

UNCLASSIFIED

AD 408 523

DEFENSE DOCUMENTATION CENTER

FOR

SCIENTIFIC AND TECHNICAL INFORMATION

CAMERON STATION, ALEXANDRIA, VIRGINIA



UNCLASSIFIED

NOTICE: When government or other drawings, specifications or other data are used for any purpose other than in connection with a definitely related government procurement operation, the U. S. Government thereby incurs no responsibility, nor any obligation whatsoever; and the fact that the Government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data is not to be regarded by implication or otherwise as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use or sell any patented invention that may in any way be related thereto.

63-4-2


MATERIEL MANAGEMENT STUDIES

CATALOGUE

AS AD N

408 523

A COMPARATIVE ANALYSIS OF THE BAYES INVENTORY POLICY



PLANNING RESEARCH CORPORATION
LOS ANGELES, CALIFORNIA



A COMPARATIVE ANALYSIS
OF THE BAYES INVENTORY POLICY

PRC R-201

30 March 1961

Herbert Scarf
John VanderVeer

Prepared for
Bureau of Supplies and Accounts
Department of the Navy
Under Contract NONr-2713(00)

This is a working paper. It may be expanded, modified, or withdrawn at any time.

PLANNING RESEARCH CORPORATION
LOS ANGELES, CALIF. WASHINGTON, D.C.

ABSTRACT

Three kinds of provisioning policies are postulated and compared: a Bayes policy, a dynamic programming policy based on initial demand estimate only, and a dynamic programming policy with periodic recomputation using revised demand forecasts based upon an average of past demands. Cost comparisons are made by simulating (in a Monte Carlo sense) the use of the different policies for several values of estimated mean demand and shortage cost to unit cost ratios, and for various values of actual demand less than and greater than the estimated demand.

Based upon the parameter values chosen, and under the limitations and assumptions of the study, the Bayes policy appears superior (less cost) when demand is underestimated, particularly for high values of the shortage cost to unit cost ratio. The dynamic programming policies are superior when demand is overestimated, with little difference between the two kinds of dynamic programming policies.

TABLE OF CONTENTS

	<u>Page</u>
ABSTRACT	ii
I. INTRODUCTION	1
A. Kinds of Provisioning Policies	1
B. Evaluation Method	3
II. TECHNICAL BACKGROUND	4
III. THE CASE RUNS	9
IV. RESULTS OF THE SIMULATION	12
REFERENCES	16

LIST OF EXHIBITS

	<u>Page</u>
1 Input Values for Case Runs	10
2 Comparison of Policy Costs--Case 1 (Estimated Mean = .12, SC/UC = 10)	17
3 Comparison of Policy Costs--Case 2 (Estimated Mean = .90, SC/UC = 10)	18
4 Comparison of Policy Costs--Case 3 (Estimated Mean = 10.2, SC/UC = 10)	19
5 Comparison of Policy Costs--Case 4 (Estimated Mean = .12, SC/UC = 100)	20
6 Comparison of Policy Costs--Case 5 (Estimated Mean = .90, SC/UC = 100)	21
7 Comparison of Policy Costs--Case 6 (Estimated Mean = 10.2, SC/UC = 100)	22
8 Comparison of Policy Costs--Case 7 (Estimated Mean = .12, SC/UC = 1,000)	23
9 Comparison of Policy Costs--Case 8 (Estimated Mean = .90, SC/UC = 1,000)	24
10 Comparison of Policy Costs--Case 9 (Estimated Mean = 10.2, SC/UC = 1,000)	25
11 Comparison of Dynamic Programming with Recomputation and Bayes (High-Confidence) Policy Costs (Estimated Mean = .12)	26
12 Comparison of Dynamic Programming with Recomputation and Bayes (High-Confidence) Policy Costs (Estimated Mean = .90)	27
13 Comparison of Dynamic Programming with Recomputation and Bayes (High-Confidence) Total Costs	28

I. INTRODUCTION

In any dynamic programming calculation of optimal inventory policies, current decisions are based on explicit assumptions as to the levels of costs and the types of probability distributions which will occur in the future. Costs and policies are insensitive to misestimates of future values for some of these qualities. On the other hand, some of the parameters required as inputs have a high degree of sensitivity, as far as the costs of actually using policies predicated on these estimates are concerned. Generally speaking, the most important parameter with respect to effect on costs is the initial guess as to the future mean demands.

For a relatively new item with limited historical experience, the initial provisioning problem is difficult because of the inability to predict future mean demands with precision. Of course, as time goes on, more and more information about the demand behavior of the part will accumulate, and this very frequently will enable better estimates to be made as to the distribution of demand for the particular part. For this reason, the problem of estimating future demand is considerably more serious in the initial stages of provisioning, as compared with the time after which data has been collected on the behavior of the item.

A. Kinds of Provisioning Policies

In two previous papers [1] [2], a procedure for coming to grips with the uncertainty in our initial estimate of mean demand has been presented. This procedure is to assume that the true distribution of demand has a density function $\phi(\xi|\omega)$ with ω an unknown parameter which might, for example, be the mean demand per period. In addition, an a priori distribution, with density $f(\omega)$, is assumed to describe initial guesses as to the possible values of the unknown parameter. The specific form of $f(\omega)$ may be obtained in many ways, including perhaps a statistical analysis of failure data for parts believed to be similar to the particular part in question. The estimate for $f(\omega)$

is incorporated very heavily in the initial provisioning decisions, but as time goes on, specific historical data on demand will play a more and more important role in the determination of stockage policies. The important feature of this procedure, which is frequently called a Bayes approach, is that the statistical analysis of demand is not isolated from considerations of the cost of initial procurements but is used jointly with purchase cost, holding cost, and shortage cost in determining the degree of conservatism to be used in deciding on initial purchases.

There are several alternative procedures that might be used in the problem of initial provisioning. One, for example, is to take some initial estimate of the mean demand for the item, and compute stockage policies based on this estimate, under the assumption that the estimate will accurately predict mean demand in the future. Then, as time goes on, the initial guess may prove erroneous, and a new estimate of the mean demand may be made on the basis of past demand experience. Policies may then be recomputed every period as if the revised estimate of the mean demand were really the correct one for the remainder of the program. In this study, the estimate of future mean demand was taken to be the average of the past observed demands, though perhaps a better procedure would have been to weight this average and the initial guess. This procedure, referred to as the dynamic programming recomputation procedure, is easy to carry out in practice and makes a certain amount of intuitive sense.

As another possibility, the initial estimate may be made and kept, regardless of the actual data which accumulates as time goes on. The use of this policy, referred to as the straight dynamic programming policy, seems to be less desirable than following the other policies in the sense that damage due to an initial misestimate of demand is not corrected.

The primary purpose of this paper is to explore the relationships among these three policies--the Bayes, dynamic programming with periodic recomputation, and the straight dynamic programming--as far as the cost implications are concerned.

B. Evaluation Method

A number of situations encompassing specific assumptions as to holding cost, purchase cost, shortage cost, and various guesses as to the future demand rate were selected. For each of these cases the three types of policies were computed and the performance of the policies tested, using a Monte Carlo simulation of demand distributions based on a wide variety of means.

Before describing the details of the study, it is perhaps worthwhile to give some general characteristics of the Bayes policy as contrasted with the straight dynamic programming policy. The Bayes policy is generally more conservative in the following sense:

Assuming the same inputs in the two computations of optimal policies--the same unit, holding, and shortage costs and the same estimated mean and variance of demands--then the Bayes procedure attempts to protect against a somewhat wider variety of possible demands than does the straight dynamic programming calculation, which assumes that the mean demand is known with complete certainty. For this reason, if the shortage cost is high, the Bayes policy will make an initial purchase which is generally higher than the straight dynamic programming initial purchase, and occasionally considerably higher. There is no reason to think, however, that because of this conservatism the Bayes approach is necessarily better. If a Bayes approach is applied in a case where the dynamic programming policy actually happens to be based on a good prediction of future demand, then the Bayes approach, being somewhat more conservative, will overbuy for that case and thereby produce a higher cost than will the straight dynamic programming policy.

On the other hand, if the shortage cost is high and the mean demand is underestimated in the straight dynamic programming calculation, then the Bayes approach will save future shortage costs, especially if there is a fairly long lead time in the delivery of items.

II. TECHNICAL BACKGROUND

Attention may now be directed to a description of the specific technical inputs to the study and a description of the procedures used to calculate optimal policies. The dynamic programming calculations are, of course, quite standard [3]. They proceed by means of the usual recursive calculation of functions of a single variable. The Bayes approach is somewhat more complicated.

In addition to the customary inputs to an inventory problem, purchase cost, holding and shortage costs, interest rate, etc., it is necessary to specify a parametric family of demand distributions $\phi(\xi|\omega)$ and an a priori distribution $f(\omega)$, and then use the system of functional equations described in [1]. The calculations generally will consist of recursive calculations of sequences of functions of two variables, one of them being current stock and the other being total past observed demands.

In a later paper [2], a procedure was described which takes advantage of some additional assumptions and permits a simplification in the Bayes calculations. This is accomplished by replacing the recursive calculations of functions of two variables by the recursive calculations of functions of a single variable. The basic ingredients in this latter simplification consist first of all in a specification of the parametric family of demand distributions to be the Gamma family of demand distributions,

$$\phi(\xi|\omega) = \frac{\omega^a \xi^{a-1} e^{-\omega\xi}}{\Gamma(a)} ;$$

secondly, an assumption that the a priori distribution is itself a Gamma distribution; i.e.,

$$f(\omega) = \frac{\lambda^b \omega^{b-1} e^{-\lambda \omega}}{\Gamma(b)} ;$$

thirdly, to insist that the purchase cost function be composed entirely of a unit cost with no setup cost; and then finally, an assumption which is customarily made anyway, namely, that the holding cost and the shortage cost are both linear functions of their arguments.

These assumptions were adopted in [2] so as to simplify the recursive calculation of Bayes policies, though it is by no means impossible to compute Bayes policies based on the calculation of functions of two variables rather than a single variable. Because of this simplification, however, the study was restricted to problems in which no setup cost appeared. Possibly a comparison of the three policies would have produced different results in the case in which a large setup cost was part of the picture.

In order to make the problems comparable, it was also assumed that the Gamma distribution was the relevant probability distribution for the straight dynamic programming calculation and the dynamic programming with recomputation.

A further specification may be made concerning the procedure that was used in the Bayes computation to translate the estimated mean demand and some notion of confidence in this estimate into numerical quantities. As was mentioned before, the parametric family for demand was selected to be the Γ -family, i. e.,

$$\phi(\xi | \omega) = \frac{\omega^a \xi^{a-1} e^{-\omega \xi}}{\Gamma(a)} .$$

The mean of this distribution is $\frac{a}{\omega}$ and the variance $\frac{a}{\omega^2}$. The a priori distribution was also selected to be a member of the Γ -family, with the specific form

$$f(\omega) = \frac{\lambda^b \omega^{b-1} e^{-\lambda \omega}}{\Gamma(b)} .$$

In all, it was necessary to specify three parameters with certainty: the parameters a , b , and λ . The procedure selected for determining these values was as follows:

As one input into the problem, a specific value for the estimated mean was taken, that is, a value that the quantity $\frac{a}{\omega}$ is expected to have. If the expectation of $\frac{a}{\omega}$ is computed with respect to the Bayes distribution,

$$\int \frac{a}{\omega} f(\omega) d\omega = \frac{a\lambda}{b-1}$$

is obtained. A specific value for this quantity was selected for the "initial guess." In addition, some allowance was made for the confidence in the initial guess as to the true value of $\frac{a}{\omega}$; that is, the variance of the mean demand estimate was another input to the study. Since, however, the true value of $\frac{a}{\omega}$ is not known, it is possible that it may deviate significantly from $\frac{a\lambda}{b-1}$. If there is considerable confidence in the initial guess as to the true mean, then the variance of

$$\frac{a}{\omega} - \frac{a\lambda}{b-1}$$

computed with respect to the Bayes distribution should be small; for little confidence in the initial guess, this variance should be large. This "variance of the initial guess" may be computed directly as

$$\begin{aligned} & \int_0^{\infty} \left(\frac{a}{\omega} - \frac{a\lambda}{b-1} \right)^2 f(\omega) d\omega \\ &= \frac{a^2 \lambda^2}{(b-1)(b-2)} - \frac{a^2 \lambda^2}{(b-1)^2} \\ &= \frac{a^2 \lambda^2}{(b-1)^2 (b-2)} \end{aligned}$$

Actually, it was more convenient to work with

$$\begin{aligned} & \frac{\text{"variance of the initial guess"}}{\text{"initial guess"}} \\ &= \frac{a^2 \lambda^2}{(b-1)^2 (b-2)} \bigg/ \frac{a \lambda}{b-1} \\ &= \frac{a \lambda}{(b-1)(b-2)} . \end{aligned}$$

In the study, two different choices were made for this quantity. Values of 3 and 10 were selected, furnishing two Bayes policies for each case: one based on fairly high confidence, and the other on fairly low confidence as to the initial guess.

Two conditions have now been placed on the problem, whereas, as was mentioned before, there are three constants which are to be selected: a , λ , and b .

The third condition that was imposed was the use of a member of the Gamma family similar to the one used for the dynamic programming study. For the latter computation, the member of the Γ -family with a ratio of variance to mean of 3 was consistently used. Since, for the Γ -family, the mean is $\frac{a}{\omega}$ and the variance $\frac{a}{\omega^2}$, it follows that

$$\frac{\text{variance}}{\text{mean}} = \frac{\text{mean}}{a} .$$

In this expression, the mean was replaced by estimated mean, and as a final condition it was assumed that

$$3 = \frac{\text{estimated mean}}{a} .$$

This has the virtue of making the value of "a" the same in the Bayes calculation as in the dynamic programming calculation.

The relationship between a, b, λ , and

$$m = \text{"initial guess,"}$$

$$a = \frac{\text{"variance of initial guess"}}{\text{"initial guess"}}$$

may be summarized as follows:

$$a = \frac{m}{3}$$

$$b = 2 + \frac{m}{a}$$

$$\lambda = 3\left(\frac{m}{a} + 1\right) .$$

III. THE CASE RUNS

In this study, nine basic cases were computed in order to investigate the cost effects of errors in estimating demand at various points in the parameter space. Three different means were used: 0.12, 0.90, and 10.20 per year. These means were the estimated means in the Bayes cases in the sense of being the expectation with respect to the a priori distribution of the mean of the parametric family, and the estimated means in the dynamic programming cases in the sense of being specific inputs to the study. Also, three different values for the ratio of shortage cost to unit cost were used: 10, 100, and 1000. Forming all the combinations of estimated mean demand and shortage cost to unit cost ratios yields the nine basic cases. These input parameters are summarized in Exhibit 1.

In all nine cases, the holding cost rate was .01 per year, the interest rate was 20 percent per year, the lead time was one year, and the program length was 8 years. The nine cases were also run for a program length of 5 years, but the results were not significantly different than for the 8-year program. Therefore, the cases for a 5-year program are omitted from the subsequent discussion of results.

For each of the nine cases, four different policies were computed: low-confidence Bayes, high-confidence Bayes, dynamic programming with periodic recomputation, and the straight dynamic programming. The distinction between the low-confidence and high-confidence Bayes policies is as defined on page 7.

For each of these 9 cases and each of the four possible policies, 11 different simulations were run, drawing random numbers to represent demand from a population with a true mean (a population given by a Gamma family with a ratio of variance to mean of 3) ranging from .06 to 100.2. For each of the eleven true means and for every one of the 9 cases, the expected cost was computed for an eight-period problem for the four different policies, using 200 replications in the simulation.

EXHIBIT 1 - INPUT VALUES FOR CASE RUNS

<u>Case Number</u>	<u>Mean (estimated)</u>	<u>Shortage Cost/Unit Cost</u>
1	.12	10
2	.90	10
3	10.20	10
4	.12	100
5	.90	100
6	10.20	100
7	.12	1000
8	.90	1000
9	10.20	1000

Program length = 8 years
Lead time = 1 year
Holding cost rate = .01 per year
Interest rate = 20%

The variance of total cost over the 200 replications of the simulation was also computed. Standard procedures indicate that the estimated mean total cost was within 5 percent of the true expected cost with probability at .95.

IV. RESULTS OF THE SIMULATION

The results of the study are shown in the form of two sets of graphs given at the end of this report. The first set (Exhibits 2 through 10) consists of nine graphs corresponding to the nine different cases. In each of these graphs, the Bayes policy costs and dynamic programming recomputation policy costs are compared to costs resulting from the use of the straight dynamic programming policies. These costs, in the form of ratios, are plotted as functions of true mean demand; the initial estimated mean corresponding to each case is chosen as a vertical line. On these graphs, the lower the ordinate, the better the cost consequence of following the given policy.

The second set (Exhibits 11 through 13) consists of three graphs corresponding to the three values of initial estimated mean. On each graph and for each of the three different values of shortage cost to unit cost ratio, the ratio of costs resulting from dynamic programming with recomputation to Bayes policy costs is plotted as a function of true mean demand. When the ordinate has a value greater than 1.0, the Bayes policies are superior, and when the ordinate is less than 1.0, the reverse is true.

From the first set of graphs, a number of general statements may be made about the worth of the Bayes policy as compared with the worth of the straight dynamic programming policy or the dynamic programming recomputation policy.

There are essentially two types of cases that can occur. It is possible that the initial guess overestimates the true mean, or the initial guess may underestimate the true mean. Different policies seem to be better, depending on whether one or the other of these cases occurs. For example, if the true mean is underestimated, the fact that the Bayes policy initially considers the possibility of misestimation means that the Bayes policy will generally buy more than the dynamic programming recomputation at the beginning of the problem. Since this saves shortage costs at the beginning of the problem, the Bayes policy appears to be better in this case, especially if the

shortage cost is high. The intuitive idea is that while the dynamic programming recomputation policy will eventually pick up the fact that the true mean has been underestimated, it will pick it up too late to avoid incurring large shortage costs at the beginning of the program. Exhibits 8, 9, and 10, all of which have shortage costs of 1000 times unit cost, demonstrate this fact very vividly, at least for moderate underestimates.

On the other hand, if the true mean is overestimated and at the same time there is a substantial shortage cost, the Bayes approach will buy considerably more at the beginning than the dynamic programming recomputation. If there is a sufficiently high interest rate or holding cost, then the Bayes policy will turn out to be considerably worse than the dynamic programming recomputation policy. And indeed this turns out to be the case, though the improvement of the dynamic programming recomputation over the Bayes in this case of overestimating is by no means as striking as the improvement of the Bayes over the dynamic programming recomputation procedure in the case of underestimating.

The second set of graphs establishes a criterion for choice between the Bayes policy and the dynamic programming with recomputation, assuming these are the only candidates for choice. If one of these policies were to be chosen for use over the entire range of true demand, it is apparent that the Bayes policy should be chosen for large values of shortage cost to unit cost ratio. The dynamic programming with recomputation should be chosen for low values of this ratio. This criterion may be interpreted in terms of unit cost, where the Bayes policy is better for low cost items and the dynamic programming with recomputation better for high cost items. This criterion is further sharpened when a more reasonable demand forecasting procedure is used in the case of dynamic programming with recomputation. In particular, this would lower the curves in Exhibit 13 to the right of the estimated mean, where demand is underestimated.

Actually another surprising point turns up, which has a simple enough explanation. It appears in the case where the dynamic programming recomputation procedure is better than the Bayes. If the first set

of graphs is examined, it is seen that in the case where demand is initially overestimated, there really is not much to choose between as far as the dynamic programming recomputation and the straight dynamic programming. This fact may be seen in virtually all of the 9 cases if the graphs are examined to the left of the estimated mean. The explanation is simple. The dynamic programming recomputation overestimates the initial mean, buys more than it should, and never buys again for the rest of the program. Precisely the same sort of thing happens with the straight dynamic programming without any recomputation, so that for this case of overestimating, the straight dynamic programming and the dynamic programming recomputation are virtually the same. Actually the straight dynamic programming is occasionally somewhat better in the neighborhood of a guess which is close to the true mean, for the reason that in this region the recomputation procedure has sampling error introduced into it as far as the computation of the appropriate mean to be used next period.

The main conclusion seems to be that the two competitors are the Bayes solution and the straight dynamic programming solution, based on an initial guess. This is, of course, somewhat surprising inasmuch as the recomputation procedure seems to make a certain amount of intuitive sense. However, this can be summarized by saying that the recomputation procedure is a better procedure for picking up systematic changes in demand than for coming to grips with initial misestimates in demand. Either the Bayes procedure is better in the case where demand is underestimated, or, in the other case where demand is overestimated, so much damage has been done already on the basis of the large initial buy that very little else can be done after that.

In summary, then, this analysis tends to show that the Bayes policy is better than the simpler policies when demand is underestimated and when applied to low-value items. However, these conclusions hold only under the restrictive assumptions of the study: a stationary true demand applied to a single echelon, a small program length to lead time ratio, and no fixed ordering cost. Different conclusions may pertain

for the case of a phase-in, phase-out demand pattern, a multi-echelon supply structure, a fixed ordering cost, or a larger ratio of program length to lead time. Also, no evaluation was made as to the worth of the Bayes policy in terms of potential savings as contrasted with increased computation costs, and increased costs of obtaining the additional input data that is required.

REFERENCES

1. H. Scarf, "Bayes Solutions of the Statistical Inventory Problem," Annals of Mathematical Statistics, Volume 30 (1959), pp. 490-508.
2. H. Scarf, "A Simplification of the Bayes Approach to the Inventory Problem." (Unpublished)
3. K.J. Arrow, S. Karlin, H. Scarf, Studies in the Mathematical Theory of Inventory and Production, Stanford: Stanford University Press (1958).




EXHIBIT 2 - COMPARISON OF POLICY COSTS--CASE 1
(ESTIMATED MEAN = .12, SC/UC = 10)




EXHIBIT 3 - COMPARISON OF POLICY COSTS--CASE 2
(ESTIMATED MEAN = .90, SC/UC = 10)




EXHIBIT 4 - COMPARISON OF POLICY COSTS--CASE 3
(ESTIMATED MEAN = 10.2, SC/UC = 10)




EXHIBIT 5 - COMPARISON OF POLICY COSTS--CASE 4
(ESTIMATED MEAN = .12, SC/UC = 100)




EXHIBIT 6 - COMPARISON OF POLICY COSTS--CASE 5
(ESTIMATED MEAN = .90, SC/UC = 100)




EXHIBIT 7 - COMPARISON OF POLICY COSTS--CASE 6
(ESTIMATED MEAN = 10.2, SC/UC = 100)




EXHIBIT 8 - COMPARISON OF POLICY COSTS--CASE 7
(ESTIMATED MEAN = .12, SC/UC = 1,000)




EXHIBIT 9 - COMPARISON OF POLICY COSTS--CASE 8
(ESTIMATED MEAN = .90, SC/UC = 1,000)




EXHIBIT 10 - COMPARISON OF POLICY COSTS--CASE 9
(ESTIMATED MEAN = 10.2, SC/UC = 1,000)




EXHIBIT 11 - COMPARISON OF DYNAMIC PROGRAMMING WITH RECOMPUTATION AND BAYES (HIGH-CONFIDENCE) POLICY COSTS (ESTIMATED MEAN = .12)




EXHIBIT 12 - COMPARISON OF DYNAMIC PROGRAMMING WITH RECOMPUTATION AND BAYES (HIGH-CONFIDENCE) POLICY COSTS (ESTIMATED MEAN = .90)




EXHIBIT 13 - COMPARISON OF DYNAMIC PROGRAMMING WITH RECOMPUTATION AND BAYES (HIGH-CONFIDENCE) TOTAL COSTS