Chapter 4
Continuous Review Perishable Inventory Models

Contents

Continuous review means that the state of the inventory system is known at all times, as opposed to periodic review where the state of the system is known only at discrete points in time labeled periods. Surprisingly, very little is known about optimal ordering policies for fixed life perishables under continuous review. The difficulty is that when an order lead time is present, aging is applied only to the units on hand, and not to the units on order. Furthermore, since units may be ordered at any point in time, there is no limit to the number of different orders that comprise the on-hand inventory at any point in time. Hence, the vector describing the on-hand inventory of each age level has an unlimited number of dimensions.

One way to circumvent this problem is to assume zero order lead time. We briefly review the two papers that take this approach. Before doing so, note that we believe that the simultaneous assumptions of continuous review and zero order lead time are unrealistic, and it is likely that these models have little practical use. There is a big difference between assuming zero order lead time in periodic review systems and in continuous review systems. In the former case, it only means that the lead time is less than a review period, as orders are assumed to be placed at the beginning of the planning period, and are assumed to arrive at the end of the planning period. In the case of continuous review, it means that orders arrive instantaneously. For that reason, the policies obtained are not likely to be useful in a practical setting.

It is worth noting that the primary driver of the continuous review heuristic (Q, R) inventory models that form the basis for most commercial inventory control systems is uncertainty of demand over the replenishment lead time. Where lead

S. Nahmias, *Perishable Inventory Systems*, International Series in Operations Research & Management Science 160, DOI 10.1007/978-1-4419-7999-5_4,

times reduced to zero, optimal order quantities would collapse to the simple EOQ formula and reorder levels would be zero. (See Nahmias (2009), Chap. 5, for example).

When there is a positive lead time for ordering, the form of an optimal order policy for perishables under continuous review is not known. Part of the problem is defining the state of the system. As noted above, the state could be an infinite dimensional vector of all previous orders placed and their ages. Given the complexity of the continuous review problem, a reasonable starting point is to find the best policy from a prespecified class of policies.

4.1 One for One (S–1, S) Policies

The only rigorous analysis of a continuous review perishable inventory system with positive order lead time of which we are aware is Schmidt and Nahmias (1985) (and an extension by Perry and Posner (1998)). They assumed the so-called ($S-1$, S) policy in which the inventory position (stock on hand plus stock on order) is maintained at a fixed level S. Assume that demands occur one-at-a-time, which would be the case, for example, if demands were generated by a stationary Poisson process. For nonperishables, a ($S-1$, S) policy places an order for one unit at each occurrence of a demand. With perishability, orders are placed at both the occurrence of demands and outdating. ($S-1$, S) policies are optimal for very high value items, and are common in military resupply and repair systems. See Nahmias (1981) for a review of models for repairable item systems. A more comprehensive and up-to-date discussion of these systems can be found in Muckstadt (2005)). It is common in both military and civilian applications that equipment goes in for maintenance either when the equipment fails, or at fixed intervals, whichever comes first. Since unplanned failures occur at random, these can be modeled as a stochastic process, and may be labeled the demand process. If an item reaches age m and has not failed (i.e., been demanded), the system for routine maintenance is taken out, and thus "outdates." Hence, the extension of traditional ($S-1$, S) policies to the case of fixed life perishable inventories, is potentially a very useful extension for modeling maintenance systems.

As noted above, what makes the problem difficult is the interaction of perishability and the order lead time. Assume that demands, (i.e., failures), occur one-at-a-time completely at random according to a stationary Poisson process with rate λ. Costs are charged in the usual way against proportional ordering at c per unit, holding at h per unit held per unit time, outdating at θ per unit, and lost sales at p per unit of unsatisfied demand. Assume a positive order lead time of τ.

Define the multidimensional stochastic process $\xi(t) = (\xi_1(t), \xi_2(t), ..., \xi_S(t))$ as the amount of time elapsed since the last S orders were placed. That is, $\xi_S(t)$ is the elapsed time since the last order was placed, $\xi_{S-1}(t)$ is the elapsed time since the second to last order was placed, etc. This process is the superposition of the demand and outdating processes. The approach is to derive the stationary distribution of $\xi(t)$

and use that to develop an explicit expression for the expected system cost rate as a function of S.

Let $p(t, x_1, x_2, ..., x_S)$ be the probability density of $\xi(t)$. Utilizing the partial differential equations of the process, the authors show that

$$\lim p(t, x_1, x_2, ..., x_S) = \begin{cases} Ke^{-\lambda\tau} & \text{for } x_1 < \tau \\ Ke^{-\lambda x_1} & \text{for } x_1 \geq \tau \end{cases}$$

Where

$$K = \left(e^{-\lambda\tau}\tau^S/S! + \int_{\tau}^{\tau+m} \left[{x_1}^{S-1} e^{-\lambda x_1}/(S-1)! \right] \, \mathrm{d}x_1 \right)^{-1}.$$

From these results one can derive the steady state distribution of the on-hand inventory at a random point in time, say $(P_0, P_1, ..., P_S)$.. The expected cost rate as a function of S is:

$$C(S) = c\lambda + (p-c)\lambda P_0 + h\sum_{j=1}^{S} jP_j + (c+\theta)\pi,$$

Where π is the expected outdating rate given by:

$$\pi = Ke^{-\lambda(\tau+m)}(\tau+m)^{S-1}/(S-1)!.$$

Although these equations appear straightforward, computation of the stationary state probabilities, $(P_0, P_1, ..., P_S)$ is complex, requiring numerical integration. Although the authors were unable to prove convexity of $C(S)$, numerical tests suggested that it is at least quasi-convex.

One of the interesting aspects of the problem revealed by numerical tests is the relationship between S and m. One would expect that S would be an increasing function of m. That is, as the lifetime of the product increases, the optimal stocking level would also increase, since the penalty for outdating decreases as m increases. Consider the following case: Fix $\tau = 1, \theta = 100, p = 600$ and $\lambda = 50$. The optimal values of S and the optimal expected cost, $C(S)$ for various values of m observed in this case are:

m	S^*	$C(S^*)$
0.005	0	30.00
0.007	7	29.98
0.008	25	29.62
0.009	38	29.01
0.01	47	28.30
0.02	72	21.79

(continued)

(continued)

m	S^*	$C(S^*)$
0.03	73	17.92
0.04	71	15.49
0.05	70	13.82
0.08	66	10.95
0.10	64	9.86
0.30	62	6.45
0.50	64	5.71
0.80	68	5.45
1.00	68	5.43

As one would expect, the expected cost rate decreases as the product lifetime increases. However, note the unusual behavior of S^* as m increases. The reason that S^* is 0 when the product lifetime is very small is that at this value of m, it is more economical not to run the system at all, since virtually all new items outdate before being able to satisfy demand. As m increases, the likelihood that a unit is able to satisfy demand before outdating increases, and so it becomes economical to hold positive stock. Because we chose a very high stockout penalty cost, it is desirable to maintain a larger inventory to decrease the likelihood of stockouts when m has a moderately low value. The value of S^* reaches a maximum in this case when $m = 0.03$. As m continues to increase, the likelihood of outdating decreases, so less stock is needed to maintain the same risk of stocking out. Note that for $m = 1$, the lifetime is essentially infinite, and the solution coincides with the optimal value of S for the nonperishable problem.

Perry and Posner (1998) developed the following extension of this model. Suppose that customers arriving to the system when it is out of stock are willing to wait a random amount of time Y for the next arrival of a unit. The rationale for this extension is that one can keep track of the times of orders, so one can determine the instance of the next arrival. If this is imminent, it makes sense that a customer would wait rather than leave the system unsatisfied.

As an example, suppose that H is the CDF of Y. They suggest the following form for H:

$$H(y) = q + p1_{\{y \geq \tau\}}$$

Where $1_{\{y \geq \tau\}}$ is the indicator function of the event $\{y \geq \tau\}$. This means that a customer is willing to wait for τ units of time for the next arrival (which must come within that time) with probability p. Since $H(y) = q$ for $y \leq \tau$, Y has mass q at 0 so that $P\{0 < Y \leq \tau\} = p$). The authors also consider several other forms of H.

Their analysis is based on computing the joint distribution of the process W given by

$$W = \{W_1(t), W_2(t), ..., W_S(t) : t \geq 0\}$$

where $W_i(t)$ is time to outdate the ith youngest item if the demand process were stopped at time t. While this extension is interesting, and could be potentially useful in some circumstances, the authors do not provide any calculations to compare the results of their model to those of Schmidt and Nahmias (1985). Until that is done, there is no way to tell if their extension leads to policies that differ significantly from those computed assuming customers do not wait when the system is out of stock.

4.2 Continuous Review Models with Zero Lead Time

As noted above, we feel the simultaneous assumptions of continuous review and zero lead time are unrealistic. Nevertheless, we briefly review the major results in this area. The first to look at this problem was Weiss (1980). Weiss assumed that demands are generated by a stationary Poisson process with fixed rate λ. His main result is that, in the case of lost sales, the form of the optimal policy is to either never order, or to order to a fixed level S when the inventory level drops to zero. Notice that the issue of perishability never really comes up, as order cycles are completely independent. One simply waits until the inventory level drops to zero (whether through demand or outdating) and replenishes to a fixed level (much as one does in the simple EOQ model).

Weiss's results were generalized by Liu and Lian (1999). They assumed full backordering of demand and generalized from a Poisson demand process to a stationary renewal process. They showed that the form of the optimal policy is (s, S) with $s \leq -1$, and provide explicit formulas for computing the optimal policy parameters. (A correction of their results appears in Gurler and Ozkaya (2003).) Gurler and Ozkaya (2008) extended their model to the case where the lifetime of a batch is a random variable. Note that for all of these models, the form of the optimal policy (namely that one only orders when the system is backordered), depends heavily on the assumption that order lead times are zero. It is unlikely that one would use such a policy in the real world, where order lead times are always positive. Hence, it would appear that these results do not provide much insight into the general continuous review perishable inventory problem.

4.3 Optimal (Q, r) Policies with Positive Lead Time

Short of finding optimal policies, a common method in inventory control is to determine the best policy from a class of policies. This is the approach taken by Berk and Gurler (2008). It is well known that for many nonperishable continuous review systems, the optimal policy is a (Q, r) policy. That is, when the inventory position hits r, an order for size Q is placed. As the form of the optimal policy is not known, it would seem natural to consider the best (Q, r) policy in the perishable inventory setting. In this version, the authors assume that $r<Q$, which means that there is at most a single order outstanding. (Future work considers the generalization to the case where multiple orders are outstanding.) They also assume throughout that

demands are generated from a stationary Poisson process, thus allowing them to make use of the memoryless property of the exponential interarrival time distribution.

Because items can perish, a slight modification of the standard policy is required. In nonperishable continuous review systems, one is guaranteed that an order can be placed at the instant the inventory level (or inventory position when more than one outstanding order is allowed) hits the reorder point, r. However, if only a single order is in inventory, the entire order will perish at the same instant, thus dropping the inventory level immediately to zero. Hence, the policy must be modified to: One orders Q units whenever the inventory level hits r or zero, whichever comes first.

This modification points out the inadequacy of this policy, however. Consider a blood bank storing a rare form of blood (AB negative, for example). The lead time for replenishment is three weeks owing to the difficulty of finding donors with this blood type. Suppose that there is a substantial amount of AB negative blood on hand due to outdate in one day. According to this policy, the blood bank would sit on their current supply until it outdates the next day, dropping their inventory to zero. During the three week lead time, there would be no AB negative blood available, precipitating a crisis situation. (Note that this problem does not arise in the model considered by Schmidt and Nahmias (1985), since all orders are for size one only, so the on-hand inventory level can only drop by one unit at a time independent of whether the drop is due to outdating or filling demand.)

Clearly, perishability fundamentally changes the form of the optimal policy. The blood bank seeing the supply of this rare blood type would soon expire and would take steps to replenish the inventory far in advance of the outdating. Hence, this type of (Q, r) policy does not make sense in this context. Before one can consider effective heuristics for this problem, a better understanding of the nature of the optimal policy is required.

A modification of the (Q, R) policy which ameliorates this problem to some extent is considered by Tekin et al. (2001). They suggest a (Q, R, T) policy. The policy is implemented in the following way: A replenishment order of size Q is placed whenever the on-hand inventory level drops to r, or when T units of time have elapsed since the last instance at which the inventory level hit Q, whichever occurs first. For this policy to make sense, they require the unusual assumption that the aging of items in a batch begins only after all of the units of the previous batch are exhausted either by demand or outdating. It seems likely that this would be true only in rare circumstances. With this assumption, epochs at which the inventory hits Q are regeneration points of the system. The analysis can then proceed using renewal reward processes. Their main finding is that when service levels are high, the value of r is not increased as much as it would be in an ordinary (Q, r) system. This policy only makes sense when their assumption about aging is accurate, however.

4.4 An Alternative Approach

As noted above, the difficulty with trying to characterize the optimal policy for a continuous review perishable inventory system is properly defining the state variable. If the state is defined in terms of on-hand inventory (as is almost always done in inventory modeling), then the optimal solution would be a function of an infinite dimensional state variable. This would correspond to all of the on-hand orders and their ages. An interesting question is whether or not there is a single dimensional surrogate variable that might provide reasonable operating policies. We speculate that there is such a surrogate variable. It is the expected remaining lifetime or virtual lifetime of the on-hand inventory. That is, given all of the on-hand orders and their ages, and knowledge of the demand process, what is the expected time that this inventory will be depleted either by demand or by outdating?

Consider the following two scenarios. Product lifetime is 10 days, and expected inter-demand time is 1 day. In case one, assume that there are 100 units of 9-day-old inventory on hand, and in case two, there are 100 units of 1-day-old inventory on hand. The expected remaining lifetime of the on-hand stock in the first case is essentially one day, and in the second case it is slightly less than 9 days (since it is possible for 100 demands to occur before the nine days have elapsed). The expected remaining lifetime measure is far more useful than just knowing the number of units on hand, and a policy based on its value should perform reasonably well. The policy we suggest is the following: Let $E(L)$ be the expected remaining lifetime of the current on-hand inventory. Place an order for Q units at the first instance that $E(L)$ drops below a trigger level, r. (This policy assumes only one outstanding order. If multiple orders are outstanding, the form of the policy would have to be modified to take into account unit on order as well as on hand.)

A problem with this idea is that computation of the expected remaining lifetime for an arbitrary stockpile of items appears to be very difficult. An alternative approach would be to use some type of virtual lifetime measure. This would operate much like the virtual waiting time in queueing theory. At each arrival of new stock, the virtual lifetime process would jump upward by an amount related to the size of the order, the lifetime of the product, and the demand process. One would reorder when the virtual lifetime crossed a critical threshold. Conceptually, this is similar to the suggestion above to base the reorder decision on the expected remaining lifetime in the stockpile, but could be easier to compute and implement.

We are not aware of any work that has explored these ideas, but they could provide a way of obtaining reasonable control policies for perishables in a continuous review environment.