

1 Stochastic Demand

In this section we discuss the problem of controlling the inventory of an item with stochastic demand. Initially we consider the single period problem known as the newsvendor problem and then extend it to multi-period, and infinite horizon problems with and without setup costs.

1.1 The Newsvendor Problem

Let D denote the one period random demand, with mean μ and variance σ^2 . Let c be the unit cost, $c(1 + m)$ the selling price and $c(1 - d)$ the salvage value. If Q units are ordered, then $\min(Q, D)$ units are sold and $(Q - D)_+$ units are salvaged. Here $x_+ = \max(x, 0)$ and $x_- = \max(-x, 0)$. Thus the expected profit is given by

$$\pi(Q) = c(1 + m)E \min(Q, D) + c(1 - d)E(Q - D)_+ - cQ$$

This can be written as

$$\pi(Q) = cm\mu - G(Q)$$

where

$$G(Q) = cmE(D - Q)_+ + cdE(Q - D)_+.$$

For convenience let $h = cd$ and $p = cm$. It is convenient to think of h as the per unit overage cost and of p as the per unit underage cost. A little algebra reveals that $G(Q)$ can be written as

$$G(Q) = h(Q - \mu) + (h + p)E(D - Q)_+.$$

It is easy to verify that the loss function $G(Q)$ is convex in Q , and that if the distribution of D is continuous, then the optimal value of Q can be obtained by setting the derivative to zero, resulting in

$$P(D > Q) = \frac{h}{h + p}. \tag{1}$$

This means that the optimal order quantity Q^* is such that the probability of having a demand larger than Q^* is equal to the ratio of the overage cost h to the overage plus underage cost $h + p$.

Frequently D is assumed to be Poisson or Normally distributed. Unfortunately there is no closed form solution to equation (1) for these distributions of demand. However, the optimal order quantity against the worst possible distribution of demand with mean μ and variance σ^2 can be expressed in closed form:

$$Q^S = \mu + \frac{\sigma}{2} \left(\sqrt{\frac{p}{h}} - \sqrt{\frac{h}{p}} \right). \tag{2}$$

This formula is due to Herbert Scarf [5], is fairly robust, and it can be shown that

$$G(Q^S) \leq \sqrt{ph}\sigma$$

with equality holding for a certain distribution of demand with mass concentrated at two points. Notice that Scarf's formula (2) suggests ordering more (resp., less) than the mean demand when $p > h$ (resp., $p < h$). Moreover, $|Q^S - \mu|$ increases linearly in σ for $h \neq p$.

It is also possible to show that if we order the mean, then

$$G(\mu) \leq \frac{1}{2}(h + p)\sigma.$$

Consequently,

$$\pi(\mu) \geq p\mu - \frac{1}{2}(h+p)\sigma.$$

Notice that

$$1 \geq \frac{\pi(\mu)}{p\mu} \geq 1 - \frac{1}{2} \frac{h+p}{p} \frac{\sigma}{\mu}.$$

This indicates that ordering the mean is guaranteed to be near optimal whenever

$$\frac{1}{2} \frac{h+p}{p} \frac{\sigma}{\mu}$$

is small, e.g., when the coefficient of variation $\frac{\sigma}{\mu}$ is small.

1.2 Multi-period Models

In this section we consider a variety of multi-period models. Initially, we discuss models without setup costs and with zero lead times. Later we extend the analysis to the case of positive setup costs and positive lead times.

1.3 Finite Horizon Models

Let D_1, \dots, D_T be the demands for the next T periods. We assume that the D_t 's are independent random variables, and that all stockouts are backordered. Let c_t denote the unit cost in period t , and let x_t denote the inventory level at the beginning of period t , where a positive x_t indicates that x_t units of inventory are carried from the previous period, and a negative x_t indicates that a backlog of $-x_t$ is carried from the previous period. Let $y_t - x_t \geq 0$ denote the size of the order in period i . Then the inventory level at the beginning of period $t+1$ is given by

$$x_{t+1} = y_t - D_t.$$

At the end of period i a per unit cost h_t is incurred for every unit of inventory carried into period $t+1$ and a per unit cost p_t is incurred for every unit backordered into period $t+1$. Thus, if $y_t = y$ the loss function

$$G_t(y) = h_t E(y - D_t)_+ + p_t E(D_t - y)_+$$

in period i is the expected holding and backorder cost charged to period i .

Let $C_{T+1}(x_{T+1})$ be an arbitrary cost function on the inventory level at the end of period T , let $0 < \alpha \leq 1$ be the one period discounted cost and let $C_t(x)$ denote the optimal expected discounted cost starting in period t with x_t units of inventory. Then,

$$C_t(x_t) = \min_{y \geq x_t} \{c_t(y - x_t) + G_t(y) + \alpha EC_{t+1}(y - D_t)\}. \quad (3)$$

It can be shown that if $C_{T+1}(\cdot)$ is convex, then $C_t(\cdot)$ is convex for all $t = 1, \dots, T$, and the optimal policy is to order $\max(0, y_t^* - x_t)$ units in period i where y_t^* minimizes

$$c_t y + G_t(y) + \alpha EC_{t+1}(y - D_t).$$

This class of policies is known as order-up-to policies. The idea is that we order up to y_t^* in period i if $x_t < y_t^*$ and not to order otherwise.

Notice that the above problem needs to be solved recursively starting with period T down to period 1. This requires a computer code that can be written in less than one hour by an experienced programmer. The quality of the solution depends on the quality of the estimates of future cost and of the demand distributions.

1.3.1 The Myopic Policy

Here we describe a myopic policy that is frequently used in practice. To develop this policy we need to write a slightly different but equivalent set of recursive equations. To this end let

$$M_t(x_t) = C_t(x_t) + c_t x_t$$

and notice that with this definition, the recursion becomes

$$M_t(x_t) = \alpha c_{t+1} \mu_t + \min_{y \geq x_t} [m_t y + G_t(y) + \alpha E M_{t+1}(y - D_t)],$$

where $m_t = c_t - \alpha c_{t+1}$. The myopic policy ignores at time i , the future discounted costs

$$\alpha E M_{t+1}(y - D_t),$$

and orders $\max(0, y_t^m - x_t)$ units in period i , where y_t^m minimizes the current costs

$$m_t y + G_t(y).$$

If the demand is continuous, then y_t^m satisfies

$$P(D_t > y) = \frac{h_t + m_t}{h_t + p_t}.$$

How is the myopic policy related to the optimal policy? The most important known result is that

$$\min\{y_t^m, \dots, y_T^m\} \leq y_t^* \leq y_t^m$$

which implies that $y_t^* = y_t^m$ when the y_t^m are non-decreasing.

Using Scarf's min-max approach, the myopic policy is to order $\max(0, y_t^S - x_t)$ where

$$y_t^S = \mu_t + \frac{\sigma_t}{2} \left(\sqrt{\frac{p_t - m_t}{h_t + m_t}} - \sqrt{\frac{h_t + m_t}{p_t - m_t}} \right).$$

1.4 Infinite Horizon, Stationary Models

If all the costs are stationary, i.e., $c_t = c$, $h_t = h$ and $p_t = p$ for all i , and the demands are independent and identically distributed (IID), then the infinite horizon cost function $C(x)$ satisfies the functional equation

$$C(x) = \min_{y \geq x} \{c(y - x) + G(y) + \alpha EC(y - D)\}.$$

In terms of $M(x) = C(x) + cx$, the functional equation can be written as

$$M(x) = \alpha c \mu + \min_{y \geq x} \{c(1 - \alpha)y + G(y) + \alpha EM(y - D)\}.$$

The myopic policy orders $\max(0, y^m - x)$ units where y^m minimizes the current cost

$$c(1 - \alpha)y + G(y).$$

If the one period demand has a continuous distribution, then y^m satisfies

$$P(D > y) = \frac{h + c(1 - \alpha)}{h + b}.$$

Surprisingly, the myopic policy is optimal if the distribution of D is over the non-negative numbers. This can be seen as follows. Suppose that $M(\cdot)$ is known and that y^* minimizes

$$c(1 - \alpha)y + G(y) + \alpha EM(y - D).$$

Then for $x \leq y^*$ we have

$$M(x) = \alpha\mu + c(1 - \alpha)y^* + G(y^*) + \alpha EM(y^* - D).$$

Notice that the right hand side of the last equation is independent of x , so there is a constant, say M^* , such that $M(x) = M^*$ for all $x \leq y^*$. Since $D \geq 0$, $y^* - D \leq y^*$ so $M(y^* - D) = M^*$. Therefore M^* satisfies

$$M^* = \alpha\mu + c(1 - \alpha)y^* + G(y^*) + \alpha M^*.$$

Solving for M^* yields

$$M^* = \frac{\alpha\mu + c(1 - \alpha)y^* + G(y^*)}{1 - \alpha}$$

so y^* must minimize the current cost

$$c(1 - \alpha)y + G(y)$$

just as y^m . Therefore $y^* = y^m$ if $c(1 - \alpha)y + G(y)$ has a unique minimizer or we can select y^* as y^m if this function admits more than one minimizer.

Thus, for the infinite horizon stationary cost model the optimal policy is to order up to y^* . This policy is also known as a base-stock policy because orders are placed in each period to restore the inventory to y^* . Let us look at how the policy works in practice. Suppose first that inventory at the beginning of period 1 is $x_1 < y^*$. Since $x_1 < y^*$ we place an order for $y^* - x_1$ units bringing the inventory up to y^* . Suppose that during period 1 $D_1 \geq 0$ units are demanded. Then the inventory at the beginning of period 2 is $x_2 = y^* - D_1$ and D_1 units are ordered to bring the inventory up to y^* . It is easy to see that the order placed in period i is exactly for D_{i-1} units for all $i \geq 2$. If, on the other hand, $x_1 > y^*$ then nothing is ordered until the first period, say i , where $y^* - x_t > 0$. After period i , we simply order to replenish the quantity demanded in the previous period.

Notice that as α increases to one, i.e., no discounting, the optimal policy is to order up to y^* where y^* satisfies

$$P(D > y) = \frac{h}{h + b}.$$

Notice also that the myopic policy is also optimal for the finite horizon stationary problem provided we set $c_{T+1}(x) = -cx$.

1.5 Positive Lead Times

Suppose that an order placed at the beginning of period t arrives at the beginning of period $t + L$. To work with positive, but deterministic, lead times, we need to add the inventory on order to the inventory level to summarize the state space at the beginning of each period. The resulting quantity is known as the inventory position and is equal to the inventory on hand plus on order minus backorders. When the lead time is zero, the inventory position is equal to the inventory level. Let x_t be the inventory position at the beginning of period t , after we receive the order placed L periods ago, but before we make the ordering decision for period t . Suppose that we order to bring the inventory position to $y_t \geq x_t$. This order will arrive at the beginning of period $t + L$. All orders placed prior to period t would have arrived by the beginning of period $t + L$. Moreover, orders placed after period t will not arrive until after period $t + L$. Consequently, the inventory level at the end of period $t + L$ is given by $y_t - D[t, t + L]$ where $D[t, t + L] = \sum_{s=t}^{t+L} D_s$. The demand $D[t, t + L]$ over periods $\{t, \dots, t + L\}$ is known as the lead time demand starting from period t . Notice that $D[t, t + L]$ contains the demand over $L + 1$ periods and reduces to D_t when $L = 0$. Since the decision made at time t determines the holding and penalty costs incurred at the end of period $t + L$ it makes sense to charge these costs to period t . This is accomplished by redefining the loss function to be

$$G_t(y) = h_t E(y - D[t, t + L])^+ + p_t E(D[t, t + L] - y)^+.$$

Let $C_t(x_t)$ be the minimal expected discounted cost of managing the system starting from period t with inventory position x_t . Then,

$$C_t(x_t) = \min_{y_t \geq x_t} \{c_t(y_t - x_t) + G_t(y_t) + \alpha EC_{t+1}(y_t - D_t)\}.$$

This formulation is equivalent to (3) except that x_t is now the inventory position and G_t is defined differently. One additional difference is that the last ordering period is $T - L$ instead of T . Other than this, the problems are mathematically equivalent. The myopic policy calls for bringing the inventory position up to y_t^m in period where y_t^m satisfies

$$P(D[t, t + L] > y) = \frac{h_t + m_t}{h_t + p_t}.$$

The infinite horizon policy calls for bringing the inventory position up to y^* where y^* satisfies

$$P(D[t, t + L] > y) = \frac{h + c(1 - \alpha)}{h + p}.$$

Let

- μ mean demand per period
- σ standard deviation of daily demand
- μ_d mean of the leadtime demand.
- σ_d standard deviation of the leadtime demand.

If we assume that the period demands are statistically independent, then $\mu_d = \mu(1 + L)$ and $\sigma_d = \sigma\sqrt{1 + L}$. Often $D[t, t + L]$ can be modeled as normally distributed with mean μ_d and standard deviation σ_d . In this case,

$$y^* = \mu_d + z\sigma_d$$

where

$$\bar{\Phi}(z) = \frac{h_t + m_t}{h_t + p_t}.$$

1.5.1 Random Lead Times

When lead times are random things become complicated because of the possibility of order crossing, i.e., a recent order arrives before an old order. There is no easy way to account for order crossings. In many practical manufacturing and distribution situations orders do not cross or they cross so rarely that it makes sense to build a model under the assumption that orders do not cross. If we are willing to assume that orders do not cross, that demand lead times over different periods are normal and identically distributed then the problem can be solved once we find the mean and the variance over the lead time.

Let L be the lead time. To simplify the notation we will let μ_l and σ_l to denote respectively the mean and the standard deviation of $L + 1$. Our objective is to write μ_d and σ_d in terms of μ , σ , μ_l and σ_l under the assumption that the period demands are statistically independent. The formula for the mean lead time demand is again $\mu_d = \mu\mu_l$. The formula for σ_d , which is what we are interested in, is given by

$$\sigma_d = \sqrt{\mu_l\sigma^2 + \sigma_l^2\mu^2}.$$

This formula is derived on page 153 in reference [3].

Numerical Example The mean daily demand for a product is $\mu = 80$ units and the standard deviation is $\sigma = 20$ units.

- Scenario 1. The leadtime is short, but unreliable: The mean leadtime is $\mu_l = 5$ days but the standard deviation is $\sigma_l = 4$ days. In this case, the standard deviation of the leadtime demand is

$$\sigma_d = \sqrt{5(20)^2 + (4)^2(80)^2} = 323.$$

- Scenario 2. The leadtime is long, but reliable: The mean leadtime is $\mu_l = 25$ days but the standard deviation is $\sigma_l = 0$ days. In this case, the standard deviation of the leadtime demand is

$$\sigma_d = \sqrt{25(20)^2 + (0)^2(80)^2} = 100.$$

Since the holding and penalty costs are proportional to the standard deviation of demand, we see that the costs are over three times higher with the shorter and more unreliable leadtime. Comparing the standard deviation of the lead time demand to the mean lead time demand shows that the incidious effect of randomness in the lead time is even worse than the direct comparison between the standard deviations would indicate.

2 Positive Ordering Costs

2.1 (Q, r) Policies

The above models assumed that decisions are made at discrete points in time. In this section we consider a continuous time model with nonzero lead times and positive setup costs. We follow the notation in Zipkin [9] in developing performance measures.

Let $D(t)$ denote the cumulative demand up to time t and define the following quantities:

- $I(t)$ inventory on hand at time t .
- $B(t)$ backorders at time t .
- $IN(t)$ net inventory at time t .
- L lead time.
- $IO(t)$ inventory on order at time t .
- $IP(t)$ inventory position at time t .

By definition $I(t) = IN(t)_+ = \max(0, IN(t))$ and $B(t) = IN(t)_- = \max(0, -IN(t))$, so the net inventory $IN(t)$ is equal to the inventory when it is nonnegative. When $IN(t) < 0$ we have $B(t) = -IN(t)$ backorders. We will assume that an order is received L time units after it is placed. The inventory on order $IO(t)$ at time t is therefore equal to the number of orders placed during the interval $(t - L, t]$. The inventory position $IP(t)$ is defined as the inventory on hand plus the inventory on order minus the number of backorders. Mathematically,

$$\begin{aligned} IP(t) &= I(t) + IO(t) - B(t) \\ &= IN(t) + IO(t). \end{aligned}$$

Notice that $IO(t) = 0$ when $L = 0$, and in that case the inventory position is equal to the net inventory $IN(t)$.

Given a stationary supply and demand processes, and a specific policy, we will determine a number of performance measures of interest including the long run fraction of time out of stock, the average number of units on inventory, the average number of units backordered, and the average frequency of orders.

We will restrict our attention to the class of (Q, r) policies that are often used in practice. Under an (Q, r) policy we monitor the inventory position continuously and place an order of size Q whenever

the inventory position is at or below the reorder point r . We will assume that demands are for one unit at a time. Under this assumption (Q, r) policies are equivalent to (s, S) policies with $s = r$ and $S = r + Q$. Under an (s, S) policy we monitor the inventory position continuously and place an order to restore the inventory position to S whenever the inventory position drops to or below s .

2.2 $Q = 1$

Initially, we will compute the performance measures for the case $Q = 1$. This mode of operation is optimal when there are no setup costs or they are small relative to the cost of holding inventory, e.g., for expensive items with low demand rates. For convenience let $S = r + Q = r + 1$. Notice that in this case the (Q, r) policy is actually a base stock policy with base stock level S . Under this policy we order to keep the inventory position equal to S .

Then, $IP(t) = S$ except at ordering epochs where the inventory position momentarily drops below S . Since

$$S = IP(t) = IN(t) + IO(t)$$

we have

$$IN(t) = S - IO(t).$$

Now $IO(t)$ is the number of units ordered during the interval $(t, t - L)$. Under a base-stock policy orders are placed to keep the inventory position constant so $IO(t)$ is equal to $D(t) - D(t - L)$ the number of units demanded over the interval $(t - L, t]$. Let $D(t|L) \equiv D(t) - D(t - L)$. As $t \rightarrow \infty$ we have

$$IN(\infty) = S - D(\infty|L)$$

where $D(\infty|L)$ is the stationary lead time demand. Thus the stationary distribution of IN is determined by the stationary distribution of the lead time demand. For example, if $D(t)$ is Poisson (λt) it follows that $D(\infty|L)$ is Poisson (λL) .

We now compute some performance measures of interest. Let $A = Pr(IN(\infty) \leq 0)$, then A is the long run probability of being out of stock. From the above, we have

$$A = Pr(D(\infty|L) \geq S).$$

Let $B = E[B(\infty)]$ then

$$B = E(D(\infty|L) - S)_+.$$

Finally, let $I = E[I(\infty)]$, then

$$I = E(S - D(\infty|L))_+ = S - ED(\infty|L) + B.$$

Thus, if we want to minimize the long run expected holding and backorder costs, we are faced with the problem of minimizing

$$bB + hI = bED(\infty|L) - S)_+ + hE(S - D(\infty|L))_+.$$

The good news is that this is a newsvendor problem, which, as we know, is solved by letting S be the smallest integer such that

$$Pr(D(\infty|L) \leq S) \geq \frac{b}{h+b}.$$

2.3 Q a Positive Integer

Now, suppose that Q is a positive integer. Then, under very general conditions on the demand process, it can be shown that the stationary inventory position is uniform between $r + 1$, and $r + Q$. That is,

$$P(IP(\infty) = j) = \frac{1}{Q} \quad j = r + 1, \dots, r + Q.$$

Moreover, it can be shown that $IP(\infty)$ is independent of $D(\infty|L)$. This allows us to compute the relevant performance measures resulting in

$$A = \frac{1}{Q} \sum_{s=r+1}^{r+Q} Pr(D(\infty|L) > s),$$

$$B = \frac{1}{Q} \sum_{s=r+1}^{r+Q} E(D(\infty|L) - s)_+,$$

and

$$I = \frac{1}{Q} \sum_{s=r+1}^{r+Q} E(s - D(\infty|L))_+.$$

If the average demand per unit time is λ , then the long run frequency of orders is

$$O = \frac{\lambda}{Q}.$$

The above performance measures can then be combined to form a cost function. Suppose that the ordering cost is K , the holding cost rate is h and the backorder cost rate is b . Then, the total average cost can be written as

$$c(Q, r) = \frac{K\lambda}{Q} + \frac{1}{Q} \sum_{y=r+1}^{r+Q} G(y),$$

where

$$G(y) = hE(y - D(\infty|L))_+ + bE(D(\infty|L) - y)_+.$$

We will discuss an algorithm to find the optimal (Q, r) pair, as well as some bounds and heuristics.

2.4 Algorithm

An algorithm to find the optimal (Q, r) policy is easily obtained after making three observations.

First, since $-G(y)$ is unimodal, the problem

$$c(Q) = \min_r c(Q, r)$$

is easily solved by finding the set of Q consecutive integers

$$\{y_1, \dots, y_Q\}$$

such that

$$y_1 = \min\{G(y) : y \in \mathcal{Z}\},$$

and, given y_1, \dots, y_k

$$y_{k+1} = \min\{G(y) : y \in \mathcal{Z}, y \neq y_i, i = 1, \dots, k\}.$$

Letting G_k denote $G(y_k)$ we can write

$$c(Q) = \frac{K\lambda + \sum_{k=1}^Q G_k}{Q}.$$

In words, the Q smallest values of G , namely G_1, \dots, G_Q are attained by Q consecutive integers, i.e., by some reordering of the set $\{y_1, \dots, y_Q\}$.

The second observation is that we can write $c(Q)$ as a convex combination of $c(Q - 1)$ and G_Q . Indeed it is easy to verify that

$$c(Q) = \frac{Q-1}{Q}c(Q-1) + \frac{1}{Q}G_Q.$$

This implies that $c(Q) < c(Q - 1)$ if and only if $C(Q - 1) > G_Q$ which implies that

$$G_Q < c(Q) < c(Q - 1).$$

The third observation is that $-c(Q)$ is unimodal, which implies that the optimal batch size is the largest value of Q for which

$$G_Q < c(Q - 1).$$

Algorithm

1. Set $Q = 1$ and find y_1 , G_1 and $c(1)$.

2. Let

$$L_Q = \min\{y_1, \dots, y_Q\} - 1$$

and

$$R_Q = \max\{y_1, \dots, y_Q\} + 1$$

If $G_{Q+1} = \min(G(L_Q), G(R_Q)) \geq c(Q)$ then stop. Otherwise compute

$$c(Q+1) = \frac{Q}{Q+1}c(Q) + \frac{1}{Q+1}G_{Q+1}$$

and set $y_{Q+1} = L_Q$ if $G(L_Q) < G(R_Q)$ and $y_{Q+1} = R_Q$ otherwise.

3. Set $Q \leftarrow Q + 1$ and return to Step 2.

This algorithm is due to Federgruen and Zheng [1]

To facilitate the use of this algorithm it is convenient to write the increment of the $G(y)$ as

$$G(y+1) - G(y) = (h+p)P(D(\infty|L) \leq y) - p.$$

For Poisson demands we can update $P(D(\infty|L) = y)$ and $P(D(\infty|L) \leq y)$

$$P(D(\infty|L) = y+1) = \frac{\lambda L}{y+1}P(D(\infty|L) = y),$$

starting from $P(D(\infty|L) = 0) = e^{-\lambda L}$, and

$$P(D(\infty|L) \leq y+1) = P(D(\infty|L) \leq y) + P(D(\infty|L) = y+1).$$

2.5 Sensitivity, Bounds and Heuristics

Let us consider again the cost function

$$c(Q, r) = \frac{K\lambda + \sum_{y=r+1}^{r+Q} G(y)}{Q}$$

that arises when the demand rate is λ , the ordering cost is K , the holding cost is h the backorder cost is b and the lead time demand is a random variable D with mean μ and variance σ^2 .

Notice that if the variance $\sigma^2 = 0$ the demand is deterministic and the resulting problem is essentially an economic order quantity where we need to balance the ordering holding and backorder costs. On the other hand, if the ordering cost $K = 0$ then the problem reduces to the newsvendor

problem where we need to decide on the stock level to minimize the holding and backorder costs. Thus, the cost function $c(Q, r)$ reduces to well known subproblems if either $\sigma^2 = 0$, or $K = 0$.

Although we have developed a fairly deep understanding of both the EOQ and the newsvendor subproblems and have an efficient algorithm to minimize the cost function $c(Q, r)$, we don't yet have a deep understanding of the cost function $c(Q, r)$. Is it more or less sensitive than the EOQ to mispecifications of the batch size or the cost parameters. Is it more or less sensitive than the newsvendor problem to the specification of the distribution of the lead time demand? Can we obtain effective bounds on the average cost without having to run the algorithm? How does the average cost behave as a function of the set up cost and the variance of the lead time demand? Can we find upper and lower bounds on Q ? Are there effective heuristics for the batch size? We now provide answers to some of these questions. The results, except as noted, are due to Gallego [2].

2.5.1 Sensitivity

It can be shown that $c(Q) = \min_r c(Q, r)$ is less sensitive than the EOQ in the sense that

$$\frac{c(Q)}{c(Q^*)} \leq \frac{1}{2} \left(\frac{Q}{Q^*} + \frac{Q^*}{Q} \right).$$

Notice that we have an inequality for the case of random demands, where we had an equality for the EOQ cost function. This result is due to Zheng [7].

2.5.2 Bounds

We have the following closed form bounds on the cost function

$$\sqrt{c_d^2 + G_1^2} \leq c(Q^*) \leq \sqrt{c_d^2 + \bar{G}_1^2}.$$

where c_d is the average cost of the EOQ subproblem,

$$G_1 = G(y_1) = \min\{G(y) : y \in \mathcal{Z}\}.$$

is the newsvendor cost, and

$$\bar{G}_1 = \sigma\sqrt{hb}$$

is Scarf's upper bound on the newsvendor cost. Recall that $c_d = \sqrt{2HK\lambda}$ where $H = \frac{hb}{h+b}$.

Closed form bounds on Q^* are given by

$$Q_d \leq Q^* \leq Q_e$$

where

$$Q_d = c_d/H$$

is the economic order quantity, and

$$Q_e = \sqrt{c_d^2 + \bar{G}_1^2}/H = \sqrt{Q_d^2 + Q_\sigma^2}$$

where

$$Q_\sigma = \frac{\bar{G}_1}{H}.$$

2.5.3 Heuristics

It can be shown that

$$\frac{c(\sqrt{2} Q_d)}{c(Q^*)} \leq 1.0625,$$

so using a batch size that is $\sqrt{2}$ times the EOQ results in a cost increase of at most 6.25%. In practice, the $\sqrt{2} Q_d$ heuristic can be improved by using the batch size

$$Q_g = \min(\sqrt{2} Q_d, \sqrt{Q_d Q_e}).$$

2.6 General Demand Sizes

When demands are not for one unit at a time an order under an (Q, r) policy consists of the number of batches of size Q that are necessary to bring the inventory position to the interval $[r, r+Q)$. In this case, (Q, r) policies are no longer optimal. Managerially (Q, r) policies are still attractive because the more restricted order size facilitates packaging, transportation, and coordination. Let X denote the random demand size. Then, the long run average cost under an (Q, r) policy is given by

$$c(Q, r) = \frac{K\lambda E \min(Q, X) + \sum_{j=1}^{r+Q} G(y)}{Q} \quad (4)$$

To see how the ordering cost arises, notice that when the inventory position is $r+j$, a demand of size X triggers an order if and only if $X \geq j$. Since the inventory position is uniform $\{r+1, \dots, r+Q\}$ the probability, and the long run average frequency, of placing an order is $\sum_{j=1}^Q P(X \geq j)/Q$. Since $X \geq 0$ and $E \min(Q, X) = \sum_{j=1}^Q P(X \geq j)$, the cost function (4) results.

3 (s, S) Policies

Under an (s, S) policy, $s < S$, the inventory manager places an order to increase the item's inventory position to the *order-up-to level* S , whenever he finds the item's inventory position to be at or below the *reorder-level* s .

Academicians have devoted a large effort to the problem of identifying single-item stochastic inventory models for which an (s, S) policy is optimal. It turns out that (s, S) policies are optimal for a large class of single-item inventory models including the one we will study in this section. Here we will take the optimality of (s, S) policies for granted and will concern ourselves with the problem of computing an optimal (s, S) policy for a model where both the demand and the relevant costs are time stationary.

We assume that

- orders may be placed at the beginning of each period,
- orders are delivered immediately,
- all stockouts are backordered,
- period demands are independent and identically distributed,
- costs are stationary over time.

The objective is to minimize the long run average cost over an infinite horizon.

Notation:

- D the one period demand,
- $p_j = Pr(D = j)$, $j = 0, 1, \dots$,
- $K > 0$ fixed cost of placing an order,
- $G(y)$ one period expected cost starting with y units.

The typical form of $G(y)$ is

$$G(y) = hE(y - D)_+ + pE(y - D)^-,$$

where h is the holding cost rate and p is the stockout penalty cost rate. However, other forms of $G(\cdot)$ may also arise. In any event, all that we will require of $G(\cdot)$ is that:

- (i) $-G(\cdot)$ is unimodal,
- (ii) $\lim_{|y| \rightarrow \infty} G(y) > \min_x G(x) + K$.

Let $c(s, S)$ denote the long run average cost of using the policy (s, S) . To obtain an expression for (s, S) we use the well known reward-renewal theorem that states that the long run average cost is equal to the expected cost per cycle divided by the expected cycle length. A cycle is interpreted as the time elapsed between the placement of two consecutive orders. We say that the system renews itself after each cycle because the item's inventory position immediately after an order is placed is equal to S .

We are now concerned with the determination of the expected cost per cycle, and the expected cycle length. For $y > s$, let $k(s, y)$ denote the total expected cost until the next order is placed when the starting inventory position is equal to y units. Our interest, of course, is in finding a formula for $k(s, S)$. Likewise, let $M(j)$ be the expected total time until an order is placed when starting with $s + j$ units. Our interest, of course, is to find a formula for $M(S - s)$. Once these formulas are obtained, we can write

$$c(s, S) = \frac{k(s, S)}{M(S - s)}.$$

It is clear that the functions $k(s, \cdot)$, and $M(\cdot)$ satisfy the discrete renewal equations

$$k(s, y) = G(y) + K \sum_{j=y-s}^{\infty} p_j + \sum_{j=0}^{y-s-1} p_j k(s, y - j), \quad y > s$$

and

$$M(j) = 1 + \sum_{i=0}^{j-1} p_i M(j - i), \quad j = 1, 2, \dots$$

Let $m(0) = 1/(1 - p_0)$, $M(0) = 0$, and

$$m(j) = \sum_{k=0}^j p_k m(j - k), \quad j = 1, 2, \dots$$

It follows that

$$M(j) = M(j - 1) + m(j - 1), \quad j = 1, 2, \dots,$$

and

$$k(s, y) = K + \sum_{j=0}^{y-s-1} m(j)G(y - j) \quad y > s.$$

Consequently,

$$c(s, S) = \frac{K + \sum_{j=0}^{S-s-1} m(j)G(S - j)}{M(S - s)}.$$

Unfortunately the cost function $c(s, S)$ is not, in general, convex. For a long time this fact precluded the development of efficient algorithms. However, Zheng and Federgruen [8] have observed that

$$c(s - 1, S) = \alpha_n c(s, S) + (1 - \alpha_n)G(s) \tag{5}$$

where

$$\alpha_n \equiv \frac{M(n)}{M(n + 1)},$$

and $n = S - s$. Based on this observation, they have derived a very effective algorithm to compute an optimal (s, S) policy. We present here some of their key results, as well as their algorithm. From (1) we see that $c(s - 1, S)$ is a convex combination of $c(s, S)$ and of $G(s)$, and consequently

$$\min \{c(s, S), G(s)\} \leq c(s - 1, S) \leq \max \{c(s, S), G(s)\}.$$

We will use property (1) to determine necessary and sufficient conditions on s^o to be the optimal reorder-level for a fixed order-up-to level S . Then, we will obtain lower and upper bounds on an optimal reorder-level and an optimal order-up-to level.

For fixed S the reorder-level s^o is optimal if

$$c(s^o, S) \leq c(s, S) \quad \forall s.$$

Consequently s^o must satisfy

$$c(s^o - 1, S) \geq c(s^o, S) \leq c(s^o + 1, S),$$

but then from (1)

$$G(s^o + 1) \leq c(s^o, S) \leq G(s^o). \quad (6)$$

Let $y_1^* = \min\{y : G(y) = \min_x G(x)\}$, and notice that $-\infty < y_1^* < \infty$.

We will now establishing lower and upper bounds on an optimal reorder-level s^* .

Proposition 1 *Let s_l^* denote the smallest optimal reorder-level, then*

$$s_l^* \leq \bar{s} \equiv y_1^* - 1.$$

Proof: Let s_l^* be the smallest optimal value of s that minimizes $c(s, S^*)$. Suppose for a contradiction that $s_l^* \geq y_1^*$, then it follows from the form of $c(s, S)$ that $c(s_l^*, S^*) \geq G(s_l^*)$ which in turn implies that $c(s_l^* - 1, S^*) \leq c(s_l^*, S^*)$ contradicting the definition of s_l^* . \square

Proposition 2 *Let s_u^* denote the largest optimal reorder-level $< y_1^*$. Then*

$$s^o \leq s_u^*$$

where s^o is the optimal order level for some arbitrary order-up-to level S .

Because s_u^* is optimal for S^* it follows that (2) must hold. In fact, we claim that $G(s_u^* + 1) < c(s_u^*, S^*)$ holds. Suppose for a contradiction that $s_u^* < y_1^* - 1$, and that $G(s_u^* + 1) = c(s_u^*, S^*)$ holds. Then $s_u^* + 1 < y_1^*$ is also optimal, contradicting the definition of s_u^* . On the other hand, if $s_u^* = y_1^* - 1$, then, by the definition of y_1^* , $G(y_u^*) = G(s_u^* + 1) < c(s_u^*, S^*)$. Now, given any S , and an optimal reorder-level s^o for S , we have

$$G(s_u^* + 1) < c(s_u^*, S^*) \leq c(s^o, S) \leq G(s^o).$$

But then because $G(s)$ is unimodal, $G(s^o) \geq G(s_u^*) \geq G(s_u^* + 1)$, so $s^o \leq s_u^*$. \square

Corollary 3 *There exists an optimal solution s^* satisfying*

$$s^o \leq s^* \leq \bar{s}. \quad (7)$$

where s^o is an optimal reorder-level for an arbitrary order-up-to level S .

We now turn our attention to bounds on S^* . To this end, let $\underline{S} \equiv \max\{y : G(y) = \min_x G(x)\}$; notice that $y_1^* \leq \underline{S} < \infty$. Let $c^* = c(s^*, S^*)$ denote the optimal average cost value, and let $\bar{S} \equiv \max\{y \geq \underline{S} : G(y) \leq c^*\}$.

Proposition 4 *There exists an optimal policy (s^*, S^*) for which*

$$\underline{S} \leq S^* \leq \bar{S}. \quad (8)$$

Proof: We start by proving the lower bound. Let (s^*, S^*) be an optimal (s, S) policy that maximizes the value of S^* . Assume for a contradiction that $S^* < \underline{S}$. Note that for $j \geq 0$, $G(S^* + 1 - j) \leq G(S^* - j)$, so $c(s^* + 1, S^* + 1) \leq c(s^*, S^*)$ contradicting the maximality of S^* .

To show the upper bound, assume for a contradiction that $G(S^*) > c^*$. Notice that from the definition of $k(s, \cdot)$ and $M(\cdot)$ we can write

$$\begin{aligned} c^* &= \frac{G(S^*) + K \Pr(D \geq S^* - s^*) + \sum_{j=0}^{S^* - s^* - 1} p_j k(s^*, S^* - j)}{1 + \sum_{j=0}^{S^* - s^* - 1} p_j M(S^* - s^* - j)} \\ &> \frac{c^* + \underline{k} \Pr(D < S^* - s^*)}{1 + \underline{M} \Pr(D < S^* - s^* - 1)}, \end{aligned}$$

where

$$\begin{aligned} \underline{k} &= \frac{\sum_{j=0}^{S^* - s^* - 1} p_j k(s^*, S^* - j)}{\Pr(D < S^* - s^*)} \\ \underline{M} &= \frac{\sum_{j=0}^{S^* - s^* - 1} p_j M(S^* - s^* - j)}{\Pr(D < S^* - s^*)}. \end{aligned}$$

Consequently,

$$c^* > \frac{\underline{k}}{\underline{M}}. \quad (9)$$

However, we can identify the right hand side of (5) as the average cost of a feasible policy! This contradicts the optimality of (s^*, S^*) so $G(S^*) \leq c^*$.

Corollary 5 *Let $c > c^*$, and $\bar{S}_c \equiv \max\{y \geq \underline{S} : G(y) \leq c\}$, then $S^* \leq \bar{S} \leq \bar{S}_c$.*

Corollary 5 can be used to identify increasingly tighter upper bounds for S^* as increasingly better (s, S) policies are found.

For any fixed order up to level S , let

$$c^*(S) = \min_{s < S} c(s, S).$$

S is said to be *improving* upon S^o , if $c^*(S) < c^*(S^o)$.

Lemma 6 *For a given order-up-to level $S^o (\geq y_1^*)$, let $s^o (< y_1^*)$ be an optimal reorder-level. Then $c^*(S) < c^*(S^o)$ if and only if $c(s^o, S) < c(s^o, S^o)$.*

Proof: Suppose $c(s^o, S) < c(s^o, S^o)$, then $c^*(S) \leq c(s^o, S) < c(s^o, S^o) = c^*(S^o)$.

Conversely, assume that $c^*(S) < c^*(S^o)$, and that $c(s^o, S) \geq c(s^o, S^o)$. To reach a contradiction it is enough to show that $c(s, S) \geq c(s^o, S^o)$ for all $s < y_1^*$. First, consider $s^o < s < y_1^*$, and notice that the optimality of s^o implies that $c(s^o, S^o) \geq G(s^o + 1)$, and since $-G(\cdot)$ is unimodal $G(S - j) \leq c(s^o, S^o)$ for $j = S - s, \dots, S - s^o - 1$. Consequently,

$$\begin{aligned} c(s^o, S) &= \frac{K + \sum_{j=0}^{S-s-1} m(j)G(S-j) + \sum_{j=S-s}^{S-s^o-1} m(j)G(S-j)}{M(S-s^o)} \\ &= \frac{c(s, S)M(s, S) + \sum_{j=S-s}^{S-s^o-1} m(j)G(S-j)}{M(S-s^o)} \\ &\leq \frac{c(s, S)M(s, S) + \sum_{j=S-s}^{S-s^o-1} m(j)c(s^o, S^o)}{M(S-s^o)} \\ &= \beta c(s, S) + (1 - \beta)c(s^o, S^o), \end{aligned}$$

where $\beta = \frac{M(S-s)}{M(S-s^o)}$. Thus for $s^o < s < y_1^*$, $c(s^o, S)$ is dominated by a convex combination of $c(s, S)$ and $c(s^o, S^o)$. But then, $c(s^o, S) \geq c(s^o, S^o)$ implies $c(s, S) \geq c(s^o, S^o)$.

Now, for $s < s^o$, the fact that $G(S-j) \geq c(s^o, S^o)$ for $j = S-s^o, \dots, S-s-1$, allow us to write

$$c(s, S) \geq \gamma c(s^o, S) + (1-\gamma)c(s^o, S^o),$$

where $\gamma = \frac{M(S-s^o)}{M(S-s)}$, and consequently $c(s, S) \geq c(s^o, S^o)$. \square .

Thus, given (s^o, S^o) , we can easily identify an improving S' by simply comparing $c(s^o, S^o)$ and $c(s^o, S')$. If S' improves on S^o , then we want to find an optimal reorder-level s' for S' . The following lemma restricts the search for s' to s^o, \dots, \bar{s} .

Lemma 7 *Assume that $s^o \leq \bar{s}$ is an optimal reorder-level for S^o and that S' improves on S^o , then there exists an optimal reorder-level s' for S' with $s' \in \{s^o, \dots, \bar{s}\}$.*

Proof: Given S' we know from Proposition 1 that there exists an optimal reorder-level $\leq \bar{s}$. Let s' be the largest optimal reorder-level ($\leq \bar{s}$) for S' . Then $G(s'+1) < c(s', S') \leq c(s^o, S') < c(s^o, S^o) \leq G(s^o)$. Since $-G(\cdot)$ is unimodal it follows that $s^o \leq s'$. \square .

We are now ready to present an algorithm to find an optimal (s^*, S^*) policy.

Algorithm.

Let y^* be a minimizer of $G(\cdot)$.

Step 0. (Initial Solution)

$S^o = y^*$;

$s = y^* - 1$;

DO WHILE $c(s, S^o) > G(s)$;

$s = s - 1$;

ENDO;

$c^o = c(s, S^o)$, $S = S^o + 1$;

Step 1 (Main Step)

DO UNTIL $G(S) > c^o$;

IF $c(s, S) < c^o$;

$S^o = S$;

DO WHILE $c(s, S^o) \leq G(s+1)$;

$s = s + 1$;

ENDO;

$c^o = c(s, S^o)$;

ENDIF;

$S = S + 1$;

ENDO;

END;

4 Multi-echelon Systems

Consider a serial system comprised of J stages arranged in series, where external demands occur only at stage J . Stage $j = 2, \dots, J$ is replenished by stage $j-1$, and stage 1 replenishes from an outside supplier. We assume that the outside supplier never runs out of stock. Let L_j and D_j denote the lead time and the lead time demand at stage $j = 1, \dots, J$. Let h'_j denote the holding cost at stage j . We assume that

$$h'_1 \leq h'_2 \leq \dots \leq h'_J.$$

For convenience define $h'_0 = 0$. Then the echelon holding cost at stage j is defined to be

$$h_j = h'_j - h'_{j-1} \geq 0, \quad j = 1, \dots, J.$$

Finally, let b be the backorder cost at stage J .

Consider now a local base stock inventory control system where stage j monitors its inventory position (inventory on hand plus on order minus backorders) and places orders to keep its inventory position at s'_j , $j = 1, \dots, J$.

In order to assess the cost of this policy, we need to determine the inventory at each stage, the inventory on transit, and the backorders at stage J .

The inventory in transit from stage j to stage $j + 1$, is equal to $\mu_{j+1} = ED_{j+1}$. If we charge this inventory at rate h'_j , the in transit inventory cost is $\sum_{j=1}^{J-1} h'_j \mu_{j+1}$. To compute the inventory at each stage and the backorders at stage J . We use the following recursion: Set $B_0 = 0$, then

$$B_j = (B_{j-1} + D_j - s'_j)_+$$

and

$$I_j = (s'_j - D_j - B_{j-1})_+$$

denote the backorders and the inventory at stage $j = 1, \dots, J$. Letting $B = B_J$, the average cost can be written as

$$c(s'_1, \dots, s'_J) = \sum_{j=1}^J h'_j E I_j + b E B + \sum_{j=1}^{J-1} h'_j \mu_{j+1}.$$

For $j = 2$ the cost can be written as

$$c(s'_1, s'_2) = h'_1(s'_1 - \mu_1) - h_2 E(D_1 - s'_1)_+ + h'_2(s'_2 - \mu_2) + (h'_2 + b)E((D_1 - s'_1)_+ + D_2 - s'_2)_+$$

which is not necessarily convex in (s'_1, s'_2) .

Consider now a central control policy where stage j monitors its *echelon* inventory position (inventory on hand at stage j plus all the inventory that has gone through stage j but has not yet left the system plus inventory on order at stage j minus backorders at stage J) and places orders to keep the echelon inventory position at s_j .

It can be shown that for serial systems both policies are equivalent, and that

$$s_j = \sum_{i=j}^J s'_i.$$

The net echelon inventory at stage j can be obtained recursively as follows: Let $IN_0 = \infty$ then

$$IN_j = \min(IN_{j-1}, s_j) - D_j.$$

for $j = 1, \dots, J$.

With this notation, the average cost can be written as:

$$c(s_1, \dots, s_J) = \sum_{j=1}^J h_j E IN_j + (b + h'_J) E B.$$

The resulting problem can be solved sequentially as follows:

Let

$$C_J(y) = E\{h_J(y - D_J) + (b + h'_J)(D_J - y)_+\}.$$

This is a newsvendor problem. Let s_J^* be the minimizer of $C_J(y)$, and set $\underline{C}_J(x) = C_J(x)$ if $x < s_J^*$ and equal to $C_J(s_J^*)$ otherwise. In other words,

$$\underline{C}_J(x) = C_J(\min(x, s_J^*)).$$

Now, suppose that you have computed s_J^* and that you have constructed the function $\underline{C}_J(x)$. Then define

$$C_{j-1}(y) = E\{h_{j-1}(y - D_{j-1}) + \underline{C}_j(y - D_{j-1})\}.$$

Let s_{j-1}^* denote the minimizer of $C_{j-1}(y)$ and repeat the procedure until $j = 1$. At this point you have obtained the numbers s_1^*, \dots, s_J^* . To construct a local control policy we can proceed as follows: If the s_j^* are non-increasing, the local base stock levels are given by

$$s'_j = s_j^* - s_{j+1}^*.$$

Otherwise, define

$$\bar{s}_j^* = \min_{i \leq j} s_i^*$$

and proceed as above to compute the local control policy. This provides an efficient procedure to solve the multi-echelon serial inventory problem. But once again, the algorithm fails to give us an intuitive understanding of the system.

Recently Shand and Song [6] have developed bounds on the optimal echelon base stock levels based on the observation that for each stage j , the optimal echelon base stock level s_j^* depends on the upstream stages $\{1, \dots, j-1\}$ only through the sum of the echelon holding cost rates at these stages $\sum_{i=1}^{j-1} h_i$ and not on the base stock levels s_1^*, \dots, s_{j-1}^* . Based on this observation they are able to show that

$$s_j^l \leq s_j^* \leq s_j^u, \quad j = 1, \dots, J \quad (10)$$

where

$$\begin{aligned} s_j^l &= F_j^{-1} \left(\frac{b + \sum_{i=1}^{j-1} h_i}{b + \sum_{i=1}^J h_i} \right) \\ s_j^u &= F_j^{-1} \left(\frac{b + \sum_{i=1}^{j-1} h_i}{b + \sum_{i=1}^j h_i} \right) \end{aligned}$$

where F_j is the cumulative distribution of $\sum_{i=j}^J D_i$. In addition, these authors show that the optimal cost is bounded by

$$\begin{aligned} \sum_{j=1}^{J-1} h'_j \mu_{j+1} + G_1^l(s_1^u) &\leq c(s_1^*, \dots, s_J^*) \\ &\leq \sum_{j=1}^{J-1} h'_j \mu_{j+1} + G_1^u(s_1^l). \end{aligned}$$

where

$$G_1^l(y) = E[h_1(y - D[1, J])^+] + (b + h[1, J])(D[1, J] - y)^+$$

and

$$G_1^u(y) = E[h[1, J](y - D[1, J])^+] + (b + h[1, J])(D[1, J] - y)^+,$$

where we have used the notation $h[j, k] = \sum_{i=j}^k h_i$ and $D[j, k] = \sum_{i=j}^k D_i$. In their experiments the authors find that

$$c(s_1^*, \dots, s_J^*) \simeq \sum_{j=1}^{J-1} h'_j \mu_{j+1} + G_1^u(s_1^l).$$

References

- [1] Federgruen, A. and Y.S. Zheng (1992) "An Efficient Algorithm for Computing an Optimal (Q, r) Policy in Continuous Review Stochastic Inventory Systems," *Operations Research*, 40, 808-813.

- [2] Gallego, G. (1998) "New Bounds and Heuristics for (Q, r) Policies." *Management Science*, 44, 219-233.
- [3] Hadley, G., and T. Whitin' *Analysis of Inventory Systems*, Prentice Hall. 1963.
- [4] Veinott, A. F. (1965) Lecture Notes Stanford University.
- [5] Scarf, H. (1958) "A Min-Max Solution of an Inventory Problem," *Ch. 12 in Studies in The Mathematical Theory of Inventory and Production, Stanford Univ. Press.*
- [6] Shang, K., and J. Song (2001) "Newsvendor Bounds and Heuristics for Optimal Policies in Serial Supply Chains. Working paper University of California at Irvine. Irvine, CA.
- [7] Zheng, Y.S. (1992) "On Properties of Stochastic Inventory Systems," *Management Science*, 38, 87-103.
- [8] Zheng, Y.S. and A. Federgruen (1991) "Finding Optimal (s, S) Policies is About as Simple as Evaluating a Single Policy," *Operations Research*, OR 39, 654-665.
- [9] Zipkin, P. (1993) Teaching notes. Graduate School of Business, Columbia University, New York, NY