Chapter 1
Preliminaries

Contents

1.1 Deterministic Demand

When demand is known with certainty, the problem of managing perishables is straightforward for the most part. Consider first the basic EOQ model. Suppose the demand rate is λ, the fixed cost of placing new orders is K, and the holding cost per unit time is h. Then, it is well known that the optimal order size is

$$Q^* = \sqrt{\frac{2K\lambda}{h}}$$

and the optimal time between placement of orders is $T^* = Q^*/\lambda$. Suppose now that the item has a usable lifetime of m. All deliveries are assumed to be of fresh units only. Then, there are two cases: (a) $T^* \leq m$ and (b) $T^* > m$. In case (a) the optimal policy remains the same, since in each order cycle, all units are consumed by demand before they expire. However, in case (b) if Q^* is ordered at the beginning of the cycle, there will be positive inventory on hand at time m, which will have outdated and must be discarded, and a new order placed at that time. Notice, however, that if we reduce the order quantity from $Q^* = \lambda T^*$ to $Q = \lambda m < Q^*$, then the cycle length will remain at m, no units will expire and holding costs will be reduced, since average inventory will be reduced from $Q^*/2$ to $Q/2$. Hence, the modification of the standard EOQ model to include perishability is straightforward.

However, not all deterministic perishable inventory problems are solved so easily. In particular, consider the deterministic nonstationary production planning

S. Nahmias, *Perishable Inventory Systems*, International Series in Operations Research & Management Science 160, DOI 10.1007/978-1-4419-7999-5_1,

problem. Demands over an n period planning horizon are known constants, say $(r_1, r_2, \ldots, r_n)$. Costs include holding, h_i, set-up, K_i, and marginal production cost, C_i in each period. Then, Wagner and Whitin (1958) showed (in the infinite lifetime case) that an optimal policy has the following structure. If starting inventory is zero, then the order quantity in each period is either zero or exact requirements – namely, the sum of requirements in the current to some future period. Furthermore, an optimal policy only orders in periods when the starting inventory is zero. As a consequence of this result, one only needs to determine the periods in which ordering takes place, thus reducing the calculations significantly.

It turns out that an exact requirements policy may not be optimal when perishability is introduced. A counterexample is presented in Chap. 7, and methods for resolving this problem is discussed there. When demand is nonstationary, finding optimal order policies for a fixed life inventory is not trivial. However, the vast majority of the research on ordering policies for perishables has focused on stochastic demand – a significantly more difficult problem.

1.2 Periodic Review Versus Continuous Review

Demand uncertainty and (fixed life) perishability combine to result in challenging and complex problems. Stochastic perishable inventory problems fall into one of the two basic categories: periodic review or continuous review.

Most of the research in inventory theory assumes inventory levels are reviewed periodically. This means that the state of the system (on hand inventory) is known only at discrete points in time. This assumption is appropriate, for example, if inventory levels are checked once a day, once a week, etc. The landmark collection of Arrow et al. (1958) assumed periodic review in every case considered, and set the stage for much of the subsequent research on inventories. From a practice point of view, it is probably true that most inventory systems were periodic review 50 years ago. Today, however, point-of-sale scanners and automated inventory control systems have made true continuous review more common.

There are two reasons why continuous review has grown in importance. First, with automated inventory control systems computers can automatically trigger orders when inventory levels hit predetermined levels. Second, continuous review models often are able to provide simple approximations to complex problems that are difficult to solve with periodic review formulations.

Perishable inventory research has also evolved along the two separate tracks of periodic review and continuous review. The periodic review track generalizes the kind of models considered by Arrow et al. (1958), among many others, to incorporate perishability. The continuous review track is largely an outgrowth of the theory of queues with impatient servers. An impatient customer is one who leaves the queue if they have not been served by a fixed time. Queueing models with impatient customers are discussed in detail in Chap. 9.

1.3 Periodic Review Preliminaries

As noted earlier, we assume that the on hand inventory level is known only at discrete points in time, which are labeled periods. Assume that periods are numbered 1, 2,.... Demands in successive periods are not known, but are assumed to be random variables, $D_1, D_2, \ldots$ with a known probability distribution. For convenience, assume that the demand distribution is continuous with cumulative distribution function (CDF) $F(x)$ and probability density function (PDF) $f(x)$. (Note that basic results have been shown to carry over to the discrete demand case as well. Also, virtually all of the results carry over to nonstationary demand. Stationarity of the demand distribution is assumed for notational convenience.)

Assume that new orders are always of fresh units that have a usable lifetime of m periods. A little reflection should convince the reader that it is necessary to track the entire age distribution of the on hand inventory in order to determine outdates each period. Hence, the system state is described by a vector $\mathbf{x} = (x_{m-1}, x_{m-2}, \ldots, x_1)$ where x_i is the number of units on hand with i useful periods of life remaining. Note that there are many notation options for the state vector. The state could be defined in terms of age rather than remaining lifetime and the vector could be numbered in order of oldest to youngest rather than vice versa, as we have done. This convention was chosen to reflect that aging occurs in the direction left to right, like the flow of the English language, and that the decision variable, y, can be equated to x_m and placed in the proper position in the vector. Note that x_0 would represent the number of units on hand that have just expired or outdated. We do not need to carry x_0 in the state vector since outdated units are assumed to leave the system. We use the convention throughout that boldface $\mathbf{x}$ is the vector of on hand inventories of each age level, and $x = \sum_{i=1}^{m-1} x_i$ is a scalar quantity representing the total on hand inventory.

The necessity to define a vector valued state variable is only one of the things that separate the perishable inventory problem from the conventional nonperishable problem. As we see, several other concerns arise as a result of perishability. One is the sequence that items are issued to meet demand. Note that there is a substantial literature on optimal issuing policies independent of the ordering problem. The appropriate assumption concerning issuing policies depends on whether the producer or consumer chooses which items satisfy demand. If the producer determines the issuing policy, it is clearly in his interest to issue items on an oldest first basis (known in accounting parlance as FIFO for first in first out). If the consumer determines the issuing policy, it is likely that the consumer will choose the freshest items, resulting in units issued in last in first out (LIFO) sequence. The vast majority of the perishable inventory literature assumes FIFO issuing, and we do so as well unless stated otherwise. Clearly, FIFO is most cost-efficient and results in minimum outdating. (A third alternative, which might be appropriate in some contexts, is to issue the items in a random order. To our knowledge, random issuing policies have not been considered in the context of optimal ordering policies for perishables.)

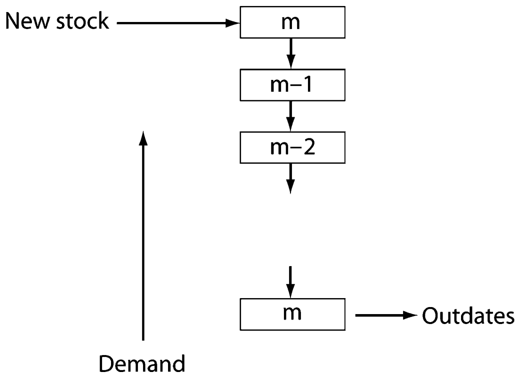

Fig. 1.1 The flow of demand and product in an FIFO fixed life perishable inventory system

If items are issued according to FIFO, then the aging and demand processes travel in opposite directions. To see what this means, consider the representation of the system state in Fig. 1.1. Bins are labeled $m-1$, $m-2$, . . ., 1, where the contents of bin i are the number of units on hand with i useful periods of life remaining. At the end of each period, all contents of a bin are moved to the next lower bin, and the contents of bin 1 outdate and must be discarded (or salvaged). Because of the FIFO assumption, demand depletes first from bin 1, then from bin 2, etc. Excess demands may be lost or backordered. If excess demands are backordered, then this is reflected in a negative value of x_{m-1}.

Consider now the ordering policy. The optimal number of units to order each period is function of the state vector, $\mathbf{x}$. We represent this function as $y(\mathbf{x})$. As we see, $y(\mathbf{x})$ is a complex nonlinear function of the state variable x. We assume that costs are assessed in the usual way for finite horizon periodic review inventory systems. At the end of each period, the total inventory is determined. If it is positive, assess a holding cost of h per unit held per period. If it is negative (which occurs when excess demands are backordered), assess a cost of p per unit of unsatisfied demand. Furthermore, we assume a marginal order cost only. That is, there is a cost of c per unit ordered. For now, assume that there is no fixed order cost. Finally, we come to the issue of how to assess the cost for items that must be discarded due to outdating. Let θ be the cost of disposing of outdated units. If D is the demand in a period, then the number of units outdating at the end of the period when starting inventories are $\mathbf{x}$, is $\max(x_1 - D, 0)$, which we represent as $(x_1 - D)^+$.

We now face the first issue. The astute reader will notice that the outdating cost, $\theta E(x_1 - D)^+$, is independent of the decision variable y. That means that any single period model ignores the effects of outdating. In fact, one would need to churn through at least m periods of a dynamic programing formulation before the outdating penalties of over ordering would be reflected in the optimal order policy. (This was, in fact, the approach taken in Fries 1975.)

Suppose, however, that one were interested in constructing a one period model that reflected the outdating penalties of over ordering. How could this be done? Let $D_1, D_2, \ldots$ represent demands in successive periods, starting with the current period. Then, the current order, y, would not outdate until m periods into the future, if it had not been consumed by demand by that time. Consider how one would determine the expected outdating of the current order y, m periods into the future.

Define the following sequence of random variables:

$$R_1 = y + \sum_{i=1}^{m-1} x_i$$
$$R_2 = [y + \sum_{i=2}^{m-1} x_i - (D_1 - x_1)^+]^+$$
$$R_3 = [y + \sum_{i=3}^{m-1} x_i - (D_2 + (D_1 - x_1)^+ - x_2)^+]^+$$

etc.

Then, R_i represents the amount of the current inventory on hand i periods into the future, assuming we start with $\mathbf{x}$ and order y. Each value of R_i is a random variable, since it is a function of the future demands. Note the use of the imbedded + functions are necessary to keep track of units lost due to outdating.

To simplify the notation, define the following sequence of random variables recursively:

$$\begin{aligned} B_0 &= 0 \\ B_1 &= (D_1 - x_1)^+ \\ &\vdots \\ B_j &= (D_j + B_{j-1} - x_j)^+ \text{ for } 1 \le j \le m-1. \end{aligned}$$

Interpret the random variable B_j as the total unsatisfied demand in period j after depleting the on hand inventory that would have outdated in period j. It follows that

$$R_m = [y - (D_m + B_{m-1})]^+$$

represents the amount of the current order y, that outdates in m periods.

The goal of the analysis is to compute the expected value of R_m and incorporate this into the one period model.

Define $G_n(t;\, \mathbf{w}_{n-1}) = P\{D_n + B_{n-1} \le t)$ where $\mathbf{w}_i = (x_i, x_{i-1}, \ldots, x_1)$. Note that $\mathbf{w}_{m-1} = \mathbf{x}$.

We present the first result without proof, which is based on a standard induction argument. Details can be found in Nahmias (1972).

Theorem 1.1. $G_n(t;\, \mathbf{w}_{n-1}) = \int_0^t G_{n-1}(v + x_{n-1};\, \mathbf{w}_{n-2}) f(t-v)\mathrm{d}v.$

Theorem 1.2. $E(R_m) = \int_0^y G_m(t; \mathbf{x})\mathrm{d}t$

Proof. It is well known that for any nonnegative random variable, X, the expectation may be computed two ways:

$$E(X) = \int_0^\infty x f(x)\mathrm{d}x = \int_0^\infty [1 - F(x)]\mathrm{d}x.$$

We have $P\{R_m \leq t\} = P\{y - (D_m + B_{m-1}) \leq t\} = 1 - G_m(y - t;\, \mathbf{x})$ for $t \geq 0$ Since R_m is a nonnegative random variable, the result follows from the second representation of the expected value above (after a change of variable). □

1.4 A One Period Newsvendor Perishable Inventory Model

Most readers should be familiar with the classic newsvendor model. A newsvendor must decide at the beginning of each day how many newspapers to purchase. Daily demand is not known, but is assumed to follow a known probability distribution. Let y be the number of newspapers purchased and D the demand. There are two penalties: overage (ordering too much) and underage (ordering too little).

Now, let us consider the perishable inventory model. The penalty for ordering too much is the future penalty of outdating, at θ per unit, and the penalty for ordering too little is penalty cost for excess demand, at p per unit. Hence, a sensible expected one period cost function for the perishable inventory problem is:

$$L(\mathbf{x}, y) = p \int_{x+y}^{\infty} [t - (x + y)]\, f(t)\mathrm{d}t + \theta \int_{0}^{y} G_m(t; \mathbf{x})\mathrm{d}t.$$

It is easy to show that $L(\mathbf{x}, y)$ is convex in y (and is strictly convex as long $f(t) > 0$ for all $t > 0$. Hence, the optimal order quantity, y, for this simple model satisfies:

$$\frac{\partial L(\mathbf{x}, y)}{\partial y} = -p(1 - F(x + y)) + \theta G_m(y;\, \mathbf{x}) = 0.$$

The optimal one period solution, say $y^*(\mathbf{x})$, is a nonlinear function of the entire state vector, $\mathbf{x}$. In this case, $y^*(\mathbf{x}) > 0$ for all positive real vectors $\mathbf{x}$. In addition, as we see in the analysis of the dynamic problem, $y^*(\mathbf{x})$ is decreasing in each component of the state vector, $\mathbf{x}$, but at less than unit rate.

Somewhat sharper results can be obtained when we add holding and marginal order costs.

Theorem 1.3. Suppose that in addition to penalty and outdate costs, we also include marginal order cost at c per unit ordered, and a unit holding cost, h, charged against each unit on hand at the end of the period. Then, the optimal solution has the following form: If $x < \bar{x}$ order $y^*(\mathbf{x})$ solving

$$c + hF(x + y) + p(1 - F(x + y)) + \theta G_m(y;\, \mathbf{x}) = 0$$

where $\bar{x}$ solves

$$c + hF(\bar{x}) - p(1 - F(\bar{x})) = 0.$$

If $x \geq \bar{x}$, no order is placed.

Proof. The expected one period cost function is now:

$$L(\mathbf{x},y) = cy + h \times \int_{0}^{x+y} (x+y-t)f(t)\mathrm{d}t + p \int_{x+y}^{\infty} (t-(x+y))f(t)\mathrm{d}t + \theta \int_{0}^{y} G_m(t;\mathbf{x})\mathrm{d}t.$$

Convexity in y is easy to show so that the minimizing value of y occurs, where the partial derivative of $L(\mathbf{x},y)$ vanishes. The partial derivative is given by:

$$\frac{\partial L(\mathbf{x},y)}{\partial y} = c + hF(x+y) - p(1 - F(x+y)) + \theta G_m(y;\, \mathbf{x})$$

thus giving the definition of the optimal ordering quantity. Notice that If $x < \bar{x}, \left.\frac{\partial L(\mathbf{x},y)}{\partial y}\right|_{y=0} < 0$ and if $x \geq \bar{x}, \left.\frac{\partial L(\mathbf{x},y)}{\partial y}\right|_{y=0} \geq 0$, thus establishing that $y^*(\mathbf{x})$ is positive in the region $x < \bar{x}$ and zero in the region $x \geq \bar{x}$. □

In the next chapter, we extend this approach to a multiperiod dynamic model.