

MEMORANDUM
RM-4720-PR
DECEMBER 1965

A SYSTEM APPROACH TO
BASE STOCKAGE OF RECOVERABLE ITEMS

G. J. Feeney and C. C. Sherbrooke

PREPARED FOR:
UNITED STATES AIR FORCE PROJECT RAND

The **RAND** *Corporation*
SANTA MONICA • CALIFORNIA

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE DEC 1965	2. REPORT TYPE	3. DATES COVERED 00-00-1965 to 00-00-1965			
4. TITLE AND SUBTITLE A System Approach to Base Stockage of Recoverable Items		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Rand Corporation, Project Air Force, 1776 Main Street, PO Box 2138, Santa Monica, CA, 90407-2138		8. PERFORMING ORGANIZATION REPORT NUMBER			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 54	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

MEMORANDUM

RM-4720-PR

DECEMBER 1965

A SYSTEM APPROACH TO
BASE STOCKAGE OF RECOVERABLE ITEMS

G. J. Feeney and C. C. Sherbrooke

This research is sponsored by the United States Air Force under Project RAND—Contract No. AF 49(633)-1700—monitored by the Directorate of Operational Requirements and Development Plans, Deputy Chief of Staff, Research and Development, Hq USAF. Views or conclusions contained in this Memorandum should not be interpreted as representing the official opinion or policy of the United States Air Force.

DISTRIBUTION STATEMENT

Distribution of this document is unlimited.

The RAND Corporation

1700 MAIN ST • SANTA MONICA • CALIFORNIA • 90406

PREFACE

This is the concluding Memorandum in a series describing a system approach to setting base stock levels on recoverable items. (See "References," p.41.) It should be of interest to anyone concerned with the problem of base stockage, whether he be in management or a member of an implementation team. The intent is to provide a readable and concise summary of the work accomplished; consequently, no derivations are included here, since the mathematical foundations have been developed in previous Memoranda. Mathematical results obtained in the previous studies are cited, but it is not necessary that the reader be familiar with those studies.

The associated computer program is described in a Memorandum by J. Y. Lu and Gabriele A. Michels (See Ref. 1).

SUMMARY

This Memorandum describes a system approach to setting base stock levels on recoverable items, and demonstrates its application.

The system approach does not replace the traditional item analysis. It translates to a system viewpoint the implications of item data, including unit cost and repair characteristics. This approach provides four major advantages:

- (1) A Bayesian procedure, using data on all items in the supply system, can provide improved demand estimates for individual items.
- (2) A system cost-effectiveness curve displaying realistic alternatives to a supply manager can be generated. This eliminates the necessity for arbitrary estimates of holding and stockout costs.
- (3) It provides a mechanism for examining sensitivity to parameter changes.
- (4) It enables us to compare different stockage policies.

An examination of the sensitivity of the stockage policy to parameter changes indicated that it is sensitive to the item variance-to-mean ratio of the assumed stuttering Poisson demand distribution but insensitive to other parameter changes.

A comparison between stock levels computed by using the stockage policy and actual stock levels from an Air Force base indicates that this policy can achieve the same performance at about one-half the investment. A second comparison between the stockage policy and the current AFM 67-1 policy indicates that for comparable performance the investment can be reduced by one-half to three-quarters.

CONTENTS

PREFACE	iii
SUMMARY	v
Section	
I. INTRODUCTION	1
II. PROBLEM DESCRIPTION	2
III. ITEM ANALYSIS	4
IV. SYSTEM APPROACH	10
V. SENSITIVITY OF SYSTEM PERFORMANCE	14
Advantages of the System Approach	14
Data Used in the Sensitivity Tests	15
Description of the Sensitivity Tests	15
Display of Sensitivity Results	16
Sensitivity to the Choice of Performance Measure	17
Sensitivity to Response Time and Level of Operation	19
Sensitivity of Fill Rate to the Item Variance- to-Mean Ratio	20
Sensitivity to the Discrete Approximation	22
VI. COMPARISON OF THE STOCKAGE POLICY WITH AIR FORCE PROCEDURES	23
Authorized and Actual Stock Levels at Oxnard AFB	23
Current Air Force Manual 67-1 Procedures	26
VII. CONCLUSION	32
APPENDIX	33
REFERENCES	41

I. INTRODUCTION

Research in inventory theory has concentrated on two areas: (1) calculation of the utility function, $U(A|\theta)$, for each inventory action, A , when the vector of item parameters, θ , is known; and (2) the estimation of the item parameters, θ , from data, x . In an earlier Memorandum⁽²⁾ the authors demonstrated the calculation of several utility functions, $U(A|\theta)$, under an $(s-1,s)$ inventory policy when demand is compound Poisson and the resupply (repair) distribution is arbitrary. The $(s-1,s)$ policy means that a reorder is placed whenever a demand occurs, and this is the optimal policy for high-cost, low-demand items. This is an important class of items including, for example, most aircraft recoverable spare parts that may cost hundreds of dollars and are typically demanded only a few times a year at an individual air base.

In a second Memorandum⁽³⁾ the authors described an objective Bayesian procedure for estimating parameters θ , from data. The procedure should be particularly significant for low-demand items, but it is not limited to any specific inventory policy.

The present Memorandum combines the results of these two papers to demonstrate a system approach to supply management. Item decisions will be made that minimize system cost for each level of system performance, generating a cost-effectiveness curve. The sensitivity of the approach to parameter changes and different criteria will be examined in detail.

II. PROBLEM DESCRIPTION

The supply process for a recoverable item operates as follows. When an item fails in the course of base operations, it is examined to determine whether repair is possible at base level. If so, the item is scheduled into base repair and, after a variable lag representing base repair-cycle length, it is returned to a serviceable condition. If base repair is not indicated, the item is either condemned or adjudged NRTS (Not Repairable This Station) and forwarded to the depot for repair action. In the latter cases the base submits a requisition to the depot for a serviceable replacement that will arrive at the base after a variable lag representing base resupply-cycle time (order and shipping time).

Present Air Force policy⁽⁴⁾ establishes that the total supply level for an item, defined as the sum of authorized quantities on hand plus on order minus back orders,* will be computed as the sum of A + B + C according to the following scheme:

Element					
Base Repair Cycle Quantity	(A)	=	Daily Demand X Rate	Fraction of Units Base X Reparable	Base Repair Cycle (Days)
Order and Shipping Quantity	(B)	=	Daily Demand X Rate	Fraction of Units NRTS X and Con- demned	Order and Shipping Time (Days)
Safety Level Quantity	(C)	=	$K \sqrt{3(A + B)}$ (normally K = 1)		

The daily demand rate for an item is simply the observed demands divided by the number of days in the observation period (the period

*Units on order and back orders include units in base repair as well as units being resupplied from depot.

must be at least 180 days). The percentage of units base reparable and its complement are also estimated from these observed data. Eventually the base repair cycle and the order and shipping time will be computed for each item, but standard times are being used at present. The order and shipping time in the continental United States is set at 8 days for XD-1 or Hi-Valu items and 30 days for XD-2 items, and over-seas at 15 and 60 days, respectively.

For example, an item with a daily demand rate of 0.0091, equivalent to one demand every 110 days, is just sufficient to establish a stock level of one (when $A + B + C = 0.5$ is rounded to the nearest integer) for an item that is always base reparable with an average repair time of seven days ($K = 1$).

In order to increase the range of items stocked, a recent addendum stipulates that when two demands have been recorded on an item, a stock level of one will be authorized if demands occur at the rate of one every 270 days (a daily demand rate of 0.0037). The item repair and resupply characteristics are ignored in this computation.

Conceptually, this new stockage policy improves on the earlier policy by considering the base repair-cycle time and the order and shipping time by item. However, this policy still ignores an important policy variable, unit cost.* Under this policy two items with the same observed demand characteristics will receive the same stock level, even if one costs \$50 and the other \$5000. An optimal policy should stock more units of a low-cost item at base level and rely on premium transportation and expediting from depot for a high-cost item.

*Unit cost enters the calculation implicitly to the extent that the standard order and shipping times are different for XD-1 and XD-2 items. However, the classification of items as XD-1 is not determined strictly by cost, and the stated objective of using actual order and shipping times means that in the future stock levels will be established independently of item cost.

III. ITEM ANALYSIS

This section develops a method for determining the optimal stock level on an item managed under an $(s-1,s)$ inventory policy. This policy simply means that a reorder (repair) is initiated whenever a demand occurs.* The following data are assumed to exist for each item:

1. Total demand during a common time period.
2. Unit cost.
3. Average response time.

The average response time for an item is the average base repair time if the item is always base reparable; it is the average base resupply time if the item is never base reparable; and if an average fraction r of the units are base reparable, the average response time is r times the average base repair time plus $(1 - r)$ times the average base resupply time. Note that the average base repair time, the average base resupply time, and the fraction r may each vary by item. The base resupply time may itself be a composite of a normal resupply time and some longer resupply time when the depot stock on hand is zero.

Any demand that cannot be met from stock on hand is backordered. Since the appropriate criterion for a supply system is subject to some dispute, we consider three obvious measures of item supply performance. Each is a function of the total supply level or spare stock, s :

- $R(s) \equiv$ Ready rate -- the probability that an item observed at a random point in time has no back orders.
- $F(s) \equiv$ Fills -- the expected number of units demanded per time period for an item that can be filled immediately from stock on hand.
- $S(s) \equiv$ Units in service -- the expected number of units in routine resupply or repair at a random point in time. This equals the expected number of units in resupply or repair minus the expected number of units in a backorder condition at a random point in time.

*The model assumes that resources are such that the time distribution for resupply (repair) of a reparable turn-in is not affected by the number of units already in resupply (repair). In theoretical terms we are assuming an infinite channel queueing facility.

Reference 2 derives the expressions for each of these three performance measurements as a function of the spare stock, s , assuming that the item's true mean demand, θ , is known. Demand is assumed to have a stuttering Poisson distribution, a generalization of the simple Poisson that appears to represent observed demand patterns far better because it has a second parameter. Reasons for considering generalizations to the simple Poisson are discussed extensively in Ref. 2.

For the sake of completeness we shall display the expressions for each of these three performance measurements. Under stuttering Poisson demand with true mean demand rate, θ , and compounding parameter, ρ , the steady-state probability of observing x demands in time T is

$$(1) \quad p(x) = \sum_{y=0}^x \frac{(\lambda T)^y e^{-\lambda T}}{y!} f^{y*}(x)$$

where $\lambda = \theta(1-\rho)$, $f^{y*}(x)$ is the y -fold convolution of $f(x)$, and

$$(2) \quad f(x) = (1-\rho)\rho^{x-1}, \quad \begin{array}{l} x \geq 1 \\ 0 \leq \rho < 1 \end{array}$$

The y -fold convolution of $f(x)$ is simply the probability distribution for the sum of y samples from $f(x)$.

When T is identified as the average item response time, Eqs. (B.7), (B.8), and (B.9) of Ref. 2 yield

$$(3) \quad R(s) = \sum_{x=0}^s p(x)$$

$$(4) \quad F(s) = \lambda R(s-1) + \rho F(s-1)$$

$$(5) \quad S(s) = S(s-1) + 1 - R(s-1).$$

Reference 2 includes a simple computer program for performing these computations. These formulas are remarkably simple because all the performance measurements are functions of the mean response time T rather than the response time distribution.

Reference 2 concludes by defining the utility function for an item as the value of performance minus the cost over a specified time

period. In this expression $G(s|\theta)$ represents an arbitrary performance measurement (ready rate, fills, or units in service), k_1 is the per-unit value of performance during the period, k_2 is the holding cost in percent for the period (obsolescence, warehousing, interest), and c is the unit cost. The utility, $U(s|\theta)$, of carrying a stock level s on an item whose known true mean demand is θ can be written

$$(6) \quad U(s|\theta) = k_1 G(s|\theta) - k_2 cs .$$

Ordering costs do not enter Eq. (6) because they are independent of the stockage, s , under an $(s-1, s)$ inventory policy.

Reference 3 was concerned with the fact that the true mean demand, θ , is seldom known exactly. The traditional approach is to estimate θ by an issue rate technique (demand observed over some past period divided by the length of the period). Such an approach is adequate if the item has relatively high demand, if a relatively long history is available, and if the history is relevant. But many of the most important items, particularly high-cost, low-demand spare parts, fail to meet one or more of these three essential requirements.

Reference 3 describes a new approach to demand analysis, based on a technique called Bayesian inference. The Bayesian approach exploits the surprising fact that we can increase our knowledge of an item by analyzing the behavior of other related items in the supply system. For example, consider the problem of spare-part stockage for an F-105 squadron. Suppose the operational program is altered and the recoverable item demand increases. Let us define a factor a , called the change in the level of operations, as the expected system demand in the prediction period divided by system demand observed in a prior period of the same length. For example, if the number of sorties is expected to double in the future, the factor a would probably be estimated as two. Since the failure rate for any specific part is very low, many spare parts will still experience no demand during any specified time period. Under the traditional approach to demand analysis, each item is treated completely independently. By contrast, the Bayesian approach enables us to adjust the true mean demand, though

unknown, for each item by the level of operations to reflect the expected change in average demand for all the F-105 spare parts. Instead of trying to make a point estimate of the item's true mean demand, θ , this approach computes the probability, $q(\theta|x)$, that an item with observed demand x has a true mean demand θ for a range of θ values. This seems to be a much more realistic way of characterizing what we can learn from an item's demand history. Now we can rewrite Eq. (6) in terms of the expected utility for an item with observed demand x :

$$\begin{aligned}
 (7) \quad \bar{U}(s|x) &= \int_{\theta} U(s|\theta)q(\theta|x)d\theta = k_1 \int_{\theta} G(s|\theta)q(\theta|x)d\theta - k_2cs \\
 &= k_1 \bar{G}(s|x) - k_2cs .
 \end{aligned}$$

The distribution of $q(\theta|x)$ is related to the stuttering Poisson demand distribution $p(x|\theta)$ by Bayes' rule,

$$(8) \quad q(\theta|x) = \frac{p(x|\theta)f(\theta)}{\int_{\xi} p(x|\xi)f(\xi)d\xi}$$

Therefore, the reformulation of inventory policy under limited data is reduced to (a) justifying the treatment of θ as a random variable, and (b) specifying the prior distribution $f(\theta)$.

Reference 3 shows how a probability distribution of true item mean values can be constructed for a system of items. In the absence of item information, the true mean value, θ , for an item can be considered as a random draw from this prior distribution. If item information is available, such as past demand, it can be incorporated with Eq. (8).

A log normal form was selected for the prior distribution and equations were derived to estimate the log normal parameters from data observed on the system of items. In other words the system of inventory items provides the information for a prior distribution

common to all items and then Eq. (8) combines this information with the demand data on a specific item to give a posterior distribution of mean demand for that item. This is what is meant by the statement that the Bayesian approach utilizes system information to improve our knowledge on an individual item.

Now we shall summarize the Bayesian procedure for the simplest case, in which the variance-to-mean ratio of the stuttering Poisson demand distribution is assumed to be a constant α for all items. The first two moments of the demand data on the set of items are computed. In Ref. 3 this was labeled the cross-section of observed demand, and the moments were denoted v_1, v_2 respectively. These moments provide estimates of the parameters for the prior distribution. However, instead of determining the log normal parameters, it is more convenient to compute the mean, μ , and variance, σ^2 , of the equivalent normal distribution. Equation (11) in Ref. 3 derived

$$(9) \quad \sigma^2 = \ln \left\{ \frac{v_2 - \alpha v_1}{v_1} \right\}$$

$$(10) \quad \mu = \ln v_1 - \frac{1}{2} \sigma^2$$

If α is unknown it can be estimated from two or more periods of cross-sectional data; the procedure is described in Ref. 3.

Our objective is to calculate a discrete approximation to Eqs. (7) and (8), and for this purpose we specify N values of the standard normal deviate* $u_1, u_2 \dots u_N$. The probability associated with u_1 is obtained from the cumulative standard normal distribution function at the point $(u_1 + u_2)/2$; that with u_2 is obtained by taking the cumulative normal distribution function at the point $(u_2 + u_3)/2$ and subtracting the previous cumulative value, etc. These values may be located in tables of the normal distribution, or a computer approximation may be used. (5)

*The standard normal deviate has mean 0 and variance 1.

In this manner we can construct an N point discrete approximation to the prior distribution, $f(\theta_i)$. To compute Eq. (8) we must convert the values of the standard normal deviates to the corresponding log normal values, i.e.,

$$(11) \quad \theta_i = e^{\mu + \sigma u_i}$$

Now we can calculate $p(x|\theta_i)$ as described in Eq. (2), and finally the utility function in Eq. (7), for any choice of performance measurement.

IV. SYSTEM APPROACH

Section III described how the item's utility function, $\bar{U}(s|x)$, in Eq. (7) could be calculated for three performance measures under an $(s-1,s)$ inventory policy. Using the traditional approach in inventory problems we would assign the cost k_1 , the per-unit value of performance during the period, and the cost k_2 , the holding cost, and then maximize the item utility in Eq. (7). The question is, why is this solution to the problem not complete?

Serious objections to the traditional approach can be raised at two levels. In the first place it is extremely difficult to estimate the critical parameters k_1 and k_2 . How much is a fill worth? How much should we charge for money tied up in stock? In most real situations we can do little better than make an arbitrary guess. A second and much more fundamental objection is that the item approach tells us nothing about what is going to happen to the system. It is all very well to say that the "proper" charge for invested funds is 15 percent per year, but we are in trouble if the resulting item decisions lead to system requirements that exceed available budgets. Similarly, a special study decreeing that a shortage costs \$357.00 is of little comfort if the resulting system performance causes widespread customer dissatisfaction.

To overcome these limitations in the conventional treatment of stock policy, we have developed what might be called a system approach to stock policy. This technique does not replace item analysis; it carries the process one step further to display the system implications of item decisions and permit system policy that is readily implemented at the item level.

Let us now contrast the item approach and the system approach. In the item approach, we measure k_1 and k_2 and maximize utility for each item from Eq. (7), noting that the maximization depends only on the ratio k_2/k_1 . In the system approach we state a system objective and then determine the item decisions. Our system objective may be

to achieve a specified level of system performance at minimum investment cost.* Here system performance is defined as the sum of the performances for the individual items. This can be formulated mathematically as

$$(12) \quad \min_{\{s_i\}} \sum_{i=1}^N c_i s_i$$

subject to

$$\sum_{i=1}^N \bar{G}_i(s_i | x_i) = \Gamma_0 ,$$

which can be rewritten with a Lagrange multiplier as

$$(13) \quad \max_{\{s_i\}} \left\{ \sum_{i=1}^N \bar{G}_i(s_i | x_i) - \lambda \sum_{i=1}^N c_i s_i \right\} .$$

Eq. (13) is formally equivalent to Eq. (7), where $\lambda = k_2/k_1$ expresses the relation between system cost and system performance. For any value of λ we can compute the optimal stock levels and the system cost and performance. As λ is decreased the optimal stock levels will increase and a cost-effectiveness curve is generated. Instead of attempting to "measure" the ratio of holding cost to stockout cost, λ , we treat this ratio as a management control variable.

The only computational problem is that the value of λ corresponding to any specified level of system performance is unknown. For this reason we employ a simple computer algorithm to generate the cost-effectiveness curve. The Appendix contains a complete mathematical description of the algorithm, but its essential elements are given below.

*The investment cost is the only relevant cost under an (s-1,s) inventory policy because the number of orders is not a variable.

For some sufficiently large positive λ , the optimal values of $\{s_i\}$ in Eq. (13) are zero. As we decrease the value of λ , it becomes optimal to stock one unit of some item j because the following quantity becomes positive:

$$(14) \quad \bar{G}_j(1|x_j) - \lambda c_j .$$

The general allocation procedure consists of a series of steps. At each step an item is selected and the stock level s_j is increased by one, where that item j gives the maximum value of

$$(15) \quad \frac{\bar{G}_j(s_j + 1|x_j) - \bar{G}_j(s_j|x_j)}{c_j} = \frac{\Delta \bar{G}_j(s_j|x_j)}{c_j}$$

where Δ denotes the first forward difference. After each step we can compute the system performance and system cost as shown in Eq. (13). When either the system performance or system cost just exceeds some target value, we print out the optimal stock levels and the corresponding values of system performance, system cost, and λ . The cost-effectiveness curve is not actually a continuous curve, because the variables $\{s_j\}$ can assume only discrete values. However, in an inventory system with several hundred items the attainable points on the cost-effectiveness curve will be closely spaced, so that for all practical purposes the curve appears to be continuous.

In the discussion above a constant value of λ was employed for all items. Using fills as our criterion, this would mean that a fill on any one item has the same value as a fill on any other item. Over the class of recoverable items this is probably a reasonable assumption, but it is not required by the mathematics. Suppose, for example, that we have an item on which a fill is considered to be three times as important as a standard item fill, and which has a holding cost that is half as large. Since $\lambda = k_2/k_1$, the correct value of λ for this special item is one-sixth of the value for the standard item. But Eq. (13) is linear in cost, and replacing c_j by $c'_j = \frac{1}{6} c_j$ for this

special item, we obtain Eq. (13) again with a constant value of λ for all items. Thus the allocation algorithm given by Eq. (13) is directly applicable on the revised cost. Of course, this cost revision procedure can be employed with any performance measure.

In summary, the system approach has two important features. First, it makes visible the best attainable system alternatives. The supply manager can select a realistic alternative that presents the most acceptable tradeoff between cost and effectiveness. Second, it provides a workable mechanism for implementing system decisions, even in a completely decentralized organization. The control ratio, λ , selected at the system level, can be used to guide item stock decisions in remote locations without intruding on detailed operations.

V. SENSITIVITY OF SYSTEM PERFORMANCE

ADVANTAGES OF THE SYSTEM APPROACH

Two advantages of the system approach over conventional techniques have been described. In Sec. III the system approach enabled us to use a Bayesian procedure in estimating the true mean demand for each item. In Sec. IV the system approach generated a cost-effectiveness curve displaying the efficient system alternatives to a supply manager. Each system alternative was related to a set of item stock levels without a requirement for arbitrary estimates of holding or stockout costs.

The system approach has a third advantage: the ability to examine the sensitivity of a supply system to changes in the various input parameters. It is not immediately obvious, however, that the system approach is more appropriate for sensitivity testing. For example, if we are interested in the effect of doubling the average item response time, it is reasonable to ask why we cannot use the item approach twice to calculate the item decisions and the system cost and performance for the same value of $k_2/k_1 = \lambda$. Mathematically, it is perfectly acceptable to do so, since an efficient set of stock levels will be computed, but for the supply manager this information is likely to be irrelevant. The matter is illustrated by Fig. 1. Let us assume that originally we were at point A of the short-response curve. By selecting the same value of the slope, λ , on the long-response curve, we obtain point B. But the supply manager is probably more interested in determining point C, the degradation in system performance at the same level of investment, or point D, the additional investment required to maintain the same level of performance. The corresponding values of λ are unknown, but the algorithm described in the Appendix provides a simple method for locating these targets, expressed in terms of either investment or performance.

DATA USED IN THE SENSITIVITY TESTS

The data consisted of six months of demand history on 651 recoverable items for the F-101 aircraft at Oxnard AFB (January through June, 1960). For each item the cost and total demand during the six-month period were known; these data were the basis for the calculations herein, as they were in Ref. 3. In actuality, transaction data were at hand for the entire year; they were used

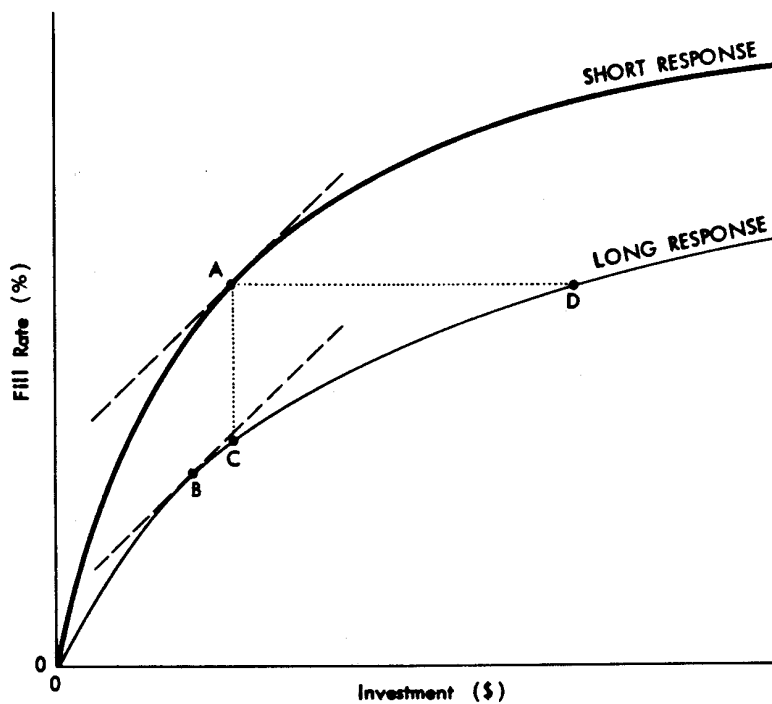


Fig. 1 -- Sensitivity Analysis of the Cost-Effectiveness Curve

in a maximum likelihood calculation to estimate the variance-to-mean ratio, α , of the stuttering Poisson at a value of two. This value was used as the standard or benchmark in the comparisons.

DESCRIPTION OF THE SENSITIVITY TESTS

Examinations will be made of the sensitivity of the cost-effectiveness curves to the choice of the performance measurement,

the average response time, the level of operations, the variance-to-mean ratio of the stuttering Poisson, and the discrete approximation of the integral in Eq. (7). Each type of sensitivity is considered separately and related to a standard case in which fill rate is the performance measurement, the average response time is seven days, the level of operations is one, the item variance-to-mean ratio $V/\theta = \alpha + \beta\theta$ has values $\alpha = 2$, $\beta = 0$, and the integral is approximated at ten points on the standard normal curve equally spaced between -2 and +3.

We did not examine sensitivity to errors in unit cost because this appears to be a data-processing problem rather than an estimation problem. Of course, a systematic error in which each cost is multiplied by a factor, a , still leads to an optimal allocation, but other types of cost errors will cause a degradation. Sensitivity to any assumptions about the distribution of errors in unit cost can be assessed easily.

The sensitivity of the cost-effectiveness curve to the length of the data period for the standard case above was considered in Ref. 2. It was introduced there to illustrate the advantages of the proposed Bayesian forecasting method over the issue-rate technique for demand prediction. We shall not display this comparison, but a modification of the standard case, assuming Poisson demand ($\alpha = 1$), will be used in the same way for comparison with the stock-leveling technique described in Air Force Manual 67-1.

DISPLAY OF SENSITIVITY RESULTS

Each performance measure is displayed as a rate. Fill rate is simply the expected number of fills during a period divided by the expected number of demands. Service rate is the expected number of units in routine resupply or repair divided by the expected number of units in resupply or repair.

Instead of displaying system investment cost, we show the average days of supply for the system. This is defined as the system investment divided by the average dollar value of daily demand for the system, the value being computed from the demands during the data

period--i.e., investment divided by daily usage. We believe that the average days of supply should be a more meaningful measure than investment, particularly in the comparison of different supply systems, because it is invariant to item cost, demand rate, and the number of items in the system. Note that this measure is quite different from the traditional Air Force procedure, in which a specified number of days of supply is carried on each item. In our case the number of days of supply will vary by item from zero days to some large positive number of days, depending on the item characteristics.

SENSITIVITY TO THE CHOICE OF PERFORMANCE MEASURE

We shall examine the three system-performance measures defined earlier -- ready rate, fill rate, and service rate -- plus a fourth measure called Operational Rate. This is simply the probability that no item in the system has a back order at a random point of time; it is the product of the item ready rates. This measure is particularly relevant in computations of the flyaway-kit type. Each of the three system-performance measures defined earlier was obtained as the sum of the item-performance measures, a requirement of the allocation algorithm. To convert this multiplicative measure, Operational Rate, to an additive measure we simply take the logarithms of the item Ready Rates. This procedure is appropriate because a function and the logarithm of a function achieve their maximum at the same point.

In Table 1 we see the effect of assuming one performance measure, calculating the corresponding set of optimal stock levels for some target expressed in days of supply, and then evaluating these stock levels against various true performance measures. For example, if we assume that ready rate is the appropriate measure, and our target is fifteen days of supply, we can expect a fill rate of 0.601. If we had assumed that fill rate was our measure and had optimized stock levels on that basis, we could have achieved the slightly higher fill rate of 0.608. By construction, the diagonal elements in each array cannot be exceeded by other elements in the same row.

The remarkable observation is that the minimum and maximum in each row are almost identical over the three levels of investment corresponding to 15, 60, and 180 days of supply. This suggests that any of the four performance criteria will lead to stock levels that

Table 1

SENSITIVITY TO CHOICE OF PERFORMANCE MEASURE

True Measure	Assumed Measure			
	Ready Rate	Fill Rate	Service Rate	Operational Rate
15 days Supply				
Ready rate	.981	.981	.981	.981
Fill rate	.601	.608	.603	.599
Service rate	.639	.636	.641	.640
Operational rate	.000	.000	.000	.000
60 Days Supply				
Ready rate	.995	.995	.995	.995
Fill rate	.882	.886	.885	.883
Service rate	.902	.900	.902	.902
Operational rate	.030	.027	.030	.030
180 Days Supply				
Ready rate	.999	.999	.999	.999
Fill rate	.988	.989	.988	.988
Service rate	.991	.991	.991	.991
Operational rate	.725	.718	.725	.725

NOTE: A day of supply has a value of \$13,816.00.

are approximately optimal with respect to the other three criteria when the item response times are all about the same. We did not examine the sensitivity when response times are different by item because we did not have appropriate data. However, if the item times vary substantially, the service rate criterion will stock more heavily than the fill rate criterion on items with long response times. Fill rate was chosen as the standard measure for the subsequent sensitivity tests here, because this is most familiar to the Air Force, and it is easy to measure in an application.

SENSITIVITY TO RESPONSE TIME AND LEVEL OF OPERATIONS

Reference 3 discussed a type of nonstationarity in which the true mean demand for each item was multiplied by a factor, a , called the change in the level of operations. The past demand data were assumed to be appropriate, but because of a change in the number of aircraft or the flying program, the mean system demand in the future is estimated to be equal to the current mean system demand multiplied by a factor a . Fill rate is calculated by dividing the expected number of fills in a period, given by Eq. (4), by the expected number of demands, λ . The result is that fill rate is a function of the product λT , but not of λ or T individually. Since λ is proportional to the true mean demand ($\lambda = (1 - \rho)\theta$), a change in the level of operations by the factor a has the same effect as multiplying the response time, T , by the factor a . Therefore, the same sensitivity analysis applies to changes in the level of operations or response time. (see Table 2).

Table 2

SENSITIVITY OF FILL RATE TO CHANGES IN RESPONSE TIME OR LEVEL OF OPERATIONS (λT) BY A FACTOR a

True Value of a	Assumed Value of a		
	0.5	1.0	2.0
15 Days Supply			
0.5	.642	.637	.610
1.0	.603	.608	.593
2.0	.532	.548	.554
60 Days Supply			
0.5	.903	.901	.893
1.0	.884	.886	.882
2.0	.837	.849	.856
180 Days Supply			
0.5	.991	.991	.990
1.0	.988	.989	.988
2.0	.977	.982	.983

NOTE: A day of supply has a value of \$13,816.00.

Table 2 considers three values of the factor, a , equal to 0.5, 1.0, and 2.0. Note that the maximum degradation in performance occurs in the first row of the table, where a 0.642 fill rate is reduced to a 0.610 fill rate by a fourfold error in the assumed value of a . Of course, if a is assumed to be 0.5 and the true value is 2.0, we will not obtain the 0.642 fill rate we might desire, but only a 0.532 fill rate with a 15-day supply of stock. In this example the fourfold increase in the factor, a , from 0.5 to 2.0 requires about a 50-percent increase in supply investment to maintain the same fill rate. However, the important observation is that if the investment level is unchanged, the 0.532 fill rate is only slightly less than the optimal fill rate (0.554) when the factor a is estimated correctly.

These sensitivity tests indicate that the factor a has a substantial impact on performance, but for a specified level of investment even a fourfold misestimate of the factor a leads to an allocation of stock that is nearly optimal.

SENSITIVITY OF FILL RATE TO THE ITEM VARIANCE-TO-MEAN RATIO

In Table 3 the objective is to assess the sensitivity of fill rate to different assumptions on the variance-to-mean ratio. Assuming this ratio is a linear function of the item mean $V/\theta = \alpha + \beta\theta$, four cases are investigated. The case $\alpha = 1, \beta = 0$ is simple Poisson; the case $\alpha = 2, \beta = 0$ is the benchmark case, in which α was estimated by maximum likelihood techniques; the case $\alpha = 3, \beta = 0$ corresponds to extreme variability, which is probably an upper bound; and the case $\alpha = 1.5, \beta = 0.5$ means that the item variability increases with the mean demand. Since variance is computed as $V = 1/T \sum (x - \bar{x})^2$, it has the dimensions of (demands)²/time whereas the mean $\theta = 1/T \sum x$ has the dimensions of demands/time. Thus the dimension of α is demands and the dimension of β is time. In our example, demand is expressed in units and a six-month period is used for the time reference. For an assumed $\alpha = 1.5$ and $\beta = 0.5$ on an item with a six-month true mean of 20 demands, the variance-to-mean ratio is $V/\theta = 1.5 + 0.5\theta = 11.5$; with a six-month true mean of 0.1, the variance-to-mean ratio is 1.55. On the other

hand, if we had used a three-month period, for example, the value of β would have to be doubled to yield the same set of variance-to-mean ratios, since β has the dimension of time; α is unaffected.

Table 3

SENSITIVITY OF FILL RATE TO THE ITEM
VARIANCE-TO-MEAN RATIO
 $V/\theta = \alpha + \beta\theta$ for Various α, β

True Values	Assumed Values			
	1,0	2,0	3,0	1.5, 0.5
15 Days Supply				
1,0	.742	.720	.679	.606
2,0	.581	.608	.599	.547
3,0	.474	.523	.532	.499
1.5,0.5	.393	.442	.458	.480
60 Days Supply				
1,0	.974	.955	.933	.901
2,0	.845	.886	.879	.857
3,0	.726	.809	.819	.806
1.5,0.5	.631	.716	.731	.741
180 Days Supply				
1,0	.999	.999	.997	.995
2,0	.976	.989	.987	.981
3,0	.920	.962	.966	.959
1.5,0.5	.839	.895	.906	.918

NOTE: A day of supply has a value of \$13,816.00

We observe that an optimization based on an incorrect assumption of the variance-to-mean ratio usually leads to a substantial degradation in performance. The common assumption of Poisson demand ($\alpha = 1$) is particularly susceptible to error even when the true variance-to-mean ratio has the modest value of two. For example, a 0.886 fill rate is reduced to 0.845 with 60 days of supply. Lest this difference in fill rate seem unimportant, the reader should recall that a few percentage points in the high-fill-rate range are equivalent to a

large investment. If Poisson demand is incorrectly assumed, 75 days of supply or a 25-percent increase in investment would be required to achieve a fill rate of 0.886 when the true variance-to-mean ratio is two. The situation is even more bleak if that ratio is three. A Poisson assumption in such a case requires a 50-percent increase in investment for the same performance. It is important to stress that we are comparing entries in the same row which correspond to the same true state of the system, and we find that an incorrect assumption on variance-to-mean ratio causes a bad allocation.

How can we attenuate this sensitivity to the variance-to-mean ratio? It is tempting to select a moderate value of the ratio, such as two rather than the Poisson extreme of one. But, in this case, if the true ratio is one, a 0.974 fill rate is decreased to 0.955. Maintaining a 0.974 fill rate entails a 25-percent increase in investment as in the complementary case described above.

Thus our conclusion must be that the correct assessment of the variance-to-mean ratio is essential for an appropriate allocation of funds. Furthermore, to the extent possible, the base operating procedures should be designed to minimize the item variance-to-mean ratio observed at base supply. As Table 3 indicates, a 0.742 fill rate can be achieved with 15 days of supply if demand is Poisson, but a 60-day supply is necessary to obtain the same fill rate with high variability ($\alpha = 1.5$, $\beta = 0.5$).

SENSITIVITY TO THE DISCRETE APPROXIMATION

Finally, to check the approximation of the integral in Eq. (7), the fill rates in the standard case were computed from a numerical integration at twenty equally spaced points on the standard normal curve between -3 and +4. The fill rates agree with those computed from ten equally spaced points on the standard normal curve between -2 and +3 within 0.001 at each investment level. Consequently, we conclude that the standard case with ten points yields an excellent approximation.

VI. COMPARISON OF THE STOCKAGE POLICY
WITH AIR FORCE PROCEDURES

AUTHORIZED AND ACTUAL STOCK LEVELS AT OXNARD AFB

This Section compares the base stockage policy with both the policy at Oxnard AFB in 1960 and the current Air Force policy.

The first comparison is interesting because it employed actual data to evaluate stock levels determined by supply personnel. The procedure described in Sec. V was applied to the first six months of data to compute a set of item stock levels for the standard case in which fill rate is the performance measurement, the average response time is seven days, and the item variance-to-mean ratio is two. This was done for several investment targets. The data also contained the authorized and actual stock levels for each item. These were evaluated by using the second six months of actual transaction data. These data specified the number of demands occurring on each day for each of the 651 line items, enabling the computation of the system fill rate for each of the two sets of levels. A seven-day response time was assumed.

The results of the simulation are displayed in Fig. 2. Compared with either the authorized or the actual stock levels, the stockage policy described here can theoretically^{*} attain the same performance with less than one-half the investment. This is because unit price is a factor in our calculations, leading to higher stock levels on the low-cost items. The results for the base stockage policy are particularly impressive because a six-month data period is not a long one when the variance-to-mean ratio of demand is two.

Table 4 gives a sample of the corresponding item detail. As mentioned above, the demand in Period 2 is used only to evaluate the stock levels. However, Table 4 displays the total demand for each item in Period 2 to indicate the high variability between the periods.

*In an implementation the performance actually attained will be slightly below the predicted performance for two reasons: 1) The model does not represent the true process exactly. 2) Since the true parameter values are unknown there are estimation errors.




Fig. 2 -- Oxnard AFB Cost-Effectiveness Curve (180-day Data Period, Stuttering Poisson Demand, Variance to Mean Ratio = 2)

Table 4

SAMPLE OF DETAIL FOR FIGURE 2

Stock Number	Cost	Demand in Period		Item Stock Levels						
		1	2	Author-ized	Actual	Computed by Stockage Policy				
4310693+2148	540.00	4	2	0	3	2	2	3	4	4
4440329+7966	137.00	3	30	39	20	3	3	4	5	6
4920504+1645	1606.00	1	0	6	6	0	0	1	2	2
4920775+6915	125.00	0	1	2	1	1	2	3	3	4
4920775+6925	400.00	0	2	3	2	0	1	2	2	3
4920776+3523	50.00	0	1	2	0	2	3	4	4	5
4920795+1364	656.93	0	1	0	0	0	0	1	2	2
5821503+3290	124.00	1	2	1	1	2	3	4	4	5
5821505+0465	254.00	2	2	1	1	2	3	3	4	5
5821505+0945	3804.00	28	49	2	14	2	4	5	6	6
5821505+1166	2152.00	4	3	3	2	0	1	2	2	3
5821505+1691	1052.00	5	1	1	0	1	2	3	3	4
5821505+3038	618.00	1	2	2	3	1	1	2	3	3
5821505+3244	225.00	0	4	2	1	1	2	2	3	3
5821505+3313	300.00	1	0	0	0	1	2	3	3	4
5821578+5722	298.00	1	6	2	2	1	2	3	3	4
5826347+7178	106.00	1	1	1	2	2	3	4	5	5
5826505+1535	500.00	1	11	0	3	1	2	2	3	3
5826505+1536	400.00	1	0	1	0	1	2	3	3	4
5826505+1676	90.00	1	2	1	0	3	3	4	5	5
5826505+2219	700.00	0	1	0	0	0	0	1	2	2
5826505+3092	441.00	1	1	0	0	1	2	2	3	4
5826505+3141	441.00	2	0	0	0	1	2	3	3	4
5826505+3510	280.00	1	1	0	1	1	2	3	4	4
5826519+5089	30.00	0	1	2	1	3	3	4	5	5
5826553+4709	2866.00	1	0	0	0	0	0	1	1	2
Days of supply				100.5	86.7	15.1	30.2	60.0	90.1	120.1
Range				.786	.697	.657	.762	.880	.948	.972
Fill rate				.733	.683	.491	.599	.780	.870	.911

for the stockage policy computation, the targets used were 15, 30, 60, 90, and 120 days of supply. Since units of stock are discrete, the targets were exceeded by a small margin as shown in Table 4. The range of items stocked is the fraction of items on which positive stock is maintained. Although the table shows a sample of only 26 items, the range and fill rate pertain to the entire system of 651 line items.

When interpreting Table 4, the reader should keep in mind that it is often possible to fill all demands during a period with a stock level that is less than the number of demands. This is because the items are recoverable. A demand is filled, and after a time lag for base repair or depot resupply, the faulty item is returned to a serviceable condition.

CURRENT AIR FORCE MANUAL 67-1 PROCEDURES

Section II described the current Air Force stock-leveling procedure for recoverable items. This procedure was used in the present study to calculate stock levels for a set of simulated items. As they did in Ref. 3, these 651 simulated items have the item costs of the Oxnard data and true mean values drawn at random from a log normal distribution whose cross-sectional moments agree with the Oxnard data. In this case, however, it was assumed that item demand has a Poisson distribution to favor the AFM 67-1 policy. This does not mean that the 67-1 policy assumes Poisson demand; but if the true demand pattern is compound Poisson, the 67-1 policy will be affected adversely for two reasons:

(1) The estimation of demand by the issue-rate technique will result in larger errors.

(2) The system fill rates will be unacceptably low, even with larger values of K.

It is reasonable to ask why simulated data were employed, rather than actual transaction data as in the previous comparison. The answer is twofold. In the first place, the Oxnard transaction data was quite nonstationary, and therefore may be atypical. But the major consideration is that with simulated data we can make meaningful comparisons of the behavior of the 67-1 policy as the base data period is extended.




Fig. 3 -- Fill Rate vs. Investment in Inventory (90-day Data Base)




Fig. 4 -- Fill Rate vs. Investment in Inventory (180-day Data Base)




Fig. 5 -- Fill Rate vs. Investment in Inventory (360-day Data Base)




Fig. 6 -- Fill Rate vs. Investment in Inventory (720-day Data Base)

These comparisons are displayed in Figs. 3 to 6. The four figures correspond to four different base data periods: 90, 180, 360, and 720 days. Each figure depicts three cost-effectiveness curves. The solid curve is calculated by the stockage policy described in this Memorandum. This is the standard case as defined above, except that demand is assumed to be Poisson. The dashed curve is calculated in the same way, except that the mean demand for each item is estimated by an issue-rate technique (observed demand divided by the length of the period) instead of the Bayesian technique. As the data period is extended to 720 days, these curves naturally converge. The dotted curve corresponds to the 67-1 policy described in Sec. II. This curve was obtained by varying the value of K and considering an average response time of 8 days for all items.

Several interesting observations can be made. (1) The base stockage policy described in this Memorandum produces the same system fill rate as the AFM 67-1 policy, with an investment that is never more than half as large, and often less than one quarter. (2) Even in this stationary system, the investment and performance under the 67-1 policy fluctuate wildly as the data period is extended. When the data period was 180 days it produced a 0.637 fill rate at an investment of \$407,000 whereas with 720 days it produced a 0.812 fill rate at an investment of \$604,000. (3) Under the stockage policy described here, it is apparent that the Bayesian cost-effectiveness curve is affected only slightly when the data period exceeds 180 days. This is, of course, related to the Poisson assumption. In Ref. 3, the authors noted that when an item variance-to-mean ratio of two was assumed, extending the data period beyond 360 days had little effect on the cost-effectiveness curve. (4) Comparing Bayes and non-Bayes solutions we find that the Bayesian procedure is significantly better until at least 360 days of data have accumulated. With higher variance to mean ratios this time period is even longer. Thus in the non-stationary world of application we must conclude that the Bayesian technique is a fundamental part of the optimal base stockage policy.

VII. CONCLUSION

The base stockage policy described in this Memorandum appears to offer significant improvement over the AFM 67-1 procedures, both past and present. Under a system approach, Bayesian procedures can be used for demand prediction and efficient system cost-effectiveness alternatives presented to management. The system approach enables us to examine the sensitivity of the base stockage policy. The policy appears sensitive to the item variance-to-mean ratio, but insensitive to other parameter changes.

The base stockage policy shows what level of performance is realistically attainable for a given budget, and furnishes a benchmark against which actual performance can be measured. Second, it provides a mechanism for keeping a system within budget constraints. Finally, it enables a manager to estimate the impact of alternative support postures for a weapon system -- for example, the impact of reducing the base's repair capability. Since the mathematical solution to the recoverable-item base stockage problem is analytic, the computer program for these calculations is extremely efficient.

APPENDIX

The objective of this algorithm is to determine a set of item stock levels, s_1, s_2, \dots, s_N such that a specified performance level Γ_0 is obtained with minimum investment or such that maximum performance is obtained for a specified investment level I_0 . This is expressed by Eq. (13). (See p. 11) The system performance and cost are the sum of the performance and cost over the individual items. Since the stock level for an item does not affect the performance or cost for any other item, we have a separable cell integer programming problem.

For any value of λ the optimal stock level for each item, s_1 , is determined by taking the first difference in Eq. (13). When \bar{G} is concave, the first difference of \bar{G} is monotone decreasing and we obtain

$$(A.1) \quad \frac{\Delta \bar{G}_i(s_i)}{c_i} \geq \lambda > \frac{\Delta \bar{G}_i(s_i + 1)}{c_i} .$$

If the appropriate value of λ is known, these item stockage decisions are straightforward. Usually, however, the value of λ corresponding to a target for system performance or system investment is unknown, and this is why the algorithm is needed.

It is important that \bar{G} be concave so that each value of λ determines a unique s_i . From Eq. (5) we find that for any probability distribution

$$\Delta S(s+1) - \Delta S(s) = - p(s+1).$$

Since the p 's are non-negative we have established that units in service is a concave function. However, for $\lambda T > 1$, fills, ready rate and operational rate are not necessarily concave functions. For this reason we must devise a variation on the well known algorithm for concave functions which generates price attainable efficient allocations.

For this application we assume that there is some maximum stock

level s_i^* for each item beyond which we will not allocate. We require that $\bar{G}(s_i) \leq \bar{G}(s_i^*) < \infty$ for $s_i < s_i^*$. But, from Eq. (3) it is obvious that this property holds for ready rate, and therefore operational rate, under any probability distribution and any value s_i^* . It is easy to show⁽²⁾ that for any compound Poisson distribution with compounding distribution f_i

$$\Delta F(s) = \lambda \left\{ p(s) + (1-f_1)p(s-1) + (1-f_1-f_2)p(s-2) \right. \\ \left. \dots + (1-f_1-f_2 \dots - f_s)p(0) \right\}.$$

Since each term is non-negative, $F(s)$ is monotone increasing and it is obviously bounded by λT .

We shall first demonstrate the allocation algorithm in which any function G_i is artificially converted to a new concave function \bar{G}_i . Then a proof of optimality will be given.

ALLOCATION ALGORITHM

Step 1: Concave Extension of G_i

Each function G_i is converted to a new function \bar{G}_i which is concave (see Figure 7)

For each G_i with s_i^* specified

$$\text{set } s_0 = 0 \\ \bar{G}_i(0) = G_i(0)$$

A. find $\max_{s_0 < s \leq s^*} \frac{G_i(s) - G_i(s_0)}{s - s_0}$

Suppose maximum is at $s=s'$
then for $s=s_0+1, s'$

$$\text{set } \Delta \bar{G}_i(s-1) = \frac{G_i(s') - G_i(s_0)}{s' - s_0}$$

if $s' < s^*$ set $s_0=s'$ and go to A
otherwise go to next function G_i .




Fig. 7 -- Concave Extension of G_i

Step 2: Initialization.

$$\begin{array}{l}
 0 \rightarrow \Gamma \\
 0 \rightarrow I \\
 0 \rightarrow s_i \\
 0 \rightarrow \Delta \bar{G}_i(s_i^*) \\
 \left. \begin{array}{l}
 \frac{\Delta \bar{G}_i(0)}{c_i} \rightarrow \lambda_i \\
 \bar{G}_i(0) + \Gamma \rightarrow \Gamma
 \end{array} \right\} i = 1, 2, \dots, N
 \end{array}$$

Step 3: Search For Best Item.

$$\begin{array}{l}
 \text{Find } j \text{ such that } \lambda_j \geq \lambda_i \text{ for all } i \\
 \lambda_j \rightarrow \lambda_{\max}
 \end{array}$$

Step 4: Increment.

$$\begin{array}{l}
 I + c_j \rightarrow I \\
 \Gamma + \Delta \bar{G}_j(s_j) \rightarrow \Gamma \\
 s_j + 1 \rightarrow s_j \\
 \frac{\Delta \bar{G}_j(s_j)}{c_j} \rightarrow \lambda_j \\
 \text{if } \lambda_j = \lambda_{\max}, \text{ repeat Step 4.}
 \end{array}$$

Step 5: Test.

$$\begin{array}{l}
 \text{If } I \geq I_0 \text{ or } \Gamma \geq \Gamma_0 \text{ terminate;} \\
 \text{otherwise, return to Step 3.}
 \end{array}$$

In Step 4 the condition $\lambda_j = \lambda_{\max}$ means that for the value of s_j just determined, $\bar{G}(s)$ was artificially increased by the concave extension procedure in Step 1. Since this value of $\bar{G}(s)$ is not attainable, we want to increment the stock level until we reach an attainable value of $\bar{G}(s)$. This is signified by the condition $\lambda_j < \lambda_{\max}$ in Step 4.

PROOF OF OPTIMALITY

It is obvious that a concave function can be constructed in Step 1 with at least two admissible points at $G_i(0)$ and $G_i(s_i^*)$. Now we can extend the concave argument of Barlow and Proschan⁽⁶⁾ as follows:

Let $\{\bar{s}\}$ be any allocation generated by the algorithm and let $\{s\}$ be any other allocation such that

$$\sum G_i(s_i) > \sum G_i(\bar{s}_i).$$

We shall show that this implies $\sum c_i s_i > \sum c_i \bar{s}_i$, or that more investment is required.

By construction the algorithm restricts any solution $\{\bar{s}\}$ to have the property

$$G_i(\bar{s}_i) = \bar{G}_i(\bar{s}_i) \text{ for all } i.$$

Also by construction, we know that for any allocation $\{s\}$ not necessarily generated by the algorithm

$$G_i(s_i) \leq \bar{G}_i(s_i) \text{ for all } i.$$

Designate the set of indices for which $s_i > \bar{s}_i$ by I_1 , and the set of indices for which $s_i < \bar{s}_i$ by I_2 . Then

$$\begin{aligned} 0 &< \sum G_i(s_i) - \sum G_i(\bar{s}_i) \leq \sum \bar{G}_i(s_i) - \sum \bar{G}_i(\bar{s}_i) \\ &= \sum_{i \in I_1} [\bar{G}_i(s_i) - \bar{G}_i(\bar{s}_i)] - \sum_{i \in I_2} [\bar{G}_i(\bar{s}_i) - \bar{G}_i(s_i)] \\ &= \sum_{i \in I_1} \sum_{h=1}^{s_i - \bar{s}_i} [\bar{G}_i(\bar{s}_i + h) - \bar{G}_i(\bar{s}_i + h - 1)] \end{aligned}$$

$$\begin{aligned}
 & - \sum_{i \in I_2} \sum_{h=0}^{\bar{s}_i - s_i - 1} \left[\bar{G}_i(\bar{s}_i - h) - \bar{G}_i(\bar{s}_i - h - 1) \right] \\
 & \leq \sum_{i \in I_1} (s_i - \bar{s}_i) \left[\bar{G}_i(\bar{s}_i + 1) - \bar{G}_i(\bar{s}_i) \right] \\
 & \quad - \sum_{i \in I_2} (\bar{s}_i - s_i) \left[\bar{G}_i(\bar{s}_i) - \bar{G}_i(\bar{s}_i - 1) \right]
 \end{aligned}$$

by concavity of each \bar{G}_i . But from Eq. (A.1) the last expression does not exceed

$$\sum_{i \in I_1} (s_i - \bar{s}_i) \lambda c_i - \sum_{i \in I_2} (\bar{s}_i - s_i) \lambda c_i.$$

Since $\lambda > 0$ this implies

$$0 < \sum c_i s_i - \sum c_i \bar{s}_i.$$

COMMENT

The concave extension algorithm generates a subset of the price attainable points which would be generated if the G_i were concave. The system target I_0 or Γ_0 is usually exceeded, but with several hundred items and different item costs, the amount of overshoot is typically small as in Figure 2. In the strict sense we have not solved the combinatorial problem of maximizing performance without exceeding I_0 (minimizing investment that just achieves Γ_0), but the distinction is irrelevant in this application.

The concave extension procedure is valid for assumptions other than finite maxima s_i^* . For example, if there exists an s_0 such that

$G(s_0) \geq G(s)$ for all $s < s_0$ and $G(s)$ is concave for all $s > s_0$, the algorithm can be applied.

REFERENCES

1. Lu, J. Y., and Gabriele A. Michels, Description of the Computer Program for Aggregate Base Stockage Policy of Recoverable Items, The RAND Corporation, RM-4395-PR, April 1965.
2. Feeney, G. J., and C. C. Sherbrooke, The (s-l,s) Inventory Policy Under Compound Poisson Demand: A Theory of Recoverable Item Stockage, The RAND Corporation, RM-4176-PR, September 1964.
3. Feeney, G. J., and C. C. Sherbrooke, An Objective Bayes Approach for Inventory Decisions, The RAND Corporation, RM-4362-PR, March 1965.
4. Air Force Manual 67-1, Vol. II, Chap. 11.
5. Hastings, C., Approximations for Digital Computers, Princeton University Press, Princeton, New Jersey, 1955, p. 185.
6. Barlow, R. E. and F. Proschan, Mathematical Theory of Reliability, John Wiley & Sons, New York, 1965, pp. 167-168.

