Chapter 8

SAM: A DECISION SUPPORT SYSTEM FOR RETAIL SUPPLY CHAIN PLANNING FOR PRIVATE-LABEL MERCHANDISE WITH MULTIPLE VENDORS

Stephen A. Smith
Department of Operations and Management Information Systems
Leavey School of Business, Santa Clara University
Santa Clara, CA 95053
ssmith@scu.edu

Narendra Agrawal
Department of Operations and Management Information Systems
Leavey School of Business, Santa Clara University
Santa Clara, CA 95053
nagrawal@scu.edu

Andy A. Tsay
Department of Operations and Management Information Systems
Leavey School of Business, Santa Clara University
Santa Clara, CA 95053
atsay@scu.edu

Abstract A number of retail firms use a "private-label" strategy in which merchandise is sold under a brand name exclusive to the retail firm, but manufactured by one or more independent vendors. While offering a number of benefits, this approach also poses a different set of supply chain challenges than manufacturer-brand-based retailing, in that the retail firm must take a more active role in organizing and coordinating the planning and materials management activities in a supply base that is often dispersed and heterogenous.

J. Geunes et al. (eds.),
Supply Chain Management: Models, Applications, and Research Directions, 163–198.

This chapter describes a methodology for planning capacity commitments, scheduling shipments, and managing inventory for an assortment of private-label retail merchandise produced by multiple vendors. The vendors differ in their lead time requirements, costs, and production flexibility. Product demand is uncertain, and fluctuates over time. We develop an optimization model to choose the production commitments that maximize the retailer's expected gross profit, given market demand forecasts and vendors' capacity and flexibility constraints. The model has been incorporated into a PC-based decision support system called the *Sourcing Allocation Manager* (SAM). This was developed in collaboration with supply chain planners at a global retailer of seasonal and fashion merchandise, and has been tested by buyers at two major retailers.

Keywords: Sourcing Strategy, Retailing, Capacity Planning, Multi-item Inventory Planning

1. Introduction

A number of retail firms use a "private-label" strategy in which merchandise is sold under a brand name exclusive to the retail firm, but manufactured by one or more independent vendors. This practice can allow a retailer to avoid the premium charged by brand-name vendors, fill gaps in its product assortment, exercise greater control over product attributes, gain leverage in the manufacturer-retailer balance of power, and convert product brand loyalty to store loyalty. For well-received products, there are additional benefits to be enjoyed from being the exclusive seller. However, this also poses a different set of supply chain challenges than manufacturer-brand-based retailing, in that the retail firm must take a more active role in organizing and coordinating the planning and materials management activities in a supply base that is often dispersed and heterogenous[1]. As a result, some such retail firms have become increasingly interested in tools and techniques for effective supply chain management and design. This is the case with the retailer (a multinational firm with several billion dollars of annual revenue from private-label sales) that approached us with the business problem motivating the research described here.

[1] Private-labeling poses a number of marketing challenges as well. The retailer takes sole responsibility for brand management tasks such as advertising and creating store displays, and foregoes manufacturer-sponsored provisions that mitigate market risks, such as return privileges and price protection. Our intent is not to address the question of when a retailer should use private-label, but to provide guidance on supply chain planning when this strategy is pursued.

We consider the problem of how to optimally plan and execute the sourcing of seasonal and fashion private-label merchandise carried by department stores and specialty retailers. For a given selling season, the sourcing decisions, typically made by the retail buyer responsible for each merchandise department, include the following components: (1) purchases of raw materials (e.g., fabric) for use by vendors, (2) supply contracts and production commitments with vendors, (3) a weekly plan for sales, shipments, and inventory, and (4) adjustments based on subsequent market information. This research develops a formal planning methodology for this decision problem that accommodates multiple products and multiple suppliers, and explicitly accounts for demand uncertainty and adjustments to the plan during the season. The resulting optimization model has been embedded within a PC-based decision support system named the *Sourcing Allocation Manager* (SAM).

A more theory-oriented treatment of this modeling research is presented in Agrawal et al. (2001). Parts of that document describing the model formulation are included here for the reader's convenience, but those who are interested in such a perspective and an extensive numerical case study should refer to that paper. This chapter focuses on the software implementation and how the business environment influenced the design of the graphical user interface.

The Business Setting

Many of the challenges of this application are due to attributes of the demand patterns and the supply base, and how they interact. Demand in this environment typically fluctuates sharply throughout the year. This is exemplified by the data in Figure 8.1, which illustrates recent sales for a men's casual slacks product.

This type of demand becomes most challenging when production capacity is constrained, which is commonly the case in this industry. Specifically, demand during the peak Fall ("Back to School") and Christmas seasons typically exceeds available manufacturing capacity, while surplus capacity tends to exist during the Spring and Summer. Producing in advance of peak periods improves the ability to meet demand, but creates inventory buildup arid requires that commitments to production and fabric purchases be made under greater uncertainty.

Sourcing strategies must also reflect the performance capabilities of the supply base. In most cases there are a variety of possible vendors that differ in costs, lead times, and flexibility of production. Vendors with the lowest cost generally offer virtually no flexibility with respect to capacity commitments. These vendors tend to have long lead times for booking capacity (e.g., nine months), shipment times of several weeks,

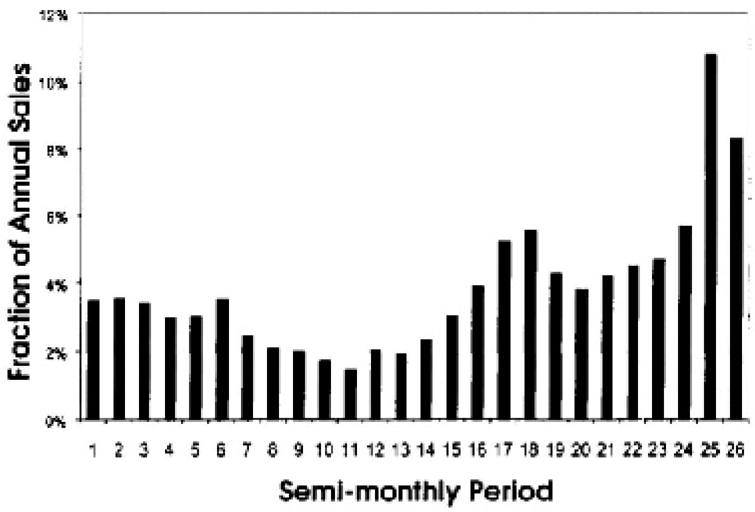

Figure 8.1. Seasonal Patterns in Demand[2]

and often require that the total production be allocated relatively evenly throughout the year. More responsive vendors may have shorter lead times and allow greater flexibility vis-a-vis production commitments. Additionally, different vendors may be willing to store limited amounts of finished product prior to delivery for a fee.

Retailers tend to leverage a portfolio of such vendors, resulting in supply chains such as that shown in Figure 8.2. The portfolio approach enables strategies such as exploiting lower cost production for the more predictable segment of demand, while sourcing the more speculative segment via the more flexible, but more costly, vendors. Operationalizing this strategy in a multi-product, multi-vendor setting is nontrivial, and is further complicated by many production and logistical constraints described later. This was our retail collaborator's motivation in sponsoring this project. In fact, our methodology is unique in its focus on designing contracts with a portfolio of vendors that simultaneously exploits the comparative advantages of each, as opposed to selecting a single most desirable vendor.

Research Contribution

Relative to previous academic research detailed in Section 2, our for-

[2]Since the retailer providing this data aspires to and usually achieves very high fill rates for this product, the difference between sales and demand is insignificant.

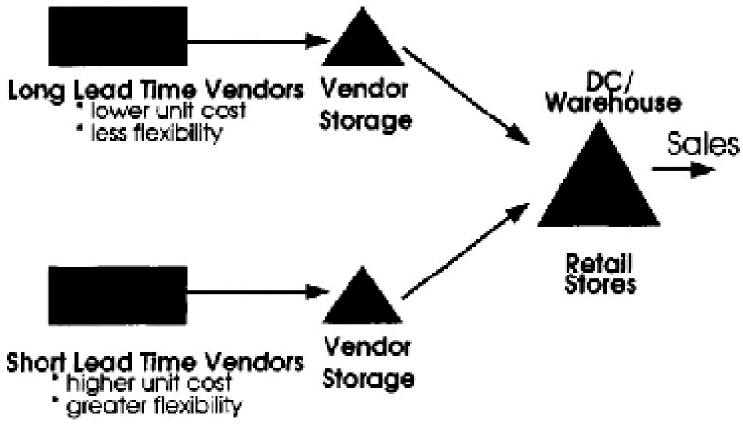

Figure 8.2. Supply Chain Structure

mulation of the multi-vendor sourcing problem is novel in representing the complex constraints and changing states of information under which different sourcing commitments must be made. We address numerous issues associated with the design of the supply chain, and provide insight into a universal question in sourcing: how to balance unit costs versus supplier attributes such as flexibility. Overall, our model builds on the key aspects of the literature described in Section 2, incorporating seasonal patterns in demand and detailed production and logistical constraints in a stochastic demand environment with forecast updating. While subsets of these issues have been treated previously, we believe our formulation to be unique in addressing all of them simultaneously.

Our formulation evolved in close collaboration with retail practitioners, whose involvement occurred at two different levels. A committee of senior executives from different firms regularly reviewed our assumptions and problem framing to ensure the broad applicability of our model to a variety of retail settings. However, the specifics were developed in collaboration with executives and buyers at a particular retailer, who confirmed that our level of detail captures the key complexities faced by retail planners. Their help was especially useful in identifying the cost tradeoffs and constraints most important for sensitivity analysis, leading to variable and constraint modifications that allowed discovery and presentation of the most critical shadow prices. Furthermore, feedback from these buyers and planners was instrumental in the incorporation of our model into a decision support software package with a graphical user interface. Given the depth and breadth of the practitioners' participation, we believe this model to be widely applicable to retail firms that manage the sourcing and production of private-label merchandise, and to certain nonretail firms as well.

Organization of This Chapter

The remainder of this chapter is organized as follows. Section 2 reviews the relevant literature. Section 3 details the mathematical formulation of the optimization model, discussing in depth the assumptions we made to capture the salient features of the particular retail environment. Section 4 describes the decision support software, the business issues motivating the design features of the user interface, and summarizes our retail collaborator's experiences with the software, and Section 5 concludes.

2. Literature Review

The use of formal decision models in aggregate production planning has a long tradition, and has been the subject of hundreds of academic studies. See Silver and Peterson (1985) for a textbook treatment and some historical background. A review of the academic literature is provided in Nam and Logendran (1992), and a survey of the usage of such models in practice is provided in Buxey (1993) and Buxey (1995). The predominant optimization approach is based on linear programming (LP), which allows for non-stationary but deterministic demand, and can handle large numbers of products simultaneously. Forecast uncertainty and information updating are usually dealt with only in an indirect fashion, by using a rolling-horizon implementation of a snapshot deterministic solution (the formal term for this is "Open-Loop Feedback Control", cf. Bertsekas (1976)), and also perhaps through the specification of safety stock levels, usually exogenously (e.g., Guerrero et al. (1986), Gunther (1982), Heath and Jackson (1994), Miller (1979)).

More direct treatment of demand uncertainty is called for in the retail setting, especially where hard-to-forecast fashion or style goods are involved. This can be provided by newsvendor-style models, but at the expense of sacrificing the dimensionality and detailed constraint structure that can be supported by LP formulations. In this type of approach, the entire selling season for a product is summarized as a small number (possibly one or two) of random variables with known joint probabilities. This allows analytic incorporation of forecast uncertainty into production planning (e.g., Crowston et al. (1973), Hartung (1973), Hausman and Peterson (1972), Murray and Silver (1966), Ravindran (1972), Wadsworth (1959), and more recently Brown and Lee, Donohue (2000), Eppen and Iyer (1997), Fisher and Raman (1996)), albeit in stylized ways. Various approaches to obtaining the probability distributions of these demand random variables, especially for fashion products, are proposed by Chang and Fyffe (1971), Hausman and Sides (1973), Hertz and Schaffir (1960), Riggs (1984), Riter (1967), and Wolfe (1968).

The efforts closest in spirit to our work are Bitran et al. (1986), Eppen et al. (1989), Kira et al. (1997), and Nuttle et al. (1991). The first three are based on mathematical programming, while the fourth takes a simulation approach. We discuss them briefly below.

In Bitran et al. (1986), the authors perform multi-period production planning for families of consumer electronics products which in turn consist of specific items. Setup costs for switchovers between families are such that each family will be run only once during the season, while switchovers between items within a family are assumed costless. Demand occurs in the last period, and estimates of this demand are revised each period. Demands for all items are assumed to be normally distributed, and the standard deviation of forecast error at each time period is known, given by an arbitrary, decreasing sequence of numbers which must be provided as data. The updated forecasts at each period also follow a joint normal distribution, with a known covariance matrix. The exact problem is a difficult-to-solve, stochastic mixed-integer program, for which the authors develop a deterministic mixed-integer approximation. While both their model and ours consider multi-product planning with forecast updates, the respective areas of emphasis differ. Whereas they take production capacity as given and then determine how to schedule the production of a variety of items, we consider as decision variables the capacities to be reserved with a variety of vendors at different points in time. They model the operations within a single factory at a greater level of detail, whereas our scope spans multiple vendors' factories as well as the retailer's distribution center, and includes the scheduling of shipments from the former to the latter. Their representation of item demand is more general but also data-intensive. We pursue a discrete simplification of forecast dynamics as part of an overall strategy of retaining a basic LP structure that allows an exact solution in real time.

In Eppen et al. (1989) a model is developed for General Motors to aid in making decisions about capacity for several lines of automobiles produced in multiple factories. A general sequence of events is considered in each of five years: (1) the available capacity is configured in terms of tooling the production lines for specific products, (2) demand occurs, and (3) a production plan is implemented that attempts to meet the demand given the capacity configuration. Demand uncertainty is represented by defining three different "scenarios" for each year that specify the demand and price for individual products. Scenario probabilities are assigned, and are assumed to be independent from year to year. The resulting optimization problem is a mixed-integer program that maps out individual sample paths of all possible scenario combinations. This scenario approach is similar to our representation of demand uncertainty. However, our production decisions are based on imperfect demand signals, while theirs assume that all uncertainty has been resolved. Further, our notions of capacity are slightly different. Their optimal capacity configuration is selected from a discrete number of predefined possibilities, hence the integer variable structure. Ours is chosen from a simplex region defined by a variety of constraints that explicitly represent features of the business relationship between the retailer and each vendor.

In Kira et al. (1997), the authors use a probability structure similar to that in Eppen et al. (1989), with a single-factory production environment that is much simpler than ours. Capacity planning is not treated, and the nuances of managing a supply chain composed of multiple, independently-managed physical nodes are not incorporated into their formulation.

In Nuttle et al. (1991) a software application called "The Sourcing Simulator" is described, which was developed by researchers at North Carolina State University in concert with the Textile/Clothing Technology Corporation and the American Textile Partnership-Demand Activated Manufacturing Architecture (AM-TEX-DAMA) project. This treats the same industry setting as we do, and makes many similar assumptions in addressing the question of how the replenishment frequency and lead time of a vendor affects a retailer's performance. This purely descriptive simulation approach allows a detailed representation of certain aspects of the setting, especially in the range of allowable replenishment strategies and consumer behavior. However, because it assumes single-sourcing (with the single vendor abstracted as simply a lead time and reorder frequency), it cannot simultaneously allocate production across a portfolio of time-phased vendors. Like the three models described previously, the scope of this formulation is largely confined to a single firm. Nevertheless, various studies based on this model (Hunter et al. (1992), Hunter et al. (1996) and King and Hunter (1996)) have validated the importance of the ability to react to improved demand information, which is a key rationale for the sourcing strategies that we model.

3. Model Specification

This section outlines the mathematical formulation of the planning problem faced by a retailer leveraging a portfolio of time-phased vendors. Our discussion uses the language of apparel retailing because this is our sponsor firm's primary line of business. However, we believe our underlying methodology to be broadly applicable to other product settings.

3.1 Timeline of Events and Information Assumptions

In chronological order, the critical time points for the retailer's sequential decision problem for a specified "selling season[3]" are as follows:

t_0 = time at which initial vendor commitments and fabric purchases[4] are made

t_1 = second time at which commitments to vendors are made, for those vendors allowing capacity decisions to be deferred to this time[5]

t_b = beginning of selling season

t_f = end of selling season, when actual demand becomes known.

We assume that our model analysis is performed at some time at or before t_0 for a selling season that spans the horizon (t_b, t_f). The retail planner's information regarding demand evolves continuously over time, shaped by economic forecasts, new fashion and color trends, and observed sales results for similar products. However, for our formulation it is only necessary to define the possible states of information at the specific points in time at which decisions are made. Evaluation of the expected profit also requires knowledge of the actual demand information at time t_f. To represent the evolving demand information, we define the following random variables[6]:

---

[3]This might correspond, for example, to the Fall season (running from roughly August through January) or the Spring season (February through July). For certain merchandise, some retailers use four or more shorter seasons per year. In some instances a season may be as short as 8 weeks.

[4]In many cases the fabric is purchased by the retailer and shipped to vendors for cutting and sewing. This provides control of raw material quality and leverages the buying power that a major retailer enjoys.

[5]Our discussions with the retailer's production planning managers indicated that two decision points (times t_0 and t_1) are adequate for a typical apparel planning decision process. However, the formulation can easily be extended to include more decision points by simply adding more variables to the model.

[6]For example, we have assumed that the initial demand information for any product at time t_0 is deterministic, i.e., X_0 has only one possible value. At time t_1, the demand information demand has three possible values based on what has been observed since t_0: High, Medium, or Low, with different probabilities. The remaining uncertainty about the actual demand is

$X_k \equiv$ a random variable corresponding to the market demand information that the retailer has at time t_k, for $k = 0, 1, f$.

At each time point, X_k has a discrete set of possible values. Finally, at time t_f, the actual demand corresponds to one of a discrete set of demand scenarios. We define the following probability distributions to describe the likelihood of observing particular sequences of demand information:

$p(\xi_1) \equiv P\{X_1 = \xi_1\}$ for each possible ξ_1 value at time t_1

$p(\xi_f|\xi_1) \equiv P\{X_f = \xi_f|X_1 = \xi_1\}$ for each possible combination of ξ_1 and ξ_f, and

$p(\xi_1, \xi_f) \equiv p(\xi_f|\xi_1)\,p(\xi_1)$ = the joint probability of X_1 and X_f.

Clearly, this structure can be generalized to characterize information that is revealed in any number of stages, but we will describe only the two-stage case since that corresponds to our particular application.

Market "scenarios" are frequently used by retailers in developing marketing plans for alternative contingencies[7]. We extend this concept to include market demand information that is revealed in stages, resulting in the sequential stochastic decision model illustrated in Figure 8.3. The underlying assumption is that as the selling season gets closer, the sales estimates in the plan improve for several reasons. For example, there is new sales information for related products. Also, updated sales estimates are at least in part the result of revisions in the merchandise plan, e.g., deciding to feature more or less of a particular type of merchandise, giving it a more or less prominent display and floor space, etc. For

then described by the conditional probability distribution , and is not completely resolved until the end of the selling season at time t_f. We also note that this modeling structure is easily generalizable to include additional stages of information and decision points.

[7]We model demand uncertainty through discrete scenarios for three reasons. The first reason is analytical tractability. Modeling uncertainty using continuous random variables would rule out certain complexities categorically declared by our corporate collaborators to be essential attributes of their business setting. The second reason is consistency with common managerial practice. Our corporate collaborators indicated that their planning methodology often requires the articulation of "worst case," "most likely," and "best case" scenarios for market uncertainties. However, in the past these scenarios have typically been used only for financial planning, due to a lack of technical know-how for translating them into contingency plans for vendor and production management. The third reason is that there is an established precedent in the literature for using scenarios to model uncertainty in a variety of contexts. As described in Section 2, Eppen et al. (1989) and Kira et al. (1997) used a scenario approach similar to ours for capacity planning. Discrete demand scenarios were used in Smith et al. (1998) to obtain optimal inventory and promotional plans for retail chains. Of course, there is a rich tradition in the financial economics literature of modeling uncertainty in the prices of stocks and securities this way (cf. Cox and Rubinstein (1985)). More recently, Huchzermeier and Cohen (1996) have used discrete scenarios to study the operations management implications of exchange rate fluctuations.

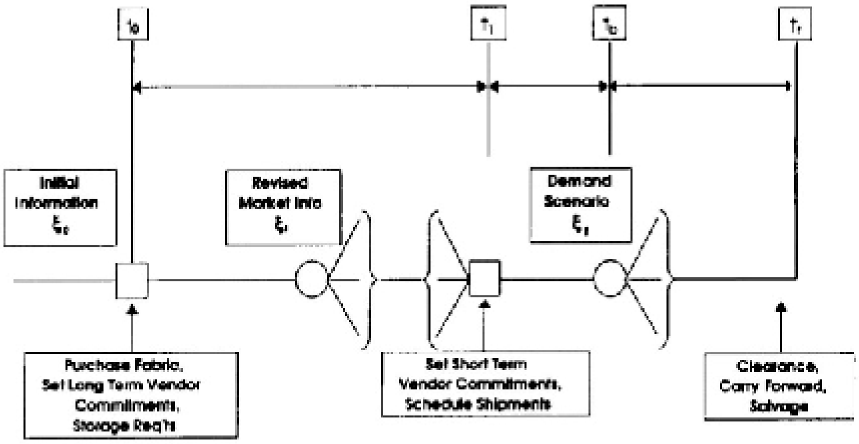

Figure 8.3. Decision Tree for Production Planning

changes of this type, there is a good base of experience for the buyers to update their subjective estimates of the demand. This determines the conditional probabilities $p(\xi_f|\xi_1)$. Note that in principle the same method can use early sales results to update the demand probabilities (after the selling season begins) and in fact, our formulation approach is compatible with Bayesian updating of the probabilities of the discrete demand levels based on early sales results. However, for this application the vendor deadlines did not permit changes in capacity commitments after the start of the selling season, other than changes in the color, style, or size mix. Since our model is meant to support capacity planning at an aggregate level, this is appropriate for our application[8]. With some assistance from the authors, the retail planners at two major retailers were able to subjectively estimate the required probabilities.

3.2 Decision Variable Definitions

The following indices will be used for variabledefinitions: j for products, i for vendors, t for the time increment used for production, ship-

[8]Most papers that consider updating of forecasts in a model of reasonably realistic detail only consider updating prior to the occurrence of any sales. This includes the mathematical-programming-based models most similar to ours, as described in Section 2. Those models that do accommodate forecast updating based on in-season sales tend to have very simplified inventory analysis that would not scale to the constraint and decision variable complexity in our decision model (e.g. Chang and Fyffe (1971), Crowston et al. (1973), Fisher and Raman (1996), Hartung (1973), Murray and Silver (1966)).

ment, and sales decisions (typically weeks), and q for the time increment used for reservation of capacity (typically quarters). In this model, we assume that the term "product" refers to an aggregation of styles, not to an individual SKU (distinguished by style/size/color). Variable names in upper case represent decision variables, while those in lower case or Greek symbols are fixed parameters.

The main basis for classifying vendors (into "short lead time" and "long lead time" types) is the time at which commitments for each product must be made. This is denoted by

$\tau_{ij} \equiv$ time at which a commitment is required by vendor i for product j,

and the corresponding state of retailer information

$X_{ij} \equiv$ demand information available at time τ_{ij}, which takes on discrete values ξ_{ij}.

For our implementation, $\tau_{ij} = t_0$ or t_1, since these are the only production commitment time points. It follows that X_{ij} is either X_0 or X_1 for every combination of i and j.

For each possible (ξ_1, ξ_f) combination, the production and inventory variables are defined as follows:

$F_j \equiv$ fabric commitment (in yards) made at time t_0 for product j

$P_{ij}(t|\xi_{ij}) \equiv$ production by vendor i of product j during period t

$Z_i(q|\xi_1) \equiv$ total production by vendor i during quarter q

$Z_j^F(\xi_1) \equiv$ yards of fabric actually used for product j

$M_{ij}(t|\xi_1) \equiv$ beginning inventory of product j stored by vendor i in period t

$S_{ij}(t|\xi_1) \equiv$ quantity of product j shipped from vendor i in period t

$U_j(t|\xi_1, \xi_f) \equiv$ retailer's unit sales of product j in period t

$I_j(t|\xi_1, \xi_f) \equiv$ retailer's beginning inventory of product j in period t

The decision variables depend on the information states in different ways, i.e., what information is known when each variable's value is specified. These dependencies determine the dimensionality of the variables. We denote this dependence explicitly in our formulation, using the "|" notation. For example, since the production schedule of an item i at a vendor j is fixed at time τ_{ij}, when the state of information is ξ_{ij}, the corresponding vendor production variables are denoted as $P_{ij}(t|\xi_{ij})$. The total production and total fabric usage depend upon ξ_1 because they are defined for both short and long lead time production decisions. Similarly, the vendors' inventory and shipment decisions depend upon ξ_1 because that is the information available to the vendor when the shipping decisions are made. However, the realized unit sales, and consequently the retailer's on hand inventory, depend on both ξ_1 and ξ_f. This is because the on-hand inventory depends on both the actual demand scenario and all the production decisions, some of which depend on ξ_1. Since the unit sales are affected by the inventory level, this depends on ξ_1 and ξ_f as well. In the LP optimization, the information states ξ_1 and ξ_f are simply treated as additional "subscripts" on variables.

3.3 Inventory Balance Equations and Production Constraints

The production, inventory, and shipping variables are related to each other by the following inventory balance equations for the retailer and vendors:

$$I_j(t+1|\xi_1,\xi_f) = I_j(t|\xi_1,\xi_f) + \sum_{i|t\geq\tau_{ij}+l_i} S_{ij}(t-l_i|\xi_1) - U_j(t|\xi_1,\xi_f),$$

$$\text{for all } i, j, \xi_1, \xi_f, t, \tag{8.1}$$

where l_i is the shipping delay for vendor i,

$$M_{ij}(t+1|\xi_1) = M_{ij}(t|\xi_1) + P_{ij}(t|\xi_{ij}) - S_{ij}(t|\xi_1),$$
$$\text{for all } i, j, \xi_1, t. \tag{8.2}$$

When the states of information "subscripts" in one constraint are different for different variables, the variable with fewer subscripts simply keeps the same value for a subset of the equations.

For simplicity, our model considers only the total inventory in the retailer's system, as opposed to inventory levels in individual stores[9]. This assumes that inventory is generally balanced across the stores, and

[9]Once the merchandise reaches the retailer's distribution center (DC), it is usually distributed to the stores and displayed for sale within two to three days. In order to maximize the productivity per square foot, there is generally little storage space in stores, and all store merchandise is placed on display for sale as quickly as possible. The only significant delays in this type of supply chain arise from production commitment lead times, which are usually several months, and shipping times, which may be several weeks for surface shipments.

is appropriate because inventory is re-balanced during the season by allocating replenishments to the stores that most need additional stock. For some merchandise, transshipments are made from one store to another to balance the inventory, but only if the repackaging and shipping costs can be justified.

Constraints on each vendor's storage space can be represented as

$$\sum_j \nu_j M_{ij}(t|\xi_1) \leq w_i(t) \equiv \text{vendor } i\text{'s maximum storage for period } t,$$

$$\text{for all } i, t, \xi_1 \qquad (8.3)$$

where $\nu_j \equiv$ storage space required per unit of product j.

A retailer may also specify an upper bound on the amount of inventory contained within its system[10]. This can be specified by

$$\sum_j \nu_j I_j(t|\xi_1, \xi_f) \leq w^R(t) \equiv \text{retailer's maximum storage for period } t,$$

$$\text{for all } t, \xi_1, \xi_f \qquad (8.4)$$

The initial and final inventories may also be required to satisfy constraints of the form:

$I_j(t_b|\xi_1, \xi_f) \geq i_j^0 \equiv$ minimum initial retailer inventory for product j for all ξ_1, ξ_f

$I_j(t_f|\xi_1, \xi_f) \geq i_j^f \equiv$ minimum final retailer inventory for product j for all ξ_1, ξ_f

$M_{ij}(t_b|\xi_1) \geq m_{ij}^0 \equiv$ minimum initial inventory of product j at vendor i for all ξ_1

$M_{ij}(t_f|\xi_1) \leq m_{ij}^f \equiv$ maximum final initial inventory of product j at vendor i for all ξ_1.

The initial inventory i_j^0 must be sufficient to create an attractive display of merchandise with which to begin the selling season. For continuing, or "basic" products, the minimum final inventory i_j^f may be set to the desired initial inventory for the subsequent season. The vendor's initial inventory m_{ij}^0 can be used to satisfy demand in the current season, while the final inventory m_{ij}^f is available for the subsequent season.

[10]This can represent either a physical or budget restriction. In the latter case, ν_j will have a different meaning.

For some aspects of aggregate production planning, managers use quarters as the appropriate increment of time. Using $q(t)$ to denote the quarter corresponding to a time period t, the following relationship tallies the total production of vendor i within a quarter y:

$$\sum_{j} \sum_{t|q(t)=y} \kappa_j P_{ij}(t|\xi_{ij}) = Z_i(y|\xi_1), \quad \text{for all } i, y, \xi_{ij} \tag{8.5}$$

where $\kappa_j \equiv$ production capacity required per unit of product j. This enables us to model quarterly constraints. For instance, to ensure diversification a vendor may be willing to commit only a fraction of its quarterly capacity to a single retailer. On the other hand, less flexible vendors may also insist on a minimum quarterly production commitment from the retailer as a condition for doing business. These can be included as follows:

$$\underline{k}_i(q) \leq Z_i(q|\xi_1) \leq \overline{k}_i(q), \quad \text{for all } i, q, \xi_1 \tag{8.6}$$

where the bounds do not depend on the demand information. To achieve the economic benefits of level production, certain vendors also permit only limited changes of total production from quarter to quarter, which can be expressed as follows:

$$(1-\alpha_i) Z_i(q-1|\xi_1) \leq Z_i(q|\xi_1) \leq (1+\beta_i) Z_i(q-1|\xi_1), \quad \text{for all } i, q, \xi_1 \tag{8.7}$$

where $0 \leq \alpha_i \leq 1$ and $\beta_i \geq 0$. In general, vendors that allow later commitments also typically allow greater quarter-to-quarter flexibility (larger α_i and β_i parameters).

Production is also constrained by the fabric procurement decision as follows:

$$\sum_{y} \sum_{j} \sum_{t|q(t)=y} \kappa_j^F P_{ij}(t|\xi_{ij}) = Z_j^F(\xi_1) \leq F_j, \quad \text{for all } j, \xi_{ij} \tag{8.8}$$

where $\kappa_j^F \equiv$yards of fabric required per unit of product j.

3.4 Modeling Product Demand

The demand pattern for each product over time is an input to the model that is conditional on the demand scenario ξ_f, denoted as follows:

$d_j(t|\xi_f) \equiv$ actual demand for product j in period t.

To specify these values, we used a forecasting model form that has been applied successfully to retail sales forecasting. Econometric marketing studies have found that multiplicatively separable models of the form

$$\begin{pmatrix}\text{Period } t \text{ demand}\\ \text{for product } j\end{pmatrix} = \begin{pmatrix}\text{Total season demand}\\ \text{for product } j\end{pmatrix} \cdot \begin{pmatrix}\text{Seasonality}\\ \text{effect at } t\end{pmatrix} \cdot \begin{pmatrix}\text{Marketing}\\ \text{effects at } t\end{pmatrix}$$

fit observed retail sales data well (Achabal et al. (1990), Kalyanam (1996)). Thus we let

$$d_j(t|\xi_f) = b_j(\xi_f) \cdot f_j(t) \cdot \rho_j(t) \tag{8.9}$$

where

$b_j(\xi_f) \equiv$ full-season demand for product j under demand scenario ξ_f

$f_j(t) \equiv$ fraction of total demand for product j that occurs in period t

$\rho_j(t) \equiv$ marketing effects for product j during period t, including price/ advertising effects.

This approach greatly reduces the model dimensionality by confining the effect of information updating to the full-season demand, which is a scalar. The full set of relative seasonality factors $f_j(t)$, such as that shown in Figure 8.1, generally do not require updating. Similar representations of demand have been used by Chang and Fyffe (1971), Crowston et al. (1973), and Hartung (1973). The specification of demand parameters and price variations due to any retail promotional strategies is exogenous to the optimization model, hence does not affect the linearity structure.

3.5 Calculating Unit Sales

Unit sales volume in period t is bounded by the period's demand, so

$$U_j(t|\xi_1, \xi_f) \le d_j(t|\xi_f), \quad \text{for all } j, t, \xi_1, \xi_f. \tag{8.10}$$

While traditional inventory models assume that lost sales occur only when inventory is fully exhausted, in retail marketing environments the amount of on-hand inventory can influence sales. In apparel, for example, sales rates can deteriorate as inventory drops because the remaining inventory consists of increasingly broken assortments with incomplete selections of sizes and colors (Smith and Achabal (1998)). Low inventory also increases the likelihood that some stores are inadequately stocked, i.e., the inventory is not "balanced." While the relationship between inventory level and sales is not necessarily linear (Smith and Achabal (1998)), a linear approximation is reasonable within the range of values of the inventory level that is expected in practice. This lends considerable analytical tractability to our formulation. Therefore, we allow unit sales to depend upon the beginning inventory according to the following constraints:

$$U_j(t|\xi_1,\xi_f) \le \eta_j I_j(t|\xi_1,\xi_f), \text{ for all } j,t,\xi_1,\xi_f \tag{8.11}$$

where $\eta_j \equiv$ maximum fraction of the beginning inventory that can be sold in one period[11]. Because of (8.1) and the production capacity constraints in (8.6) and (8.7), it is also possible that neither (8.10) or (8.11) will be binding for a given t.

Constraints (8.10) and (8.11) assume that the unfilled demand is lost (to competitors, for example), which is more common than backordering for most retail merchandise. Backordering, which is actually more straightforward to model, can easily be accommodated within our formulation by modifying the inventory balance equations.

3.6 The Objective Function

The objective function will be defined in terms of the following economic parameters:

$\pi_j(t) \equiv$ average selling price for product j in period t

$c_{ij} \equiv$ unit procurement + shipping cost ("landed cost") for product j purchased from vendor i

$r_j \equiv$ residual value per unit of product j at the end of the selling season

$c_j^F \equiv$ cost per yard of fabric for product j

$r_j^F \equiv$ residual value per yard of fabric for product j at the end of the selling season

$h_j \equiv$ retailer's unit holding cost per period for product j

[11] Retailers typically track the "sell-through" rate, i.e., the fraction of the beginning on-hand inventory that is sold in each time period. If the sell-through rate is too high, it is assumed that some sales have been lost due to insufficient inventory (see Smith et al. (1998) for further discussion).

$v_{ij} \equiv$ vendor i's unit storage charge per period for product j

The average selling price $\pi_j(t)$ may vary by time period to allow periodic price markdowns during the season. The value of r_j has different interpretations for seasonal and fashion items. For a seasonal item, it corresponds to the unit value of this product in the next selling season (i.e., the avoided replacement cost minus any holding cost). For fashion items it describes a "salvage value." At the selling season's end, any remaining fashion items may be sold through outlet stores or in bulk to discounters, resulting in markdowns to prices possibly below the original cost.

The expected revenue and cost for each product, denoted as R_j and C_j, respectively, are:

$$\begin{aligned} R_j &= \sum_{t,\xi_1,\xi_f} p(\xi_1,\xi_f)\Big\{\pi_j(t|\xi_f)\,U_j(t|\xi_1,\xi_f) \\ &\quad + r_j I_j(t_f|\xi_1,\xi_f) + r_j^F\left(F_j - Z_j^F(\xi_1)\right)\Big\} \qquad (8.12) \\ C_j &= \sum_{i,t,\xi_1} p(\xi_1)\left\{c_{ij}P_{ij}(t|\xi_{ij}) + v_{ij}M_{ij}(t|\xi_1)\right\} \\ &\quad + \sum_{t,\xi_1,\xi_f} p(\xi_1,\xi_f)\,h_j I_j(t|\xi_1,\xi_f) + c_j^F F_j \qquad (8.13) \end{aligned}$$

where $p(\xi_1,\xi_f)$ and $p(\xi_1)$ are the previously defined joint and marginal probabilities, respectively. The total objective to maximize is then $\sum_j \{R_j - C_j\}$

The fabric commitments, production capacity commitments, and shipping schedules that optimize this objective function correspond to a sequence of decisions under uncertainty, where the demand information changes at each decision point. In general, this can be viewed as a stochastic dynamic programming problem (with linear constraints). Unfortunately, the size of the resulting state space and the complexity of the objective make this solution approach impractical. However, as long as the states of information are restricted to a discrete set of values, the equations for R_j and C_j are linear in the decision variables, so that this optimization problem is a linear program[12].

[12]This approach for handling uncertainty within an LP formulation was first suggested by Dantzig (1955). Including decision variables whose values may be chosen after the resolution of the uncertainty leads to what is generally termed as a stochastic linear program with recourse. See Hansotia (1980) and Infanger (1994) for discussion of various technical aspects of solving such models and extensive reviews of the literature.

3.7 Model Extensions for Sensitivity Analysis

Important insights from an optimization analysis are often derived from shadow prices and other sensitivity outputs. In vendor sourcing, this information can identify the most critical vendor production and storage constraints, and therefore guide the retailer in negotiating these limits or in identifying alternative vendors with appropriate capabilities. The retailer's storage limits or end–season inventory requirements may also be opportunities for performance improvement.

Because of the multitude of variables and constraints associated with the specific time periods and information states, most individual shadow prices in our model are not directly meaningful. However, useful sensitivity information can be obtained by introducing additional variables. For instance, since increases in production and storage capacity would typically be made for the entire horizon rather than by individual periods, it is appropriate to introduce a single variable that increments a given vendor's capacity uniformly in all periods and information states. If this variable is then constrained to be 0, the corresponding shadow price will reveal the marginal benefit of increasing the vendor's capacities in all periods at once. We add variables for these aggregate constraints as follows:

$\overline{\Delta}_i \equiv$ increase in quarterly production capacity (000's) for vendor i for all quarters

$\underline{\Delta}_i \equiv$ decrease in quarterly minimum production (000's) for vendor i for all quarters

$\omega_i \equiv$ increase in storage capacity at vendor i (cartons)

The appropriate constraint equations ((8.6) and (8.3)) are then replaced with the following:

$$\underline{k}_i(q) - \underline{\Delta}_i \leq Z_i(q|\xi_1) \leq \overline{k}_i(q) + \overline{\Delta}_i, \text{ for all } i, q, \xi_1 \tag{8.14}$$

$$\sum_j \nu_j M_{ij}(t|\xi_1) \leq w_i(t) + \omega_i, \text{ for all } i, t, \xi_1 \tag{8.15}$$

$$\underline{\Delta}_i, \overline{\Delta}_i, \omega_i = 0. \tag{8.16}$$

This enhancement was made for components of the formulation deemed most important by the retail planners: vendor production capacity, vendor flexibility, vendor storage, end-season retail inventory, and product demand.

3.8 Positioning This Model in the Retailer's Planning Process

Our discussions with executives at our retail sponsor highlighted two key issues relevant to the implementation of our methodology. The first deals with the timing of the analysis. Even though our planning model formulates the demand and supply dynamics over a finite horizon, like many other such models it would actually be used on a rolling horizon basis. (As noted earlier, this approach can be termed "open-loop feedback control.") Thus, the production planning actions recommended by each run of the model will serve as important inputs to the subsequent run[13]. The second issue deals with the level of product aggregation at which the analysis is performed. The retail executives envisioned this model being used for analysis at the product category level (e.g., T-shirts, denim jackets, or denim pants) as well as at a lower product type level (e.g., Pocket Tees, V-Neck Tees, and Crew Neck Tees). The former analysis will typically be of interest to product managers who are responsible for the profitability of separate categories. The latter will be of primary interest to buyers who devise procurement plans for product types.

4. The Decision Support System

With extensive input from sourcing managers at the retail chain, the optimization model described above was implemented as a PC-based decision support system (DSS) named the *Sourcing Allocation Manager* (SAM). The user interface screens were programmed in Visual Basic and the optimization engine is LINGO, supplied to us by LINDO Systems. For test problems with four products, four vendors, a nine-month planning horizon, and 27 distinct sample paths of information realizations, the LP has several thousand decision variables and constraints. It was solved on a 300 MHz Pentium II PC in approximately 3-5 minutes.

The DSS development was a "proof of concept" exercise with several goals: (1) to provide a context for defining the user inputs and outputs of the model, (2) to test the practical viability of the optimization algorithm, (3) to demonstrate to the sourcing managers the potential benefits of the system, and (4) to identify through experience the cost tradeoffs

[13] For example, within the context of our formulation, at time t_0 one could be planning for a six month selling season that begins six months hence (i.e., $t_b - t_0 = 6$ months, and $t_f - t_b = 5$ months). The entire planning horizon thus consists of 4 quarters, with planning decisions being revised at a weekly level. In this case, the previously committed production, which might be the result of a prior run of the model, can serve as input constraints to the current run of the model.

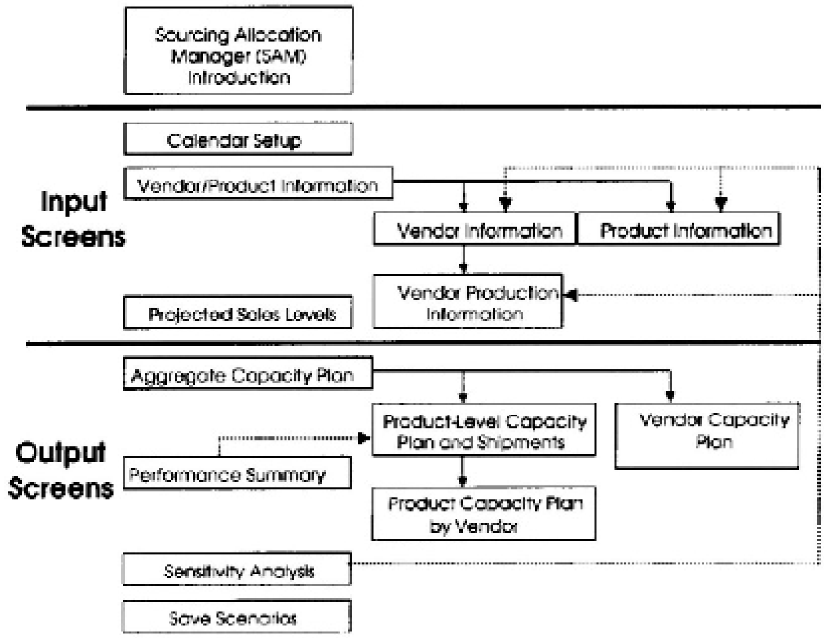

Figure 8.4. SAM Screen Flowchart

most important for planning. These goals were largely achieved, and the extensive involvement of retail planners profoundly influenced the resulting system in numerous ways. For instance, over the course of its development, the DSS evolved from a batch-processing application that generated the optimal sourcing plan for a particular text file consisting of all relevant parameters, to a system that enables beginning users to perform sensitivity and cost tradeoff analyses interactively. Feedback from users over the course of their cumulative experience with the DSS led to a number of key enhancements of the core mathematical formulation as well.

4.1 Graphical User Interface

The logical flow of the DSS screens is illustrated in Figure 8.4. In general, all input screens must be completed before any output screens can be viewed, although input scenarios can be stored for subsequent analyses. Data for the input screens can either be keyed in manually or read from a Microsoft Excel spreadsheet file, which a user can view and modify interactively. In a full-scale implementation most of these values would likely be fed directly from other applications or databases.

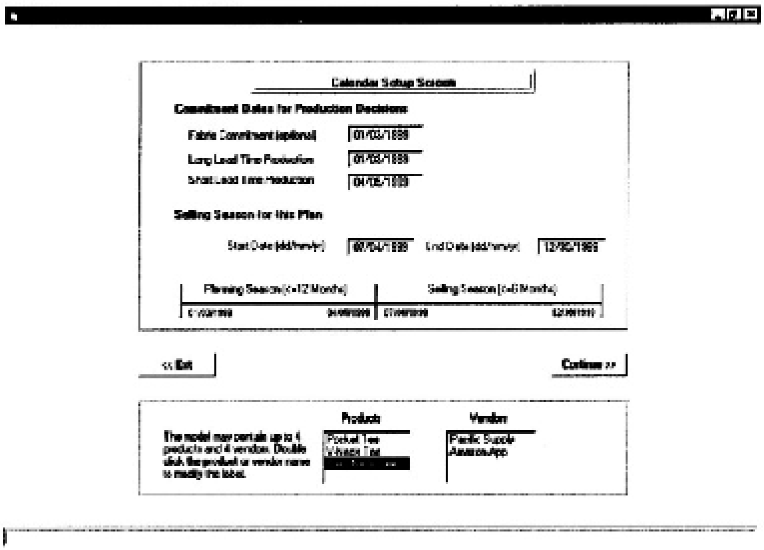

Figure 8.5. Calendar Setup Screen

Below we will describe the main screens, although space limitations preclude the inclusion of all screen views.

Input Screens

The Calendar Setup Screen in Figure 8.5 allows the user to specify dates delineating the timetable for planning. The first date field is for Fabric Commitment, indicating when the fabric must be ordered for all products under consideration. The second and third dates are the commitment times for the long and short lead-time vendors, respectively. For reasons discussed in Section 3.1 and Section 3.2, all commitments with vendors are modeled as being made at one of these two dates. However, a single vendor is allowed different commitment dates for different products. The Selling Season corresponds to the retailer's season for this set of products or the time frame for which this set of production commitments is in effect, whichever is shorter. (For continuing products, linkage to selling periods beyond this season is achieved by requiring end of season inventory, as described in Section 3.3.) This screen also allows the specification of vendor and product names (up to 4 of each).

Figure 8.6 shows the screen displaying vendor and product attributes. These are organized into a matrix with column headings (vendor names), row headings (product names), and interior cells that each provide click-

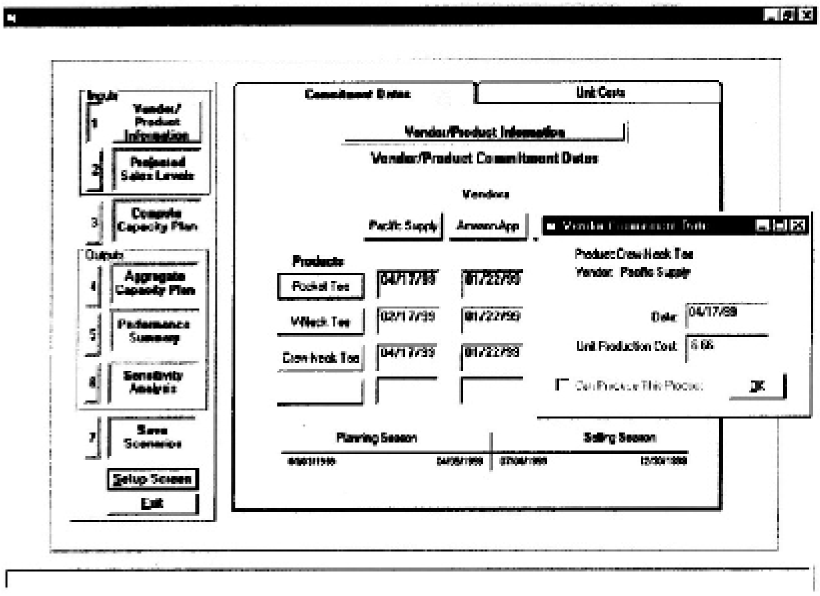

Figure 8.6. Vendor/Product Information Screen

through access to the appropriate type of information. This matrix framework persists throughout the DSS.

This screen has two different views accessible via the folder "tabs" at the top of the screen. These present the most significant attributes of each vendor-product sourcing combination – Commitment Date and Unit Cost. In the Commitment Date view depicted in Figure 8.6, the interior cells report the commitment deadlines required by each vendor for supplying each product. Short and long lead-time vendors are differentiated by color coding of these cells (although this is not apparent in a black-and-white graphic). Clicking on an interior cell calls up the dialog box presented in the foreground, in which a user can view or alter the Commitment Date or the Unit Cost. Toggling to the Unit Cost view presents the matrix of unit costs for all vendor-product combinations. Reflecting the richness of detail which our model can accommodate, separate input screens are required to fully specify the attributes of each vendor and each product. These may be accessed by clicking the appropriate column or row heading buttons, as described below.

Clicking a column heading button in the Figure 8.6 screen calls up the vendor information shown in Figure 8.7. The allowable quarter-to-quarter production volume adjustment (see equation (8.7)) and shipping

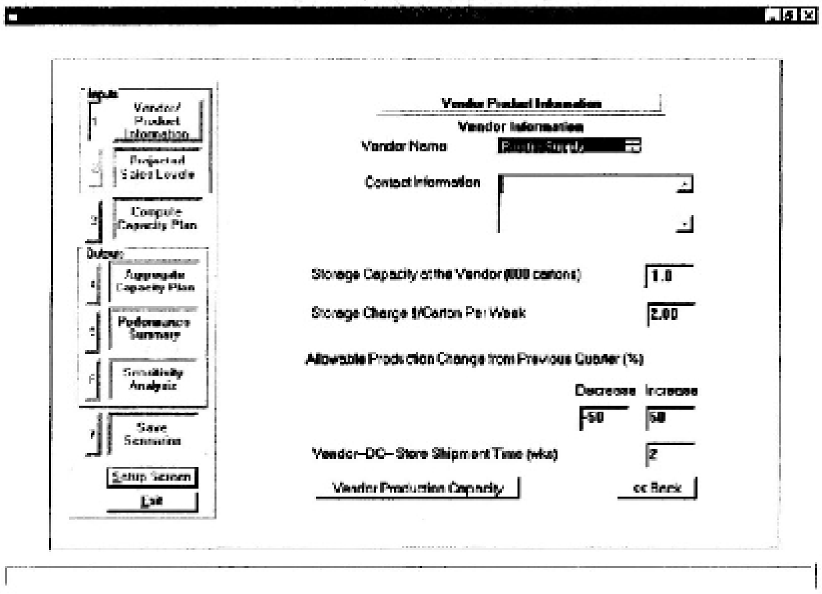

Figure 8.7. Vendor Information Screen

lead time dictate the relative flexibility of this vendor. The Vendor Production Capacity button allows access to a screen detailing each vendor's total production capacity by quarter, shown in Figure 8.8. The storage capacity and total quarterly production capacity are shared across all products made by this vendor.

Clicking a row heading button in the Figure 8.6 screen calls up the Product Information screen shown in Figure 8.9. Here each product's sales forecasts (for the "Most Likely" case, as described in Section 3.1), retail prices (week by week to accommodate frequent price changes if dictated by the retailer's promotional strategy), inventory costs and requirements, and fabric information are entered (or taken from a spreadsheet using an embedded interface accessible from the "Show Spreadsheet" button) and displayed. The inventory constraints and costs on this screen apply only to inventory held in the retailer's distribution system and stores.

The Projected Sales Levels button on the left-hand menu calls up the screen shown in Figure 8.10, which solicits the retail planners' beliefs about demand uncertainty. The "Most Likely" total season forecast for each product is automatically computed by summing the estimated weekly sales shown in Figure 8.9. The user specifies what a "Low"

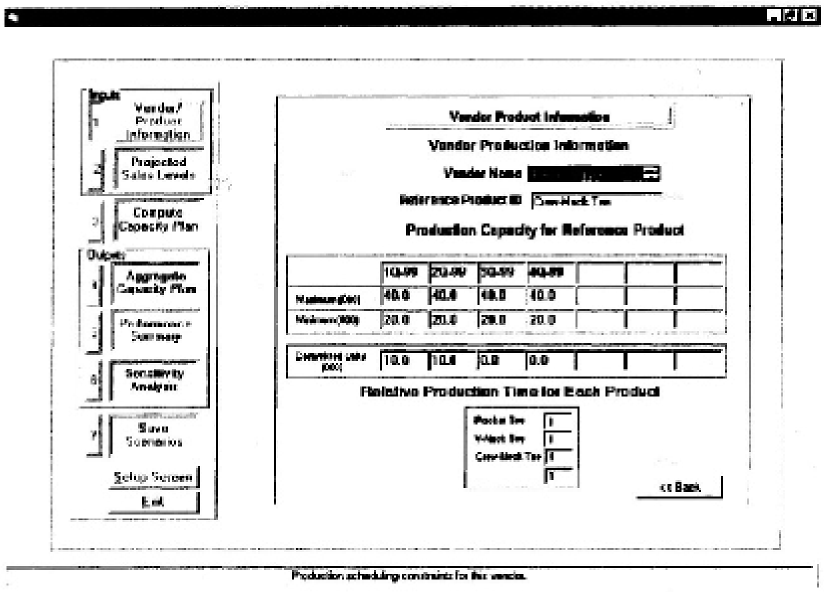

Figure 8.8. Vendor Production Information Screen

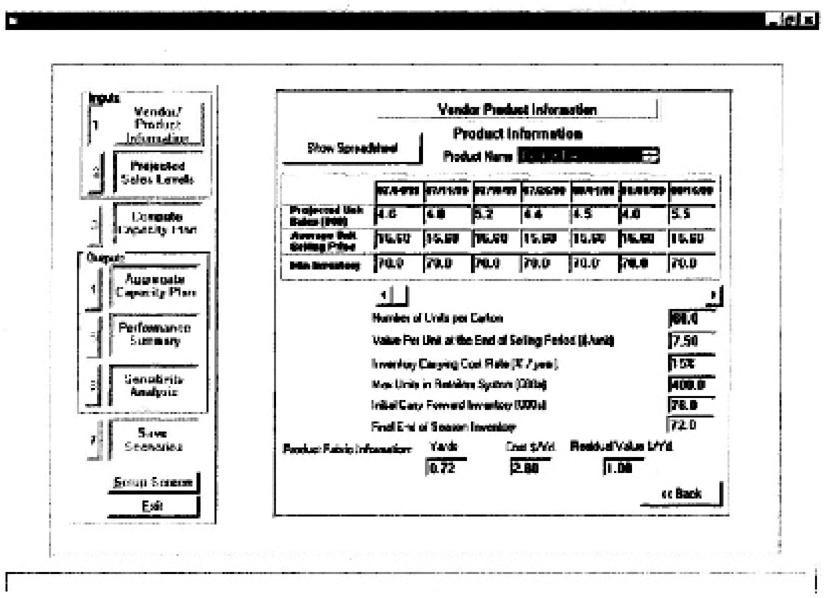

Figure 8.9. Product Information Screen

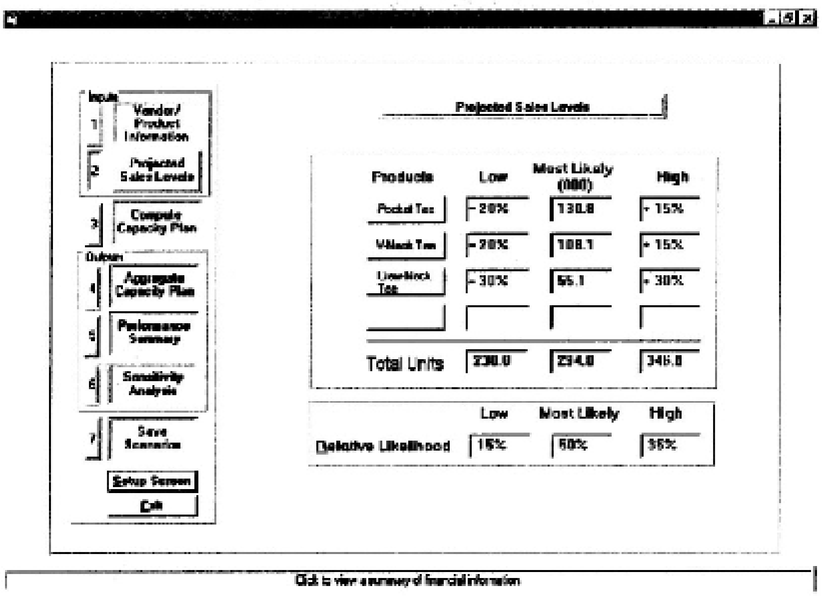

Figure 8.10. Projected Sales Levels Screen

and "High" forecast update would mean for each product in terms of a percentage deviation from the "Most Likely" volume. The percentage changes input here for a product capture the uncertainty about its demand, i.e., the extent to which the projections about that product's demand might change between t_0 and t_1. Stable products will tend to have more narrow ranges than newer or fashion-oriented products. These parameters are used to scale the weekly sales according to the demand model described in Section 3.4. At the bottom area of the screen the user must specify relative likelihoods for each of the three scenarios. After considerable discussion and experimentation, this input format was preferred by the sourcing managers because they are accustomed to developing strategies for three scenarios (cf. footnote 7).

Output Screens

A complete set of inputs allows the optimal sourcing plan to be determined by the LP solver engine. Since this plan contains considerable detail as well as contingency plans, the output is summarized across several screens. The main output screen is shown in Figure 8.11, which reports the total amount of each product that should be committed to each vendor under the three scenarios. (The buttons along the bottom

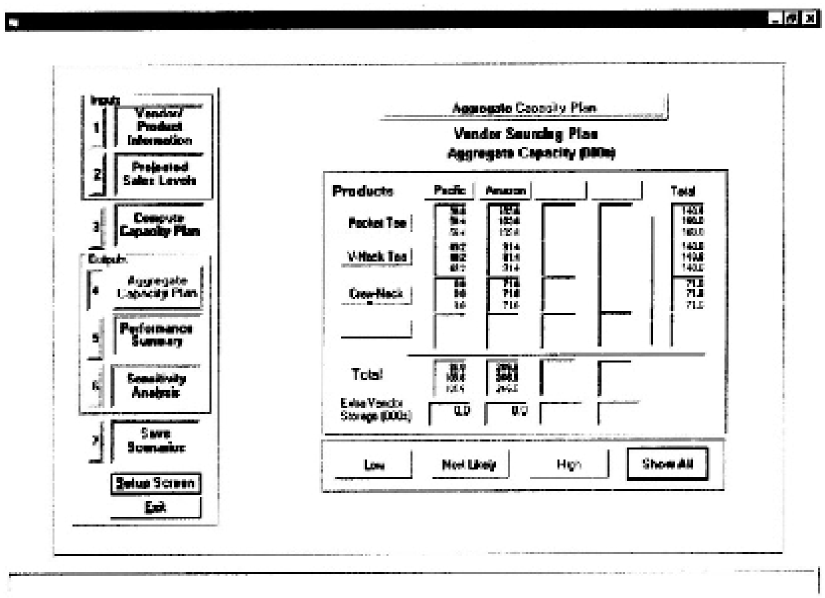

Figure 8.11. Aggregate Capacity Plan

of the screen allow the user to toggle through the plans for each individual scenario, or to juxtapose all three as shown in the figure. Each scenario button has a different color, which is used to display the corresponding plan in the "Show All" view.) By definition, long lead-time vendors must have the same commitments in all scenarios, while short lead-time vendors may receive commitments that depend on the scenario. This flexibility justifies any cost premium the latter vendors may charge. The summary provided by this screen can give each vendor a reasonable picture of how its total volume of business might vary. The buttons on the row and column headings are analogous to those in Figure 8.6, in that they provide paths to further details by vendor or product.

To eliminate any potential LP infeasibility due to the vendor storage constraints, the formulation was modified to allow unlimited auxiliary storage (at some very high price, to discourage the pursuit of this option). The Extra Vendor Storage cells at the bottom of the screen report the additional space (in thousands of cartons) required by the modified formulation's optimal plan. The sourcing managers considered this relaxation to be reasonable since, in spite of formally stated vendor storage limits, additional storage can almost always be obtained at some price.

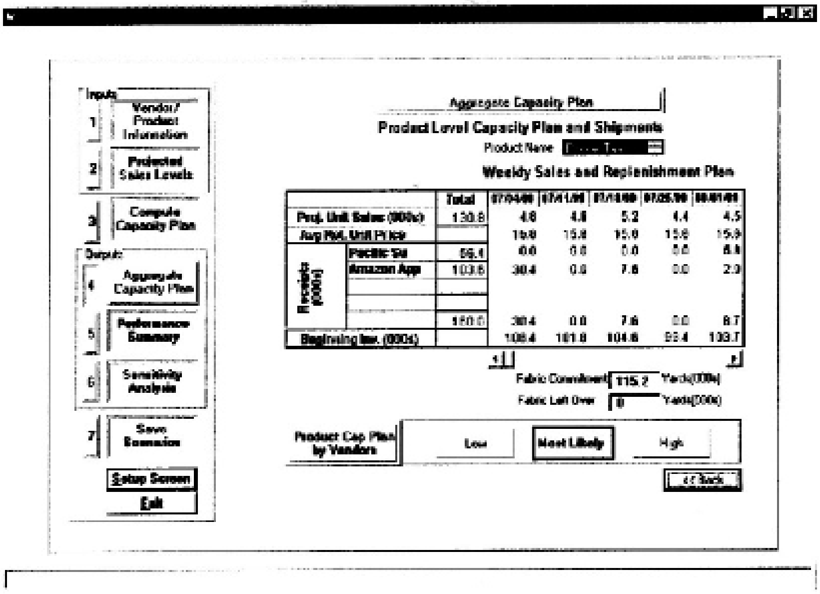

Figure 8.12. Product-Level Capacity Plan and Shipments Screen

The Product-Level Capacity Plan and Shipments Screen in Figure 8.12 shows week-by-week sourcing plans by product, including shipping receipts and fabric commitments for each scenario (displayed, as always, by toggling via the buttons near the screen's bottom). Figure 8.13 provides the vendor perspective on quarterly production and shipments for all products under each scenario. Barcharts are presented beneath the numerical table to visually illustrate the changes in vendor commitments for each product across the scenarios.

Figure 8.14 presents a summary of key retail performance metrics for each of the sales scenarios. GM represents gross margin dollars, inventory turnover is the total annualized sales divided by the average inventory level during the season, and GMROI (Gross Margin Return On Investment) is the product of the inventory turnover and the gross margin per unit (cf. Berman and Evans (1998)). Average cost is scenario-dependent because salvage revenues vary with the amount of residual fabric and end-of-season inventory. Again, the metrics take different values for each scenario, and the barcharts compare the values for each metric across the three scenarios.

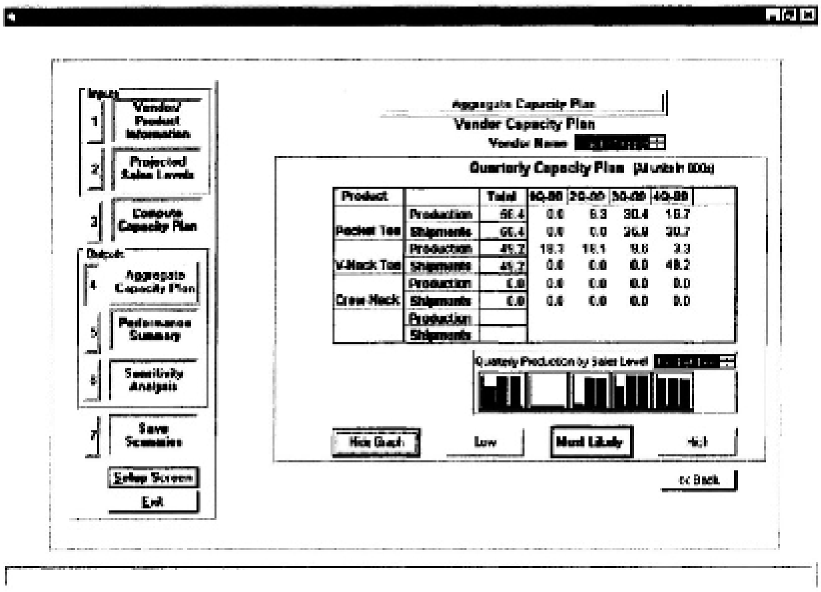

Figure 8.13. Vendor Capacity Plan Screen

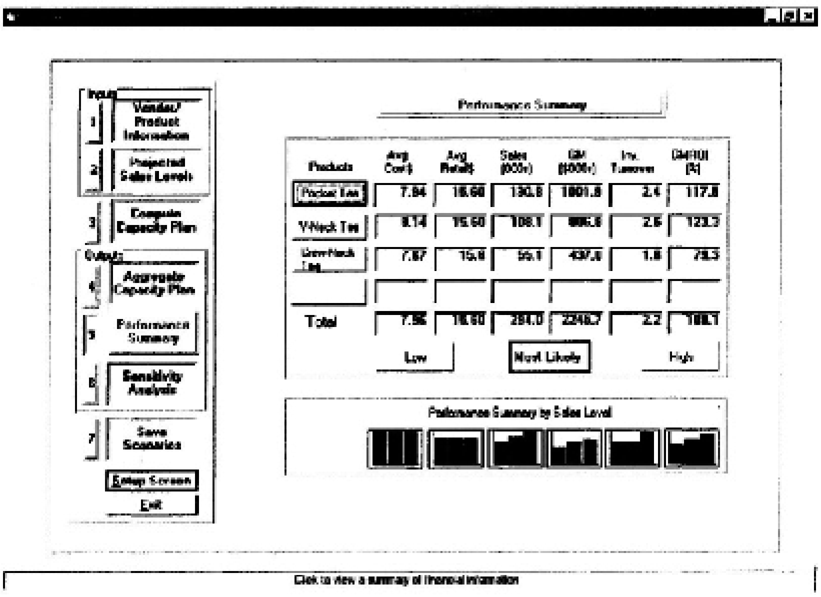

Figure 8.14. Performance Summary Screen

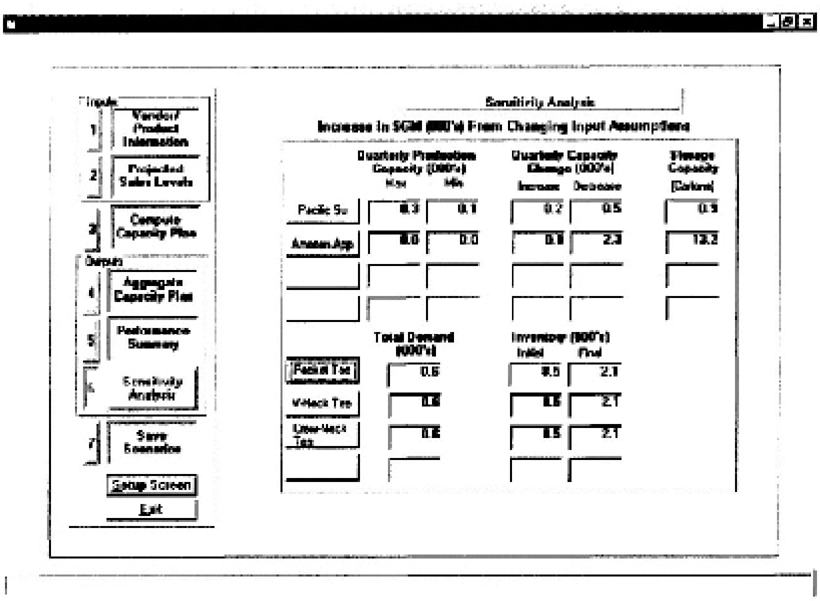

Figure 8.15. Sensitivity Analysis Screen

Sensitivity Analysis

Figure 8.15 shows the screen interface to the sensitivity analysis framework described in Section 3.7. Clicking on the vendor or product name buttons takes the user back to the appropriate input screen to view the current parameter value and make any desired modifications. This allows "What if?" analyses to be conducted quickly and easily for a wide range of assumptions. Added at the request of our retail sponsor, this screen has had the longest evolutionary path of any of the SAM screens and is considered to have the greatest strategic importance. By running the DSS, the retailer can determine which of each vendor's constraints are the most significant obstacles to improved profitability and can also estimate the value of relaxing any particular constraint. This can play a valuable role in negotiating with vendors by providing information that was not previously quantifiable.

Figure 8.16 shows the screen associated with the Save Scenarios button on the left menu. Because of the large number of input parameters, and to facilitate what-if analysis, SAM allows the capture of inputs and outputs for up to five problem scenarios. The user can annotate these in the text window near the bottom of the screen, where the various scenarios can be saved and retrieved.

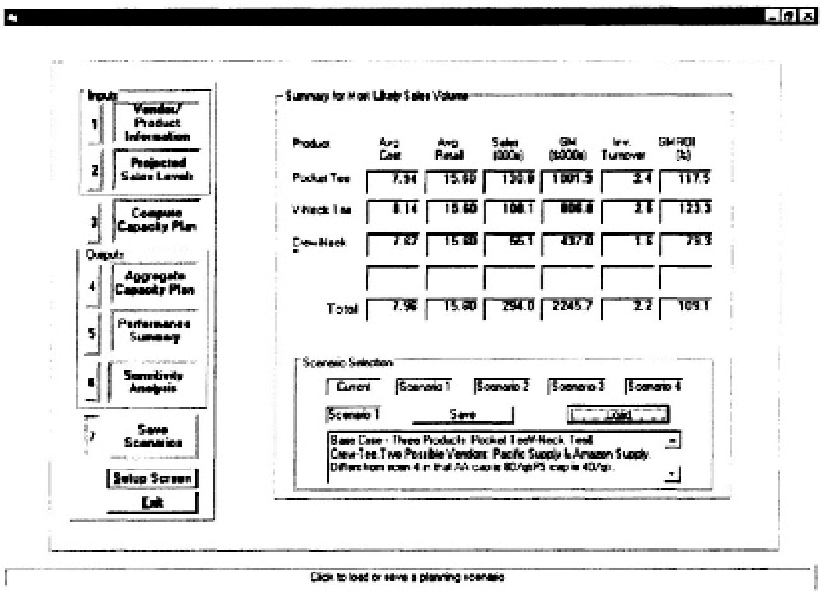

Figure 8.16. Save Scenarios Screen

4.2 User Experiences with the DSS

Experience with the SAM DSS was obtained through an analysis conducted with our sponsor firm, using representative but disguised data. The retailer's goal for this analysis was to gain experience with the model and develop an understanding of the key tradeoffs between vendor capabilities and unit costs. The details are presented as a case study in Agrawal et al. (2001). Some of the resulting insights are as follows:

1 Building a stochastic model allows presentation of distributional information about any system performance metric. This can provide valuable insight about the extent of intrinsic risk to which a decision-maker is exposed. The sourcing managers were able to connect this to their strategies for managing risk, including those related to the selection of product assortments. This would not be possible under any deterministic planning methodology.

2 Under reasonable assumptions, SAM's recommendations can improve expected profits by several percentage points relative to typical sourcing practices. Since net profit margins in retailing tend to be very small, this suggests that our methodology can offer very meaningful gains.

3 There can be value in using a portfolio of vendors with differing production flexibility. In practice, while buyers know at an intuitive level that flexibility has value, the inability to quantify this has left them biased toward vendors quoting the lowest unit costs. Our model easily demonstrates that additional vendor flexibility can indeed be worth a price premium when demand uncertainty exists, and we provide a means to evaluate this tradeoff with a realistic level of detail.

4 Not all production capacity is created equal. Capacity cannot be properly valued independently of its flexibility constraints, such as commitment lead time and allowable production change from quarter to quarter. The effect of these conditions is a function of the attributes of the type of merchandise, in particular the predictability of demand and the cost of obsolescence.

5 While conventional wisdom suggests that inventory turnover is determined by the replenishment policies adopted at the store level, tension between demand seasonality and the vendors' desire to maintain stable production schedules profoundly affects retailer inventory levels. Thus, efforts to increase turnover should also consider negotiations with vendors to seek greater production flexibility.

6 From an organizational point of view, our methodology can provide a vehicle for facilitating cross-functional communication and negotiation. Specifically, in a retail firm the merchandising, sourcing, and finance organizations typically have somewhat conflicting objectives with respect to inventory management strategy. (In mathematical terms, each group typically perceives a different segment of the overall objective function.) An early insight for us and our corporate sponsors was that our DSS could serve as a tool for brokering the concerns of these groups by solving the global optimization problem, explicitly quantifying tradeoffs, and, most importantly, defining a common vocabulary for discussion.

5. Conclusion

Estimating the value of adding or dropping a vendor, renegotiating the terms of a supply contract, or improving forecast capability requires the respecification of the production schedule in ways that may differ dramatically from past plans. The complexity of such decisions renders the subjective selection of optimal or even near-optimal plans extremely difficult or impossible. While many retail buyers and merchandise planners rely on extensive databases and query tools for decision support, there are few computer-based methods for optimal decision making or sensitivity analysis regarding these decisions. Our model and the associated decision support software provide retail planners with the power to identify and evaluate a wide variety of potential supply chain improvements that they are not currently able to consider.

Capturing market uncertainty through discrete scenarios is a familiar mechanism that simplifies the required user inputs and allows the application of linear programming optimization. Because of the many types of production and sales constraints that may apply in a retail environment, simplicity of use is essential to the practicality of a decision support tool. Tests of our model by buyers and planners within a major retail organization indicate that our framework is compatible with the production commitment decisions they face.

Acknowledgments

The authors are grateful to the corporate sponsors of the Retail Workbench at Santa Clara University for their financial support, for suggesting the business problem, and for their comments and suggestions in the course of this research. We are especially grateful to Professor Dale Achabal, Retail Workbench Director, who provided valuable guidance throughout the project, and to executives from the retail firm sponsoring this project for working with us in developing the modeling assumptions and designing the user screens for the software tool. The software owes much of its existence to Jerry Currie, who wrote and debugged the hundreds of lines of Visual Basic code that went into it. We would also like to acknowledge LINDO Systems for providing a special version of its LINGO optimization software for integration into our system. The authors are fully responsible for the opinions expressed and any remaining errors.

References

Achabal, D.D., S.H. McIntyre, and S.A. Smith (1990). Maximizing Profits from Periodic Department Store Promotions. *Journal of Retailing,* 66, 4, 383-407.

Agrawal, N., S.A. Smith, and A.A. Tsay (2001). Multi-Vendor Sourcing in a Retail Supply Chain. Working Paper, Leavey School of Business, Santa Clara University, Forthcoming, *Production & Operations Management.*

Berman, B. and J.R. Evans (1998). *Retail Management: A Strategic Approach,* Prentice Hall, Upper Saddle River, NJ.

Bertsekas, D.P. (1976). *Dynamic Programming and Stochastic Control,* Academic Press, New York, NY.

Bitran, G.R., E.A. Haas, and H. Matsuo (1986). Production Planning of Style Goods with High Setup Costs and Forecast Revisions. *Operations Research,* 34, 2, 226-236.

Brown, A.O. and H.L. Lee (1998). Optimal 'Pay to Delay' Capacity Reservation with Application to the Semi-conductor Industry. Working Paper, Owen Graduate School of Management, Vanderbilt University, Nashville, TN.

Buxey, G. (1993). Production Planning and Scheduling for Seasonal Demand. *International Journal of Operations & Production Management,* 13, 7, 4-21.

Buxey, G. (1995). A Managerial Perspective on Aggregate Planning. *International Journal of Production Economics,* 41, 1-3, 127-133.

Chang, S.H. and D.E. Fyffe (1971). Estimation of Forecast Errors for Seasonal Style-Goods Sales. *Management Science,* 18, 2, B89-B96.

Cox, J.C. and M. Rubinstein (1985). *Options Markets,* Prentice-Hall, Englewood Cliffs, NJ.

Crowston, W.B., W.H. Hausman, and W.R. Kampe (1973). Multistage Production For Stochastic Seasonal Demand. *Management Science,* 19, 8, 924-935.

Dantzig, G.B. (1955). Linear Programming Under Uncertainty. *Management Science,* 1, 3, 197-206.

Donohue, K.L. (2000). Efficient Supply Contracts for Fashion Goods with Forecast Updating and Two Production Modes. *Management Science,* 46, 11, 1397-1411.

Eppen, G.D. and A.V. Iyer (1997). Backup Agreements in Fashion Buying - The Value of Upstream Flexibility. *Management Science,* 43, 11, 1469-1484.

Eppen, G.D., R.K. Martin, and L. Schrage (1989). A Scenario Approach to Capacity Planning. *Operations Research,* 37, 4, 517-527.

Fisher, M. and A. Raman (1996). Reducing the Cost of Demand Uncertainty Through Accurate Response to Early Sales. *Operations Research,* 44, 1, 87-99.

Guerrero, H.H., K.R. Baker, and M.H. Southard (1986). The Dynamics of Hedging the Master Schedule. *International Journal of Production Research,* 24, 6, 1475-1483.

Gunther, H.O. (1982). A Comparison of Two Classes of Aggregate Production Planning Models Under Stochastic Demand. *Engineering Costs and Production Economics,* 6, 89-97.

Hansotia, B.J. (1980). Stochastic Linear Programming With Recourse: A Tutorial. *Decision Sciences,* 11, 1, 151-168.

Hartung, P.H. (1973). A Simple Style Goods Inventory Model. *Management Science,* 19, 12, 1452-1458.

Hausman, W.H. and R. Peterson (1972). Multiproduct Production Scheduling For Style Goods With Limited Capacity, Forecast Revisions And Terminal Delivery. *Management Science,* 18, 7, 370-383.

Hausman, W.H. and R. Sides (1973). Mail-order Demands For Style Goods - Theory And Data Analysis. *Management Science,* 20, 2, 191-202.

Heath, D.C. and P.L. Jackson (1994). Modeling The Evolution Of Demand Forecasts With Application To Safety Stock Analysis In Production/Distribution Systems. *IIE Transactions,* 26, 3, 17-30.

Hertz, D.B. and K.H. Schaffir (1960). A Forecasting Method for Management of Seasonal Style-Goods Inventories. *Operations Research,* 8, 1, 45-52.

Huchzermeier, A. and M. A. Cohen (1996). Valuing Operational Flexibility Under Exchange Rate Risk. *Operations Research,* 44, 1, 100-113.

Hunter, N.A., R.E. King, and H.L.W. Nuttle (1992). An Apparel-supply System for QR Retailing. *Journal of the Textiles Institute,* 83, 3, 462-471.

Hunter, N.A., R.E. King, and H.L.W. Nuttle (1996). Evaluation of Traditional and Quick-response Retailing Procedures by Using a Stochastic Simulation Model. *Journal of the Textiles Institute,* 87, 1, 42-55.

Infanger, G. (1994). *Planning Under Uncertainty: Solving Large-Scale Stochastic Linear Programs,* Boyd & Fraser, Danvers, MA.

Kalyanam, K. (1996). Pricing Decisions Under Demand Uncertainty: A Bayesian Mixture Model Approach. *Marketing Science,* 15, 3, 207-221.

King, R.E. and N.A. Hunter (1996). Demand Re-estimation and Inventory Replenishment of Basic Apparel in a Specialty Retail Chain. *Journal of the Textiles Institute,* 87, 1, 31-41.

Kira, D., M. Kusy, and I. Rakita (1997). A Stochastic Linear Programming Approach to Hierarchical Production Planning. *Journal of the Operational Research Society,* 48, 2, 207-211.

Miller, J.G. (1979). Hedging the Master Schedule. In *Disaggregation Problems In Manufacturing And Service Organizations.* L.P. Ritzman, et al. (Eds.) Martinus Nifhoff, Boston, MA, 237-256.

Murray, G.R. and E.A. Silver (1966). A Bayesian Analysis of the Style Goods Inventory Problem. *Management Science,* 12, 11, 785-797.

Nam, S. and R. Logendran (1992). Aggregate Production Planning–A Survey of Models and Methodologies. *European Journal of Operational Research,* 61, 3, 255-272.

Nuttle, H.L.W, R.E. King, and N.A. Hunter (1991). A Stochastic Model of the Apparel Retailing Process for Seasonal Apparel. *Journal of the Textiles Institute,* 82, 2, 247-259.

Ravindran, A. (1972). Management of Seasonal Style-Goods Inventory. *Operations Research,* 20, 2, 265-275.

Riggs, W.E. (1984). A Short-Term Forecasting Model for Producers of Seasonal Style Goods. *Production & Inventory Management,* 25, 4, 42-49.

Riter, C. (1967). The Merchandising Decision Under Uncertainty. *Journal of Marketing,* 31, 44-47.

Robison, L.J. and P.J. Barry (1987). *The Competitive Firm's Response to Risk,* Macmillan, New York, NY.

Silver, E.A. and R. Peterson (1985). *Decision Systems for Inventory Management and Production Planning,* 2nd Edition, Wiley, New York.

Smith, S.A. and D.D. Achabal (1998). Clearance Pricing and Inventory Policies for Retail Chains. *Management Science,* 44, 3, 285-300.

Smith, S.A., N. Agrawal, and S.H. McIntyre (1998). A Discrete Optimization Model for Seasonal Merchandise Planning. *Journal of Retailing,* 74, 2, 193-221.

Wadsworth, G.P. (1959). In *Notes on Operations Research,* assembled by the Operations Research Center, M.I.T., Technology Press, Cambridge, MA.

Wolfe, H.B. (1968). A Model for Control of Style Merchandise. *Industrial Management Review,* 9, 2, 69-82.