Design of component-supply contract with commitment-revision flexibility

by Y. Bassok
A. Bixby
R. Srinivasan
H. Z. Wiesel

In this paper, we study a type of supply contract that is frequently used in the electronics industry. A common feature of these supply contracts is that at the beginning of the contract, the buyer makes purchasing commitments to the supplier for each period. The buyer may have some flexibility to purchase quantities that actually deviate from the original commitments. Moreover, as time passes and more information about the actual demand is collected, the buyer may update the previous commitments, in a way that is described. We develop a heuristic that is easy to implement and that determines nearly optimal commitments and purchasing quantities. We show that in many cases of practical interest, the heuristic results in solutions that are close to the optimal.

Introduction

Considering the very competitive business environment, the executives of the IBM Printer System Company recognized that in order to maintain profitability, significant improvements must be made throughout the business. Focus was placed on reduced cycle times, marketing, brand management, and fulfillment. (Brand management deals with the formulation and execution of strategies related to product offering, pricing, etc. Fulfillment represents the coordination of internal production and distribution units and external suppliers to satisfy customers' demands for products.) The fulfillment process and the related asset-management activities became cornerstone initiatives and prime candidates for improvement. Focusing on these areas produces significantly lower inventory levels while maintaining responsiveness to customers and achieving an improved cash position for the company. Although many of the improvements were the result of "traditional" asset-management techniques, many unique techniques were applied as well. One of these is an analytical technique that determines the optimal terms and conditions for supply contracts. We describe this technique in the following.

The relationship with suppliers is emerging as a key determinant to the competitiveness of production enterprises like IBM, which manufacture products with purchased materials that range from basic raw materials to complex assemblies. This fact has been repeatedly demonstrated in the last few years in the arena of printers and personal computers, where more than 50% to 70% of the total manufacturing cost of a personal computer is added by the suppliers. Another aspect of the relationship with suppliers that requires consideration is the dynamic

0018-8646/97/$5.00 © 1997 IBM

and uncertain nature of demand for the components and subassemblies provided by the suppliers. Rapid advances in hardware technology and computer applications and the unpredictability of customer interest in these advances affect the demand. For instance, the sale of personal computer units in 1993 by some of the biggest manufacturers in the industry increased by as much as 35% to 140% over that in 1992 [1]. As a result of the sharp increase in demand, these manufacturers were severely constrained by the supply of some of the components, resulting in higher production costs.

In order that the products be cost-competitive in such a dynamic environment, the relationship between the end-product manufacturer, hereafter referred to as the buyer, and the component supplier has to be financially rewarding to both in the long run. The situation is similar in the consumer electronics and automobile industries. This relationship should begin at the time the end product is conceived by the designer and continue until the after-sale service of the product—in short, over the life cycle of the product. The stages in between are product development, product engineering, volume planning/forecasting, supply contracting, parts ordering, parts delivery, manufacturing, distribution, and sale.

A crucial aspect of this interaction, which is the subject of this study, is supply contracting. Key elements of a typical supply contract include a) capacity, the resources the supplier agrees to dedicate to the manufacture of the buyer's components—that is, dedication of manufacturing lines; b) flexibility, the fluctuation in the buyer's demand that the supplier is willing to accommodate at no extra cost (a higher price may be charged for increased flexibility); c) liability, the cost the buyer incurs in the event the purchase is canceled or the volume is below the agreed-upon limit; d) quality, the upper limit on the percentage of defective units in the supply that is acceptable to the buyer; e) warranty, the terms under which defective components will be compensated for by the supplier during a specified period following the shipment from the supplier; and f) currency-exchange fluctuations, the terms under which the fiscal effects of fluctuations in exchange rates will be shared by the buyer and the supplier.

The issue focused upon in this paper is flexibility. (Liability and capacity are implicit in the commitment and flexibility agreed upon by the buyer and supplier.) We discuss two types of flexibility: the flexibility to update previous commitments and the flexibility to purchase quantities that deviate from previously made commitments. The trade-offs for the buyer and the supplier are clear: the supplier usually incurs additional cost by providing flexibility, because of the need for additional production capacity and/or inventory of raw materials and finished components to meet the buyer's uncertain demand in a timely manner. Owing to the higher flexibility, the buyer is able to meet, with greater probability, the uncertain demand and thus realize savings. It is important to notice that although providing flexibility is costly for the supplier, the benefits for the buyer may be high enough that the total system (supplier and buyer) realizes savings. These savings may be allocated between the supplier and the buyer so that both benefit and realize savings due to the high flexibility in the system. As a means of allocating the savings due to flexibility, the supplier may be willing to offer different levels of flexibility at appropriate costs to the buyer. The buyer, therefore, has to choose the level of flexibility consistent with the buyer's objectives, in order to minimize cost and provide a certain level of service to the marketplace. In the case of electronic components such as memory and logic chips, it is customary in the industry to have a single contract cover several components of the same technology produced by the same supplier. The supply flexibility in this case is typically stated for both the total volume for the technology as well as the volumes of the individual components. The scope of our work in this paper is limited to considering flexibility and liability in a contract that deals with a single component.

There are numerous studies in the area of inventory control (e.g., Veinott [1]) that analyze and determine optimal or efficient purchasing (or production) policies. Common to most of these studies is the notion that purchasing decisions are performed dynamically—that is, at the beginning of each planning period or interval (periodic review), or when a certain event happens (continuous review). In contrast to these studies, we assume that some, but not necessarily all, decisions are performed and negotiated between the buyer and the supplier at the beginning of the horizon (the contract duration). Those decisions have a major impact on the dynamic decisions. For example, flexibility limits and dedication of capacity are determined at the beginning of the horizon and affect the dynamic decisions of the actual purchasing quantity.

Models in the production-management literature that utilize volume-based price discounts may be viewed as special cases of supply contracts. Most of the studies in this area assume a deterministic environment with known and fixed demand rate. It is well known that in this environment, the economic order quantity (EOQ) is the optimal purchasing quantity, minimizing the buyer's cost (see Tersine [2]). The EOQ might not be optimal from the point of view of the supplier or the system (the total costs of the buyer and the supplier). In these cases, it might be optimal for the supplier to provide the buyer with a discount to induce the buyer to order quantities that differ from the EOQ but are closer to the quantities optimal for the supplier to deliver. The discount may result in a situation in which both the buyer and the supplier benefit and reduce their costs compared with the EOQ solution. Lee and Rosenblatt [3] and Rosenblatt and Lee [4] are examples of this line of research.

Sadrian and Yoon [5] deal with a supply-management problem with deterministic demand and multiple suppliers, in which discounts are a function of the volume purchased. The main objective is to determine the set of suppliers from whom to purchase components, so that the total purchasing cost is minimized.

Jucker and Rosenblatt [6], assuming a single-period problem with random demand, study different types of price discounts. Bassok and Anupindi[1] assume a finite-horizon problem with random demands, in which the buyer makes a commitment at the beginning of the horizon to purchase a minimum quantity. The supplier, in return for this commitment, provides the buyer with a discount. They present an algorithm to determine the optimal purchasing quantities for each period, on the basis of the original commitment, and determine the optimal commitment. Anupindi and Bassok [7] present a similar problem with a minimum-budget commitment. Here, at the beginning of the horizon, the buyer makes a commitment to purchase components of a technology for a minimum dollar volume. Bassok and Anupindi [8] present and solve a model in which, at the beginning of the horizon, the buyer provides the supplier with commitments for each period. The buyer has the flexibility to purchase quantities that deviate somewhat from the original commitment. They show that their solution is open-loop-feedback-control optimal. Computational time to determine the optimal policy is relatively extensive. In contrast, we concentrate in this paper on developing a heuristic that is easy to implement and understand and requires a short time to compute the nearly optimal policy. Tsai and Lovejoy [9] present a supply contract similar to the one presented here. The main difference is that they are concerned with a multi-echelon environment. They define two types of supply nodes (or production stages): a flex node and a semiflex node. A flex node deals with the internal demand from a downstream production stage, while a semiflex node deals with external demand. Their main interest is in the flex nodes, while ours is in what they define as a semiflex node.

In this paper, we consider the problem from the point of view of the buyer, who determines his optimal purchasing and commitment policy on the basis of the terms of the contract provided by the supplier. We do not consider the cost of the flexibility to the supplier, although a richer model would also consider this cost and determine the equilibrium for the entire system. Such an approach is beyond the scope of this paper. A good example of such a model can be found in Moses and Seshadri [10].

One of the objectives of this study is to develop a method for enabling the buyer to compare different contracts with different terms. By solving a large number of contracts with different terms, we are able to construct break-even curves that indicate to the buyer when a contract is acceptable and when it is not. Also, because the computation time of the solution procedure that we present is very small, we are able, if necessary, to compare different contracts and to choose the optimal one.

The rest of the paper is organized in the following way: In the following section, we present the main assumptions and notations. In the next section, we present the basic model and solution methods. In the last major section, we present computational results.

Assumptions and notation

We are now ready to describe the problem that we discuss in this paper in detail: a supply contract problem with periodic commitments and update flexibility (PCUF). At the beginning of the planning horizon (the length of the horizon is usually a year), the buyer makes purchasing commitments to the supplier for every period (usually a month) in the horizon. The actual purchasing quantities and the commitments may be modified as time passes and more information about the demand is collected. The fashion in which the purchasing quantities and the commitments are modified will be discussed shortly. The objective of this study is to determine "good" initial commitments, purchasing quantities for each period, and a mechanism to update the commitments.

- *Horizon*

We assume a finite horizon comprising T periods. Although the periods may not have identical durations, they are considered identical here, with no loss of generality. (**Table 1** presents the notation used in this paper.)

- *Demands*

For a component, the demand in each period is a random variable with known probability density function, and the probability density functions of the different periods are not necessarily identical. The demand in each period is independent of the demands in the other periods. Of course, the demand for a component depends upon the demand for the buyer's products that include the component.

Unsatisfied demand in one period is carried forward to the next period as "backlog."

[1] Y. Bassok and R. Anupindi, "Analysis of Supply Contracts with Forecasts and Flexibility," Northwestern University, March 1995. Submitted to *Management Science* for publication.

Table 1 Notation for variables and functions.

T	Number of periods in the horizon.
c	Purchasing cost per component unit.
h	Holding cost per unit (cost incurred by the buyer in carrying one unit of the component for one time period).
p	Shortage cost per unit (cost incurred by the buyer for being short by one unit of the component).
x_t	Initial component inventory at period t.
Q_t	Purchasing quantity for period t.
Q_t^*	Optimal purchasing quantity for period t.
q_t^i	Commitment made at the beginning of period i.
$q{*}_t^i$	Commitment made at the beginning of period i for period t, according to the heuristic.
S_t	Base-stock level obtained by solving the news-vendor problem.
α	Purchasing flexibility.
β	Update flexibility.
$f(\cdot)$	Probability density function of the demand. (This may vary from period to period; we omit the subscript used to designate the period, however, for the sake of simplicity of notation.)
D_t	Demand at period t. (Although it is a discrete random variable, we treat it as a continuous random variable.)
$L(\cdot)$	Expected holding and shortage cost, defined as $L = h \int_0^{x_t+Q_t} (x_t + Q_t - D_t) f(D_t)\, dD_t + p \int_{x_t+Q_t}^{\infty} [D_t - (x_t + Q_t)] f(D_t)\, dD_t$.

$C_t(Q_t, q_{t+1}^t, \cdots, q_T^t | x_t, q_t^{t-1}, \cdots, q_T^{t-1}$ is the expected total cost for periods t through T, under the assumption that the initial inventory in period t is x_t and the commitments made at period $t-1$ for periods t through T are $q_t^{t-1}, \cdots, q_T^{t-1}$.

$C_t^*(x_t, q_t^{t-1}, \cdots, q_T^{t-1})$ is the optimal expected total cost for periods t through T, under the assumption that the initial inventory in period t is x_t and the commitments made at period $t-1$ for periods t through T are $q_t^{t-1}, \cdots, q_T^{t-1}$.

- *Costs*

The following costs are considered:

- Purchasing cost, proportional to the quantity purchased.
- Inventory-holding cost, proportional to the quantity carried.
- Shortage cost, proportional to the unsatisfied demand.

- *Sequence of events*

1. At the beginning of the horizon, the buyer makes a multiperiod commitment to the supplier for the entire horizon, $q_1^1, \cdots, q_T^1$, where q_t^i is the commitment made at the beginning of period i for period t; i.e., the buyer agrees to purchase q_t^i units during period t. This may be subsequently adjusted according to the flexibility permitted, as described below.
2. At the beginning of period t, the buyer observes the inventory on hand and the orders for finished products and then makes the following two decisions:
   - The purchasing quantity for period t, Q_t, which must satisfy the constraint
   $$(1-\alpha)q_t^{t-1} \le Q_t \le (1+\alpha)q_t^{t-1} \qquad 0 \le \alpha < 1,$$
   where α is the purchasing flexibility ≥ 0.
   - The commitments for future periods, which must satisfy the set of constraints
   $$(1-\beta)q_{t+i}^{t-1} \le q_{t+i}^t \le (1+\beta)q_{t+i}^{t-1} \qquad 0 \le \beta < 1, \quad i = 1, \cdots, T-t,$$
   where β is the update flexibility, defined similarly to α.
3. Components are delivered instantaneously.
4. Excess inventory (components remaining after production) is stored and used in the next period, and excess demand is backlogged. The following equation describes excess inventory:
   $$x_t = x_{t-1} + Q_{t-1} - D_{t-1}.$$
   When $x_t > 0$, there is excess inventory; when $x_t < 0$, there is demand backlog.

It is important to observe that the PCUF problem has a different structure than the well-known finite-horizon "news-vendor" problem (see Karlin [11]). In solving the news-vendor problem, the purchasing decisions are made sequentially, which means that at the beginning of every period, the decision maker observes the state of the system (inventory on hand and in the pipeline and orders for finished products) and then makes a purchasing decision. Solving the PCUF problem, the decision maker first makes a static decision (that is, determines the optimal commitments for each period of the horizon) and makes a purchasing decision for the first period. In addition, the decision maker may dynamically update the commitments. Note that when the flexibility provided by the supplier is infinite ($\alpha = \beta = \infty$), the update process is not constrained, and the two problems are identical. On the other hand, when $\alpha = \beta = 0$—i.e., there are no purchasing and update flexibilities—all decisions are static and are made at the beginning of the horizon.

We assume instantaneous delivery of components. This is not a limitation, since it is well known (see Karlin [11]) that a problem with known lead time and back-ordering of excess demand can be transformed (by changing the demand distribution) into an equivalent problem with zero lead time.

It is important to remember that as time passes and more information is collected, the buyer has the flexibility to change his or her previously made commitments, but this flexibility is limited. The magnitude of the flexibility is greater for periods further away, because the further away the period is, the more times the commitments may be updated. This captures the notion that once the supplier begins to produce the component, the supplier can offer very limited flexibility. On the other hand, the supplier has a considerably greater flexibility dealing with orders to be delivered in the distant future. In many instances, suppliers are unable or unwilling to provide very large long-term flexibility; in such instances, they limit the sum of the changes that may be made by the buyer.

Model and solution method

It is possible to formulate the problem by means of the following recursion, which is a standard recursion for inventory-optimization models (see Karlin [11]):

$$C_1(Q_1, q_2^1, \cdots, q_T^1 | x_1)$$
$$= cQ_1 + L(x_1 + Q_1) + C_2(x_2, q_2^1, \cdots, q_T^1) \quad (1a)$$

and, in general,

$$C_{t-1}(Q_{t-1}, q_t^{t-1}, \cdots, q_T^{t-1} | x_{t-1})$$
$$= cQ_{t-1} + L(x_{t-1} + Q_{t-1}) + C_t(x_t, q_t^{t-1}, \cdots, q_T^{t-1})$$
$$t = 2, \cdots, T. \quad (1b)$$

The optimal expected total cost is

$$C_t^*(x_t, q_t^{t-1}, \cdots, q_T^{t-1})$$
$$\equiv \min_{Q_t, q_{t+1}^t, \cdots, q_T^t} C_t(Q_t, q_{t+1}^t, \cdots, q_T^t | x_t, q_t^{t-1}, \cdots, q_T^{t-1})$$
$$t = 1, \cdots, T, \quad (2)$$

where c, L, C_t, and C_t^* are as described in Table 1.

We are now in a position to define problem P1:

Minimize $C_1(Q_1, q_1^1, \cdots, q_T^1 | x_1)$ so that

$$(1 - \alpha)q_t^{t-1} \le Q_t \le (1 + \alpha)q_t^{t-1} \qquad t = 2, \cdots, T \quad (3)$$

and

$$(1 - \beta)q_{t+i}^{t-1} \le q_{t+i}^t \le (1 + \beta)q_{t+i}^{t-1}$$
$$t = 2, \cdots, T - 1, \quad i = 1, \cdots, T - t. \quad (4)$$

Constraint (3) requires that the actual purchasing quantity for period t be within the flexibility bounds. Notice that at period t, the commitment q_t^{t-1} is known, since it was determined at period $t - 1$. Similarly, constraints (4) require that the commitment updates be within the update flexibility constraints. Note that q_t^i, in general, is clearly a function of x_i, the inventory on hand at the start of period i. A more accurate notation would be $q_t^i(x_i)$. For simplicity, we write q_t^i instead of $q_t^i(x_i)$.

The above model assumes that the flexibilities are the same for all periods—i.e., that α and β do not change. However, this is done only in the interest of simplicity and clarity. The solution method presented later can be modified to account for nonconstant flexibilities.

Observe that the number of decision variables at period t is equal to $T + 1 - t$. For example, at period 1, the decisions are the optimal purchasing quantity for period 1 and the commitments for periods 2 to T. When the number of decision and state variables is large, obtaining the optimal solution with the use of standard dynamic programming is not feasible. In what follows, we present a heuristic solution to problem P1. We first develop lower and upper bounds on the total cost and then show that the gap between the upper and lower bounds is small enough to justify the use of the upper bound as an approximate solution to problem P1.

- *Lower bound*

The optimal expected total cost in problem P1 decreases as the purchasing and update flexibilities increase, because constraints (3) and (4) become more relaxed. When flexibilities are large (approaching ∞), we actually face a standard, finite-horizon news-vendor problem; thus, the expected total cost obtained by solving the news-vendor problem is a lower bound on the cost obtained by solving problem P1.

Noting that a "base-stock" policy, in which the inventory is raised to the base-stock level S_t at every period, is known to be optimal for the news-vendor problem (for details, see Karlin [11]), we use the structure of the base-stock policy in designing the proposed heuristic solution to problem P1.

- *Upper bound*

A feasible upper bound is obtained by our heuristic approach. The basic idea of the heuristic is to choose feasible purchasing quantities and commitments so that the probability of being able to raise the inventory to the base-stock levels S_t (obtained by solving the news-vendor problem) is maximized.

For example, if the initial inventory at the first period is S_1, the maximum possible inventory level at period 2 is $S_1 - D_1 + (1 + \alpha)q_2^1$, and the minimum possible inventory level is $S_1 - D_1 + (1 - \alpha)q_2^1$. Since our objective at period 2 is to raise the inventory level to S_2, we choose the commitments for period 2 to maximize the probability

$$\Pr[S_1 - D_1 + (1 + \alpha)q_2^1 \ge S_2 \ge S_1 - D_1 + (1 - \alpha)q_2^1] \quad (5)$$

that S_2 is within the above range. In general, the same arguments lead to determining the commitment at period 1 for period t by maximizing the probability

$$\Pr[S_1 - (D_1 + \cdots + D_{t-1}) + (1 + \alpha)q_2^1 + \cdots + (1 + \alpha)q_t^1$$
$$\ge S_t \ge S_1 - (D_1 + \cdots + D_{t-1}) + (1 - \alpha)q_2^1 + \cdots + (1 - \alpha)q_t^1].$$

Observe that at the beginning of period 1, it is necessary to determine the $T - 1$ period commitments, which we do by sequentially maximizing the $T - 1$ expressions for periods $2, \cdots, T$. We begin by determining q^{*1}_2, according to the maximum of Expression (5), and using this value to determine q^{*1}_3. In general, we use $q^{*1}_2, \cdots, q^{*1}_{t-1}$ to determine q^{*1}_t. It should be noted that we use the symbol q^{*i}_t to represent the commitment according to the heuristic—not the actual optimal commitment.

At period t, we repeat the same process by solving the following constrained problem P2. The optimal purchasing quantity is determined by solving the following problem:

$$\text{Minimize } Q_t + x_t - S_t \text{ over } Q_t \tag{6}$$

such that

$$(1-\alpha)q^{*t-1}_{t} \le Q_t \le (1+\alpha)q^{*t-1}_{t}. \tag{7}$$

Observe that the objective is to purchase a quantity Q_t that will bring the initial inventory x_t as close as possible to the base-stock level S_t. At period t, the commitments for periods $t+1, \cdots, T$ are updated by solving the $T-t$ problems; i.e., maximize the following over q^{t}_{t+k}:

$$\begin{aligned}\Pr\Bigg[&S_1 - \sum_{i=1}^{t-1} d_i + \sum_{i=1}^{t} Q_i^* - \sum_{i=t}^{t+k-1} D_i + \sum_{i=1}^{k-1}(1+\alpha)q^{*t}_{t+1} \\ &+ (1+\alpha)q^{t}_{t+k} \ge S_{t+k} \ge S_1 - \sum_{i=1}^{t-1} d_i + \sum_{i=1}^{t} Q_i^* \\ &- \sum_{i=t}^{t+k-1} D_i + \sum_{i=1}^{k-1}(1-\alpha)q^{*t}_{t+i} + (1-\alpha)q^{t}_{t+k}\Bigg]\\ &k = 1, \cdots, T-t \end{aligned} \tag{8}$$

such that

$$(1-\beta)q^{*t-1}_{t+k} \le q^{t}_{t+k} \le (1+\beta)q^{*t-1}_{t+k}. \tag{9}$$

Here, Q^* denotes the optimal purchasing quantity determined by the heuristic. Note that constraints (7) and (9) ensure that the purchasing quantity and commitment updates are feasible; thus, the solution obtained by solving problem P2 is a feasible solution to problem P1. As a result, the cost obtained by solving P2 is an upper bound on the cost obtained by solving the original problem, P1.

The above heuristic is clearly not optimal. It does capture the dynamic nature of the problem by maximizing the probability of reaching the base-stock levels that are optimal for the news-vendor problem and provides a mechanism to determine the static commitments. Morton [12] uses a similar approach to solve finite-horizon, nonstationary inventory problems.

We can rewrite Expression (5) as

$$\Pr[S_1 - S_2 + uq^1_2 \ge D_1 \ge S_1 - S_2 + lq^1_2] = \int_{S_1-S_2+lq^1_2}^{S_1-S_2+uq^1_2} f(D_1)\,dD_1,$$

where f is the probability density function of the demand D_1, $u \equiv 1+\alpha$, and $l \equiv \max(0, 1-\alpha)$.

Then problem P2 can be expressed as follows:

$$\text{maximize} \int_{S_1-S_2+lq^1_2}^{S_1-S_2+uq^1_2} f(D_1)\,dD_1$$

over q^1_2 so that $q^1_2 \ge 0$.

If we assume that the demands are normally distributed, with mean μ_1 and standard deviation σ_1, problem P2 can be expressed as maximizing the following over q^1_2:

$$\frac{1}{\sqrt{2\pi\sigma^2}} \int_{S_1-S_2+lq^1_2}^{S_1-S_2+uq^1_2} e^{-\frac{(D_1-\mu_1)^2}{2\sigma_1^2}}\,dD_1, \tag{11}$$

such that $q^1_2 \ge 0$.

It is easy to show that the probability integral in expression (11), a function of q^1_2, is not unimodal, which means that potentially problem P2 has more than one optimum. It will be shown shortly that the cost function has two stationary points: one for $q^1_2 < 0$, which is a minimum and clearly not feasible, and the other for $q^1_2 > 0$, which is a maximum. As a result, in the feasible region, the cost function is unimodal, and problem P2 has a unique feasible solution.

Taking the derivative of the cost function with respect to q^1_2 and equating it to zero, we obtain

$$ue^{-\frac{(S_1-S_2+uq^1_2-\mu_1)^2}{2\sigma_1^2}} - le^{-\frac{(S_1-S_2+lq^1_2-\mu_1)^2}{2\sigma_1^2}} = 0, \tag{12}$$

which is equivalent to

$$\log\left(\frac{u}{l}\right) = \frac{(S_1-S_2+uq^1_2-\mu_1)^2}{2\sigma_1^2} - \frac{(S_1-S_2+lq^1_2-\mu_1)^2}{2\sigma_1^2}. \tag{13}$$

Equation (13) is a quadratic function and has at the most two solutions. It is trivial to show that only one of these solutions is positive and that the cost function has a maximum at q^{*1}_2. Thus, problem P2 has a unique feasible and optimal solution. To determine the optimal commitments for periods $2, \cdots, T$, we proceed sequentially, as described above. For every period, we need to solve only a quadratic equation. At period 1, the commitment for period t is determined by solving the following equation:

$$\begin{aligned}\log\left(\frac{u}{l}\right) &= \frac{[S_1 - S_t + u(q^{*1}_2, \cdots, q^{*1}_{t-1}, q^1_t) - (t-1)\mu]^2}{2(t-1)\sigma^2} \\ &- \frac{[S_1 - S_t + l(q^{*1}_2, \cdots, q^{*1}_{t-1}, q^1_t) - (t-1)\mu]^2}{2(t-1)\sigma^2}.\end{aligned} \tag{14}$$

For ease of notation, the expression is written for stationary normal distributions; however, it could easily be written for the nonstationary case in which μ_i and σ_i vary from period to period. As before, Equation (14) is a quadratic equation, and all arguments made for Equation (13) hold here too.

Table 2 Comparison of approximate solutions TC^H and TC^N to problem P1.

σ	$\alpha=\beta$								
	0.05	*0.1*	*0.2*	*0.05*	*0.1*	*0.2*	*0.05*	*0.1*	*0.2*
250	522,982	517,099	514,560	510,756	510,237	510,643	0.98	0.98	0.99
500	582,380	566,387	554,249	541,702	544,543	541,679	0.93	0.96	0.98
1000	774,438	702,955	663,562	605,395	602,955	605,884	0.78	0.85	0.91
	TC^H			TC^N			TC^H/TC^N		

In this section we have described the way in which the original commitments are determined. At period t, the problem of determining the updated commitments is very similar. The main difference is that when updating the commitments at period t, we must consider the updating constraints, which are dependent upon the previous commitments. As before, finding the maximum value for expression (8) is equivalent to solving the resulting quadratic equation. Let the maximum value be denoted by $\bar{q}^{t}_{t+k}$; then, because of the concavity of the positive region of expression (8), we see that the optimal constrained solution is

$$q^{*t}_{t+k} = \min \{\max [(1-\beta)q^{t-1}_{t+k}, \bar{q}^{t}_{t+k}], (1+\beta)q^{t-1}_{t+k}\}.$$

The above equation holds because of the concavity of the cost function and the constraints

$$q^{*t}_{t+k} \geq (1-\beta)q^{t-1}_{t+k} \text{ and } q^{*t}_{t+k} \leq (1+\beta)q^{t-1}_{t+k}.$$

Computing the bounds

We now discuss the quality of the above heuristic. Its appeal lies in its simplicity, ease of implementation, and very low computation time. In addition, it does not require that the demands in different periods be identically distributed. To use this heuristic, we must be convinced that it produces results that are close to the optimal solution. To evaluate the heuristic, we compare the expected total cost obtained by the heuristic TC^H to the expected total cost TC^N obtained by solving the news-vendor problem. Recall that TC^N is a lower bound on the cost of the original problem P1. Note that the expected total cost is the sum of holding, shortage, and purchase costs totaled over all periods. Getting closed-form expressions for TC^N and TC^H is impossible; thus, we must use numerical procedures to obtain these costs. In what follows, we describe our computational results. In all cases, we assume (unless we specifically mention otherwise) the following parameter values: $h = 1$, $c = 40$, $p = 100$, and $T = 12$. D_t is defined by $f(\cdot)$, a normal density function with $\mu = 1000$ and σ as specified in the tables below. (We truncated the normal distribution at demand = 0.) These parameters indicate that the holding cost per period is 2.5% of the component cost and that the service level, i.e., the probability that the buyer satisfies his demand, is around 60% (calculated by means of the simple news-vendor approximation). We repeated the cost computations for $p = 400$ and $p = 800$, to ensure service levels of 90% and 95%. The quality of the approximation is not affected by this parameter.

We now present the computational results in **Table 2**, which demonstrates the performance of the heuristic for different flexibilities and demand uncertainties. The larger the fraction TC^N/TC^H, the greater the accuracy of the heuristic; i.e., when TC^N/TC^H is very close to 1, the gap between the bounds is very small. When σ/μ is small (low variability of demand), the heuristic is very close to optimal, even when the flexibility is limited (i.e., α and β are small). When $\sigma/\mu = 0.25$, for example, the largest gap between TC^H and TC^N is around 2%. This can be explained as follows. For a given flexibility, the lower the variability of demand, the lower the likelihood of having to update the forecast. This is because the likelihood of demand having a value outside of the flexibility limits is smaller. Naturally, the gap between the two solutions increases as the uncertainty increases. When $\sigma/\mu = 1$ and flexibility is 5%, the gap is around 22%. The fact that the gap between the lower bound and the heuristic is high when the coefficient of variation (σ/μ) is high is not surprising; it means that the lower bound is not tight when the uncertainty in the demand is large. For the high-variation, low-flexibility cases, we compare the heuristic solution to the optimal solution with zero flexibility, using the algorithm in Bassok and Anupindi (see footnote 1). Clearly, the problem with zero flexibility is an upper bound to problem P1 with flexibility. We expect that as the flexibility decreases to zero, the heuristic solution converges to the optimal solution of the problem with zero flexibility. A Monte Carlo simulation was performed to compute the total expected cost for both solutions. Demand for 12 periods was simulated, and 2000 samples were used. The simulation results demonstrate that the heuristic solution indeed converges to the optimal solution

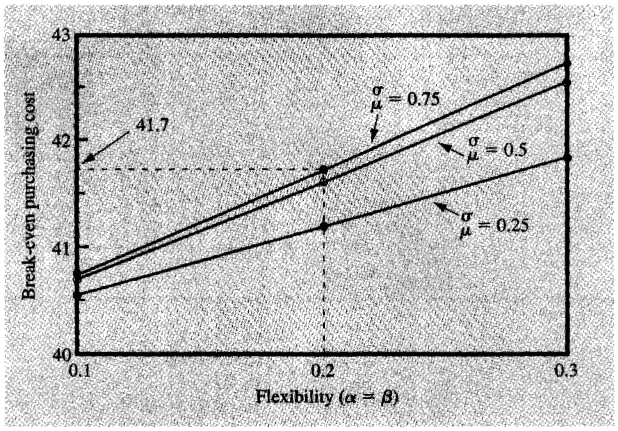

Figure 1

Break-even purchasing cost vs. flexibility under stationary demands ($c = 40$, $h = 1$, $p = 100$, $T = 12$, $\mu = 1000$).

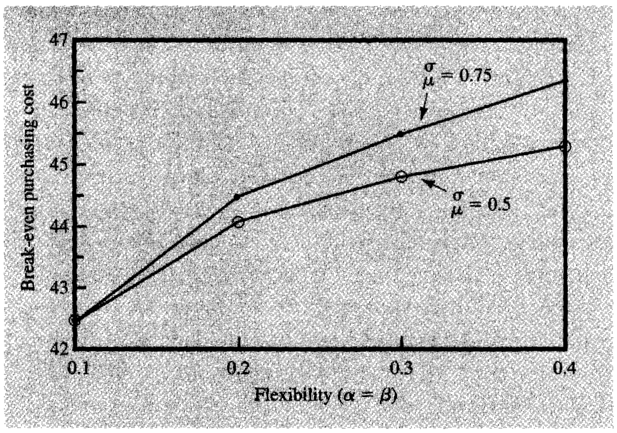

Figure 2

Break-even purchasing cost vs. flexibility under nonstationary demands ($c = 40$, $h = 1$, $p = 100$, $T = 12$, μ as in Table 2).

Table 3 Nonstationary demands.

Period	*Mean demand* (μ)	*Period*	*Mean demand* (μ)
1	333	7	1000
2	666	8	1000
3	1000	9	800
4	1000	10	600
5	1000	11	400
6	1000	12	200

of the problem with zero flexibility. A procedure for calculating the optimal period commitments with zero flexibility is presented in Bassok and Anupindi (see footnote 1). We observe that for all parameter combinations, as the flexibility approaches zero, the commitments obtained by using the heuristic approach the optimal solution. This increases our confidence that the heuristic performs well.

The worth of flexibility

In the previous section, we presented a heuristic that enables us to easily compute near-optimal purchasing quantities and period commitments for the problems presented. We use this heuristic to evaluate the worth of flexibility. There are many different ways to define and calculate the worth of flexibility. Here, we determine the iso-curves for flexibility and purchasing cost per unit; i.e., we determine all the pairs of flexibility and purchasing cost per unit that result in the same total expected cost. Such curves enable a buyer to negotiate for flexibility, providing an upper bound on the additional amount per unit that the buyer should pay for a given flexibility. **Figure 1** demonstrates this relationship for three different variabilities. On the basis of the top curve, for example, the buyer might be willing to pay 41.7 cost units for each unit purchased in order to have purchase and update flexibilities of 0.2. Any cost higher than 41.7 units would make it not worthwhile for the buyer to receive flexibilities of 0.2. This cost is referred to here as break-even cost. As intuition would suggest, the higher the variability, the higher the worth of flexibility.

In many real-life situations, the demands are not stationary. We have modified the above heuristic to deal with nonstationary demands. We calculate the base stock levels for the nonstationary case (using dynamic programming) and proceed in exactly the same way as described above. **Figure 2** is based on the data presented in **Table 3.** These data are typical of the mean demand for a product. In the initial periods (months), the demand increases; then, it levels; and toward the end of the life cycle of the product, the mean demand decreases. Figure 2 demonstrates the worth of flexibility for this nonstationary case. We make calculations for two constant values of coefficient of variation σ/μ, 0.5 and 0.75. It seems that the curves in Figure 2 have a different form than the curves in Figure 1 and are not upward concave. This is because of the range of the flexibility used in presenting these graphs. When the flexibility is increased beyond 0.3, the graphs in Figure 2 have exactly the same form as those in Figure 1.

Supply contracts in the IBM Printer System Company

Since delivery cycle time and vendor flexibility are important factors affecting responsiveness to customers (serviceability), the above approach helps the buyer to negotiate terms and conditions for a given supplier and product. IBM's printer business is very heavily dependent on external suppliers as sources of parts and assemblies. It is a business with worldwide procurement activities performed under contracts between IBM and the suppliers. The responsiveness of the suppliers to IBM is included in these contracts, as well as levels of inventory that must not be exceeded. What has not been understood well is the relationship between inventories and conditions such as lead times, frozen zones (duration prior to shipping in which the supplier has no flexibility), and demand variability. This complex set of relationships has been dealt with according to perceived and intuitive sets of assumptions. The application of the heuristic described here has demonstrated the relationships among these elements and helped to establish nearly optimum terms and conditions. By knowing and understanding these relationships and the effects on inventory of the various terms and conditions, the buyer has been able to articulate the requirements to the supplier in a precise manner. (For example, lower flexibility provided by the supplier results in lower ability for the buyer to revise purchasing commitments. When demand in a period is low, the buyer may be left with inventory yet still may need to purchase from the supplier in order to satisfy the terms and conditions in the contract.) Negotiations have been able to proceed with knowledge of the trade-off between the costs and benefits of additional flexibility.

Following is an example of a situation in which the heuristic was used by the IBM Printer System Company to arrive at the terms and conditions during the negotiation process for a new contract.

The random demand for a certain product is assumed to be stationary and normally distributed, with mean demand of 833 units per period (month) and standard deviation of 501. The cost per unit is $1707, and the holding cost per unit per period is $28. In our analysis, we assume that the Company will continue to sell the product for a long period of time; thus, the salvage value of the product (value of the leftover inventory at the end of the contract horizon) is assumed to be equal to the purchasing cost. Clearly, we could perform the analysis assuming different salvage values. The salvage value in the last period is used in the calculation of base-stock levels $S_1, \cdots, S_{12}$ by solving the news-vendor problem by means of the dynamic programming approach.

Two different supply contracts are evaluated and compared. The old supply contracts (currently in place) have a horizon of six months and provide the following flexibility: The buyer may change the size of previously made orders to be delivered three, four, five, and six months in the future by no more than 20% ($\beta = 0.2$). The size of the orders to be delivered one or two months in the future cannot be changed ($\beta = 0$). The Company considered a new type of supply contract in which orders to be delivered in the next period have an update flexibility of 30% ($\beta = 0.3$), orders to be delivered two periods in the future have 50% update flexibility ($\beta = 0.5$), and orders to be delivered three, four, five, and six months in the future have a very high update flexibility (β close to 1).

Table 4 Initial commitments for the old and new contracts.

Month	*Old contract*	*New contract*
1	1684	1386
2	1068	1074
3	1070	939
4	772	
5	803	
6	816	

It is clear why the Company found the new contract more attractive than the old one, and why the suppliers preferred the old contract. In the negotiation process between the buyer and the supplier, it is crucial to understand and know the exact worth of the additional flexibility in the new contract.

Since a major objective of the Company is to provide its customers with 98% serviceability, we ensured that in both contracts the Company orders quantities such that the 98% serviceability objective is achieved. We first calculate the initial commitments for both contracts (the actual orders may differ from the commitments because of the flexibility). In **Table 4**, we present the values of the initial commitments. (For the new contract, we do not calculate commitments for months four, five, and six because the conditions of the new contract provide unlimited flexibility for those months; thus, there is no need to determine these quantities ahead of time.)

We compare the expected holding cost per month resulting from the two contracts. The old contract results in an expected holding cost per month of $34509, while the new contract results in $20600, a reduction of approximately 70%.

The above results suggest that the worth of the added flexibility (when one considers only the reduction in holding cost) is around $15000 per month. When the supplier is willing to accept the terms and conditions of the new contract but requires compensation of less than $15000 per month, the contract is advantageous to the Printer Company; on the other hand, when the supplier requires more than $15000 per month as compensation, the old contract is more attractive from the Company's point of view.

Conclusions

We have presented a technique that can assist a component buyer such as IBM in negotiating contracts with suppliers. The features of this technique are the following:

1. It permits specification of the advance commitments that the buyer may make to the supplier.
2. It captures the buyer's flexibility to adjust commitments and actual orders on the basis of the observed demands.
3. It specifies how to calculate the worth of the buyer's flexibility to adjust the commitments and orders according to the changing conditions of the marketplace.

We have provided an actual example in which this model was used to modify the terms and conditions of contracts currently in place. We have demonstrated the potential benefits that will result from using this technique.

In our research, we extended the above model in order to consider a family of components purchased from a single supplier. In this case, the flexibility is stated for the family as a whole as well as individually for each component in the family. In our current research, we consider the following extensions to the work presented in this paper:

1. Supply contracts for multiple components that are used in different products assembled by the buyer are planned simultaneously, requiring coordinated design of multiple contracts.
2. The supply contracts are not restricted to having the same terms and conditions.

References

1. A. F. Veinott, "State of Mathematical Inventory Theory," *Manage. Sci.* **12,** 747–777 (1966).
2. R. J. Tersine, *Principles of Inventory and Materials Management*, North-Holland, New York, 1988, pp. 91–94.
3. H. L. Lee and M. Rosenblatt, "A Generalized Quantity Discount Pricing Model to Increase Supplier's Profits," *Manage. Sci.* **32,** 1177–1185 (1986).
4. M. Rosenblatt and L. Lee, "Improving Profitability with Quantity Discount and Fixed Demand," *IIE Trans.* **17,** 388–395 (1985).
5. A. Sadrian and Y. S. Yoon, "A Procurement Decision Support System in a Business Volume Discount Environment," *Oper. Res.* **42,** 14–22 (1994).
6. James V. Jucker and M. Rosenblatt, "Single Period Inventory Models with Demand Uncertainty and Quantity Discount: Behavioral Implication and New Solution Procedure," *Naval Res. Logist. Quart.* **32,** 537–550 (1985).
7. R. Anupindi and Y. Bassok, "Analysis of Supply Contracts with Total Minimum Commitment and Flexibility," *Proceedings of the 2nd International Symposium on Logistics*, University of Nottingham, England, July 11–12, 1995, pp. 123–128.
8. Y. Bassok and R. Anupindi, "Analysis of Supply Contracts with Total Minimum Commitment," *IIE Trans.* **29,** 373–381 (May 1997).
9. A. Tsai and W. S. Lovejoy, "Supply Chain Control with Quantity Flexibility," working paper, University of Michigan Business School, Ann Arbor, December 1994.
10. M. Moses and S. Seshadri, "Policy Mechanisms for Supply Chain Coordination," working paper, Operations Management area, New York University, New York, 1996.
11. S. Karlin, "Optimal Inventory Policy for the Arrow–Harris–Marschak Dynamic Model," *Studies in the Mathematical Theory of Inventory Production*, K. J. Arrow, S. Karlin, and H. Scarf, Eds., Stanford University Press, Stanford, CA, 1958, Ch. 9.
12. T. Morton, "The Near Myopic Nature of the Lagged Proportional Cost Inventory Problem with Lost Sales," *Oper. Res.* **19,** 1708–1716 (1971).

Received February 23, 1996; accepted for publication June 26, 1996

Yehuda Bassok *Department of Management Science, School of Business Administration, University of Washington, Seattle, Washington 98195.* Prof. Bassok received his B.Sc. degree in mathematics from the Hebrew University in Jerusalem, Israel, and his Ph.D. degree in engineering and public policy from Carnegie Mellon University in Pittsburgh. At present, he is a visiting scholar at the University of Washington in Seattle. Prof. Bassok studies problems in production, distribution, and supply chains. Lately he has been studying the effects of different structures of distribution systems, supply agreements, and incentive mechanisms on the performance of suppliers, buyers, and the entire supply chain. He has been developing optimal or close-to-optimal, easy-to-implement policies that benefit all of the participants in the supply chain. Usually, these policies have some elements of cooperation and risk-sharing among the different participants in the supply chain.

Ann Bixby *Industrial Engineering and Management Sciences Department, Northwestern University, 2225 North Campus Drive, Evanston, Illinois 60201 (abixby@iems.nwu.edu).* Ms. Bixby is a Ph.D. candidate in the Industrial Engineering Department of Northwestern University. She received a B.A. degree in German and mathematical sciences from Rice University in 1993 and an M.S. degree in industrial engineering from Northwestern in 1996. Her doctoral research is focused on the effective solution of vehicle-routing problems. She is a member of the Mathematical Programming Society and the Institute for Operations Research and the Management Sciences.

Ramesh Srinivasan *IBM Storage Systems Division, 5600 Cottle Road, San Jose, California 95193 (sramesh@us.ibm.com).* Dr. Srinivasan received a Ph.D. degree in industrial engineering from Purdue University in 1986. From 1986 until 1996, he worked at the IBM Thomas J. Watson Research Center on modeling and optimization of the IBM production, procurement, and distribution networks. In particular, he has developed new approaches for materials planning under uncertainty, in collaboration with colleagues in IBM and academia. Recently he led an IBM Research team to develop and deploy decision support tools for demand and supply planning in the IBM integrated supply chain. In 1996, he joined the IBM Storage Systems Division in San Jose, where he is involved in the reengineering of demand and supply planning processes. He has received IBM awards for technical achievement. Dr. Srinivasan is an active participant in INFORMS meetings. In 1993, he was co-organizer of the Multi-Echelon Inventory Systems Conference at the Thomas J. Watson Research Center.

Harry Z. Wiesel *IBM Product Support Services Global Parts Logistics, Amsterdam, The Netherlands (hwiesel@vnet.ibm.com).* Mr. Wiesel received a B.S. degree in electromechanical engineering from the City College of New York and an M.S. degree in computer science from the State University of New York; he is currently a candidate for an M.B.A. degree at the University of Phoenix. He joined IBM in 1974 as a Quality Engineer and since then has held assignments in a variety of manufacturing and development positions. He is currently on assignment in Amsterdam, The Netherlands, coordinating parts logistics projects for the Product Support Services organization. Prior to this assignment he was responsible for asset management and logistics in Somers, New York. His most recent line assignment was as the Material and Procurement Manager for the IBM Printer Systems Company. Mr. Wiesel received Division Excellence Awards for contributions in the areas of new product introductions, quality, and logistics management. He has also received an Authors Recognition Award for professional publications.