Chapter 3
Extensions of the Basic Multiperiod Dynamic Model

Contents

The basic model treated in Chap. 2 is the building block for several extensions. In many cases, these extensions are suggested by specific applications. Other extensions provide a wide range of applications of the basic theory.

3.1 Random Lifetime

A natural extension of the basic model is to allow for uncertainty in the product lifetime. The basic model is applicable for products with a known expiration date, since the time of outdating is known in advance. However, there are many classes of items whose useful lifetime cannot be predicted in advance. Fresh produce, meat, fowl, and fish would fall into this category. Even in blood banking, local blood banks may receive transfers from other banks that are not completely fresh, making the remaining lifetime uncertain. Perishable inventory models with uncertainty in the product lifetime could be used to model some classes of items subject to obsolescence as well.

Providing an accurate model of lead time uncertainty has long been difficult problem in inventory management (see Hadley and Whitin (1963), pp. 200–204, for a discussion of the issues involved.) The main issue is order crossing. That is, are orders placed on Tuesday allowed to arrive before orders placed on Monday? If successive lead times are independent random variables, then order crossing is inevitable. However, if one is placing orders with the same supplier, then order crossing does not make sense. If we do not allow order crossing, then successive lead times are dependent random variables, making the analysis very complex.

The simpler results for (Q,R) models with stochastic lead times (such as the formula presented in Hadley and Whitin on page 153), are based on simultaneously assuming independence of lead times and no order crossing.

As noted, if lead times are random and orders do not cross, then successive lead times are dependent random variables. Allowing for dependency is difficult, and is rare in operations research. A very innovative approach to this problem was developed by Kaplan (1970). Let $A_1, A_2, \ldots$ be a sequence of IID discrete random variables defined on the set $\{1, 2, \ldots, \tau\}$ where τ is the maximum value of the lead time in any period. In any period, when the event $\{A_i = k\}$ is realized, all orders k or more old periods arrive in the next period.

Nahmias (1977c) adapted this approach to the perishable inventory problem. Here, one assumes that if $\{A_i = k\}$ is realized, then all on hand inventory remaining in stock (after satisfying demand) that is at least k periods old at the end of period i will outdate at that time. The random variables $A_1, A_2, \ldots$ are *not* realizations of the lifetimes of successive orders, but the lifetime distribution can be derived from them.

Suppose that $P_k = P\{A_n = k\}$ independent of n. If the event $\{A_n = k\}$ is realized, then the one-period transfer function $\mathbf{s}(y, \mathbf{x}, t, k) = (s_{m-1}(y, \mathbf{x}, t, k), \ldots, s_1(y, \mathbf{x}, t, k))$ is:

for $k = 1$:

$$s_j(y, \mathbf{x}, t, 1) = 0 \qquad \text{if } 1 \leq j \leq m - 2$$
$$s_{m-1}(y, \mathbf{x}, t, 1) = -(t - x - y)^+$$

and for $2 \leq k \leq m$:

$$s_j(y, \mathbf{x}, t, k) = \begin{cases} 0 & \text{if } 1 \leq j \leq m - k \\ (x_{j+1} - (t - w_j)^+)^+ & \text{if } m - k + 1 \leq j \leq m - 1 \\ y - (t - x)^+ & \text{if } j = m - 1 \end{cases}$$

where $w_j = \sum_{i=1}^{j} x_i$, $x = \sum_{i=1}^{m-1} x_i$.

The sequence of events in a period is given as follows. First, the current state is observed, and based on that an order is placed for fresh stock. After the order arrives, the demand for the current period is realized, which then is satisfied to the extent possible on an FIFO basis. After demand is either satisfied or backordered (or lost), the aging process random variable, A_n is realized, and outdating determined. To understand the mechanics of the transfer function, consider the following example. Fix the current period at period 10. Suppose that $m = 4$, and the starting state in period 10 is $\mathbf{x} = (18, 12, 20)$ (i.e., 18 units remain from the order delivered one period ago, 12 units remain from the order delivered two periods ago, and 20 units remain from the order delivered three periods ago). Assume that the order quantity of fresh stock in period 10 is 23 units and the realization of demand is 28 units.

Consider how the system state evolves given all of the possible realizations of the aging process random variable, A_{10}. By FIFO, all of the 20 units are consumed first, followed by 8 units of the 12 units of two-period old stock. The vector of on hand inventories in period 10 after the arrival of fresh stock and the realization of the demand is (23, 18, 4, 0). If $A_{10} = 4$, there is no outdating, since all inventories of age 4 was depleted by demand, and the starting state in period 11 is (23, 18, 4). If $A_{10} = 3$, then the on hand inventory ordered three periods ago will outdate (4 units), and the starting state in period 11 is (23, 18, 0). If $A_{10} = 2$, then there will be $18 + 4 = 22$ units outdating at the end of period 10, and the starting state in period 11 is (23, 0, 0), and finally if $A_{10} = 1$ is realized all 45 units on hand outdate, and period 11 starts with zero inventory.

The functional equations defining an optimal policy are similar to those in Chap. 2, and are now given by:

$$C_n(\mathbf{x}) = \min_{y \geq 0} \Big\{ cy + L(x+y) + \theta \sum_{k=1}^{m} P_k \int_0^{w_{m-k+1}} (w_{m-k+1} - t) f(t) \mathrm{d}t$$

$$+ \alpha \sum_{k=1}^{m} P_k \int_0^{\infty} C_{n-1}\, \mathbf{s}(y, \mathbf{x}, t, k) f(t) \mathrm{d}t \Big\}.$$

The optimal policy when n periods remain in the horizon, $y_n(\mathbf{x})$, possesses virtually all of the same properties of the optimal policy discussed in Chap. 2 for the deterministic problem. The full statement of the appropriate theorem for analyzing the optimal policy can be found in Nahmias (1977c). Note that these functional equations charge the outdating cost in the current period rather than m periods into the future when the current order outdates (similar to Fries 1975). This was done to make the dynamic problem tractable, for as we see, computing the expected outdating of the current order is a very complex problem in itself.

If the lifetime of successive orders were independent, it would not be possible to compute the expected outdating of the current order, since future orders could outdate before the current order. However, due to our assumption that orders outdate in the same sequence that they enter stock, we are guaranteed that this will not occur. Hence, given a current order y and state vector $\mathbf{x}$, it should be possible to determine the expected outdating of the current order (which could be 1, 2,..., m periods into the future).

To understand the mechanics of the outdating process when the product lifetime is uncertain, consider another simple example. Let $j_1, j_2, \ldots$ correspond to successive realizations of the lifetime process, A_n. Suppose that $m = 5$, $j_1 = 4$, $j_2 = 4$, $j_3 = 2$, and the starting state in period 1 is (x_4, x_3, x_2, x_1). The goal is to compute the expected outdating of the order placed in period 1, y. Based on the realizations of the aging process, the remaining amounts of x_1 and x_2 (after satisfying demand in period 1) outdate at the end of period 1. The remaining amount of x_3 on hand at the end of period 2 after satisfying demand outdates at that time, and the remaining

amounts of y and x_4 outdate at the end of period 3. The expected amount of $y + x_4$ outdating at the end of period 3 is given by $\int_0^{y+x_4} G_3(u, x_3, x_1 + x_2)du$ so that the expected outdating of y only at the end of period 3 is $\int_0^{y+x_4} G_3(u, x_3, x_1 + x_2)du -$ $\int_0^{x_4} G_3(u, x_3, x_1 + x_2)du = \int_{x_4}^{y+x_4} G_3(u, x_3, x_1 + x_2)du$. This simple example should give the reader a glimpse of the complexities in the general calculation.

Define $r_k = \sum_{i=1}^k P_i$. Then, it can be shown that the expected outdating of the current order for y units is

$$\sum_{k=1}^m r_k \sum_{j_1=2}^m \cdots \sum_{j_{k-1}=k}^m \prod_{i=1}^{k-1} P_{j_i} \int_{v_k}^{y+v_k} G_k(u, \mathbf{v}(k-1))du,$$

where the vector terms $\mathbf{v}(k-1)$ are appropriate partial sums of the state vector $\mathbf{x}$ based on the realizations of the aging process, A_n. This is a very tedious calculation and it is unlikely that one would resort to trying to find an optimal policy except for small values of m. However, an explicit representation of the expected outdating of the current order can be used to construct simple approximations, similar to the manner that approximations were constructed for the deterministic lifetime problem.

The essence of the calculation is to find an expression for the unconditional probability that any new order has a lifetime of exactly k periods for all $1 \le k \le m$. We call this probability q_k. Based on the definition of the aging process, it follows that

$$q_k = \sum_{k=1}^m r_k \sum_{j_1=2}^m \cdots \sum_{j_{k-1}=k}^m \prod_{i=1}^{k-1} P_{j_i}.$$

It is not difficult to show that this expression can be simplified to:

$$q_k = r_k \prod_{i=1}^{k-1} (1 - r_i).$$

The idea is to use an approximate expression for the expected outdating of the current order, y, and utilize this expression to approximate the expected one-period cost function as well as the one-period transfer function, much as was done in Chap. 3 for the fixed life case. Again, one can use either the Nahmias's bounds or the Chazan and Gal bounds. Calculations reported in Nahmias (1977c) indicate that the Chazan and Gal bounds give slightly better results, and appear to improve as m increases while Nahmias's bounds degrade as m increases.

Using the Chazan and Gal bounds, the myopic expected one-period cost function one obtains is:

$$W(z) = c(1 - \alpha)z + L(z) + \sum_{k=1}^{m} q_k(\alpha^{k-1}\theta + \alpha c)\mu_k(z).$$

Since $W(z)$ is quasi-convex in z, the minimizing value occurs where the first derivative is zero so that z^* satisfies:

$$c(1 - \alpha) + L'(z) + \sum_{k=1}^{m} q_k(\alpha^{k-1}\theta + \alpha c)\left\{\frac{F(z/k) + F^{k^*}(z)}{2k}\right\} = 0.$$

Comparisons of this approximation with the optimal policy for lifetimes of 2 and 3 periods, and a variety of costs and demand distributions shows close agreement with the optimal.

3.2   Inclusion of a Set-Up Cost

All of the work reviewed thus far on the periodic review perishable inventory problem assumes there is no fixed cost (set-up cost) for placing an order. Fixed order costs can rarely be ignored in practice. In rare occasions, when deliveries are pre-scheduled on a daily or weekly basis, fixed order costs are sunk costs. However, it is more common that fixed costs determine the optimal frequency of orders, and cannot be ignored.

The results of this section are based on Nahmias (1978). Both optimal and approximate ordering policies are considered. The optimal policy is derived only for the one-period problem, but it is likely the multiperiod version possesses the same policy structure. The optimal policy requires calculation of two functions of the starting state, $s(\mathbf{x})$ and $y(\mathbf{x})$ with $s(\mathbf{x}) \le y(\mathbf{x})$. These functions determine the optimal policy as follows: if $s(\mathbf{x}) > 0$, the optimal policy is to place an order for $y(\mathbf{x})$, and if $s(\mathbf{x}) \le 0$, then no order should be placed. As in the case of no set-up cost, $y(\mathbf{x})$ minimizes the expected one-period cost function, $B(\mathbf{x}, y)$. The function $s(\mathbf{x})$ solves $B(\mathbf{x}, s(\mathbf{x})) = B(\mathbf{x}, y(\mathbf{x})) + K$, and $s(\mathbf{x}) \le y(\mathbf{x})$. Since $B(\mathbf{x}, y)$ is convex in y, these equations uniquely define $s(\mathbf{x})$. Note that Nahmias is able to derive only first period results, but given the fact that the structure of the optimal policy is the same for the first period problem as the multiperiod problem when there is no set-up cost present, it is reasonable to speculate that this is the case here as well.

Except for product lifetimes of two or three periods, the optimal policy is too complex to be practical. Hence, approximations are of particular interest. The natural form of an approximate policy is (s, S), since this policy is known to be

optimal for the nonperishable problem, and is easy to implement in practice. (An (s, S) policy is one in which an order is placed up to S if and only if the starting inventory position is less than s.) The approach to constructing an approximation follows closely the method outlined in Chap. 3 for the case of $K = 0$. That is, the expected outdate function is approximated by a function of the total on hand inventory using either the Nahmias or Chazan and Gal bounds, and this approximation is then used to construct a surrogate problem whose optimal solution is used to approximate the original problem. The surrogate problem has only a single dimensional state variable, and is solved by dynamic programing.

Extensive computations reported in Nahmias (1978) indicate that the approximation generally gives costs within 1% of the optimal, although because of the difficulty of computing an optimal policy, the comparisons were done for $m = 2$ only. However, the results would seem to indicate that this is a valid approach for approximating this problem.

3.3 Multiproduct Models of Perishables

Blood bank applications, in particular, have given rise to several multiproduct perishable inventory models. The earliest was explored by Nahmias and Pierskalla (1976). Blood banks store both whole blood and frozen packed red blood cells. Since frozen red blood cells have a lifetime of a year or longer, one can assume that the usable lifetime of frozen blood is infinite. Although frozen blood can be used in place of fresh blood when necessary, one would only want to use the frozen blood as a last resort. This is because the cost of extracting the red blood cells, freezing them, and then subsequently thawing them and combining them with plasma is considerably more expensive than the cost of extracting and storing fresh blood.

The model developed is a direct extension of the approach applied by Nahmias (1975a). Assume that at the beginning of each planning period prior to the realization of demand, orders for both product 1 (the perishable product) and product 2 (the nonperishable product) must be placed. There is a single demand source that depletes first from product 1 according to FIFO, and then product 2 until the demand is either satisfied or backordered. (Note: in the context of blood banking, it is probably more accurate to assume that excess demand is lost rather than backordered. This would mean that excess demand is made up by emergency shipments from other blood banks. The backorder assumption is not crucial to the analysis.)

Costs are charged separately for holding at h_1 for product 1 and h_2 for product 2. Outdating of product 1 is charged at θ and there is a penalty cost of p for unsatisfied demand. Let $\mathbf{x}$, x, and y be defined as earlier for the single product case. Define x^2 as the on hand inventory of product 2 at the start of any period, and let z be the on hand inventory of product 2 after ordering in any period so that

the order quantity of product 2 is $z - x^2$. Then, the one-period expected cost function has the form:

$$L(\mathbf{x}, y, z) = c_1 y + h_1 \int_0^{x+y} (x + y - t) f(t) dt + \theta \int_0^y G_m(u; \mathbf{x}) du + c_2 (z - x^2)$$

$$+ h_2 z F(x + y) + h_2 \int_{x+y}^{x+y+z} (x + y + z - t) f(t) dt + p \int_{x+y+z}^{\infty} (t - x - y - z) f(t) dt.$$

Assume the following relationships among the cost parameters:

(i) $0 \le h_2 \le h_1$
(ii) $0 < c_1 < c_2$
(iii) $p > (1 - \alpha) c_2$
(iv) $0 \le (1 - \alpha)(c_2 - c_1) + (h_2 - h_1) < \theta.$

The rationale for these assumptions is given as follows. Assumption (1) says that the cost of storing the nonperishable product is at least as large as the cost of storing the perishable product. Assumption (2) says that the cost of acquiring the nonperishable product exceeds the cost of acquiring the perishable product. Assumption (3) is common for single product inventory systems. It says that the unit cost of excess demand exceeds the one-period discounted cost of acquiring new nonperishable product. This guarantees that it is not more economical to incur stockouts than purchase new stock – otherwise, there would be no motivation to operate the system. The expression in the final case corresponds to the cost of acquiring one unit of product 2 and salvaging it one period later minus the cost of acquiring one unit of product 1 and salvaging it one period later plus the difference in the holding costs. If this term were negative, one would never order product 1. If this term exceeded the outdate cost, it would never be optimal to order product 2.

The functional equations defining an optimal policy for the multiperiod version of the problem are

$$C_n(\mathbf{x}, x^2) = \min_{y \ge 0, z \ge x^2} \left\{ L(\mathbf{x}, y, z) - c_2 x^2 + \alpha \int_0^{\infty} C_{n-1}(\mathbf{s}_1(\mathbf{x}, y, t), s_2(x + y, z, t)) f(t) dt \right\}.$$

In this case, the transfer function is m dimensional. The first $m - 1$ components, $\mathbf{s}_1(\mathbf{x}, y, t)$, are essentially the same as the single product form. The mth component is $s_2(x + y, z, t) = z - (t - (x + y))^+$ in the backorder case and $s_2(x + y, z, t) = [z - (t - (x + y))^+]^+$ in the lost sales case.

The structure of the optimal policy is shown in Fig. 3.1. When $m = 2$, the state vector is two dimensional, so the optimal policy can be represented graphically.
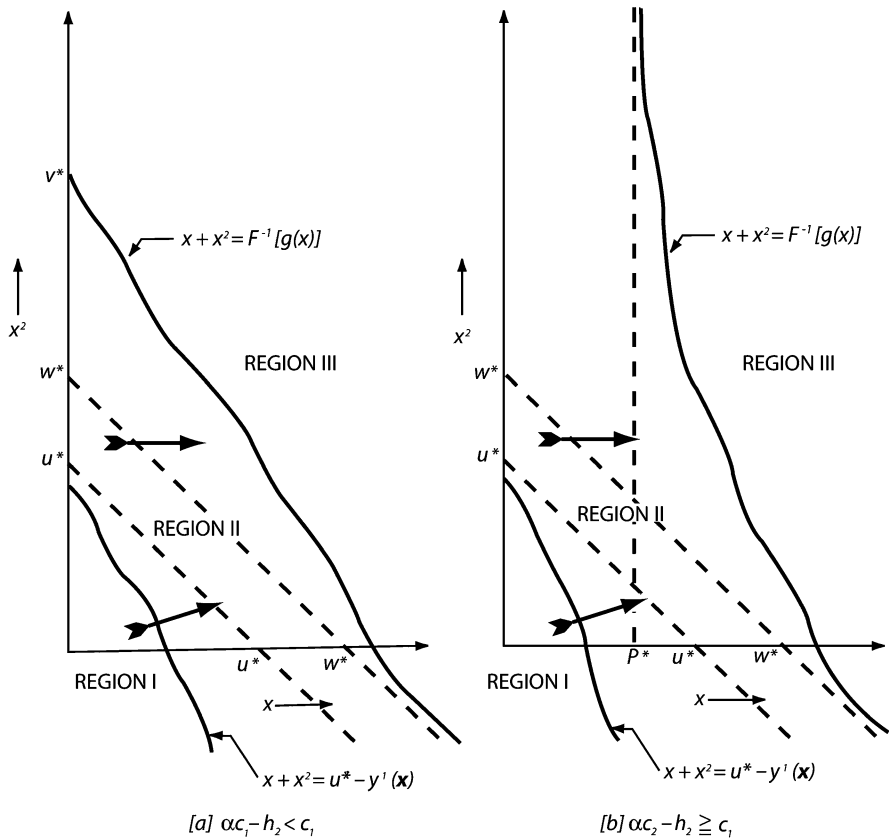
Fig. 3.1 The optimal ordering regions for a two-product perishable/nonperishable problem

Note that figure includes two cases: (a) where $\alpha c_1 - h_2 < c_1$ and (b) where $\alpha c_1 - h_2 \geq c_1$. The structure of the optimal policy is essentially the same in both cases, but the boundary between Regions II and III differ slightly. Region I corresponds to ordering both the perishable and nonperishable products. In Region II, one orders only the perishable product, and in Region III, one does not order. Note that in Region I, one orders so that the total stock on hand after ordering is u^* where

$$u^* = F^{-1}\left[\frac{p - c_2(1 - \alpha)}{p + h_2}\right].$$

What is unusual is that the line $x + x^2 = u^*$ lies completely outside of Region I. Typically, when there are ordering regions in a multiproduct inventory problem, one orders to the boundary of the region. In Region I, however, one orders to the

interior of Region II. Just the opposite is true in Region II. When ordering only the perishable product, one orders so that one remains short of the boundary between Regions II and III. This unusual behavior of the optimal policy is a direct result of the perishability of product 2. If both products were nonperishable, then the boundary between Regions I and II would be the line $x + x^2 = u^*$. However, the perishability adds a penalty to orders of product 2, which has the effect of tightening the boundary between Regions I and II. In the same way, when one is in Region II, the additional penalty of perishability results in a decreased order size so that one does not order to the boundary of Regions II and III. Note that there an apparent instability in Region II. Since one orders short of the boundary, it implies that one would want to order again, since after ordering the total stock on hand still lies within Region II. However, note that this is not exactly the case, since the state consists of one-period old inventory, and the order quantity consists of fresh inventory.

Note that Nahmias (1982) shows how the approximation techniques discussed in Chap. 3 can be applied to this problem. Although these approximations have never been tested, the success of the approximations in the single product case suggests that this approximation should be effective as well.

Deuermeyer (1979) considered a multiproduct perishable inventory model also suggested by a blood banking application. Consider two production processes, A and B. Process A produces two products (1 and 2) while Process B produces only product 2. Both products are perishable, and both face independent stochastic demands. In the blood banking application, product 1 is platelets (a blood component) and product 2 is whole blood. Process A separates some of the whole blood into platelets and packed red cells while Process B just produces whole blood.

Suppose that the lifetime of product 1 is m periods and the lifetime of product 2 is l periods. The state vector consists of the starting stocks of each product at each age level, represented as $(\mathbf{x}, \mathbf{w})$ having dimension $(m-1) + (l-1)$. The decision variable is the pair (a, b), which corresponds to the production quantities for Process A and Process B, respectively. Assume that a units of Type A production yields $\eta_1 a$ of product 1 and $\eta_2 a$ of product 2 while b units of Type B production yield b units of Product 2.

Deuermeyer develops a one-period model using the expected future outdating of both products in the one-period cost function as we did in Chap. 2. This makes it likely that the structure of the optimal one-period policy is the same as the structure of the optimal multiperiod policy. (Unfortunately, due to the complexity of the problem, he was able to obtain first period results only.)

The optimal policy structure is pictured in Fig. 3.2 when $m = l = 2$. The policy is defined by four regions in the $(\mathbf{x}, \mathbf{w})$ space. In Region I, both processes are used. The graph pictures the increase in the inventories of both products first due to Process A and then the increase in product 2 inventory due to Process B. In Region II, one only applies Process A (which increases the inventories of both products) while in Region III one applies only Process B so that only product 2 inventory is increased. Finally, in Region IV, neither process is applied.
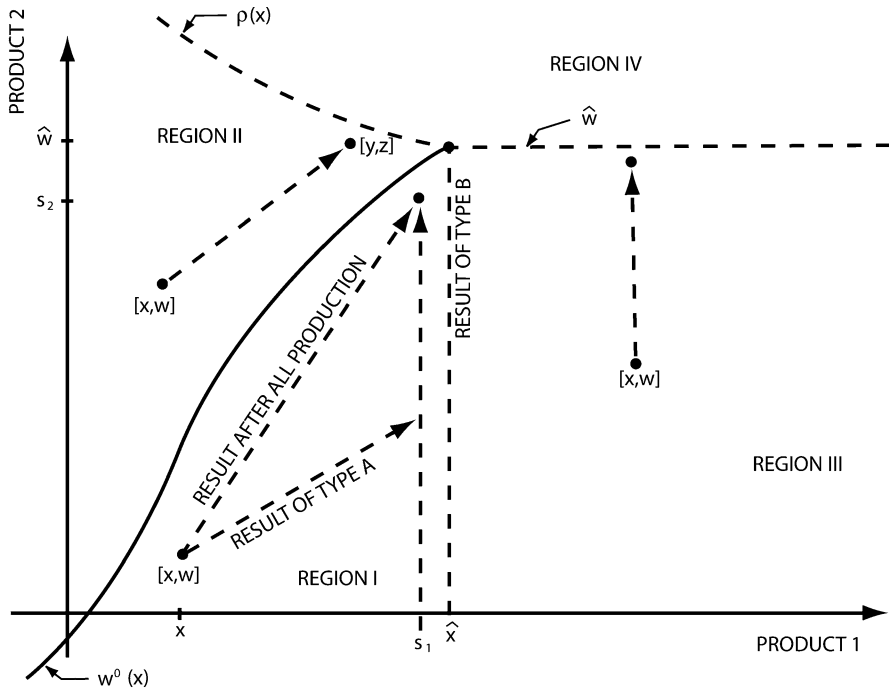
Fig. 3.2 Characterization of the optimal policy for the Devermeger two-product model

Deuermeyer (1980) considers a system consisting of n products all of which are perishable with respective lifetimes of l_i for $1 \le i \le n$, and assumes that the system has the substitution property. That is, the order quantity of each product is a nonincreasing function of the on hand inventories of the other products. Again, he is able to characterize the optimal policy for a single-period problem only.