1 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.

1.1 Finite Horizon Models

Let $D_1, \ldots, D_T$ 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 $-x_t$ units is carried form the previous period. Let $y_t - x_t \geq 0$ denote the size of the order in period t resulting in a procurement cost $c_t(y_t - x_t)$ and an increase of the inventory level to y_t. Since D_t units are demanded during the period, the inventory level at the beginning of period $t+1$ is given by

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

If $y_t = y$ the loss function in period t is given by

$$G_t(y) = h_t E(y - D_t)_+ + b_t E(D_t - y)_+ \tag{1}$$

where h_t is the overage or holding cost and b_t is the underage or backorder penalty cost. Notice that the interpretation of h_t and b_t for period $t \leq T$ is different from that of period $T+1$ in that in period $T+1$ we would typically salvage remaining items and either produce or reimburse customers if there are backlogs.

Before continuing with the formulation, we remark that the cost function G_t is convex and also coercive. A function is coercive if it grows to ∞ as $|y|$ goes to ∞. The results that we are about to obtain would continue to hold for general convex and coercive functions G_t, and are not limited to specific form of Equation (1).

Let $C_{T+1}(x_{T+1})$ be an arbitrary cost function on the inventory level at the end of period T (beginning of period $T+1$), 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,

$$C_t(x_t) = \min_{y \geq x_t} \{c_t(y - x_t) + G_t(y) + \alpha E C_{t+1}(y - D_t)\} \tag{2}$$

represents a recursive, Dynamic Programming equation that can be solved backwards starting with period T.

It can be shown that if $C_{T+1}(\cdot)$ is convex and coercive, then $C_t(\cdot)$ is convex and coercive for all $t = 1, \ldots, T$. Let y_t^* be a minimizer of

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

Then an optimal policy is to order $(y_t^* - x_t)^+$ units in period t. This class of policies is known as order-up-to policies. The idea is that we order up to y_t^* in period t if $x_t < y_t^*$ and do not to order otherwise.

As an example of a convex terminal cost functions C_{T+1}, consider the case where left over stock, $x_{T+1} > 0$, is salvaged at a unit price s_{T+1} and and backlogged sales, $x_{T+1} < 0$ are cancelled by at a unit cost p_{T+1}. Then, $c_{T+1}(y) = -s_{T+1}y_+ + p_{T+1}y_-$ is convex as long as $p_{T+1} \geq s_{T+1}$. Another form of C_{T+1} that is often used is $C_{T+1}(x) = -c_{T+1}x$. This is a situation where excess units are salvaged at c_{T+1} and excess demand is satisfied by producing at unit cost c_{T+1}.

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 the data and the demand distributions.

1.1.1 The Myopic Policy

We now describe a myopic policy that is frequently used in practice. The advantage of the myopic policy is that the computations reduce to that of solving one Newsvendor problem for each period $t = 1, \ldots, T$, and thus avoid the computational effort of solving the Dynamic Programming problem (2).

To develop this policy we need to write a slightly different but equivalent set of recursive equations. To this end let

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

$M_t(x_t)$ is the expected cost-to-go $C_t(x_t)$ starting with x_t units of inventory plus the period t value of the x_t units of inventory. With this definition, the recursion becomes

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

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

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

and orders $(y_t^m - x_t)^+$ units in period t, where y_t^m minimizes the *current* expected cost

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

If the demand is continuous, then y_t^m satisfies

$$P(D_t \leq y) = \frac{b_t - m_t}{h_t + b_t}.$$

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

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

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

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

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

Another way of deriving the myopic policy is to write down the total cost over the entire horizon and then separate the terms that depend on the decision made in period t. To this end, let

$$C_1(x_1) = \min_{y_t \geq x_t} E[\sum_{t=1}^{T} \alpha^{t-1}\{c_t(y_t - x_t) + G_t(y_t)\} + \alpha^T C_{T+1}(x_{T+1})]$$

Notice that y_t appears in the sum as $\alpha^t\{(c_t - \alpha c_{t+1})y_t + G_t(y_t)\}$. If $C_{T+1}(x) = -c_{T+1}x$, then can write

$$C_1(x_1) = -c_1 x_1 + \min_{y_t \geq x_t} E \sum_{t=1}^{T} \alpha^{t-1}\{(c_t - \alpha c_{t+1})y_t + G_t(y_t)\} + \sum_{t=1}^{T} \alpha^t c_{t+1} E D_t.$$

We can now see that the myopic policy minimizes the cost function term-by-term, but ignores the possible interactions among the terms. However, if the myopic solution is such that $y^m_{t+1} \geq y^m_t - D_t$ with probability one, then the decisions in one-period do not preclude us from achieving the minimum cost in the next period, so the myopic policy is optimal in this case. This occurs, for example, if the y^m_t's are non-decreasing in t, so a natural question to ask is when can we guarantee that the y^m_t's are non-decreasing in t. One such case is when the ratios $(b_t - m_t)/(h_t + b_t)$ is independent of t and $Pr(D_t \leq y)$ is non-decreasing in t, or equivalently if the sequence of random variables is *stochastically* increasing.

1.2 Infinite Horizon, Stationary Models

If all the costs are stationary, i.e., $c_t = c, h_t = h$ and $b_t = b$ for all t, and the demands are independent and identically distributed (IID), then finite-horizon discounted costs (when $\alpha < 1$) converge, so the DP (2) can be written as

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

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 $(y^m - x)^+$ 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

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

Surprisingly, the myopic policy is optimal, under the mild assumption that D takes only non-negative values. 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

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

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

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

Solving for M^* yields

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

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.

Finally, notice that for $x \leq y^m = y^*$,

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

and this can be interpreted as the cost of the safety stock $c(y^* - \mu)$ minus the cost of the inventory already available cx, plus the discounted purchasing and inventory related costs $(c\mu + G(y))/(1-\alpha)$.

The policy of ordering up to y^* works as follows. If x is initially greater than y^* we do nothing until x drops below y^*. Once x drops below y^* and we place the initial order $y^* - x$, all subsequent orders will be equal to the previous period demand. To see this, suppose that we order up to y^* at the beginning of period t. Then $x_{t+1} = y^* - D_t$, so $y^* - x_{t+1} = D_t$ is the amount to be ordered at the beginning of period $t + 1$. This policy is also known as a base-stock policy because orders are placed in each period to restore the inventory to y^*.

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

$$P(D \leq y) = \frac{b}{h+b}. \tag{3}$$

Also, as α increases to one, the discounted cost goes to infinity and it makes more sense to talk about the average cost per period. It can be shown, e.g., by using the vanishing discount cost method, that the policy that sets y^* as prescribed in equation (3) is indeed an optimal solution for the average cost case.

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

1.3 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 $\{t, \ldots, t + L\}$ 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

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

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,

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

This formulation is equivalent to (2) 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

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

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

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

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

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

1.3.1 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 although this assumption may be inconsistent with the assumption that demands are time-independent. If we are willing to assume that orders do not cross, then the problem can be solved, at least approximately, 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

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

These results are direct applications of the well known formulas:

$$E[\sum_{i=1}^{N} X_k] = E[E[\sum_{i=1}^{N} X_k|N]]$$

and

$$\text{Var}[\sum_{i=1}^{N} X_k] = \text{Var}[E[\sum_{i=1}^{N} X_k|N]] + E[\text{Var}[\sum_{i=1}^{N} X_k|N]].$$

and can be found on on page 153 in reference [**?**].

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 insidious effect of randomness in the lead time is even worse than the direct comparison between the standard deviations would indicate.