DYNAMIC MODELING OF INVENTORIES SUBJECT TO OBSOLESCENCE

George W. Brown
John Y. Lu
Robert J. Wolfson

November 1963

P-2825

DYNAMIC MODELING OF INVENTORIES SUBJECT TO OBSOLESCENCE*

George W. Brown

University of California, Los Angeles, California

John Y. Lu

IBM Federal Systems Division, Bethesda, Maryland

Robert J. Wolfson

The RAND Corporation, Santa Monica, California

INTRODUCTION

The models presented here arose in connection with studies on the problem of obsolescence as met in the supply system, with particular reference to spare parts for Naval aircraft. Inventory control procedures in current usage ordinarily allow for obsolescence by making an arbitrary blanket charge, to be distributed over all supply items. The models to be discussed are concerned with optimization of ordering and disposal decisions, leading to policies of the (s,S) type, incorporating specific elements descriptive of the uncertainty of future demand for the particular item.

---

*Any views expressed in this paper are those of the authors. They should not be interpreted as reflecting the views of The RAND Corporation, or the official opinion or policy of any of its governmental or private research sponsors. Papers are reproduced by The RAND Corporation as a courtesy to members of its staff.

The research reported herein was performed while the first author was consultant to, and the second and third authors were staff members of C-E-I-R Inc., Los Angeles, California. Work was performed with the support of United States Navy Bureau of Supplies and Accounts, under Contract Nonr 3333(00) with C-E-I-R, Inc. This work is a joint effort of the three authors, with the mathematical development primarily by G. W. Brown.

Obsolescence may occur, for a particular item, because the function served by that item is no longer required, because units are replaced as they are consumed by a substitute item which performs similar or identical functions, or because of a program of systematic replacement by a substitute item. While the processes which bring about obsolescence are themselves of considerable interest, the emphasis here is on evaluation of the threat of obsolescence, present and future, as it pertains to optimum ordering and disposal policies.

The simplest treatment of problems of this type corresponds to the situation in which demand parameters are assumed to vary slowly as compared with the ordering cycle, so that a steady state inventory model serves as the basis of the optimization, while demand forecasting takes place outside the model, for the purpose of estimating the parameters of the model.

Parameter estimation in the steady state case may be handled by assuming a form for the probability distribution of demand together with an a priori distribution for the parameters of the demand distribution, with the result that conditional demand distributions may be derived on the basis of experience, using Bayes' formula as discussed by Raiffa and Schlaifer.[5][6] To accommodate obsolescence, or other types of varying demand, one possible generalization of the steady state model assumes demand distributions which vary over time, but with parameters specified in advance as functions of time, in which case the corresponding dynamic programming problem can be solved as in References 1, 2 and 3. In this case the assumption is made implicitly that the probability distribution of demand at time t is always the same, as seen at all times $t' < t$, or, in other words, that all the knowledge that will ever exist about the demand at time t, prior to observations of the actual demand at t, already exists at the outset.

The models discussed below were devised primarily to incoroprate a Bayesian process of dynamic adjustment of future demand distributions as one advances in time toward the future, so that the latest knowledge of actual demands, and perhaps other ancillary information which may have become available, will be reflected at all times in the present description of future demand, rather than simply at the outset. This approach requires a description of the statistical nature of the variation in demand parameters which might be encountered. The next section is devoted to a special class of problems, in which demand parameters are assumed constant until obsolescence, with time of obsolescence following a given distribution (i.e. the useful period of the item follows a mortality distribution). In these cases it turns out that it is useful to consider mortality distributions with a particular kind of reproducibility. The remainder of the paper is concerned with a more general Markov model of underlying demand-generating states, permitting consideration of more complicated patterns of demand variation other than simple obsolescence, admitting the possibility of simultaneous consideration of items with linked demands, and allowing incorporation of ancillary information other than past observed demand.

Some policies have been computed under this type of model. They are of the (s,S) type and are, in general suprisingly insensitive to small changes in a priori state probability. The results of these computations are presented in this paper.

"SUDDEN DEATH" OBSOLESCENCE

Although the main body of the paper does not depend on this section, this material has been included for its own interest, in view of the fact that the reproducibility property to be introduced may have other applications. This section also provides a simple introduction to the philosophical orientation of the material presented later.

This section is concerned with the situation in which a particular supply item is supposed to follow a given probability distribution of demand until some unknown future date, at which time all demand will cease thenceforth. Before that date the item is not obsolescent, after that date it is obsolescent, that is, all demand permanently ceases. Hence the term "sudden death". At the outset there is given a quantity P_o, the probability that the item is not obsolete, and a cumulative distribution function $F_o(\tau)$, the conditional probability that obsolescence takes place by time τ assuming the item not obsolescent at t=0, with $F_o(0) = 0$. Let Y_t be a vector which describes the inventory status at time t, let x_t be the observed demand in the inventory period beginning at time t, and let A_t be the set of available actions at time t, so that α is a possible action if $\alpha \in A_t$.

For the purposes of this section it is unnecessary to formalize the constraints which relate the sequences $\{Y_t\}$, $\{x_t\}$, $\{\alpha_t\}$, or to complete the model by relating the costs and rewards to the sequences. Given all these constraints the optimum action α at $t = 0$ will be a function of P_o, Y_o, and the function F_o. If $\hat{\alpha}_o$ denotes the optimum α at t = 0 then, symbolically, $\hat{\alpha}_o = \Psi\left[P_o , Y_o , \left\{F_o(\tau)\right\}\right]$, given A_o and all other constraints. In problems of this type the determination of $\hat{\alpha}_o$ involves implicitly the optimum α_t for later t's. It becomes of interest, therefore, to examine the recursive specification of the quantities P_t and the functions $F_t(\tau)$ which, together with Y_t will characterize the situation at later time t, corresponding to P_o , $F_o(\tau)$ and Y_o which characterize the initial situation.

The recursive specification of P_{t+1} and $F_{t+1}(\tau)$ (taking the inventory period as the unit) in terms of P_t , $F_t(\tau)$, and x_t, takes into account the information added by x_t and the passage of time from t to t+1. In this simple model an observed $x_t > 0$ implies the item was not obsolete at time t, but, of course, the item may still become obsolete in the interval (t, t+1), with probability $F_t(1)$. On the other hand $x_t = 0$ will modify the probability that the item was already obsolescent at time t, by Bayes formula, and the passage of time from t to t+1 modifies further the probability of obsolescence. The detailed calculation should make the situation clear.

Let p be the probability that $x_t = 0$, if the item is not obsolescent, whence 1 - p is the probability that $x_t > 0$, conditioned similarly. If the item is obsolescent then $x_t = 0$ necessarily. Denoting the events of obsolescence - non-obsolescence as of time t by O_t and O_t' respectively, then

$$\Pr\left\{x_t=0 \mid O_t\right\} = 1, \ \Pr\left\{x_t > 0 \mid O_t\right\} = 0, \ \Pr\left\{x_t=0 \mid O_t'\right\} = p,$$

$$\text{and } \Pr\left\{x_t > 0 \mid O_t'\right\} = 1 - p.$$

Also $\Pr\{O_t\} = P_t$, from which

$$\Pr\{O_t \mid x_t = 0\} = \frac{\Pr\{O_t , x_t = 0\}}{\Pr\{x_t = 0\}}$$

$$= \frac{\Pr\{O_t\} \Pr\{x_t=0 \mid O_t\}}{\Pr\{O_t\} \Pr\{x_t=0 \mid O_t\} + \Pr\{O_t'\} \Pr\{x_t=0 \mid O_t'\}}$$

$$= \frac{P_t}{P_t + (1 - P_t)p}$$

and $$\Pr\{O_t \mid x_t > 0\} = \frac{\Pr\{O_t , x_t > 0\}}{\Pr\{x_t > 0\}} = 0.$$

In other words, the a posteriori probability of obsolescence at t is either $\dfrac{P_t}{P_t + (1 - P_t)\, p}$

or 0, according to whether $x_t = 0$ or $x_t > 0$.

Allowing now for the passage of time from t to t+1, we have

$$P_{t+1} = \Pr\{O_{t+1} \mid x_t\} = \Pr\{O_t \mid x_t\} + \Pr\{O_t' \mid x_t\}\, F_t\,(1).$$

Substitution of the value derived above leads to the result

$$x_t = 0 \Rightarrow P_{t+1} = \frac{P_t}{P_t + (1-P_t)p} + \frac{(1 - P_t)p}{P_t + (1-P_t)\, p}\, F_t\,(1)$$

$$= \frac{P_t + (1-P_t)p\, F_t(1)}{P_t + (1-P_t)p}$$

and

$$x_t > 0 \Rightarrow P_{t+1} = F_t\,(1).$$

Turning now to the function $F_{t+1}(\tau)$, which represents the conditional probability of $O_{t+1+\tau}$ given O'_{t+1} , it is seen that

$$F_{t+1}(\tau) = \Pr\left\{O_{t+1+\tau} \mid O'_{t+1}\right\} = \frac{F_t(1+\tau)-F_t(1)}{1-F_t(1)} ,$$

independently of x_t. The recursion has now been established, for the determination of P_{t+1} and $F_{t+1}(\tau)$.

As was noted earlier, the optimum action at t=0 is $\alpha_0 = \Psi\left[P_0 , Y_0 , \left\{F_0(\tau)\right\}\right]$. Subject to the obvious stationarity assumptions, such as that A_t should depend on t only through its dependence on Y_t, and corresponding assumptions about the cost and reward functions over time, it can then be concluded that the optimum action at time t is $\alpha_t = \Psi\left[P_t , Y_t , \left\{F_t(\tau)\right\}\right]$, where Ψ is the same function which applied to α_0. Possible identification of the same policy Ψ as an optimum policy for all t has important implications for the computational problem of approximating to optimal policies.

Examining further the computational problem, it is clear that in certain approaches it will be necessary to adopt tabular representations of the arguments of Ψ. P_t is simply a number between 0 and 1, Y_t is a vector, hopefully with not too high dimension, but if $F_t(\tau)$ can be an arbitrary mortality distribution its representation may be extraordinarily expensive. Exploration of this question led to consideration of special families of mortality distributions which would be closed under the transformation

$$G(\tau) = \frac{F(1+\tau) - F(1)}{1 - F(1)} .$$

If a suitable family of distribution functions can be found, represented by a small parameter set, and closed under the above transformation, then for all t $F_t(\tau)$ will be a member of the family as long as $F_o(\tau)$ is so chosen, with the result that the policy ψ may be represented as $\psi\left[P_t, Y_t, \theta_t\right]$, where θ_t is the current parameter set of the distribution $F_t(\tau)$. Such a family is said to be conditionally reproducible.

The simplest example of such a family is given by a distribution of the form $F(\tau) = 1 - e^{-a\tau}$, $a > 0$, $\tau \geq 0$, the simple exponential distribution. In this case

$$G(\tau) = \frac{F(1+\tau) - F(1)}{1-F(1)} = \frac{1 - e^{-a(1+\tau)} - 1 + e^{-a}}{e^{-a}}$$

$$= 1 - e^{-a\tau} = F(\tau) ,$$

and the family is zero-dimensional. This is, of course, a trivial example. Non-trivial examples are furnished by the following families, in which the parameters are properly restricted to yield proper mortality distributions.

family 1: $F(\tau) = 1 - e^{-a(e^{b\tau} - 1)}$

family 2: $F(\tau) = 1 - (1 + b\tau)^{-c}$

family 3: $F(\tau) = 1 - e^{-a\left[(1+b\tau)^{c}-1\right]}$

The first two are both limiting cases of family 3; the exponential distribution may also be interpreted as a limiting case of family 3. Verification of the closure property also yields the corresponding transformations on the parameters.

For family 1,

$$G(\tau) = \frac{F(1+\tau)-F(1)}{1-F(1)} = \frac{1-e^{-a(e^{b+b\tau}-1)} - 1 + e^{-a(e^{b}-1)}}{e^{-a(e^{b}-1)}}$$

$$= 1 - e^{-a(e^{b+b\tau}-e^{b})}$$

$$= 1 - e^{-ae^{b}(e^{b\tau}-1)}$$

$$= 1 - e^{-a'(e^{b\tau}-1)}, \quad \text{with } a' = a\, e^{b} .$$

Thus, for fixed b, the family is a one-parameter family. For family 2,

$$G(\tau) = \frac{F(1+\tau)-F(1)}{1-F(1)} = \frac{1-(1+b+b\tau)^{-c}-1+(1+b)^{-c}}{(1+b)-c}$$

$$= 1 - (1 + \frac{b}{1+b}\tau)^{-c}$$

$$= 1 - (1 + b'\tau)^{-c}, \quad \text{with } b' = \frac{b}{1+b}$$

In this case, for fixed c the family is a one-parameter family.

Finally, for family 3,

$$G(\tau) = \frac{F(1+\tau)-F(1)}{1-F(1)} = \frac{1-e^{-a[(1+b+b\tau)^{c}-1]}-1+e^{-a[(1+b)^{c}-1]}}{1-e^{-a[(Hb)^{c}-1]}}$$

$$= 1 - e^{-a[(1+b+b\tau)^{c}-(Hb)^{c}]}$$

$$= 1 - e^{-a(1+b)^{c}[(1+\frac{b}{1+b}\tau)^{c}-1]}$$

$$= 1 - e^{-a'[(1+b'\tau)^{c}-1]}$$

with $a' = a(1+b)^{c}$ and $b' = \frac{b}{1+b}$.

Family 3, for fixed c, is a two-parameter family.

In each case, therefore, if $F_o(\tau)$ is chosen as a member of one of these families, the step from $F_t(\tau)$ to $F_{t+1}(\tau)$ will always yield a member of the same family. As a consequence the problem of finding optimum policies Ψ is simplified by permitting consideration of the simultaneous solution of all problems derivable from the initial $F_o(\tau)$ by iteration of the transformation.

A MARKOV MODEL FOR VARIATION OF DEMAND PARAMETERS

The remainder of the paper is devoted to discussion of an alternative model, chosen in preference to the model of Section II, for further exploration, in view of its greater inherent richness for representing demand variation, for considering items with linked demands, and for accommodation of ancillary information. The key to this approach lies in the introduction of a set of system states, each of which generates its own characteristic set of distributions of demand and other observables. Assumption of initial probabilities for the several states, together with associated transition probabilities from state to state permit recursive evaluation of a posteriori state probabilities, based on all observables available, as time proceeds from one inventory period to the next. Items with linked demands are represented by joint distributions of demand, conditioned on each system state, and optimum policies may be found separately, item by item, whenever the available actions and the associated costs and rewards are independent of one another. For computational purposes policies may be represented simply as functions of the present probabilities of the various states and of the relevant inventory status Y_t , with the result that optimization may be performed simultaneously for all sets of state probabilities, provided that reasonable stationarity assumptions are met by the set of available actions and the cost and reward functions, as well as the conditional demand distributions. It is not necessary that the transition probabilities permit reversal of states; indeed the stochastic process of induced demand may be very far from stationary, even though the particular construction permits a stationary type of analysis.

To introduce specific notation, suppose that n different demand-generating states S_r , $r=1,2, \cdots ,n$ are distinguished. Let P_{rs} represent the customary Markov transition probability, the probability of a transition to S_s in one inventory period, conditioned on the system being in S_r , so that $\sum_s P_{rs} = 1$, $r=1,2, \cdots ,n$. Associated with each state at time t is the corresponding state probability $\pi_r(t)$, as of the beginning of the inventory period $(t,t+1)$, with $\sum_r \pi_r(t)=1$.

Represent by X_t the vector of all observables, demand or otherwise, to become available between t and t+1, and associate with each state S_r a joint probability $g_r(X)$, representing the conditional probability $\Pr\left\{X_t=X \mid S_r\right\}$. It is assumed that $\left\{g_r(X)\right\}$ are not time-dependent; the assumption that X_t has discrete distribution will not be particularly restrictive. As in the preceding section elementary calculations will yield a recursive evaluation of $\left\{\pi_r(t+1)\right\}$, incorporating the vector of observables X_t and allowing for the passage of time from t to t+1.

To illustrate the applicability of the model to variable demand parameters, consider the following simple example of a possible transition matrix.

$r \backslash s$	1	2	3
1	.9	.1	0
2	.1	.8	.1
3	0	0	1

P_{rs}

State S_1 may be a normal beginning state, with substantial demand for a set of items, S_2 may correspond to an intermediate period with somewhat reduced demand for some or all of the items, and S_3 may correspond to a period in which some or all of the items in question are not demanded at all. With the particular matrix chosen S_1 may give way only to S_2 , S_2 may be succeeded by either S_1 or S_3 , and S_3 is necessarily a terminal state, corresponding perhaps to obsolescence. Accompanying the different demand distributions associated with the separate states there may also be different distributions of additional observables which yield ancillary statistical information about the state of the system. The introduction of additional states, some of which may differ only with respect to the distributions of ancillary information, allows almost arbitrary complication, in principle.

Proceeding now to the recursive process, the a posteriori probabilities corresponding to time t, conditional on X_t , must be obtained first. Bayes' formula yields

$$\Pr\left\{S_r \mid X_t\right\} = \frac{\Pr\left\{S_r\ ,\ X_t\right\}}{\Pr\left\{X_t\right\}}$$

$$= \frac{\Pr\left\{S_r\right\} \Pr\left\{X_t | S_r\right\}}{\sum_r \Pr\left\{S_r\right\} \Pr\left\{X_t | S_r\right\}}$$

$$= \frac{\pi_r\ (t)\ g_r\ (X_t)}{\sum_{r=1}^{n} \pi_r\ (t)\ g_r\ (X_t)}$$

Introducing now the transition probabilities for the change of state which may occur in the period $(t, t+1)$ leads to the result

$$\pi_s\ (t+1) = \sum_r P_{rs} \Pr\left\{S_r | X_t\right\}$$

$$= \frac{\sum_r P_{rs}\ \pi_r\ (t)\ g_r\ (X_t)}{\sum_r \pi_r\ (t)\ g_r\ (X_t)} \quad , \ s=1,2,\cdots,n.$$

This result provides the basis for representing policies as functions $\psi\left[Y\ ,\ \left\{\pi_r\right\}\right]$; with ψ independent of t, except through the identification of $\left\{\pi_r\right\}$ with $\left\{\pi_r\ (t)\right\}$ and Y with Y_t .

OPTIMIZATION IN THE MARKOV MODEL

Determination of optimum inventory policies may be achieved by constructing the functional equation corresponding to the optimality principle of dynamic programming as discussed in References 1 and 2, and approximating the solution by suitable iterative methods. To present in one place the required notation, denote by

Y_t the inventory status vector at time t, the beginning of an inventory period;

X_t the vector of observables during the period which begins at t;

$\left\{g_r(X)\right\}$ the conditional distributions of observables, one for each state of the demand-generating system;

$\left\{P_{rs}\right\}$ the matrix of transition probabilities from demand-generating state r to demand-generating state s;

$\pi(t)$ the vector of a priori state probabilities $\left\{\pi_r(t)\right\}$ at time t;

α an available action;

A the set of available actions;

Ψ a policy, $\Psi\ (Y,\ \pi)\ \epsilon\ A$;

$L(Y,X,\alpha)$ current costs immediately associated with action α when $Y_t=Y$ and $X_t=X$;

$C_\Psi(Y,\pi)$ expected total present and future cost, referred to time t, associated with policy Ψ, when $Y_t=Y$ and $X_t=X$;

ρ a discount factor, $0 < \rho < 1$, for relating a cost at time t+1 to an equivalent cost at time t.

In this formulation all time-dependence is referred to the Markov process of demand-generating states; through the transition probabilities and the initial a priori state probabilities the observables X_t cause $\pi(t)$ to be updated in each time period, as shown in the preceding section. The cost function L depends on time only through Y_t and X_t, and is intended to reflect benefits as well as costs, taking costs to be positive. L is intended to reflect only the immediate consequences of the action α, given the circumstances described by Y, and reflecting the actual demands implicit in X. The function $C_\Psi(Y,\pi)$, used to evaluate the policy Ψ, is the expectation of the present value of the total of all L costs from time t onward, each future value of L being discounted to the time t, viewed as the present time.

The point of view is that the process is to continue, without finite termination. If the evaluation criterion is to be the minimum average L, without discounting, the discount factor may be permitted to approach $\rho = 1$. Optimization calls for finding Ψ to minimize $C_\Psi(Y, \pi)$, for all combinations (Y, π).

For any policy Ψ the evaluation function $C_\Psi(Y, \pi)$ satisfies the following recursion relation:

$$C_\Psi(Y,\pi) = E\left\{L(Y,X,\Psi(Y,\pi)) + \rho C_\Psi(Y'(Y,X,\Psi(Y,\pi)),\ \pi'(\pi,X))\right\},$$

where $Y'(Y,X,\Psi(Y,\pi))$ represents the terminal inventory status resulting from the beginning inventory status Y, acted upon by both the inventory action $\Psi(Y,\pi)$ and the actual demand implicit in X, and $\pi'(\pi,X)$ represents $\pi(t+1)$ computed from $\pi(t)$ and X_t, as in the previous section. The expectation operation is with respect to the marginal distribution of X, hence the recursion relation may be written

$$C_\Psi(Y,\pi) = \sum_{r,X} \pi_r\, g_r(X) \left\{L(Y,X,\Psi(Y,\pi)) + \rho C_\Psi(Y'(Y,X,\Psi(Y,\pi)),\pi'(\pi,X))\right\}$$

The above relation simply reflects the fact that the total expected cost is the sum of the expected current cost and the discounted expected cost associated with the new position to be encountered next period. Note that the policy Ψ affects the action taken now, which shows up in both the calculation of L and the calculation of the new beginning status Y'. If Ψ is to be an optimal policy $C_\Psi(Y,\pi)$ must also satisfy the functional equation derived from the familiar optimality principle.

$$C_\Psi(Y,\pi) = \min_{\alpha\in A} \sum_{r,X} \pi_r\, g_r(X) \left\{L(Y,X,\alpha) + \rho C_\Psi(Y'(Y,X,\alpha),\pi'(\pi,X))\right\},$$

corresponding to the statement that no α different from $\Psi(Y,\pi)$, if chosen now, can do better than $\Psi(Y,\pi)$, when followed by the optimal policy Ψ thereafter. Note that if C_Ψ were known, the optimality principle would yield $\Psi(Y,\pi)$ as the minimizing α in the functional equation.

The functional equation immediately above leads one to consider an iterative method based on an initial function $K_o(Y,\pi)$, more or less arbitrary, with

$$K_{n+1}(Y,\pi) = \operatorname*{Min}_{\alpha} \sum_{r,X} \pi_r \, g_r(X) \, \left\{L(Y,X,\alpha)+\rho K_n(Y'(Y,X,\alpha),\pi'(\pi,X))\right\} .$$

The minimizing α, as a function of (Y,π), yields a corresponding sequence of policies $\Psi_n(Y,\pi)$.* In numerical studies of an iterative solution along these lines it was found that convergence of K_n was often quite slow, particularly with ρ close to unity. Convergence of the policies Ψ_n was, however, quite rapid.

A slight modification of the iteration yielded rapid convergence, by fixing at 0 the value of $K_n(Y^o,\pi^o)$, for some arbitrarily chosen origin (Y^o,π^o), then subtracting after each iteration a constant from all values, chosen to keep $K_n(Y^o,\pi^o)$ fixed. More precisely, set

$$F_{n+1}(Y,\pi) = \operatorname*{Min}_{\alpha} \sum_{r,X} \pi_r \, g_r(X) \, \left\{L(Y,X,\alpha) + \rho K_n(Y'(Y,X,\alpha),\pi'(\pi,X))\right\}$$

and set $K_{n+1}(Y,\pi) = F_{n+1}(Y,\pi) - F_{n+1}(Y^o,\pi^o)$. When convergence is complete, $C_\Psi(Y^o,\pi^o)$ will be simply

$$\frac{\lim F_n(Y^o,\ \pi^o)}{1-\rho} \quad ,$$

*This approach is very similar to the "value iteration" approach of Howard [4].

and all values of $C_\psi(Y, \pi)$ may be obtained by adding the same constant to line $K_n(Y, \pi)$.

For illustrative purposes we present a few computed policies. In this numerical example, a demand-generating system consisting of two states, one having a high demand rate, and the other a low demand rate, is considered. Orders are made at the beginning of each regularly spaced period and a delivery lag of one period is considered. Any unfilled demand for current period becomes demand for next period. The conditional distribution of demand is asssued to be Poisson.

Definition of symbols used and parameter values assumed are given below:

y = initial inventory level. Both positive and negative values are allowed. If y assumes a negative value, it means that there are unfilled demands.

α = desired inventory level. The difference α-y represents the optimal order quantity, if it is positive. If the difference is negative, it represents the optimal disposal quantity. In the first example, we are considering $\alpha \geq \max(y, 0)$, i.e., no disposal activities are allowed. Of course, it presents no problem if one wishes to allow disposal activities as in the second example; all that needs to be done is to set the domain of α to be $\alpha \geq 0$. If we allow α to assume negative values it means that the stock can be returned to the suppliers and credit received for it.

x = demand per period

π = a priori probability that the system is in state 1. Then $(1-\pi)$ is a priori probability that the system is in state 2.

$g_1(x)$ = probability distribution of demand when the system is in state 1

$$e^{-2}\left(\frac{2^x}{x!}\right)$$

$g_2(x)$ = probability distribution of demand when the system is in state 2

$$e^{-0.4}\ \frac{(0.4)^x}{x!}$$

$h_\pi(x)$ = probability distribution of x weighted by a priori state probabilities

$$h_\pi(x) = \pi\ g_1(x) + (1-\pi)\ g_2(x)$$

d_1 = unit storage cost, 0.5

d_2 = unit out-of-stock cost, 5.0

d_3 = fixed order cost, 1.0

d_4 = unit order cost, 0.5

$||P_{rs}||$ = transition probability matrix, r,s, = 1,2

r \ s	1	2
1	0.7	0.3
2	0.1	0.9

π'' = a priori probability that the system will be in state 1 in the next period.

$$\pi'' = \frac{p_{11}\ \pi\ g_1(x) + p_{21}\ (1-\pi)\ g_2(x)}{\pi\ g_1(x) + (1-\pi)\ g_2(x)}$$

ρ = discount factor, 0.99

With the above notation, we now can proceed to specify the expected total cost function which is composed of the expected current costs and discounted future cost.

Expected holding and shortage costs, $L_1(y,\pi)$, to be charged during the current period, assuming that an order will not be delivered until next period is:

$$L_1(y,\pi) = \begin{cases} d_1 \sum_{x=0}^{y} (y-x)\ h_\pi(x) + d_2 \sum_{x=y+1}^{\infty} (x-y)\ h_\pi(x) & y \geq 0 \\ \\ d_2(-y) + d_2 \sum_{x=0}^{\infty} x\ h_\pi(x) & y < 0 \end{cases}$$

Suppose the ordering cost, $L_2(y,\alpha)$, is charged when orders are placed, it can be expressed as follows:

$$L_2(y,\alpha) = \begin{cases} d_3 + d_4\ (\alpha-y) & \alpha > y \\ \\ 0 & \alpha = y \end{cases}$$

Then the expected current cost function is:

$$L(y,\ \pi,\ \alpha) = L_1(y,\ \pi) + L_2(y,\alpha)$$

Let $C_\psi\ (y,\pi)$ represent the total expected cost if the provisioning is done optimally.

$$C_\psi(y,\pi) = \operatorname*{Min}_{\alpha > \max(y,0)} \left[L(y,\pi,\alpha) + \rho \sum_{x=0}^{\infty} h_\pi(x)\ C_\psi\ (\alpha-y,\pi'') \right]$$

The solution to this functional equation was obtained by means of the iterative procedure described above and is presented in Table 1. Note that the condition $\alpha > \max(y,0)$ restricts the policy from carrying out disposal activities.

Another set of calculations was performed by introducing disposal activities in the above problem and by reducing the unit stockout cost. (The remainder of the parameters were unchanged). More specifically,

TABLE 1: OPTIMUM POLICY TABLE WITHOUT DISPOSAL*

Initial Inventory Level y \ π (a priori state probability)	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
7	7	7	7	7	7	7	7	7	7
6	6	6	6	6	6	6	6	6	6
5	5	5	5	5	5	5	5	5	5
4	4	4	4	4	4	4	6	6	6
3	3	5	5	5	5	5	6	6	6
2	4	5	5	5	5	5	6	6	6
1	4	5	5	5	5	5	6	6	6
0	4	5	5	5	5	5	6	6	6
-1	4	5	5	5	5	5	6	6	6
-2	4	5	5	5	5	5	6	6	6

*For a given a priori state probability and initial inventory level, we read off an appropriate entry in the above table. The difference between this entry and the initial inventory level is the optimum order quantity. When the difference is zero, no action should be taken. For all $y < -2$, the optimum policy repeats itself.

we set the domain of minimization $\alpha > 0$ and set $\alpha_2 = 25$. Results are tabulated in Table 2.

TABLE 2: OPTIMUM POLICY TABLE WITH DISPOSAL

$y \setminus \pi$	*a priori* state probability								
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
7	3	7	7	7	7	7	7	7	7
6	3	6	6	6	6	6	6	6	6
5	5	5	5	5	5	5	5	5	5
4	4	4	4	4	4	4	4	4	4
3	3	3	3	3	3	3	5	5	6
2	2	4	4	4	4	4	5	5	6
1	3	4	4	4	4	4	5	5	6
0	3	4	4	4	4	4	5	5	6
-1	3	4	4	4	4	4	5	5	6
-2	3	4	4	4	4	4	5	5	6

It is of interest to note that by making a few special assumptions a model of the sort discussed in the section headed "sudden death" obsolescence falls out as a special case of the Markov process model:

One of the states is assumed to have zero demand. Once the system is in this state, it is not possible to make a transition to another state. Hence, we consider that obsolescence has occurred when the system has entered the zero demand state.

It illustrates how those failure or mortality distributions with the property of "conditional reproducibility" which have been discussed earlier enter into the general Markov inventory model. With the introduction of such a distribution into the model, it is appropriate to look for optimum policies which depend on the inventory level and on some function of the past observable demands but not explicitly dependent on time.

Consider a system with two underlying states. When the system is in State 1, the item demand follows the conditional distribution $g_1(x)$ for x=0, 1,2... When it is in State 2, the conditional demand function $g_2(x)$ is defined to be

$$g_2(x) = \begin{cases} 1 & \text{for } x = 0 \\ 0 & \text{for } x >= 0 \end{cases}$$

If we let the matrix of transition probabilities be

r \ s	1	2
1	e^{-k}	$1-e^{-k}$
2	0	1

the a priori probability that the system will be in State 1 in the next period given that the current demand = x is

$$\pi'' = \begin{cases} \dfrac{e^{-k}\pi\, g_1(x)}{\pi\, g_1(x) + (1-\pi)} & \text{for } x = 0 \\ e^{-k} & \text{for } x > 0 \end{cases}$$

With these definitions of $h_\pi(x)$ and π'', we can formulate the total expected cost function in a fashion similar to that followed with the general Markov demand model.

Note that e^{-k} is the probability that the system survives in any given period given that it did not terminate during the immediately preceding period if we assume that the system's survival follows an exponential distribution.

As described, the optimization procedure apparently applies to the simultaneous optimization of inventory actions for all items whose demands are included in X. In practice, any item or group of items whose inventory decision constraints and cost function L may be separated from the remaining items, by virtue of independent warehousing, procurement, storage, and the like, may be treated separately in the optimization process, since the overall minimum cost will simply be the sum of the separately obtained minimum costs. Thus, a single calculation may link the computation of $\pi'(\pi, X)$ for a group of items, while the optimum policy Ψ may be obtained separately for each one, each time solving the functional equation with the appropriate cost function L and derived inventory status Y'.

Although the states S_r, which play a key role in the development, have been described as "demand-generating states," note that they may also be interpreted as information-generating states. For example, two different states may generate exactly the same demands, resulting in the same cost values L and the same derived inventory positions Y', but may be distinguished with respect to the type of ancillary information implicit in X, thus leading to appropriate modifications of $\pi'(\pi, X)$.

REFERENCES

1. Bellman, R.E., Dynamic Programming, Princeton University Press, Princeton, 1956

2. Bellman, R.E., and Dreyfus, Stuart, Applied Dynamic Programming, Princeton University Press, Princeton, 1962

3. Brown, G.W., Lu, J.Y., and Wolfson, R.J., Dynamic Modeling of Inventories Subject to Obsolescence, C-E-I-R Report to BUSANDA, U.S.N., 1961

4. Howard, R.A., Dynamic Programming and Markov Processes, The Technology Press of MIT, and John Wiley and Sons, Inc., New York

5. Raiffa, H., and Schlaifer, R., Applied Statistical Decision Theory, Division of Research, Graduate School of Business Administration, Harvard University, Boston, 1961

6. Schlaifer, R., Probability and Statistics for Business Decisions, McGraw Hill, New York, 1959