manufacturers operate as job shops, but price their services dynamically in response to the level of their capacity utilization. We are interested in the volatility of prices in such a market, and we recognize that this volatility cannot be properly understood without recognizing the intertemporal nature of demand price elasticity. That is, the demand for manufacturing capacity in such markets typically comes from consumer products assemblers who have the opportunity to modulate their demand for capacity over time in response to the current price and anticipated price changes. By carrying inventories or enduring backorders with their customers (the end consumers), these assemblers can shift their demand forward or backward in time as the price of capacity changes. We investigate this phenomenon using classical economic equilibrium analysis under the assumption of perfect markets and rational expectations.

The Martingale Model of Forecast Evolution (MMFE) was developed independently by Graves et al. (1986), and Heath and Jackson (1994). In this paper, we derive a continuous time analog of the additive MMFE model. In our model, we assume that the forecast for the rate of consumer demand at any instant t, $0 < t < \infty$ evolves as a continuous Martingale process over $[0, t]$. In the second half of the paper, we apply the continuous-time MMFE model to study the evolution of the equilibrium price in a market for capacity. Our goal is to study how the resolution of consumer demand uncertainty translates into the resolution of uncertainty of the equilibrium price of capacity. We also study the impact of supply chain cost parameters on the the rate of resolving price uncertainty.

2. Continuous-Time Martingale Model of Forecast Evolution

Let $f(s, t)$ denote the forecast at time s of the rate of consumer demand for time t and, let $d_s f(s, t)$ denote the forecast update at s of rate of consumer demand at t. Thus, $f(t, t)$ represents the actual rate of consumer demand at t. Let $\mathbf{W}(\cdot)(\equiv(W_1(\cdot), W_2(\cdot), \ldots, W_n(\cdot))$ be an n-dimensional Weiner's process defined on a complete probability space $(\Omega, \mathcal{F}, \mathcal{P})$ and let $\mathcal{F}_t$ be the filtration. The model developed in this section is similar to the Futures Prices model in Health and Jara (2003).

Assumption 1. *$\mathcal{F}_t = \sigma\{\mathbf{W}(\mathbf{s})\colon 0 \le s \le t\}$. Further, $\mathcal{F}_0$ is $\mathcal{P}$-complete and $\mathcal{P}$-degenerate.*

Assumption 2. *$f(\cdot, t)\colon \Omega \times [0, t] \to \mathcal{R}$ is a Martingale.*

Using Theorem 4.15 (Karatzas and Shreve 1988, p. 182), if $f(t, t)$ is in $\mathcal{L}^2$ and is measurable with respect to $\mathcal{F}_s$, then there exist progressively measurable $\sigma_i \in \mathcal{L}^2$, $i = 1, \ldots, n$ such that:

$$f(s, t) = f(0, t) + \sum_{i=1}^{n} \int_0^s \sigma_i(u, t)\, dW_i(u),$$

which implies:

$$d_s f(s, t) = \sum_{i=1}^{n} \sigma_i(s, t)\, dW_i(s) = \sigma(s, t)\, dW(s)^T$$

where $\sigma \equiv (\sigma_1, \sigma_2, \ldots, \sigma_n)$.

2.1. Additive Model

For the rest of this section, we shall assume $\sigma(\cdot, t)$ is deterministic. Let $F(s, t)$ be the forecast at s of the cumulative consumer demand until $t \ge s$. Thus, $F(t, t)$ represents the actual cumulative consumer demand until t. The state equation for the realized cumulative consumer demand may be written as,

$$dF(t,t) = f(t,t)\,dt + \sum_{i=1}^{n} \sigma_i(t,t)\,dW_i(t).$$

Lemma 1. *$F(\cdot, t)\colon \Omega \times [0,t] \to \mathscr{R}$ is a Martingale process.*

3. Application of MMFE: The Market Model for Capacity

We consider a simple, continuous-time market for the commitment of production capacity to the actual production of a single homogeneous product. The market has numerous agents of two types: owners of capacity and consumers of capacity. The equilibrium price at every instant in the market is determined by the balance of supply and demand for capacity. We assume that the price paid by a buyer in the capacity market does not affect the price charged by the buyer in the consumer market. Therefore, the revenue earned by the buyer in the consumer market is not part of the objective function for the buyer.

Each buyer forecasts his or her share of consumer demand over the remaining horizon. While estimating the instantaneous consumer demand at any instant, the buyers take into account the resolution of uncertainty from the beginning of the horizon until that instant. Let $F_j(t,t)$ denote the cumulative demand from end consumers for the sales of buyer j through time t, and let $F_j(s,t)$ denote the forecast of this cumulative demand made at time s. Any buyer's decision regarding the order quantity is dependent on future prices in the capacity market. We invoke the theory of Rational Expectations (Muth 1961) to understand how the prices evolve in the capacity market.

Let C_k, $Y_k(t)$, and $y_k(t)$ denote the capacity, cumulative production through time t and instantaneous rate of production at time t, respectively, for seller k. We assume $Y_k(0) = Y_{0k}$. Let $X_j(t)$, X_{0j}, and $x_j(t)$ denote the cumulative orders for production placed by buyer j through time t, the initial inventory of buyer j, and the instantaneous order rate placed by buyer j at time t, respectively.

There is no lead time between order placement and order delivery: Production is instantaneously distributed from sellers to buyers. A market price $P(t)$ per unit is paid by buyers and received by sellers for each unit of production. All agents in this market are assumed to be pricetakers. We consider this price to be a premium or discount from some exogenously determined price that considers such factors as unit production costs and consumer demand price sensitivity for the final good. In this way, we focus on the role of price in managing the evolution of inventory in the economy. Because it can be either a premium or a discount, we do not constrain the sign of $P(t)$.

The finite-horizon version of the buyer's problem is:

$$\begin{aligned}
\min_{x_j(\cdot)\in\mathcal{U}_j[0,T]} \quad & \mathrm{E}\int_0^T \left\{\pi I_j(t)^2 + P(t)x_j(t)\right\} dt \\
\text{s.t.} \quad & dI_j(t) = \Bigg(x_j(t) - D_j - \alpha_j\gamma\cos\gamma t \\
& \qquad - \int_0^t \sigma_j(u,t)\,\mathbf{dW}(u)^T\Bigg)\,dt \\
& \qquad - \sigma_j(t,t)\,\mathbf{dW}(t)^T, \quad t\in[0,T] \\
& X_j(0) = X_{0j},
\end{aligned}$$

where π denotes the net inventory penalty coefficient, D_j is the average demand rate, α_j is the amplitude of seasonal variation, γ is the period of seasonal demand, and $\sigma_j(s,t)\,\mathbf{dW}(s)$ is the forecast update at s for the rate of demand at time $t > s$. We assume that the forecast updates are the same for all buyers. $I_j(t)$ is the net inventory at time t (on-hand inventory less backorders). We use the following model of instantaneous demand for buyer j:

$$\begin{aligned}
dF_j(t,t) \equiv{} & \left(D_j + \alpha_j\gamma\cos\gamma t + \int_0^t \sigma_j(u,t)\,\mathbf{dW}(u)^T\right)dt \\
& + \sigma_j(t,t)\,\mathbf{dW}(t)^T.
\end{aligned}$$

We assume that the value of D_j is relatively large compared to the randomness in the demand and, therefore, the probability of occurrence of negative order rate or production rate is negligible. The same assumption ensures that the demand rate is nonnegative with a high probability.

The finite-horizon version of the seller's problem is:

$$\begin{aligned}
\min_{y_k(\cdot)\in\mathcal{U}_k[0,T]} \quad & \mathrm{E}\int_0^T \left\{\kappa(C_k - y_k(t))^2 - P(t)y_k(t)\right\} dt \\
\text{s.t.} \quad & dY_k(t) = y_k(t)\,dt, \quad t\in[0,T] \\
& Y_k(0) = Y_{0k},
\end{aligned}$$

where $T > 0$ denotes the length of the horizon, and κ is the penalty coefficient on the over- or under-utilization of capacity. We model the equilibrium price process as

$$P(t) = a(t) + \int_0^t \mathbf{b}(u,t)\,\mathbf{dW}(u)^T$$

for suitable (equilibrium) values of $a(\cdot)$ and $\mathbf{b}(\cdot,\cdot)$. Every agent knows $a(\cdot)$ and $\mathbf{b}(\cdot,\cdot)$ and makes decisions accordingly. According to the Rational Expectations theory, this would indeed result in the realization of $P(\cdot)$ in the market. Our goal in this paper is to compute the equilibrium values of $a(\cdot)$ and $\mathbf{b}(\cdot,\cdot)$ such that

$$\sum_j x_j(t) = \sum_k y_k(t)$$

at all times $t \in [0,T]$. We assume that $a(\cdot)$ and $\mathbf{b}(\cdot,\cdot)$ are deterministic functions.

Assumption 3. *$E\int_0^T P(t)^2\,dt < \infty$ and $P(t)$ is $\mathcal{F}_t$-adapted.*

3.1. Solution Approach

We next summarize the solution approach to obtain the equilibrium market quantities.

(1) Obtain the necessary and sufficient conditions for optimality of the Buyer's and the Seller's models using the Stochastic Maximum Principle (Cadenillas and Karatzas 1995).

(2) Aggregate the two sets of conditions over all the buyers and sellers, respectively, and combine them using the balance of supply and demand condition.

(3) Solve the resulting stochastic differential equations to obtain the equilibrium market production and price.

By adopting this approach, we are able to compute closed-form expressions for the functions $a(\cdot)$ and $\mathbf{b}(\cdot,\cdot)$ in equilibrium. Define the forecast at s of market price at t as

$$P(s,t) = E(P(t)\mid\mathcal{F}_s) = a(t) + \int_0^s \mathbf{b}(u,t)\,\mathbf{dW}(u)^T.$$

Clearly, the price forecast process evolves as a martingale. Now, the forecast of price at t gets updated by

$$d_sP(s,t) = \mathbf{b}(s,t)\,\mathbf{dW}(s)^T$$

at time $s < t$. The function $\mathbf{b}(s,t)$ can be interpreted as the forecast update coefficient at s for the price of capacity at t. The behavior of $\mathbf{b}(\cdot,\cdot)$ describes the resolution of price uncertainty in the capacity market.

4. Evolution of Market Price as a Martingale

We explore the impact of various parameters on the price forecast update, $d_sP(s,t)$, for a special case, $n=1$. We numerically analyze the impact of the rate of resolution of consumer demand uncertainty and the supply chain cost parameters on the resolution of the price uncertainty. We find that as the rate of resolution of demand uncertainty increases, so does the rate of resolution of price uncertainty. Thus, the learning behavior with regard to consumer demand is translated to learning behavior for the price of capacity.

We also find that the supply chain cost parameters do affect the resolution of price uncertainty. As κ/π (that is, the relative cost of changing production with respect to holding inventory) increases, the resolution of the price uncertainty occurs more uniformly in time. On the other hand, for low values of κ/π, most of the resolution of the price uncertainty is concentrated close to time t.

References

Cadenillas, A., I. Karatzas. 1995. The stochastic maximum principle for linear convex optimal control with random coefficients. *SIAM J. Control Optim.* **33**(2) 590–624.

Graves, S. C., H. C. Meal, S. Dasu, Y. Qiu. 1986. Two-stage production planning in a dynamic environment. E. S. S. Axater, C. Schneeweiss, eds. *Multi-Stage Production Planning and Control. Lecture Notes in Economics and Mathematical Systems*, Vol. 266. Springer-Verlag, Berlin, Germany, 9–43.

Heath, D. C., P. L. Jackson. 1994. Modeling the evolution of demand forecasters with applications to safety stock analysis in production/distribution systems. *IIE Trans.* **26**(3) 17–30.

Heath, D. C., D. Jara. 2003. Term structure models based on future prices. Working paper, Carnegie Mellon University, Pittsburgh, PA.

Karatzas, I., S. E. Shreve. 1988. *Brownian Motion and Stochastic Calculus*. Springer-Verlag, NY.

Muth, J. F. 1961. Rational expectations and the theory of price movements. *Econometrica* **29**(3) 315–335.

Strategic Investments, Trading, and Pricing Under Forecast Updating

Jiri Chod
The Simon School, University of Rochester, Rochester, New York 14627, chodji@simon.rochester.edu

Advisor: Nils Rudi
The Simon School, University of Rochester, Rochester, New York 14627

When deciding how much to invest in resources such as capacity or inventory, firms must typically rely on imperfect forecasts of future market conditions, such as demand, price, and exchange rates. Often, however, firms have the opportunity to correct such quantity decisions and reduce the associated uncertainty cost by trading these resources later, when more accurate market forecasts become available. Capacity trading or subcontracting is particularly widespread in industries with high investment costs and low production asset specificity, such as pharmaceuticals, telecommunications, or electronics manufacturing. Excess inventory trading is an equally prevalent practice in most industries. According to an AMR Research press release in 2000, the global market with excess inventory of consumer goods alone is on the order of a hundred billion USD.

While the existing subcontracting and transshipment literature (e.g., Kouvelis and Gutierrez 1997, Van Mieghem 1999, Rudi et al. 2001, Lee and Whang 2002) does not capture the forecast evolution over time, the extant literature on resource pooling under forecast updating (e.g., Lee and Whang 1998, Aviv and Federgruen 2001) is limited to centralized systems. We consider two firms that must make resource investments well ahead of the selling season, based on their forecasts of future market conditions. After the investment decisions are made and the selling season approaches, the firms update their forecasts. The forecast revisions may be based on observing signals such as market trends, sales of related products, etc. In general, as the forecasting horizon becomes shorter, the forecasts become more accurate, reflecting the reduction of uncertainty over time. Given the revised forecasts, the firms have the option to engage in a resource trade. When all uncertainty is finally resolved, the firms realize their operating profits.

While some of the excess inventory trading has moved to anonymous online auctions and exchanges, because of the risks of diluting the brand, selling to competitors, or purchasing defective or counterfeited goods, a majority of deals with excess inventory are still negotiated privately between firms with a good knowledge of each other (Pickering 2001). Similarly, capacity trading and subcontracting typically require face-to-face cooperative negotiations between the trading firms because of the inherent asset specificity. We assume that our two firms negotiate a trade contract that is a Nash bargaining equilibrium. A solution concept for cooperative games commonly used in economics, the Nash bargaining equilibrium is a Pareto-optimal contract that allocates the expected value created by cooperation (i.e., trading) between the firms in proportions that depend on their relative bargaining powers (see, e.g., Van Mieghem 1999).

As opposed to trading, the firms' investment decisions will typically be made unilaterally. Thus, in our base case, investment decisions result from a noncooperative stochastic game in which the payoffs depend on an embedded Nash bargaining stochastic trading game. We characterize the Nash equilibrium investment levels and prove existence and uniqueness of this equilibrium for a fairly general functional form of the operating profits while considering two types of trading frictions: transaction cost and resource specificity. At the equilibrium, the expectation of each firm's marginal investment payoff must be equal to its unit investment cost. We find that a firm's marginal investment payoff is the weighted average of its marginal investment payoffs in the cooperative investment scenario and in the case of no trading, weighted by the bargaining powers of the firm and its trading partner, respectively. Even though the firms know that they will have the opportunity to trade, each firm strategically considers the scenario of no trading to improve its alternative to trading and, hence, its bargaining position. The more powerful the firm, the more of the value created by the trade this firm expects to extract, and the less strategically it behaves in the investment stage.

To gain additional insights into the equilibrium, we then focus on price-setting firms facing stochastic demand functions. Specifically, we assume that each firm faces an isoelastic demand curve that is subject to a multiplicative random shock. Reflecting the relative flexibility of pricing versus quantity decisions in many applications, we allow the firms to set prices in response to the realized demand functions. Given this demand model and an arbitrary distribution of the initial as well as the updated forecasts, we prove that the option to trade leads to higher investments, prices, profits, and consumer surplus, and that investment cooperation further increases investments, profits, and consumer surplus, but decreases prices.

With responsive pricing, trading enables the firms to achieve higher expected prices, which lead to a higher marginal profitability of the investments and induce the firms to invest more. The higher initial investments then result in a higher total output, which, in turn, moderates the increase in expected prices. To be able to charge a higher price, the firms must sell to consumers who are willing to accept the higher price because they value the product more. This and the larger output are the reasons why the option to trade increases the expected consumer surplus. Because in the noncooperative investment game the firms weigh the trading and nontrading scenarios, the Nash equilibrium investment levels lie between the optimal cooperative investment levels and the optimal investment levels of the firms without the option to trade. Because the option to trade induces higher investments, the optimal cooperative investment levels must exceed the Nash equilibrium investment levels. Higher investments then lead to a higher total output, which, in turn, results in lower prices and higher consumer surplus.

To study the main driving forces of trading, we further assume that the forecast evolution process follows a two-dimensional geometric Brownian motion. This model of a forecast evolution process was initially proposed by Hausman (1969), and it has been supported by statistical studies of actual forecasting systems. A forecast evolution process asymptotically follows a geometric Brownian motion if forecast revisions result from observing a large number of information signals, and the forecast adjustments induced by these signals are proportional to the existing forecast value. The quality of the forecast revision is measured by the fraction of uncertainty, in terms of the variance-covariance matrix of the log-forecast, which is resolved by the time the trade is negotiated.

Given this model of forecast evolution, we prove that equilibrium investments, expected prices, profits, and consumer surplus are (i) increasing in the quality of the forecast revision, (ii) increasing in market variability, and (iii) decreasing in market correlation. A higher quality of the forecast revision leads, in expectation, to more efficient resource allocation and, hence, to higher marginal profitability of investments, inducing higher investment levels. A more efficient resource allocation means that more output is sold in the market where a higher price is acceptable, i.e., where consumers value the product more. This and the higher total output are the reasons why a better quality of the forecast revision leads to a higher expected consumer surplus.

The higher the market variability and the lower the market correlation, the more resource, in expectation, the firms trade. As trading leads to higher investments, prices, profits, and consumer surplus, it is not surprising that all of these measures are increasing in market variability and decreasing in market correlation. The opposite effect of market variability on the expected profits compared to the traditional newsvendor-type models stems from two factors. Firstly, the ability to set prices in response to the actual market conditions allows the firms to manage market variability more efficiently. Secondly, the higher the market variability, the higher the expected price differential between the markets and, hence, the higher the value of the option to trade.

Finally, we consider resource trading in a global context, when the firms generate their revenues in different currencies. In this case, in addition to demand curve uncertainty, the firms face uncertainty of exchange rate, which is assumed to follow a geometric Brownian motion. Exchange rate volatility has an effect similar to that of demand curve uncertainty, as it increases the expected real price differential and, hence, the value of the option to trade. Furthermore, the exchange rate volatility increases the marginal profitability of the resource investments and, thus, induces higher equilibrium investments. The closer to the actual selling season the trade takes place, the better forecast of the selling season exchange rate is available, and the more efficient, in expectation, the trade will be. Thus, the relative timing of the trade between the investment stage and the selling season has an effect similar to the forecast revision quality, increasing the equilibrium investments and expected profits.

References

AMR Research. 2000. AMR Research predicts the excess consumer goods inventory market will reach $60 billion by the end of the year. *AMR Research Press Release* (June 19) http://www.amrresearch.com.

Aviv, Y., A. Federgruen. 2001. Design for postponement: A comprehensive characterization of its benefits under unknown demand distributions. *Oper. Res.* **49**(4) 578–598.

Hausman, W. H. 1969. Sequential decision problems: A model to exploit existing forecasters. *Management Sci.* **16**(2) B93–B111.

Kouvelis, P., G. J. Gutierrez. 1997. The newsvendor problem in a global market: optimal centralized and decentralized control policies for a two-market stochastic inventory system. *Management Sci.* **43**(5) 571–585.

Lee, H., S. Whang. 1998. Value of postponement. T. Ho, C. Tang, eds. *Product Variety Management: Research Advances*. Kluwer Academic Publishers, Boston, MA.

Lee, H., S. Whang. 2002. The impact of the secondary market on the supply chain. *Management Sci.* **48**(6) 719–731.

Pickering, C. 2001. The price of excess. *Business 2.0.* (January) http://www.business2.com.

Rudi, N., S. Kapur, D. F. Pyke. 2001. A two-location inventory model with transshipment and local decision making. *Management Sci.* **47**(12) 668–680.

Van Mieghem, J. A. 1999. Coordinating investment, production, and subcontracting. *Management Sci.* **45**(7) 954–971.

Mass Customization vs. Mass Production: Variety and Price Competition

Aydın Alptekinoğlu
Anderson Graduate School of Management, University of California, Los Angeles, Los Angeles, California 90095, aalpteki@anderson.ucla.edu

Advisor: Charles J. Corbett
Anderson Graduate School of Management, University of California, Los Angeles, Los Angeles, California 90095

1. Introduction

A quintessential feature of modern commerce is the proliferation of product variety in virtually every industry. As a result, many companies have begun formulating their marketing and operations strategies around customization, while others still use mass production technologies that limit them to a predetermined set of products. Hence, we witness customizing firms competing against noncustomizing firms in many industries ranging from health care to cosmetics.

The business press provides many examples of customization: multivitamins (Acumins), industrial detergents (ChemStation), pants and shirts (Lands' End), lighting systems (Lutron), bicycles (National Bicycle), sneakers (Nike), beauty-care products (Procter & Gamble's Reflect.com), golf clubs (TaylorMade), messenger bags (Timbuk2), plastic food containers (Ultra Pac), and candles (Yankee Candle). This list is by no means exhaustive, but reflects the diversity of industries in which customization is gaining ground.

Many scholars and practitioners believe that, due to advances in manufacturing and information technologies, very high levels of product variety—high enough to be termed "mass customization"—can now be achieved without a prohibitive cost penalty from reduced efficiency. The main market appeal of customization is the promise of products that fit better with customers' needs or desires. Customization tends to be employed by new entrants (e.g., Dell) rather than established mass production firms (e.g., Compaq). We can thus expect mass customizers and mass producers to compete in many industries, at least in the short run, as new entrants adopt customization before established firms respond.

In this paper we study competition between two multiproduct firms with distinct production technologies, characterized by different fixed and variable costs. The *mass customizer* (MC) has a perfectly flexible production technology, and thus can offer any variety within a product space, represented by Hotelling's (1929) linear city. The *mass producer* (MP) has a focused production technology, and thus can only offer a finite, predefined set of products. Customers are heterogenous on a single taste attribute, so firms compete on price and proximity to each customer's "ideal" product. The firms first simultaneously decide whether to enter the market. Then MP designs its product line, determining the number and location of its products. Finally, both firms simultaneously set prices. We characterize the subgame perfect Nash equilibria in this three-stage game.

2. Literature Review

We draw from the parts of economics, marketing, and operations management literatures that study product variety, focusing on multiproduct firms that practice horizontal rather than vertical differentiation, and emphasizing research inspired by Hotelling's (1929) model. (This extended abstract covers only the most relevant parts of our literature review; for a more in-depth coverage, please refer to our working paper.)

Competition between multiproduct firms remains an under-studied subject in economics (Klemperer 1992, Lancaster 1990). Klemperer (1992) explores price competition under shopping costs, when two multiproduct firms offer identical or differentiated product lines with the same given number of products. Martinez-Giralt and Neven (1988) analyze a location-price game between two firms that can each open up to two outlets with no fixed or variable cost. In one equilibrium, firms open a single outlet at the opposite ends of a linear city (or diametrically opposed in a circular city); but other equilibria may exist. Heal (1980) explores spatial distribution of retail outlets in a circular city with a producer located at the center. Heal's producer is analogous to our MC, and his collection of retail outlets to our MP. However, his central producer is passive (does not set its own price), in contrast to MC in our model.

The product strategy literatures of marketing and economics are closely interrelated. Balasubramanian (1998) studies price competition between a direct marketer and multiple retailers that are evenly spaced on a circular market. His model can also represent competition between multiple specialized brands and a mass customizer. Our model differs in several important ways: His retailers are analogous to independent single-product firms, while our "retailers" are managed by a single decision maker, MP. His retailers are evenly spaced on a circular market, while we use a linear market with no ex ante restrictions on positioning. His reservation price is sufficiently high that no consumer is priced out, an assumption we do not make.

Product variety has been handled in the operations literature almost exclusively from a supply-side perspective (Ramdas 2003, Ho and Tang 1998). De Groote (1994) focuses on coordination between marketing's product-line design (number, location, and price of products) and manufacturing's choice of process flexibility (setup cost). The market is modeled using Hotelling's framework, while the production technology is modeled using a multiproduct EOQ framework. His main finding is that decentralization of product-line design and process flexibility decisions typically leads to a serious global suboptimality. Dewan et al. (2000) study product strategy-location-price competition of two firms that can offer a single standard product or a range of custom products on a circular space. The unique equilibrium of this three-stage game involves both firms choosing to customize, because the first stage poses a Prisoner's Dilemma. The assumption of zero marginal production cost is critical, as the customizing firm can then always replicate (hence dominate) a standard-product firm by setting the range of customization to zero. Our model allows more general cost parameters, giving either MC or MP a cost advantage; besides, MP *can* offer multiple products.

3. Model

In this section we present a model of competition between two multiproduct firms with fundamentally different production technologies. The *mass producers* (MP) offers a finite set of standard products by using a focused production technology. The *mass customizer* (MC) offers an infinite variety of custom products by using a perfectly flexible production technology. MP-MC competition revolves around price and ability to match each customer's single "ideal" product among a set of an infinite number of horizontally differentiated products. Each product is fully characterized by a single taste attribute, quantified by a real number between zero and one, which we refer to as the product's *location* on the *product space* $[0, 1]$.

MP makes three decisions: *entry*, *variety*, and *price*. MP first decides whether to enter the market. If it does, it pays a fixed cost F_p for a focused technology. Second, MP decides on variety, i.e., the number and location of standard products. For each member of its product line, MP pays another fixed cost f. Let $n \in \{1, 2, 3, \ldots\}$ denote the number of products and $x_i \in [0, 1]$ the location of product $i \in \{1, 2, \ldots, n\}$. The variety decision then involves setting the dimension n and the elements of the vector $\mathbf{x} \equiv (x_1, \ldots, x_n)$. Third, MP sets the prices $\mathbf{p} \equiv (p_1, \ldots, p_n)$ for products $1, \ldots, n$, respectively. The variable unit cost of production c_p is the same for all products.

MC makes two decisions: *entry* and *price*. MC also pays, if it enters, a fixed cost F_c for a perfectly flexible technology capable of custom-producing any product within the product space. Second, MC sets a uniform price p_c for its products. The variable unit cost of production c_c is again independent of location.

Each customer buys one unit (if any) from either MP or MC. Customers are heterogenous in their tastes: They each have a *location* on $[0, 1]$, representing their ideal product. The total number of customers is λ, their locations uniformly spread between zero and one. They pay a *transportation* cost (or disutility) d per unit distance between their location and the purchased product; there is no transportation cost for custom products. Each customer buys the product with the smallest *delivered price*, the list price plus the transportation cost, as long as this does not exceed a reservation price $\bar{p}$. A customer located at $z \in [0, 1]$ receives a utility of $U(z, x, p) = \bar{p} - p - d|z - x|$ by purchasing a product at x for price p. This customer will therefore purchase the custom product located at z if $p_c < p_i + d|z - x_i|$ for all $i \in \{1, 2, \ldots, n\}$, and $p_c \le \bar{p}$; the standard product $i \in \{1, 2, \ldots, n\}$ located at x_i if $i = \arg\min_{j \in \{1, 2, \ldots, n\}}[p_j + d|z - x_j|]$ and $p_i + d|z - x_i| \le \min(p_c, \bar{p})$; and will make no purchases if $p_c > \bar{p}$, and $p_i + d|z - x_i| > \bar{p}$ for all $i \in \{1, 2, \ldots, n\}$. Ties can be broken arbitrarily. We assume that an indifferent customer will prefer the standard product over the custom product, and a purchase over no purchase.

The interaction between MP and MC unfolds in three stages. In the first stage, the *entry game*, firms simultaneously decide whether to enter the market. In the second stage, MP—if it enters—determines its product variety, i.e., sets n and designs its product line $\mathbf{x}$. No decision by MC is needed in the second stage. Finally, in the third stage, firms in the market simultaneously set their prices, p_c and $\mathbf{p}$. This sequence is consistent with the hierarchy of decisions in marketing and economics models of product strategy. Note that we do not let the firms choose their production technology upon entry. This is because endogenous technology choice would require a treatment of competition between two MPs, which appears to have no equilibrium (Klemperer 1992). Moreover, we believe the MP-MC competition is the most interesting focus of analysis.

MC's profit is given by $\pi_c(p_c, n, \mathbf{x}, \mathbf{p}) = (p_c - c_c) \cdot \lambda y_c(p_c, n, \mathbf{x}, \mathbf{p}) - F_c$, where $y_c(\cdot)$ is the total market share captured by the custom products. MP's profit is given by $\pi_p(p_c, n, \mathbf{x}, \mathbf{p}) = \sum_{i=1}^{n}[(p_i - c_p) \cdot \lambda y_i(p_c, n, \mathbf{x}, \mathbf{p}) - f] - F_p$, where $y_i(\cdot)$ is the length of the market segment captured by the standard product $i \in \{1, 2, \ldots, n\}$. y_c and $\sum_{i=1}^{n} y_i$ are MC's and MP's total market share.

The entry game has four possible outcomes. If both firms decide to enter, then a *duopoly competition* ensues. If only one firm enters, then either MP or MC is a *monopoly*. Finally, if neither firm enters, *market breakdown* occurs. Throughout, we assume that $c_c < \bar{p}$ and $c_p < \bar{p}$; otherwise, one or both firms would never enter.

We develop a full analysis of this three-stage game to address the following research questions: How would the firms set prices as monopolists or duopolists? How would MP design its product line when competing with MC, as opposed to when acting as a monopolist? Does duopoly lead to higher or lower variety of standard products than monopoly? Under what conditions can MP profitably coexist with MC? What conditions result in duopoly, or in a monopoly of either firm?

4. Summary of Analysis

In this section we provide a summary of our analysis, with all but one of the propositions omitted (please refer to our working paper for a full treatment). Our objective is to characterize the subgame perfect Nash equilibria of the three-stage game defined in §3. We first consider the two monopoly cases, where only one of the firms enters the market. In the case of MP monopoly, we show that MP's optimal pricing and positioning decisions dictate that standard products be priced uniformly and each capture an equal-length market segment. This structural result considerably reduces the solution space, as MP's second- and third-stage decisions reduce to setting the number of products n, and a uniform price p_p. We thus derive the optimum price p_p^* for any given n, and the optimum variety n^* (by relaxing the integrality of n). The case of MC monopoly is trivial: MC charges the reservation price, $p_c^* = \bar{p}$, yet still captures the entire market.

Secondly, we analyze the duopoly case. Analogous to the MP monopoly case, a structural result that reduces MP's second-stage decision to choosing n, and its third-stage decision to setting a uniform price p_p holds. We first derive the Nash equilibrium in prices $(\hat{p}_c, \hat{p}_p)$ for a given n. We then solve for MP's optimal variety, $\hat{n}$, under the presumption that the price equlibrium will follow in the third stage.

Having obtained the profits under the monopoly $(\pi_c^{monop}, \pi_p^{monop})$ and duopoly $(\pi_c^{duop}, \pi_p^{duop})$ cases, we conclude the analysis by identifying all the equilibria of the entry game, which can be stated in normal form as follow:

MC\MP	enter (e)	do not enter (w)
enter (e)	π_c^{duop}, π_p^{duop}	π_c^{monop}, 0
do not enter (w)	0, π_p^{monop}	0, 0

The most interesting outcome of the entry game, duopoly competition, can be characterized as follows: Entry by both firms (e, e), is the unique equilibrium of the entry game *iff* the exogenous parameters of the model satisfy one of the two cases below.

Case 1. $\bar{p} > d/3 + (2c_c + c_p)/3$, $[(0 < c_c - c_p \le d$ and $f \ge \bar{f})$ or $(-d/2 \le c_c - c_p \le 0)]$, $F_c \le (2\lambda/9d)[d - (c_c - c_p)]^2$ and $F_p \le (2\lambda/9d)(c_c - c_p + d/2)^2 - f$, where $\bar{f}$ uniquely satisfies $((2\lambda(c_c - c_p))/9d)(c_c - c_p - d/4) \le \bar{f} < ((\lambda(c_c - c_p)^2)/2d)$ and $\sqrt{2\lambda d\bar{f}} - \bar{f} = (2\lambda/9d)[-(c_c - c_p)^2 + (7d/2)(c_c - c_p) - d^2/4]$.

Case 2. $\bar{p} \le d/3 + (2c_c + c_p)/3$, $[(0 < c_c - c_p < d$ and $\check{f} \le f < ((\lambda(\bar{p} - c_p)^2)/2d)$ or $-d/2 < c_c - c_p \le 0$ and $f < ((\lambda(\bar{p} - c_p)^2)/2d)]$, $F_c \le (2\lambda(\bar{p} - c_c)^2)/(3\bar{p} - 2c_c - c_p)$ and $F_p \le (d/(3\bar{p} - 2c_c - c_p))[((\lambda(\bar{p} - c_p)^2)/2d) - f]$, where $\check{f}$ uniquely satisfies $((\lambda(c_c - c_p))/6d)(2c_c - c_p - \bar{p}) \le \check{f} < ((\lambda(c_c - c_p)^2)/2d)$ and $(3\bar{p} - 2c_c - c_p)\sqrt{2\lambda d\check{f}} - d\check{f} = 2\lambda(\bar{p} - c_c)(c_c - c_p) + (\lambda/2)(\bar{p} - c_p)(2c_c - c_p - \bar{p})$.

These two cases yield the following decisions by the firms: Case 1: $\hat{n} = 1$, $(\hat{p}_c, \hat{p}_p) = (d/3 + ((2c_c + c_p)/3), d/6 + ((c_c + 2c_p)/3))$; Case 2: $\hat{n} = d/(3\bar{p} - 2c_c - c_p)$, $(\hat{p}_c, \hat{p}_p) = (\bar{p}, ((\bar{p} + c_p)/2))$. Note that MP offers multiple products only when the reservation price is low enough (Case 2).

5. Discussion

5.1. Coexistence or Dominance?

A key question is whether a finite-variety MP can survive competition from an infinite-variety MC. We offer three observations. First, contrary to basic intuition, MP *can* profitably compete even if MC has a variable production cost advantage, as long as that advantage is not too large; i.e., $c_p - d/2 < c_c < c_p$. Second, MP *can* profitably compete even if MC also has a lower fixed cost; i.e., $F_p > F_c$. Third, for duopoly to occur, the difference between c_c and c_p must be small relative to d: $-d/2 < c_c - c_p < d$. These bounds are skewed in favor of MC: MP drops out with a unit cost disadvantage of $d/2$, whereas MC can withstand a unit cost disadvantage of up to d.

It is surprising that MP can profitably coexist with MC despite fixed and variable cost disadvantages. The underlying reason is that even if MC has a lower variable cost, it may choose not to lower its uniform price enough to deter entry by MP. Revenue losses in the already captured portions of the market might outweigh the gains from the market segments that would switch from MP to MC upon a price cut.

5.2. Monopoly vs. Duopoly

Comparing an MP monopoly (high F_c) with a duopoly (low F_c), we observe the following. First, as expected, prices are lower under duopoly than under monopoly; i.e., $\hat{p}_c \le p_c^*$ and $\hat{p}_p \le p_p^*$. Second, in duopoly, MC always charges a higher price and captures a larger market share than MP; i.e., $\hat{p}_c \ge \hat{p}_p$ and $y_c \ge \hat{n} y_p$.

Finally, and most importantly, MP offers *lower* product variety under duopoly competition than under monopoly; i.e., $\hat{n} \le n^*$. Naively, one might expect MP to increase its product variety to better compete against MC. However, the results indicate that higher variety can lead to increased price competition, hence lower prices. Therefore, MP prefers to cut product variety in order to soften the price competition at the third stage. This result extends a recurring theme from the economics and marketing literatures to the OM literature: Incentive to soften price competition often trumps other considerations (e.g., Martinez-Giralt and Neven 1988, Klemperer 1992, Balasubramanian 1998).

5.3. When Does Higher Variety Imply Higher Prices?

The prevailing view in the OM community seems to be that higher product variety always allows charging higher prices, because perceived value to customers supposedly increases as product variety grows. Indeed, in a widely cited study based on the PIMS (Profit Impact of Marketing Strategies) database, Kekre and Srinivasan (1990, p. 1217) hypothesized that "a broader product line may lead to higher relative prices," but found only marginal support in both industrial and consumer markets. Our results may help explain this lack of strong support.

For MP monopoly, higher variety does imply (weakly) higher prices. However, under duopoly, a higher variety offered by MP may result in a lower, equal, or higher equilibrium price. One reason why Kekre and Srinivasan (1990) found a much weaker link than expected between variety and price may be that their sample included a mix of all the above cases. Our analysis suggests that empirical studies of the link between variety and price must consider the prevailing competitive structure (monopoly versus duopoly) and the relative costs of different production technologies.

6. Concluding Remarks

We model competition between two multiproduct firms on the basis of variety and price, where one firm is a *mass customerizer* (MC) with infinite variety, while the other is a *mass producer* (MP) with finite variety. We construct a three-stage game using Hotelling's (1929) model of horizontal differentiation. First, both firms decide whether to enter the market; next, MP designs its product line, and then both firms set prices. We fully characterize the subgame perfect Nash equilibria of this game. We find that MP will *decrease* variety upon entry by MC. This reflects the firms' desire to soften the price competition under duopoly. Also, the finite-variety MP *can* profitably compete even if it has a disadvantage against the infinite-variety MC in fixed cost of production technology, variable cost of production, or both.

We contribute to the product variety literature in several ways. First, we introduce an important supply-side element (production technology) to the study of product differentiation, conducted so far mostly from a demand perspective. Also, MC in our model has perfect product-mix flexibility, so our analysis is informative on the strategic value of a production technology that can provide such flexibility. Second, within Hotelling's (1929) model of product differentiation, we introduce a mass customizing firm that can produce and market any variety with no penalty, in line with recent advances in manufacturing and information technologies. Third, we partly contradict a prevailing view in OM, by finding that MP in a duopoly needs to *reduce* product variety (compared to an MP monopoly) in order to soften price competition with MC.

References

Balasubramanian, S. 1998. Mail versus mall: A strategic analysis of competition between direct marketers and conventional retailers. *Marketing Sci.* **17**(3) 181–195.

De Groote, X. 1994. Flexibility and marketing/manufacturing coordination. *Internat. J. Production Econom.* **36**(2) 153–169.

Dewan, R., B. Jing, A. Seidmann. 2000. Adoption of Internet-based product customization and pricing strategies. *J. Management Inform. Systems* **17**(2) 9–30.

Heal, G. 1980. Spatial structure in the retail trade: A study in product differentiation with increasing returns. *Bell J. Econom.* **11**(2) 565–583.

Ho, T.-H., C. S. Tang. 1998. *Product Variety Management.: Research Advances*. Kluwer Academic Publishers, Boston, MA.

Hotelling, H. 1929. Stability in competition. *Econom. J.* **39**(153) 41–57.

Kekre, S., K. Srinivasan. 1990. Broader product line: A necessity to achieve success? *Management Sci.* **36**(10) 1216–1231.

Klemperer, P. 1992. Equilibrium product lines: Competing head-to-head may be less competitive. *Amer. Econom. Rev.* **82**(4) 740–755.

Lancaster, K. 1990. The economics of product variety: A survey. *Marketing Sci.* **9**(3) 189–206.

Martinez-Giralt, X., D. J. Neven. 1988. Can price competition dominate market segmentation? *J. Indust. Econom.* **36**(4) 431–442.

Ramdas, K. 2003. Managing product variety: An integrative review and research directions. *Production Oper. Management* **12**(1) 79–101.

Dynamic Bayesian Quantity Control when Some Customers Will Accept a Substitute

Li Chen
Department of Management Science and Engineering, Stanford University, Stanford, California 94306, skychen@stanford.edu

Advisor: Erica L. Plambeck
Graduate School of Business, Stanford University, Stanford, California 94306

1. Introduction

If a desired product is not available, a customer may accept a substitute, or choose not to make any purchase. The substitution probability is crucially important in revenue management and inventory management, but is not known with certainty. In two simple but qualitatively insightful models, this paper explores the managerial implications of simultaneous quantity control and estimation of the unknown substitution probability. In particular, we model the customers' decision to substitute or not as i.i.d. Bernoulli trials, and we update beliefs about the substitution probability using a beta-binomial conjugate prior. The beta-binomial is a very natural and surprisingly tractable formulation.

The first problem under consideration is multi-period inventory management of a perishable good with unknown demand distribution. Following the partial lost-sales model of Nahmias and Smith (1994), we assume that when a stockout occurs during a period, each excess demand generates a Bernoulli trial. "Success" indicates that the customer will wait for the product to be delivered at the end of the ordering period, and "failure" corresponds to a lost sale. We assume the demand distribution has an unknown parameter (or vector of parameters) and the Bernoulli probability of customer's willingness to wait is unknown as well. Both the demand parameter and Bernoulli wait probability are subject to some prior beliefs, with the latter one following a beta distribution. Moreover, we assume that the total demand and the amount of lost sales are observed in each period. This assumption holds when, for instance, orders arrive through a call center or over the Internet and detailed order data are recorded. Under the Bayesian sequential decision framework, we show it is optimal to stock less than a "myopic" optimal inventory level, so that the inventory manager can learn more about customers' willingness to wait as stockout is more likely to occur. This "stock less" result when the customer's decision to wait versus walk is observed contrasts nicely with the "stock more" result by Lariviere and Porteus (1999) and Ding et al. (2002) in the setting with unobserved lost sales.

The second model stems from an airline revenue management problem. Starting with the two-fare model of Lee and Hersh (1993), we introduce the probability that a discount-fare customer, denied a booking request, will "buy up" to the full fare. Assuming the buy-up probability is unknown, but subject to a beta prior belief, we again characterize the optimal booking limit policy using the Bayesian sequential decision framework. We prove that it is optimal to set the discount-fare booking limit lower than a "myopic" optimal level (no Bayesian updating). Essentially, by reducing the quantity of discount-fare seats available, the decision maker learns more about the buy-up probability and can better control the booking limit in subsequent periods.

2. Bayesian Inventory Management with Partial Lost Sales

In this section, we consider inventory management for a perishable good with unknown demand distribution. There is a total of N periods. We number the periods in a reverse order. At the beginning of each period, the inventory manager selects an inventory level for a perishable good. The inventory level is achieved immediately after the decision. (We are assuming a negligible delivery leadtime from the supplier.) For each unit of the good, the production cost is c and selling price is r. At the end of each period, any leftover stock is salvaged for a value of s per unit. If the demand during a period exceeds the inventory level, every additional customer is asked to wait for the product. Each customer is willing to wait with probability p and becomes a lost sale with probability $1-p$. If a customer agrees to wait, then the product is expedited at a cost of e and delivered to him at the end of the period. Hence, no back orders are carried over to the next period. The objective of the inventory manager is to maximize discounted expected profits; δ denotes the discount factor. For nontriviality, we assume $s<c<e<r$.

The demand in each period is independently and identically generated by a general nonnegative discrete demand distribution. The probability mass function of the demand distribution is denoted as $f(\xi \mid \theta)$, $\xi = 0, 1, 2, \ldots,$ where θ is an unknown parameter (or vector of parameters) with $\theta \in \Theta$. For period i, given a prior distribution density $\pi_i(\theta)$ and a demand observation ξ_i, the posterior distribution density $\pi_{i-1}(\theta \mid \xi_i)$ can be obtained by Bayes' rule (note that the periods are numbered in reverse order). Correspondingly, the sequence of predictive demand densities $\{f_i(\xi), i = N, \ldots, 1\}$ satisfies $f_i(\xi) = \int_\Theta f(\xi \mid \theta)\pi_i(\theta \mid \xi_{i+1})d\theta$. The predictive demand density $f_i(\xi)$ is the Bayesian estimate of the true demand distribution density $f(\xi \mid \theta)$ in period i.

Suppose the prior belief of p at the beginning of period i follows a beta distribution with parameters (α_i, β_i). Denote $v_i(\alpha_i, \beta_i, \pi_i(\theta))$ as the value function when there are i periods to go. Let y_i be the inventory-level decision. The profit-maximization problem can be recursively formulated into the following dynamic program. For $i = N, \ldots, 1$,

$$\begin{aligned}
&v_i(\alpha_i, \beta_i, \pi_i)\\
&\quad = \max_{y_i \ge 0}\Bigg\{(s-c)y_i + (r-s)\bar{\xi}_i - \left[r-s-\frac{\alpha_i}{\alpha_i+\beta_i}(r-e)\right]\\
&\qquad \cdot \sum_{\xi_i=y_i+1}^{\infty}(\xi_i-y_i)f_i(\xi_i) + \delta\sum_{\xi_i=0}^{y_i} v_{i-1}(\alpha_{i-1}, \beta_{i-1}, \pi_{i-1})\\
&\qquad \cdot f_i(\xi_i) + \delta\sum_{\xi_i=y_i+1}^{\infty} E_{(\alpha_i, \beta_i, \xi_i-y_i)}\\
&\qquad \cdot [v_{i-1}(\alpha_{i-1}, \beta_{i-1}, \pi_{i-1})]f_i(\xi_i)\Bigg\},\\
&v_0(\alpha_0, \beta_0, \pi_0) = 0,
\end{aligned} \tag{1}$$

where $\bar{\xi}_i$ is the demand mean under density $f_i(\cdot)$, and $E_{(\alpha, \beta, n)}[\cdot]$ denotes the expectation taken under a beta-binomial distribution with parameters (α, β, n). The updating schemes of prior-posterior parameters are given below:

$$\begin{cases} \alpha_{i-1} = \alpha_i \\ \beta_{i-1} = \beta_i \end{cases} \quad \text{if demand } \xi_i \le y_i,$$

$$\begin{cases} \alpha_{i-1} = \alpha_i + x_i & \text{if demand } \xi_i > y_i, \\ \beta_{i-1} = \beta_i + \xi_i - y_i - x_i & \text{and backorder is } x_i, \end{cases}$$

$$\pi_{i-1} = \frac{f(\xi_i \mid \theta)\pi_i(\theta)}{\int_\Theta f(\xi_i \mid \theta')\pi_i(\theta')d\theta'}.$$

We have the following proposition:

Proposition 1. *Denote* $V^{(\pi_i, \xi_i)}_{(\alpha_i, \beta_i, n)}$ *as*

$$E_{(\alpha_i, \beta_i, n)}\big[v_{i-1}(\alpha_i + X_i,\, \beta_i + n - X_i,\, \pi_{i-1})\big].$$

The necessary conditions for an optimal inventory level y_i^* *for Dynamic Program* (1) *are given by*

$$\text{(1)}\quad s-c+\sum_{\xi_i=y_i^*}^{\infty}\Bigg[r-s-\frac{\alpha_i}{\alpha_i+\beta_i}(r-e) + \delta\Big(V^{(\pi_i, \xi_i)}_{(\alpha_i, \beta_i, \xi_i-y_i^*)} - V^{(\pi_i, \xi_i)}_{(\alpha_i, \beta_i, \xi_i-y_i^*+1)}\Big)\Bigg]f_i(\xi_i) \ge 0,$$

and

$$\text{(2)}\quad s-c+\sum_{\xi_i=y_i^*+1}^{\infty}\Bigg[r-s-\frac{\alpha_i}{\alpha_i+\beta_i}(r-e) + \delta\Big(V^{(\pi_i, \xi_i)}_{(\alpha_i, \beta_i, \xi_i-y_i^*-1)} - V^{(\pi_i, \xi_i)}_{(\alpha_i, \beta_i, \xi_i-y_i^*)}\Big)\Bigg]f_i(\xi_i) < 0.$$

By exploiting the structure properties of the beta-binomial distribution, we can show for any $n \geq 0$, $V^{(\pi_i,\xi_i)}_{(\alpha_i,\beta_i,n)} \leq V^{(\pi_i,\xi_i)}_{(\alpha_i,\beta_i,n+1)}$. Therefore, by comparing the results of this full updating case with those of the "myopic" case (where only the prior for demand distribution is updated), we have the following theorem:

Theorem 1. *In an N-period inventory management problem with unknown demand distribution and unknown probability of customer's willingness to wait, the payoff of a full Bayesian updating policy is always greater than that of the "myopic" updating policy. Furthermore, the optimal inventory level under full updating is always less than the "myopic" optimal inventory level.*

Theorem 1 confirms the intuition that in order to gain information about the probability that a customer will wait, the manager should initially reduce the inventory level to increase the number of data points (customers forced to wait or walk). It is striking that this result holds for any nonnegative discrete demand distribution with unknown (or known) parameters.

3. Bayesian Revenue Management with Unknown Buy-Up Probability

Now consider another problem that concerns managing bookings for a single-leg flight. For simplicity, we assume that no customers cancel reservations, and overbooking is not allowed. We adopt similar notations of Lautenbacher and Stidham (1999).

Customers may purchase a full-fare ticket for r_1 or a discount-fare ticket with booking restrictions for r_2, where $r_1 > r_2$. Multiple-seat bookings are not considered. The seating capacity is C. The booking horizon, $[0, T]$, is divided into N "decision periods" in such a way that, during each period, the probability of two or more requests is negligible. Decision periods are numbered in reverse order, with Period 0 denoting the scheduled departure time. For $i = 1, \ldots, N$, we define $P^i_1 \doteq Pr\{$one full-fare class arrival during period $i\}$, $P^i_2 \doteq Pr\{$one discount-fare class arrival during period $i\}$, and $P^i_0 \doteq Pr\{$no arrival during period $i\}$, with $P^i_0 + P^i_1 + P^i_2 = 1$.

Observe that a request for the full-fare seat will always be accepted as long as seats remain available. The problem is to decide whether to accept or deny customer requests for a discount-fare ticket. When a customer is denied a discount-fare seat request, he may choose to purchase a full-fare ticket. This type of substitution is commonly called a "buy-up." We assume that a customer will choose to buy up with probability p, or to leave the system with probability $1-p$.

Suppose the prior belief of p at the beginning of period i is beta distributed with parameters (α, β). Denoting $U_i(x, \alpha, \beta)$ to be the optimal i-period-to-go revenue function when x seats are booked under a Bayesian updating policy, we have the following recursive dynamic program:

$$\begin{aligned} U_i(x,\alpha,\beta) &= P^i_0 U_{i-1}(x,\alpha,\beta) + P^i_1(r_1 + U_{i-1}(x+1,\alpha,\beta)) \\ &\quad + P^i_2 \cdot \max\Big\{ r_2 + U_{i-1}(x+1,\alpha,\beta),\ \frac{\alpha}{\alpha+\beta} \\ &\qquad \cdot r_1 + \frac{\alpha}{\alpha+\beta}\cdot U_{i-1}(x+1,\alpha+1,\beta) \\ &\qquad + \frac{\beta}{\alpha+\beta}\cdot U_{i-1}(x,\alpha,\beta+1)\Big\}, \end{aligned} \tag{2}$$

$$U_0(x,\cdot,\cdot) = \begin{cases} 0 & \text{if } x \leq C \\ -r_1(x-C) & \text{if } x > C. \end{cases}$$

In this formulation, denying the customer's request for a discount fare generates data on the buy-up probability. Hence, the accept/deny decision depends upon the value of this information, as well as the marginal expected revenue from booking a seat.

Proposition 2. *The optimal solution for the dynamic model under full Bayesian updating is to accept the discount-fare request in period i if and only if* $0 \leq x < b^{(\alpha,\beta)}_i$, *where*

$$\begin{aligned} b^{(\alpha,\beta)}_i \doteq \min\Big\{ x \geq 0 : &\frac{\alpha}{\alpha+\beta}\cdot U_{i-1}(x+1,\alpha+1,\beta) \\ &+ \frac{\beta}{\alpha+\beta}\cdot U_{i-1}(x,\alpha,\beta+1) \\ &- U_{i-1}(x+1,\alpha,\beta) > r_2 - \frac{\alpha}{\alpha+\beta}\cdot r_1 \Big\}. \end{aligned}$$

Comparing the optimal booking limit of this full updating case with the "myopic" case (where the beta prior is not updated) under the conditions of $r_2 - (\alpha/(\alpha+\beta))r_1 \leq 0$ or $P^{i-1}_0 = 0$, we have the following theorem:

THEOREM 2. *In an N-period, two-fare, dynamic booking control problem with unknown buy-up probability. The optimal booking limit under full Bayesian updating is always less than the "myopic" optimal booking limit under no updating.*

Theorem 2 supports the intuitive idea that a decision maker should lower the booking limit to gain information about customers' buy-up potential and use this to make better decisions in future periods. Compared with the multiperiod inventory model of §2, our dynamic booking control model is capacitated. As a result, the value function has properties very different from those of §2. For example, it is not true that the expected revenue under full updating is always greater than the "myopic" case. When few seats remain, the booking control decisions with full updating are identical to those for the case without updating. However, in the case with full updating, when a discount customer fails to buy up, the ongoing expected revenue is reduced relative to the case without updating. We have learned that the plane is less likely to be filled.

4. Further Discussion

The Dirichlet-Multinomial conjugate prior is a generalized version of the two-dimensional beta-binomial conjugate prior. It is useful in studying Bayesian control problems that involve estimating a discrete probability mass function. For example, in the dynamic airline reservation model of §3, we can use it to capture learning about all three of the customer arrival probabilities, P_0^i, P_1^i, and P_2^i. We conjecture that the Dirichlet-Multinomial conjugate prior will prove valuable in more general revenue management and inventory management problems with multiple products or dynamic pricing.

References

DeGroot, N. H. 1970. *Optimal Statistical Decisions.* McGraw-Hill Book Company, New York.

Ding, X., N. L. Puterman, A. Bisi. 2002. The censored newsvendor and the optimal acquisition of information. *Oper. Res.* **50**(3) 517–527.

Johnson, N. L., S. Kotz, A. W. Kemp. 1992. *Univariate Discrete Distributions*, 2nd ed. John Wiley and Sons, New York.

Lariviere, N. A., E. L. Porteus. 1999. Stalking information: Bayesian inventory management with unobserved lost sales. *Management Sci.* **45**(3) 346–363.

Lautenbacher, C. J., S. Stidham Jr. 1999. The underlying Markov decision process in the single-leg airline yield-management problem. *Transportation Sci.* **33**(2) 136–146.

Lee, T. C., M. Hersh. 1993. A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Sci.* **27**(3) 252–265.

Nahmias, S., S. A. Smith. 1994. Optimizing inventory levels in a two-echelon retailer system with partial lost sales. *Management Sci.* **40**(5) 582–596.

Demand Estimation and Assortment Optimization Under Substitution: Methodology and Application

A. Gürhan Kök
Fuqua School of Business, Duke University, Durham, North Carolina 27708, gurhan.kok@duke.edu

Advisor: Marshall L. Fisher
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104, fisher@wharton.upenn.edu

Assortment planning at a retailer involves choosing the products to be carried and setting the inventory levels for each product. With a broader assortment (i.e., more variants in a category), the retailer carries the favorite product of more customers, but due to limited shelf space, average inventory per product is lower, resulting in lower in-stock rate and higher lost sales. The degree of importance of the trade-off between breadth of assortment and in-stock rates of products depends on the customers' willingness to accept another product as a substitute if their favorite product is not available.

Although there has been significant research on inventory and assortment decisions under substitution, there is little research on the estimation of input parameters. We introduce a relatively simple model of demand substitution and present a novel methodology for demand estimation under substitution. We estimate a substitution rate common to all stores in a chain. Also, for each store, we obtain demands of products carried in the store as well as products previously not carried in that store.

The assortment optimization problem is a nonlinear, nonseparable, discrete resource allocation problem. We describe an iterative heuristic that finds solutions within 0.5% of the optimal on average. Based on our solution procedure, we establish new structural properties that relate the product selection and their inventory levels to product characteristics such as demand rate and variability, batch size, etc.

We applied our methodology at Albert Heijn, BV, a large Dutch supermarket chain. Comparing results with the existing assortments suggests more than 50% increase in profits.

Problem Structure

A subcategory is defined as a group of substitutable variants such that the difference between products within a subcategory is minimal, but the difference between subcategories is significant. Let $N_i = \{1, \ldots, J_i\}$ denote the set of potential variants in subcategory i. A product is stocked on the shelves in the allocated slots called *facings*. Let f_{ij} denote the number of facings allocated to a product j in subcategory i. Each facing has a fixed capacity c_{ij} (implying a maximum inventory level of $c_{ij} f_{ij}$) and a fixed metric width w_{ij}.

The original demand rate to each product is denoted d_{ij}. A customer who cannot find her favorite product available due to assortment decisions or stockouts may choose to accept one of the available products as a substitute (similar to Smith and Agrawal 2000). Hence, the effective demand rate of a product is

$$D_{ij}(\mathbf{f}_i, \mathbf{d}_i) = d_{ij} + \sum_{k|f_{ik}=0} \alpha_{ikj} d_{ik} + \sum_{k|f_{ik}>0} \alpha_{ikj}\, LostSales_{ik}(d_{ik}, f_{ik}), \qquad (1)$$

where $\mathbf{f}_i = (f_{i1}, f_{i2}, \ldots, f_{iJ_i})$, $\mathbf{d}_i = (d_{i1}, d_{i2}, \ldots, d_{iJ_i})$, and α_{ikj} is the probability that a customer of product k in subcategory i is willing to substitute product j (also in subcategory i) when k is not available. We define $\alpha_{ikj} = \delta_i d_{ij} / \sum_{l \in N_i \setminus \{k\}} d_{il}$. The parameter δ single-handedly determines the intensity of substitution in the subcategory of interest, hence we are able to differentiate between subcategories with low and high substitution rates. This model of substitution has features similar to those of utility-based choice models, and it can be used consistently both in estimation of substitution behavior and optimization of the assortment. The first sum in Equation (1) is the redistributed demand from *assortment-based* substitution, and the second sum is from *stockout-based* substitution. The *LostSales* function is the average unmet demand of a product depending on the specifics of the replenishment process.

The decision process involves allocating a discrete number of facings to each product in order to maximize total expected gross profits from one or more subcategories subject to a shelf space constraint.

$$\max_{f_{ij}, \forall i,j} \left\{ \sum_i \sum_j G_{ij}(f_{ij}, D_{ij}(\mathbf{f}_i, \mathbf{d}_i)), \text{ s.t.}, \sum_i \sum_j f_{ij} w_{ij} \le ShelfSpace,\ f_{ij} \in \{0,1,2,\ldots\} \right\} \qquad \text{(AP)}$$

We assume that we can compute G_{ij} given f_{ij} and D_{ij}. The unknown input to this optimization problem are the demand rates d_{ij} and the function δ_i. When there is no substitution, $D_{ij} = d_{ij}$, and d_{ij} can be estimated easily. When there is substitution, however, we do not observe d_{ij}, but observe the value of the D_{ij} function given the current assortment $\mathbf{f}_i$.

Estimation of Assortment-Based Substitution

Consider stores h and B, with the same demand rates for all products in a subcategory (subscript i is dropped for brevity). Suppose that all stores have a 100% service rate for all the products they carry. Store B carries every possible variant, assortment N; hence, no substitution takes place at store B (i.e., $D_j = d_j$). Store h carries assortment $S_h \subset N$ with subcategory sales $T_h(S_h)$. Let $T_B(S_h)$ denote the total sales of products in S_h at store B. If $T_h(S_h) > T_B(S_h)$, then it means that at store h, some of the customers whose favorite products are in $N \backslash S_h$ purchased products from S_h as substitutes. We estimate $T_h(S_h)$ using data from store h. We estimate $\mathbf{d}$ and $T_B(S_h)$ using data from all stores with assortment N. For given $\mathbf{d}$ and δ, we can theoretically compute how much products in S would sell at store h, denoted $\widehat{T}_h(S_h, \delta)$. Combining these estimates (i.e., observed and theoretical subcategory sales at a store) for multiple stores, we find the substitution rate δ that minimizes $\sum_h (\widehat{T}_h(S_h, \delta) - T_h(S_h))$. We then reverse-engineer the true demand of each product (d_j) at each store using the estimated substitution rate and the observed demands (D_js) at that store. As a result, for each store we obtain demands of products carried in the store as well as products previously not carried in that store. We test the substitution estimation scheme on different data sets and find it to be a reliable methodology. Anupindi et al. (1998) describes a method for estimating stockout-based substitution that requires inventory transactions data. We use the estimated assortment-based substitution rate for both types of substitution.

Assortment Optimization Under Substitution

AP is a knapsack problem with a nonlinear and nonseparable objective function. We propose the *Iterative Heuristic* (IH) to solve this mathematical program. We set $D_{ij} = d_{ij}$ for all i and j and solve *AP* with the original demand rates resulting in a particular facings allocation $\mathbf{f}^0$. At iteration t, we replace the G_{ij} in the objective function with $G_{ij}(f_{ij}^t, D_{ij}(\mathbf{f}_i^{t-1}, \mathbf{d}_i))$. This makes *AP* additively separable in f_{ij}^t, which is solved using an enhanced Greedy Heuristic. We keep iterating until $\mathbf{f}_i^t$ converges for all i. The algorithm concludes with a local search for improvement. Iterations in IH is a way to seek consistency between the input demand vectors and the optimal solution. Based on our solution procedure we are able to prove the following structural results.

Proposition 1. *Consider two products A and B in the same subcategory under a periodic base stock replenishment policy. Let m, ρ, and b denote per-unit gross margin, coefficient of variation of demand, and batch size of a product, respectively.*

(i) *If, ceteris paribus,* $(d_A \geq d_B)$ *or* $(m_A \geq m_B)$ *or* $(w_A \leq w_B)$, *then* $f_A \geq f_B$ *in the final solution of the IH.*

(ii) *If, ceteris paribus,* $(\rho_A \geq \rho_B)$ *or* (b_A *is an integer multiple of* b_B), *then in the final solution of the Iterative Heuristic product B is included in the assortment only if product A is included* (*i.e.,* $f_B > 0 \Rightarrow f_A > 0$).

The result regarding the demand rates is similar to the property of optimal assortments in van Ryzin and Mahajan (1999). These results imply that products with higher demand rate should be assigned the initial space, but, given sufficient space, it may be more profitable to assign more inventory to products with more constraining product characteristics.

In a numerical study, we observe that the IH solutions are within 0.5% of the optimal solution on average, with a maximum of 12.1%. We find that the value of explicitly modeling substitution is 2.5% of the total profit on average.

Application

Albert Heijn, BV is the leading supermarket chain in the Netherlands, with more than 600 stores and \$5 billion in sales.[1] The estimation and optimization system is tested on data for seven merchandise categories (114 subcategories with 880 SKUs) for a 20-week period from 37 stores. Now the system is being implemented at three pilot stores in the Netherlands. We compare the gross profit figures for the existing and suggested assortments for two categories and 10 stores. The improvement in the gross profits optimized over the peakload periods is 12.3%. Using the financial ratios from the financial statement of the company, we project that the impact of our methodology is a 50% increase in profits. Based on Albert Heijn's \$5 billion yearly sales, the bottom line impact is an additional \$75 million.

References

Anupindi R., M. Dada, S. Gupta. 1998. Estimation of consumer demand with stockout based substitution: An application to vending machine products. *Marketing Sci.* **17** 406–423.

Smith, S. A., N. Agrawal. 2000. Management of multi-item retail inventory systems with demand substitution. *Oper. Res.* **48** 50–64.

van Ryzin, G., S. Mahajan. 1999. On the relationship between inventory costs and variety benefits in retail assortments. *Management Sci.* **45** 1496–1509.

[1] Albert Heijn, BV is a subsidiary of Ahold Corporation, which owns many supermarket chains around the world, with about 8,500 stores and \$50 billion in sales.

Manufacturing & Service Operations Management 6(1), pp. 92–112, ©2004 INFORMS

A Portfolio Approach to Procurement Contracts

Victor Martínez-de-Albéniz
Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, albeniz@mit.edu

Advisor: David Simchi-Levi
Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

1. Introduction

A recent trend for many industrial manufacturers has been outsourcing; firms are considering outsourcing everything from production and manufacturing to the procurement function itself. Of course, the increase in the level of outsourcing implies that the procurement function becomes critical for a manufacturer to remain in control of its destiny. As a result, many OEMs focus on closely collaborating with the suppliers of their *strategic components*. In some cases, this is done using *private exchange* marketplaces and/or effective supply contracts; both of these try to coordinate the supply chain.

A different approach has been applied for *nonstrategic components*. In this case, products can be purchased from a variety of suppliers, and flexibility to market conditions is perceived as more important than a permanent relationship with the suppliers. Indeed, *commodity products*, e.g., electricity, steel, grain, cotton, or computer memory, are typically available from a large number of suppliers and can be purchased in spot markets. Because these are highly standard products, switching from one supplier to another is not considered a major problem.

Despite the nonstrategic nature of commodity products, it is critical to identify effective procurement strategies for these components, since manufacturers may be completely dependent on them. For instance, production costs might be very sensitive to the cost of some commodity products, e.g., electricity for automobile manufacturers, or memory for computer manufacturers. At the same time, uncertainty in supply and customer demand raises the question of whether to purchase supply now or wait for better market conditions in the future.

Thus, an effective procurement strategy for commodity products has to focus on both driving costs down and reducing risks. These risks include both *inventory* and *price* risks. By inventory risks, we refer to inventory shortages or unsold products; while price risks refers to the purchasing price, which is uncertain if the procurement strategy depends on spot markets.

The objective of this work is to develop a framework that provides *manufacturers* with the ability to select multiple contracts at the same time in order to optimize their expected profit. This approach is particularly meaningful for commodity products, since a large pool of suppliers is available, each offering a different type of contract. Thus, the manufacturer may be interested in selecting several different complementary contracts so as to reduce expected procurement and inventory holding costs.

For this purpose, we define a new type of contract, called a *portfolio* contract, that is in fact a combination of many traditional contracts, such as long-term contracts, options, and flexibility contracts. Of course, in the design of such a portfolio contract one may take into account the spot market.

Interestingly, current trends in industry are also directed toward the analysis and optimization of supply contracts. Most relevant to this research is the work done at Hewlett-Packard (HP) on the use of a portfolio approach for the procurement of electricity or memory products; see Billington (2002). Specifically, 50% of HP's procurement cost is invested in long-term contracts, 35% in option contracts, and the remaining is left to the spot market, see Carbone (2001).

We analyze the general case of a portfolio of contracts taking into account the presence of the spot market. Our objective is to design effective portfolio contracts and identify the replenishment policy so as to maximize the manufacturer's (i.e., the buyer) expected profit.

2. The Model

Consider a manufacturer of a single item who faces demand during several time periods. Future demands are unknown, but as time goes by, more precise information on the distribution of demand becomes available. The manufacturer's objective is to optimally manage its supply by buying well-chosen capacity for every time period. Specifically, the objective is to maximize expected profit by effectively managing the supply process.

For this purpose, the manufacturer needs the right trade-off between price and flexibility. That is, the manufacturer needs to find the appropriate mix of low price yet low flexibility (i.e., long-term) contracts; reasonable price, but better flexibility (i.e., option) contracts, or unknown price and quantity supply, but no commitment (i.e., the spot market). Once these decisions are made at the beginning of the horizon, the manufacturer has to manage its inventory effectively.

At the beginning of the planning horizon, i.e., in period $t=0$, the manufacturer decides on the type of contact it will buy from its suppliers for the entire planning horizon. This planning horizon is finite and has T time periods indexed from 1 to T in an increasing order. The contract specifies for every time period the cost of receiving any amount of supply.

The contract purchased at time period $t=0$ will be denoted by $\mathbf{r}(\cdot)=(r_1(\cdot),\ldots,r_T(\cdot))$. That is, at time period t, the contract allows the manufacturer to buy an amount $q_t \geq 0$ at a total cost of $r_t(q_t)$. Since we are interested in studying portfolios of contracts, we make the following assumption.

Assumption 1. *For $t=1,\ldots,T$, the contract $r_t(\cdot)$ is made up of n_t options with capacities $x_t^{i_t}$, reservation unit price $v_t^{i_t}$, and execution unit price $w_t^{i_t}$, $i_t=1,\ldots,n_t$. Without loss of generality, we assume that $w_t^1 \leq \cdots \leq w_t^{n_t}$. Therefore, the total purchase cost associated with these options is determined by solving*

$$r_t(q)=\sum_{i=1}^{n_t} v_t^i x_t^i+\min\sum_{i=1}^{n_t} w_t^i q_t^i$$
$$\textit{subject to}\ \begin{cases}\sum_{i=1}^{n_t} q_t^i = q\\ 0\leq q_t^i\leq x_t^i\ \forall\, i=1,\ldots,n_t.\end{cases}$$

This formulation is appropriate, for instance, when a manufacturer receives offers from different suppliers offering different option terms, or when a single supplier offers a broad spectrum of option contracts. It captures a wide class of contracts, including long-term or flexibility contracts.

Once the contract is signed, we move to Period 1. At the beginning of period t, $t=1,\ldots,T$, the customer demand D_t and the spot market price $s_t(\cdot)$ become known. This function $s_t(\cdot)$ describes the pricing of the spot market as a function of the amount purchased. That is, at period t, the manufacturer obtains an amount $q_t\geq 0$ of supply at a total cost $s_t(q_t)$, which is random and will only be known at the beginning of period t. Typically, $s_t(\cdot)$ is a convex function. In the context of portfolio contracts, we make the following assumption on spot market conditions, which is common in industry.

Assumption 2. *The spot market unit cost is constant and equal to S_t, and the market only offers a limited supply κ_t, for $t=1,\ldots,T$. That is, the structure of $s_t(q)$ is as follows.*

$$s_t(q)=\begin{cases}S_t q & \text{for } 0\leq q\leq \kappa_t\\ +\infty & \text{else.}\end{cases}$$

After demand and spot market conditions are revealed, the manufacturer orders from the suppliers and from the spot market. The inventory carried from previous periods plus the incoming supply can be used to satisfy this period's demand, at an exogenous price p_t. The remaining inventory is carried over to the next period. Finally, unsatisfied demand is lost to the competition.

We denote by I_t, for $t=1,\ldots,T+1$, the inventory level at the beginning of period t. We assume $I_1=0$, although the model can handle any initial inventory level. Because this is a lost-sales model, we impose $\mathbf{I}\geq 0$. As is common in traditional inventory models, a convex inventory holding cost $h_t(I_t)$ is incurred at the beginning of period t.

The decisions the manufacturer must take can be analyzed using dynamic programming. The decision space at a given time period $t\geq 1$ includes the following quantities: the quantity $q_t^r\geq 0$ that the manufacturer executes from the supply contract at a cost $r_t(q_t^r)$; the quantity $q_t^s\geq 0$ that the manufacturer purchases from the spot market at a cost $s_t(q_t^s)$; the amount of demand not satisfied q_t^n, $0 \le q_t^n \le d_t$.

At a given time period t, each one of the $n_t + 2$ sources (options plus spot market plus demand rejection) offers supply at fixed marginal cost (execution price) for a given maximum capacity. These supply sources can be represented as a pair (unit cost, capacity level): (w_t^i, x_t^i) for $i = 1, \ldots, n_t$, (S_t, κ_t) and (p_t, D_t). Define $\overline{n}_t = n_t + 2$ and sort these pairs by increasing unit cost. Let $(\overline{w}_t^j, \overline{x}_t^j)$ be the pair with jth smallest unit cost. Thus, $\overline{w}_t^1 \le \cdots \le \overline{w}_t^{\overline{n}_t}$. With this notation, the optimal replenishment policy has a simple structure: The manufacturer should *follow a modified base-stock policy for every option, the spot market, and the demand source.*

Theorem 1. *For $t = 1, \ldots, T$, there exist inventory target levels $b_t^1, \ldots, b_t^{\overline{n}_t}$ such that*

(i) *for $i = 1, \ldots, \overline{n}_t - 1$, $b_t^{i+1} \le b_t^i - \overline{x}_t^i$;*

(ii) *for $i = 1, \ldots, \overline{n}_t$, it is optimal to order from source i*

- 0 *if* $I_t \ge b_t^i$;
- $b_t^i - I_t$ *if* $b_t^i - \overline{x}_t^i \le I_t \le b_t^i$;
- $\overline{x}_t^i$ *otherwise.*

3. Portfolio Selection

The analysis so far took as an input the structure of the portfolio contract and derived the replenishment policy effectively. Thus, the fixed (reservation) cost associated with the contracts could be ignored. It remains to establish the structure of the optimal contract that will minimize both replenishment and reservation cost.

Theorem 2. *At period $t = 0$, the problem of choosing the capacities $\{x_t^{i_t}\}$ is a concave maximization problem.*

This result is key from a practical standpoint. Indeed, an extensive class of methods from nonlinear programming can handle this type of concave maximization problems.

We can provide conditions and guidelines that help establish whether a specific option is attractive for the manufacturer, i.e., such an option should be part of the optimal portfolio. This is presented in the following theorems.

Theorem 3. *At time period t, option i_t should not be included in an optimal portfolio when*

(i) *there is an option k_t such that $v_t^{i_t} > v_t^{k_t}$ and $v_t^{i_t} + w_t^{i_t} > v_t^{k_t} + w_t^{k_t}$;*

(ii) *there are options j_t and k_t such that $w_t^{j_t} < w_t^{i_t} < w_t^{k_t}$ and*

$$\frac{w_t^{i_t} - w_t^{j_t}}{w_t^{k_t} - w_t^{j_t}} v_t^{j_t} + \frac{w_t^{k_t} - w_t^{i_t}}{w_t^{k_t} - w_t^{j_t}} v_t^{k_t} < v_t^{i_t};$$

Theorem 4. *Assume that $\kappa_t = \infty$ with probability one. Then, at time period t, option i_t should not be included in an optimal portfolio when*

$$\mathbb{E}\left[S_t - w_t^{i_t}\right]^+ \le v_t^{i_t}.$$

4. Conclusion

The ability to establish effective supply contracts has been recently recognized, both in industry and academia, as an important contributor to the success of the firm. Indeed, effective supply contracts play an important role in manufacturers' abilities to reduce cost and ensure adequate supply of components. This is especially important for products for which customer demand is highly uncertain and hence supply flexibility is necessary.

Identifying effective contracts for commodity products is even more challenging since the manufacturer typically has a number of options to choose from: many suppliers willing to sign long-term or option contracts as well as spot market purchasing. Thus, portfolio contracts may be appropriate, as they provide flexibility over long-term contracts, and increase expected profit by effectively selecting the various options taking into account uncertainty in the spot market.

We thus develop a general framework for the design of supply contracts in a multi-period environment. This framework is especially suited for the analysis of a portfolio of different options. In this case, the optimal replenishment policy is particularly simple: *Every option's optimal execution policy is a modified base-stock policy*. Moreover, the contract design problem has a concave structure that makes it tractable in practice. More results can be found in Martínez-de-Albéniz and Simchi-Levi (2002).

At this point it is important to point out an important limitation of our model. So far, we have focused on the problem of maximizing expected profit. One might wonder, however, if it is possible to modify the model to take risk explicitly into account. This issue is addressed in Martínez-de-Albéniz and Simchi-Levi (2003b).

Finally, this work allows an important extension to modeling supplier competition. That is, assuming that buyers behave in the way presented here, suppliers will compete in price and flexibility to capture market share. This is presented in Martínez-de-Albéniz and Simchi-Levi (2003a).

References

Billington, C. 2002. HP cuts risk with portfolio approach. *Purchasing.com* (February 21).

Carbone, J. 2001. HP buyers get hands on design. *Purchasing.com* (July 19).

Martínez-de-Albéniz, V., D. Simchi-Levi. 2002. A portfolio approach to procurement contracts. Submitted.

Martínez-de-Albéniz, V., D. Simchi-Levi. 2003a. Competition in the supply option market. Submitted.

Martínez-de-Albéniz, V., D. Simchi-Levi. 2003b. Mean-variance trade-offs in supply contracts. Submitted.