


MEMORANDUM
RM-5078-PR
NOVEMBER 1966

AD 645684

METRIC: A MULTI-ECHELON TECHNIQUE
FOR RECOVERABLE ITEM CONTROL

Craig C. Sherbrooke

PREPARED FOR:
UNITED STATES AIR FORCE PROJECT RAND



The **RAND** Corporation
SANTA MONICA • CALIFORNIA

ARCHIVE COPY

MEMORANDUM

RM-5078-PR

NOVEMBER 1966

METRIC: A MULTI-ECHELON TECHNIQUE
FOR RECOVERABLE ITEM CONTROL

Craig C. Sherbrooke

This research is supported by the United States Air Force under Project RAND—Contract No. AF 49(638)-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

The analytic model described in this Memorandum is the culmination of several years of research instigated by a request from Logistics Plans, Headquarters USAF. References 3, 4, 6 and 9 are some of the earlier publications from this project.

The Memorandum is intended for personnel engaged in implementation of stockage policies, and for management scientists.

SUMMARY

METRIC is a mathematical model of a base-depot supply system in which item demand is compound Poisson with a mean value estimated by a Bayesian procedure. When a unit fails at base level there is a probability r that it can be repaired at the base according to an arbitrary probability distribution of repair time, and a probability $1-r$ that it must be returned to the depot for repair according to some other arbitrary distribution. In the latter case the base levies a resupply request on depot. No lateral resupply between bases is considered in the model. For high-cost, low-demand items the appropriate policy is $(s-1,s)$, which means that items are not batched for repair or resupply requests. This problem has a simple analytic solution which is a function of the mean repair times rather than the repair time distributions.

A practical and efficient computer program has been designed to show the cost-effectiveness tradeoff for a large group of recoverable items. In addition to minimizing expected backorders any system investment, the program can evaluate any distribution of stock and it can compute the optimal redistribution of stock. No arbitrary estimates of backorder cost or holding cost are required.

The study includes numerical examples, and discusses the application of the model in Air Force contexts.

CONTENTS

PREFACE	iii
SUMMARY	v
Section	
I. INTRODUCTION	1
II. GENERAL DESCRIPTION	2
Purposes	2
Features	2
Data Requirements	3
By System	3
By Item	3
By Item and Base	4
III. STRUCTURE OF THE MULTI-ECHELON PROBLEM	6
Mathematical Assumptions	6
System Objective of Minimizing the Expected Number of Backorders	6
Compound Poisson Demand	8
Demand is Stationary over the Prediction Period	9
Decision on Where Repair Is to Be Accomplished Depends on the Complexity of the Repair Only	10
Lateral Resupply is Ignored	10
System is Conservative	10
The Depot Does Not Batch Units of a Recoverable Item Unless There is an Ample Supply of Serviceable Assets ...	11
Recoverable Items May Have Different Essentialities	11
Demand Data from Different Bases Can Be Pooled	12
Multi-Echelon Theory	12
Generalization of the Objective Function	16
Limited Capability to Employ the Generalized Objective Function in METRIC	20
Demand Prediction	22
Characteristics of Initial Estimates	24
Estimation of the True Mean Demand Distribution	24
Interim Procedure for Estimating True Mean Demand	26
Estimation of Parameters	28
Combination of Demand Data and Initial Estimates	29
Debiasing the Initial Estimates	31
Estimation of Compound Poisson Parameters	31
Point Estimates of Demand	32

Preceding Page Blank

IV. APPLICATION	33
Cost-Effectiveness Decisions	33
Minimum Stock Levels	33
Maximum Stock Levels	34
Choice of Cost-Effectiveness Target	34
The Case of Demand Data and No Initial Estimates	35
Common Items	36
Estimation of the Fraction of Units Base Repairable	36
Estimation of the Average Base Repair Time	36
V. DETAILED DESCRIPTION OF METRIC COMPUTER PROGRAM, WITH AN EXAMPLE	38
Organization of the Computer Program	38
Example	39
Uncertainty Ratio of the Initial Estimate	40
Variance to Mean Ratio of the Logarithmic Poisson Process ...	45
Item Detail for the Standard Case	47
Summary Statistics for the Example	49
VI. GENERALIZATIONS OF THE RECOVERABLE ITEM MODEL	50
Bayesian Estimation for the Fraction of Units Base Reparable	50
Inclusion of Condemnations	51
APPENDIX	53
REFERENCES	55

I. INTRODUCTION

This Memorandum employs the same mathematical approach of the earlier work on a base stockage model (4) and applies it to the more complex base-depot supply system.

Consideration of this larger problem resolves many of the practical difficulties of applying any base stockage model. For example, items in long supply or short supply at the bases and depot are handled routinely as a redistribution problem. Depot stock levels are computed that are consistent with the base policy. In other words, the multi-echelon point of view is a more appropriate and simpler perspective for supply management.

Recoverable items typically have high cost and low demand. They constitute an important management problem in the Air Force, accounting for about 78 percent of the total investment in spares, or approximately 5 billion dollars. Yet from a mathematical perspective they are particularly simple because a one-for-one replacement policy is optimal at base level.

Compared to current Air Force policy our technique has the advantage that unit cost is considered in the calculation. But even more important in our view is the system approach, which displays a range of optimal cost-effectiveness alternatives to management. Instead of computing stock levels on the basis of artificial estimates of holding cost rate and backorder cost, this approach focuses management attention on the entire weapon system so that an appropriate combination of system effectiveness and system cost can be selected.

The base stockage policy has been tested successfully at Hamilton Air Force Base and George Air Force Base. To our knowledge the model described in this Memorandum is the first multi-echelon, multi-item model ever proposed for implementation.

II. GENERAL DESCRIPTION

METRIC is a mathematical model translated into a computer program, capable of determining base and depot stock levels for a group of recoverable items; its governing purpose is to optimize system performance for specified levels of system investment. METRIC is designed for application at the weapon-system level, where a particular line item may be demanded at several bases and the bases are supported by one central depot. The support depot may vary by item as in the item-manager system or it may be fixed as in the weapon-system storage site concept.

PURPOSES

1. Optimization. A major purpose of the model is to determine optimal base and depot stock levels for each item, subject to a constraint on system investment or system performance. Optimization is of prime interest in the early weapon acquisition phase.

2. Redistribution. The model can take fixed stock levels on each item and optimally allocate the stock between the bases and depot. Redistribution is a major concern when items are in long or short supply.

3. Evaluation. The model provides an assessment of the performance and investment cost for the system of any allocation of stock between the bases and depot. Evaluation is important throughout the life of a weapon system.

FEATURES

The technique embodies a number of advantageous features, outlined below.

1. Generalization of the RAND Base Stockage Policy. Because METRIC has the same mathematical foundation as the Base Stockage Policy (4), implementation experience gained with that policy is directly relevant to METRIC. In fact METRIC can also be operated as a single-echelon base stockage model, in which case it supersedes the earlier computer program.

2. Incorporation of anticipated program requirements. The model uses past data, but combines them with estimates of future program requirements to anticipate buildups or phase-outs.

3. Use of initial-estimate data with or without demand data. METRIC enables a smooth transition from initial support planning to follow-on provisioning, by allowing for the incorporation of a general demand-prediction procedure. The procedure uses initial estimates in combination with demand data when available.

4. Simplification in analyzing alternative support postures. The impact of different maintenance policies or pipeline times on the supply effectiveness for the weapon system can be readily evaluated.

5. Management capability to provide different levels of support effectiveness depending on the weapon system. Management can examine the effect of varying degrees of support depending on the mission importance of different weapon systems.

DATA REQUIREMENTS

The list of METRIC input data below includes symbolic notation for parameters that are referred to in subsequent mathematical derivations. The subscripts indicate Item i and Base j , but will be omitted in contexts where the meaning is clear.

By System

q -- Variance to Mean Ratio of Demand. This will be discussed in Sec. III under the heading, "Demand Prediction."

Optimization Targets. These are expressed as system investment in dollars or expected backorders per item. Alternatively, a ratio of backorder cost to holding cost rate can be used if a reduction in computer time is desirable.

By Item

Stock Number or Identification.

D_i -- Average Depot Repair Time. This is the average of the time required for a reparable carcass to be shipped from base to depot

and repaired at depot. (To be precise we should also subtract the average of the time between placement of an order at base and receipt of the order at depot.)

c_i -- Unit Cost.

M_i -- Initial Estimate of Mean Demand.

w_i -- Uncertainty of the Initial Estimate. This will be discussed in Sec. III under the heading, "Demand Prediction."

n_i -- Number of Squadron-Months of Demand Data (or Aircraft-Months).

u_i -- Demand Observed Over the Time Period n_i .

Total Stock for Redistribution.

By Item and Base

r_{ij} -- Average Fraction of Units that are Base Repairable.

A_{ij} -- Average Base Repair Time.

O_{ij} -- Average Order and Shipping Time. This is the average time between placement and receipt of an order at the base when the depot has serviceable units on hand.

E_{ij} -- Essentiality. This is the relative cost of a backorder on Item i at Base j compared to a backorder on some standard item.

a_{ij} -- Program Element. This reflects the anticipated change in the level of operations for each item and base; e.g., if the flying hours are expected to double over the next six months, say, this factor would probably be estimated as two.

Stock for Evaluation. This is a set of stock levels, one level for depot and one for each base.

Minimum Stocks for Optimization and Redistribution. The program will accept minimum levels for each base, the depot, and the system before optimization or redistribution is performed. In the standard case where the minimum levels are zero, input of minimum stock levels is not necessary. The section entitled "Application" describes how the minimum stock feature can be utilized.

Maximum Stocks for Optimization and Redistribution. In the standard case where the maximum levels are unconstrained, input of maximum levels is not necessary.

The computer program has been designed to minimize data-input requirements and to facilitate sensitivity testing. On the item data cards there are codes from 0-9 for the average base repair fraction, the average base repair time, the average order and shipping time, the essentiality, and the program elements. A particular code identifies a set of numbers, one for each base. Thus each parameter can be made to vary by base and item. Similarly, there are codes from 0-9 for average depot repair time and uncertainty. Each code identifies a number so that each parameter can vary by item. Sensitivity testing is simplified because only the meaning of the codes need be changed, not the item data cards themselves.

III. STRUCTURE OF THE MULTI-ECHELON PROBLEM

MATHEMATICAL ASSUMPTIONS

We shall state the general assumptions of the model, realizing that none is exactly true, but believing that all are reasonable approximations.

System Objective of Minimizing the Expected Number of Backorders

The objective will be to minimize the sum of backorders on all recoverable items at all bases pertinent to a specific weapon system. Thus, unless all bases are identical, the expected number of backorders will vary by base.

An alternative objective would be to specify an expected number of backorders for each base. The METRIC computer program does not provide this alternative objective as an option, but the mathematical procedure is described under "Multi-Echelon Theory," below, and a limited capability to employ this alternative is included in METRIC. We should emphasize that depot backorders are not explicitly considered under either objective. Depot backorders are of interest only insofar as they affect base backorders.

Let us define the backorder objective. Take a fixed period of time and add together the number of days on which any unit of any item at any base is backordered. Dividing this number by the length of the period and taking the expected value of the statistic yields a number that is independent of the period length. This is the value we seek to minimize. Under this criterion, a base backorder lasting ten days is as serious as ten backorders lasting a day apiece.

It is important to stress that our definition of backorders differs from Air Force parlance. Under our definition, a backorder exists at a point in time if and only if there is an unsatisfied demand at base level, e.g., a recoverable item is missing on an aircraft. Note that this condition can arise even when the base has a positive authorized spare stock for the item, because at a point of time all spares may be in the base repair process or in the depot resupply process.

In an earlier RAND Memorandum (4) which dealt with base stockage only, we considered other objective functions, viz., fill rate, service rate, ready rate, and operational rate. Under conditions of identical average repair times for each item, we displayed an example in which the resulting stockage policy was nearly independent of the choice of objective function.

This is not true, however, in a base-depot supply system. For example, fill rate -- defined as the fraction of demands that are immediately filled by supply when the requisitions are received -- concentrates nearly all stock at the bases. The result is that when a nonfill occurs, the backorder lasts a very long time.

Similarly, fill rate behaves improperly in allocating investment at a base when the item repair times are substantially different. Consider two items with identical characteristics except that one is base-reparable in a short time, and the other is depot-reparable with a much longer repair time. Assume that our investment constraint allows us to purchase only one unit of stock. In that case, the fill rate criterion will select the first item, and the backorder criterion the second.

Fill rate possesses an additional defect. A fill is normally defined as the satisfaction of a demand when placed. But if we allow a time interval τ to elapse, such as a couple of days, on the grounds that some delay is acceptable, the policy begins to look substantially different. As longer delays are explored, the policy begins to resemble the minimization of expected backorders.

Ready rate -- defined as the fraction of items which are not in a backorder condition -- seems to be inappropriate because it does not measure the number of units backordered on an item. Operational rate -- defined as the probability that k_j or less aircraft have a supply shortage at Base j at a random point of time -- is not very flexible. For example, it is difficult to give essentiality an economic interpretation here. Operational rate also requires the analyst to supply a set of k_j values.

In summary, the backorder criterion seems to be the most reasonable. The penalty should depend on the length of the backorder and the number of backorders; linearity is the simplest assumption. This is the criterion function most often employed in inventory models.

Finally, experimentation seems to indicate that a policy which minimizes the expected number of backorders provides good results with respect to other criteria, but the converse is not necessarily true. Expected number of backorders does have one drawback when compared to fill rate in that it is harder to measure objectively in an application. An estimate of backorders can be obtained by counting the number of backorders at a random point in time, but it may be difficult to prevent nonrandom phenomena from contaminating the estimate.

In our earlier single-echelon work we used service rate defined as $\left\{ 1 - \frac{\sum_i (\text{expected backorders on item } i \text{ with stock } s_i)}{\sum_i (\text{expected backorders on item } i \text{ with stock } 0)} \right\}$. The analogous multi-echelon measure would be obtained by substituting in the numerator for the base stock level s_i , a vector of stock levels for each base and the depot; in the denominator for the base stock level 0, a vector of 0 stock levels for each base and the depot. Since the denominator is independent of the stockage policy and the numerator is the expected backorders, the allocation obtained from service rate is identical with that resulting from expected backorders. We prefer backorders to service rate (or backorder rate) because the service rates of interest are so near one. With zero stock levels at each base and depot, the denominator is extremely large and not particularly relevant.

Compound Poisson Demand

We assume that demand for each item is described by a logarithmic Poisson process, a member of the compound Poisson family. Compound Poisson processes are generalizations of Poisson processes; the compound processes allow the flexibility of incorporating more parameters, yet retain the simple analytic properties of the Poisson.

The logarithmic Poisson is obtained by considering batches of demand where the number of batches follows a Poisson process and the number of demands per batch has a logarithmic distribution. Our earlier work employed the geometric Poisson, still another member of the compound Poisson family. The switch is occasioned by the fact

that the state probabilities -- the probability of n demands in a time-interval of specified length t -- for the logarithmic Poisson process are negative binomial, and these are particularly convenient to compute. Furthermore, for variance to mean ratios less than three (our range of interest in this application), the state probabilities for the two distributions are almost identical. A thorough discussion of discrete compound Poisson processes is given in Sherbrooke (9). In particular, it is shown that if customers follow a Poisson process with mean λ per time period, and each customer can place a number of demands that are independently and identically distributed as a logarithmic distribution so that the compound Poisson demand process has variance to mean ratio q , then the probability of x demands in the time period has the negative binomial distribution

$$(1) \quad p(x|\lambda) = \frac{(k+x-1)!}{(k-1)!x!} \frac{(q-1)^x}{q^{k+x}} \quad \begin{array}{l} x=0,1,2,\dots, \\ q>1, k>0 \end{array}$$

where $\lambda = k(\ln q)$. Defining the mean as θ , we find $\theta=k(q-1)$, where q is the variance to mean ratio. Although the negative binomial distribution is a function of two parameters, we shall only indicate the conditioning of $p(x|\cdot)$ on the customer rate λ , because q maintains a constant value on any particular item.

Demand is Stationary over the Prediction Period

It is assumed that the distribution of demand over some future period of interest, such as six months, is stationary. Of course, program element data can be used to reflect a change in the level of operations from the data period.

Although this assumption may seem very restrictive, we should point out that the compound Poisson demand model in conjunction with a Bayesian probability distribution for true mean demand enables us to represent fairly complex demand patterns over a prediction period of arbitrary length.

Decision on Where Repair Is to Be Accomplished Depends on the Complexity of the Repair Only

The assumption is that the decision to repair a unit at base level or send it to the depot is a function only of the type of malfunction and the base maintenance capability. Whenever possible, repair is accomplished at base level irrespective of maintenance workload. In a study by Weifenbach (10), shipments to depot because of shop backload accounted for only .3 percent of 10,965 reparable generations.

Lateral Resupply is Ignored

When a unit is shipped from base to depot for repair, a serviceable replacement will be resupplied from the depot if possible. If the depot has no unit on the shelf, the base must wait until a unit emerges from depot repair.

For the purposes of determining base and depot stock levels, the model ignores the possibility of lateral resupply between bases in this case. This appears to be appropriate for setting levels because the number of lateral shipments is typically small and they are apt to induce special costs of expediting. When lateral resupply is ignored, transportation costs between bases and between bases and depot are not needed because the total transportation cost is not a function of the stockage policy.

System is Conservative

Consider a particular stock item demanded at Base j . We assume that a unit of stock has a probability r_j of being reparable at Base j , with a repair time drawn at random from a base repair distribution with mean A_j ; a probability $(1-r_j)$ of being depot reparable, with an order and shipping-time distribution having mean O_j and a depot repair distribution having mean D . This implies that there are no condemnations or that the system is conservative, as the name "recoverable item" suggests. Of course, our data do show some condemnations, but these are typically less than five percent of the reparable generations.* A higher condemnation rate usually indicates that the item should be redesigned. The condemnation rate must be

* Condemnations were found to account for 4.1 percent of 10,965 reparable generations in a study by Weifenbach (10).

considered for procurement purposes, but the procurement process is not considered in the METRIC optimization. If the procurement process were included explicitly, the cost of placing a procurement at depot level would be needed as well as the condemnation rate to determine the procurement frequency and quantity.

The Depot Does Not Batch Units of a Recoverable Item for Repair Unless There is an Ample Supply of Serviceable Assets

The model assumes that depot repair begins when the reparable base turn-in arrives at the depot. This appears to be a reasonable approximation to current depot scheduling practice. Furthermore, an earlier RAND study (8) concluded that such a depot repair policy leads to minimum cost.

In the absence of priority information, the stockage model suggests a natural repair scheduling policy. Since METRIC economizes by buying fewer high-cost items, these are the items which are most likely to be in short supply. Therefore, the higher-cost items should be scheduled into repair first.

In those few cases where setup cost is an important factor and demand is reasonably high, so that some batching is indicated, the estimate of depot repair time should include the average waiting time before depot repair is initiated.

Recoverable Items May Have Different Essentialities

METRIC will accept relative backorder costs or essentialities by base and item. Suppose we define the standard essentiality on Item i at Base l as 1 . Then if Item i at Base j is estimated to have an essentiality relative to the standard of k , the correct allocation is determined when backorders at Base j for Item i are multiplied by k . This is shown at the conclusion of the section entitled "Multi-Echelon Theory," where essentiality variations by base are discussed in some detail. Essentialities may vary by item, but as a first approximation over the class of recoverable items it seems reasonable to assume that all items have equal essentiality.

Demand Data from Different Bases Can Be Pooled

We assume that demand from the several bases can be pooled in some manner so that a composite initial estimate of demand per flying hour (or any other program element) can be obtained. The pooled estimate can be obtained by a simple averaging technique or a more sophisticated procedure such as exponential smoothing. METRIC multiplies this number by the flying hours per month for each item. Of course, the number of flying hours per month will vary from base to base for which purpose the program elements, a_{ij} , are provided. Note that if on a particular Item i , a_{i1} is twice as large as a_{i2} , say, then the estimate of mean demand is twice as large at Base 1. If we believe that a doubled flying program will produce less than a total of twice as much demand, then a_{i1} should be less than twice a_{i2} .

The program element is allowed to vary by item as well as by base to cover cases where some items are applicable only to a portion of the aircraft fleet.

MULTI-ECHELON THEORY

We shall limit our discussion to the two-echelon, base-depot problem, since the extension to more echelons is an obvious consequence. Demand is assumed to be represented by a compound Poisson process. For purposes of visualization, the compound Poisson can be thought of as a series of customers following a Poisson process, each of whom can demand an amount that is independently and identically distributed according to a compounding distribution. Until further notice we consider one item stocked at J bases, with known mean customer arrival rates, $\lambda_j, j=1,2,\dots,J$. The next section considers the case in which the λ_j are unknown.

When a customer arrives at a base to place one or several demands, he turns in a like number of reparable units. It is assumed that with probability r_j these units can all be repaired at base level and with probability $1-r_j$ they must all be shipped to the depot for repair.*

*This assumption leads to the analytic simplification that the steady state probabilities are compound Poisson. For compound Poisson

Under these assumptions the customers from Base j who arrive at the depot are described by a Poisson process whose mean is $1-r_j$ times the mean of the Poisson customer arrival process at Base j . Therefore, the total demand at the depot is compound Poisson, with mean customer arrival rate $\Sigma \lambda_j (1-r_j)$. Letting \bar{f}_j be the mean demand per customer at Base j , we find that the mean depot demand rate is $\Sigma \lambda_j \bar{f}_j (1-r_j) = \Sigma \theta_j (1-r_j)$ where θ_j is defined as the mean demand at Base j .

Although the depot demand process is compound Poisson, the depot compounding distribution, the distribution of demands placed by a customer, is in general a complicated synthesis of the base compounding distributions. There is no need to specialize our argument at this time, but we note that METRIC employs a logarithmic Poisson process. By assuming that all $\bar{f}_j = \bar{f}$, which implies that demand at each base has the same variance to mean ratio, though different means, we obtain a logarithmic Poisson process at depot with that variance to mean ratio. This is an important computational advantage. The preceding assertions are proved in Sherbrooke (9).

The essence of the multi-echelon solution is the following fact derived from Eq. (12) of Feeney and Sherbrooke (3). Let s be the spare stock* for an item where demands are compound Poisson with mean customer arrival rate λ , and the resupply (repair) time is an arbitrary distribution $\Psi(t)$ with mean T . Assume that when a customer arrives, a resupply time is drawn from $\Psi(t)$ that is applicable to all demands placed by that customer. In the case where excess demand is backlogged, the expected number of backorders at a random point in time is

variance to mean ratios which are not too large we can justify this mathematical assumption on the grounds that multiple demands by a single customer occur infrequently. Thus the resulting state probabilities are nearly unaffected. For example, in a logarithmic Poisson process with a variance to mean ratio of two, the probabilities that a customer picked at random will place exactly one, two, or more than two demands are .72, .18, and .10, respectively.

* Spare stock is defined as the sum of stock on hand plus on order plus in repair minus backorders. It retains a constant value under the one-for-one replacement policy.

$$(2) \quad B(s) = \sum_{x=s+1}^{\infty} (x-s)p(x|\lambda T),$$

where $p(x|\lambda T)$ is the compound Poisson probability density for a mean customer rate λT . In the special case where the compound Poisson process is logarithmic Poisson, $p(x|\lambda T)$ is given by Eq. (1).

We note that the expected number of backorders is a convex function since

$$(3) \quad B(s+1) - B(s) = - \sum_{x=s+1}^{\infty} p(x|\lambda T).$$

The computation of the multi-echelon solution consists of five stages. (A mathematical justification will be given afterwards.)

1) We must compute, for each base, the average time which elapses between a base request for a resupply from depot and base receipt of the unit. This average response time is obviously a function of the depot spare stock. With infinite depot spare stock, this time is the average order and shipping time O_j ; with zero depot spare stock, this time is $O_j + D$, where D is the average depot repair time. Therefore, the delay at depot due to the fact that there is not always a serviceable unit on the depot shelf when a resupply request is received must be between zero and D . To compute the delay as a function of depot spare stock, s_0 , we recall that if the expected number of customers who arrive at depot during a fixed time period is λ , then $\lambda = \sum_j \lambda_j (1 - r_j)$. If there are s_0 or less units in depot repair, no resupply is being delayed. But if there are more than s_0 units in depot repair, the resupply on $x - s_0$ units is being delayed. From Eq. (2), the expected number of units on which delay is being incurred at a random point in time is

$$(4) \quad B(s_0 | \lambda D) = \sum_{x=s_0+1}^{\infty} (x-s_0)p(x|\lambda D).$$

Eq. (4) is the expected number of units backordered at the depot, but we choose not to use that terminology to avoid confusion with our objective function, which is the sum of backorders only at the bases.

The total expected system delay over any time period is simply the expected number of units on which delay is being incurred at a random point in time multiplied by the length of the time period. Since we are interested in the average delay per demand, we must then divide by the expected number of demands over that time period. For example, if λ and D have the dimensions of customers/month and months respectively, the expected number of demands per month is $\lambda \bar{f}$. Thus, the average delay per demand expressed in months is

$$(5) \quad \frac{\sum(x-s_0)p(x|\lambda D)}{\lambda \bar{f}} = \delta(s_0)D \quad ,$$

where

$$\delta(s_0) = \frac{\sum(x-s_0)p(x|\lambda D)}{\lambda D \bar{f}} = \frac{B(s_0|\lambda D)}{B(0|\lambda D)} \quad .$$

2) For each level of depot stock, s_0 , and each base we must compute the expected backorders as a function of the base stock s_j . This is accomplished by Eq. (2), with the specification $s=s_j$, $\lambda=\lambda_j$, and $T=r_j A_j + (1-r_j)(0_j + \delta(s_0)D)$.

3) For each level of depot stock, s_0 , we must determine the optimal allocation of the first, second, ... units of stock to the several bases so as to minimize the sum of expected backorders at all bases. This is accomplished by a simple marginal allocation. At each step the next unit of stock is added to that base where it will cause the largest decrease in expected backorders. From Eq. (3) backorders is a convex function when mean demand is known. The Bayesian procedure described in the next section yields a linear combination of such convex functions with positive coefficients which, of course, is also convex, thus guaranteeing that the marginal allocation technique produces optimal solutions.

4) The two-dimensional table showing expected backorders as a function of depot stock, s_0 , and total base stock s under optimal allocation, must be collapsed to one dimension. For each level of constant total system stock, $s_0 + s$, represented by a diagonal in the table, we select the minimum expected system backorders. For item decisions in Stage 5 we also record the actual allocation of stock between bases and depot corresponding to each system optimum.

5) Now we consider the multi-item problem. Marginal analysis is again employed. Using the item backorder functions computed in Stage 4, the next investment is allocated to that item which produces the maximum decrease in expected backorders divided by unit cost. This is similar to the procedure we used in Stage 3, except that there we were dealing with a single item so that unit cost was not a variable.

However, before marginal analysis is used here a preliminary convex extension of the functions must be performed because the item backorder functions are not necessarily convex.* This procedure leads to an optimal solution as shown in the Appendix. After each allocation, the system investment and system backorders are computed. Allocation terminates when the investment target is just exceeded or, alternatively, when the expected backorders are just less than a target value.

Generalization of the Objective Function

We have described a computational procedure. Now we give a mathematical justification of that procedure while generalizing the objective function to allow constraints on the expected number of backorders at each base.

The problem is to minimize system cost subject to the constraints that expected backorders at Base j not exceed b_j . Let $\beta_{ij}(s_{i0}, s_{ij})$ be the expected backorders for Item i at Base j when the depot stock is s_{i0} and the stock at Base j is s_{ij} . Then the problem can be stated:

*The possible nonconvexity may surprise the reader. After all, we know that for a specified level of depot stock the item backorder functions are convex. However, after Stage 4 is completed, nonconvexity as a function of system stock may appear whenever the depot stock levels change.

$$\min_{\{s_{ij}\}} \sum_{i=1}^I c_i \sum_{j=0}^J s_{ij} \quad ,$$

subject to

$$\sum_{i=1}^I \beta_{ij}(s_{i0}, s_{ij}) \leq b_j \quad j=1, 2, \dots, J.$$

According to Everett (2) this problem can be reformulated with Lagrange multipliers, ν_j , as:

$$(6) \quad F = \min_{\{s_{ij}\}} \left\{ \sum_{i=1}^I c_i \sum_{j=0}^J s_{ij} - \sum_{j=1}^J \nu_j \left[\sum_{i=1}^I \beta_{ij}(s_{i0}, s_{ij}) - b_j \right] \right\} .$$

The simpler objective function described earlier is obtained from Eq. (6) by setting all ν_j equal.

Since this is a separable cell integer programming problem with $F = \sum F_i$, we restrict attention to a single item. Consider a fixed set of ν_j , and assume that the optimal policy on Item i is to allocate m units according to the vector $(s_{i0}^m, s_{i1}^m, \dots, s_{iJ}^m) = S_i^m$. A necessary condition on S_i^m is then

$$(7) \quad F_i(S_i^{m+1}) - F_i(S_i^m) = c_i - \sum_{j=1}^J \nu_j \left[\beta_{ij}(s_{i0}^{m+1}, s_{ij}^{m+1}) - \beta_{ij}(s_{i0}^m, s_{ij}^m) \right] \geq 0 ,$$

where $S_i^{m+1} = (s_{i0}^{m+1}, s_{i1}^{m+1}, \dots, s_{iJ}^{m+1})$ is chosen to minimize

$$- \sum_{j=1}^J \nu_j \beta_{ij}(s_{i0}^{m+1}, s_{ij}^{m+1}) \quad ,$$

subject to the constraint

$$\sum_{j=0}^J s_{ij}^{m+1} = 1 + \sum_{j=0}^J s_{ij}^m = m+1 .$$

In other words for S_i^m to be optimal, any allocation of $m+1$ units must increase F_i . We have denoted the best such allocation by S_i^{m+1} and this then forms the binding constraint. With the choice of sign in Eq. (6) each z_j is negative. Note that z_j must have dimensions of cost. Thus, the absolute value of z_j is the imputed ratio of (back-order cost/holding cost rate) at Base j where numerator and denominator are defined for a time period of the same length.

Similarly, the allocation of one less unit of stock results in another necessary condition on S_i^m

$$(8) \quad F_i(S_i^m) - F_i(S_i^{m-1}) = c_i - \sum_{j=1}^J z_j \left[B_{ij}(s_{io}^m, s_{ij}^m) - B_{ij}(s_{io}^{m-1}, s_{ij}^{m-1}) \right] < 0 ,$$

where $S_i^{m-1} = (s_{io}^{m-1}, s_{i1}^{m-1}, \dots, s_{iJ}^{m-1})$ is chosen to minimize

$$- \sum_{j=1}^J z_j B_{ij}(s_{io}^{m-1}, s_{ij}^{m-1}) ,$$

subject to the constraint

$$\sum_{j=0}^J s_{ij}^{m-1} = m-1 .$$

In the case where

$$\sum_{j=0}^J s_{ij}^{m-1} < 0 ,$$

we define

$$B_{ij}(s_{io}^{m-1}, s_{ij}^{m-1}) = \infty \text{ for all } j.$$

A sufficient condition for optimality is that the following function be a convex function of m :

$$\Xi_i(m) = - \sum_{j=1}^J z_j B_{ij}(s_{io}^m, s_{ij}^m) .$$

Convexity assures us that any fixed set of ν_j (each ν_j is non-positive) defines a unique value of m for item i . Unfortunately, this function is not necessarily convex. Let us define a new function $\Xi'_i(m)$ which lies on the boundary of the convex hull of $\Xi_i(m)$. In other words at any point where $\Xi_i(m)$ is not convex the function is decreased to the point of convexity in forming the new function $\Xi'_i(m)$. At all other points $\Xi_i(m)$ coincides with $\Xi'_i(m)$. The optimality of the resulting allocation is shown in the Appendix. We note that $\Xi_i(m)$ is a monotonic decreasing function. This is obvious because the backorders resulting from the best allocation of m units plus an arbitrary allocation of an $(m+1)$ st unit cannot exceed the backorders from the best allocation of m units.

In order for m to be a solution, $\Xi'_i(m)$ must equal $\Xi_i(m)$ and therefore Eqs. (7) and (8) can be rewritten as

$$F_i(S_i^{m+1}) - F_i(S_i^m) = c_i + \Xi'_i(m+1) - \Xi_i(m) \geq 0$$

(9)

$$F_i(S_i^m) - F_i(S_i^{m-1}) = c_i + \Xi_i(m) - \Xi'_i(m-1) < 0 .$$

As a computational matter we note that these constraints can be incorporated in the five-stage procedure described earlier if each base backorder function is multiplied by the corresponding known ν_j before Stage 3 is undertaken. Stage 5 is replaced by the comparison indicated by Eq. (9).

The Lagrange Multiplier Method generates the set of price attainable efficient solutions. Let us relate this to our inventory problem in Eq. (6). For any set of base backorder constraints, b_j , there exist Lagrange Multipliers, ν_j , such that investment cost is minimized subject to the constraints that base backorders are $b'_j \leq b_j$. With a large number of items each b'_j can be made almost, but not quite, equal to b_j . The solution generated by the Lagrangian procedure has the property that no b'_j can be decreased without increasing investment. On the other hand it may be possible by combinatorial methods to increase at least one b'_j slightly without violating the constraint

so that the resulting investment is decreased. This solution cannot be obtained by Lagrangian methods.

With a large number of items the Lagrangian solution will differ only slightly from the combinatorial solution. Furthermore, the constraints in our problem are not rigid. They are only approximations. Our real interest is to display a set of efficient cost-effectiveness alternatives for which the Lagrangian method is sufficient. A thorough treatment of generalized Lagrange methods may be found in Everett (2).

The only difficulty with this generalized objective is that the set of z_j corresponding to the specified base backorder targets b_j are unknown, and therefore an iterative procedure is required to find the correct set of z_j . For this purpose a linear programming formulation has been proposed by Brooks and Geoffrion (1). The solution to the resulting $J+1$ by $J+1$ linear program is obtained by the Simplex Method with an additional "column generating feature."

Although we have been describing optimization, the calculations above (before convexification) indicate how any fixed set of stock levels can be evaluated. Since evaluation is not a function of the targets, no set of z_j 's is needed.

Redistribution of a fixed total quantity between the bases and depot is accomplished with the output from Stage 4 after convexification. For the generalized objective function, an appropriate set of z_j 's is again needed in Stage 3.

Limited Capability to Employ the Generalized Objective Function in METRIC

To employ the generalized objective function, we would like to be able to specify targets at each base for the expected number of backorders on all units of all items and have METRIC determine the set of z_j and the corresponding stock levels. Obviously if the appropriate set of z_j is known, the problem is trivial. Now we cannot estimate the cost of a backorder at each base, but suppose that we can estimate

the cost of a backorder at each base relative to the cost of a backorder at a particular base. If the holding cost rate is assumed to be constant across bases, then each z_j in the solution set is simply a constant multiplied by the relative backorder cost at Base j . From the computational point of view we have reduced the problem of finding J numbers to the problem of finding one. In our earlier terminology these relative backorders costs are essentialities, E_{ij} . As noted earlier we can have essentialities by item as well as by base, though let us restrict our attention here to essentialities by base.

For example, in a three-base optimization the user could specify .5, 1, 2 as the base essentialities on all items and then optimize for some investment target. METRIC will then allocate that investment and find the actual Lagrange multipliers, subject to the constraint that the backorder cost per unit at Base 1 is half that at Base 2, which in turn is half that at Base 3. The resulting set of z_j will be a constant multiplied by .5, 1, 2 where each z_j is independent of the multiplicative scale factor in the original choice of essentialities. The user must specify an investment target, not a backorder target, whenever one or more $E_{ij} \neq 1$. This is because the backorder functions, beginning in Stage 2 of the five-stage procedure, are multiplied by the essentialities to produce the correct stock allocations. After optimization has been completed, METRIC can evaluate the resulting stock levels to compute the expected backorders.

We are not providing the capability to specify base backorder targets, though by a judicious experimentation with essentialities, the same effect can be obtained by the conscientious experimenter. On the other hand, it is not unreasonable to suppose that relative backorder costs (essentialities) can be estimated for different bases though the actual values are unknown.

If essentialities are quite different from base to base on a given item, the METRIC assumptions may become unrealistic. For example, METRIC assumes that demands on depot are resupplied on a first come, first served basis. However, if essentialities vary widely from base to base, a priority system is obviously in order. In most planning applications essentiality should probably not vary by base.

DEMAND PREDICTION

We are primarily interested in computing the probability distribution of demand over a specified future time period for an item with an initial engineering estimate of mean demand per flying hour and perhaps some demand data. Our assumption is that true mean demand for the period, θ , is stationary but unknown. Demand is assumed to follow a compound Poisson process where the parameters of the compound Poisson are known functions of the unknown variable θ . For example, the METRIC computer program considers a logarithmic Poisson process. The probability distribution of demand over any fixed time period is then negative binomial, specified by two parameters, the variance to mean ratio q and mean θ . We assume that $q = \alpha + \beta\theta$ for all items where α and β are estimated by maximum likelihood techniques described later.

Our first objective is to show that a Bayesian procedure is of fundamental importance for all items, not merely low-demand items. Now, current Air Force procedures produce an initial engineering estimate of true mean demand for a new item. For the purpose of discussion we shall assume that this estimate is the expected value or mean of the possible values for true mean demand as viewed by the estimator. Actually there is a good deal of ambiguity as to what is being estimated, but we shall defer this question until later in the Section.

The naive approach would be to compute the expected backorders by using the initial estimate. For example, let us suppose that demand is Poisson and the initial estimate of demand over a unit time interval is 1. Then the expected backorders under a stock level of two is, from Eq. (2):

$$B(2) = \sum_{x=2}^{\infty} (x-2) \frac{\lambda^x e^{-\lambda}}{x!} = .1036.$$

On the other hand, we might believe there is a probability .5 that the item has a low demand of .5, and a probability .5 that the item has a relatively high demand of 1.5. Note that the mean initial

estimate is still 1. When we compute backorders under this assumption we find:

$$B^*(2) = .5 \sum_{x=2}^{\infty} (x-2) \frac{(.5)^x e^{-.5}}{x!} + .5 \sum_{x=2}^{\infty} (x-2) \frac{(1.5)^x e^{-1.5}}{x!} = .1485 .$$

For the case of Poisson demand and positive spare stock, it is easy to show that backorders computed from a point estimate of demand always understate the correct value when demand can assume more than one value. This is shown by differentiating $B(s)$ twice with respect to λ , yielding:

$$(10) \quad \frac{\partial^2(B(s))}{\partial \lambda^2} = \frac{\partial^2}{\partial \lambda^2} \left(\sum_{x=s+1}^{\infty} (x-s) \frac{\lambda^x e^{-\lambda}}{x!} \right) = \frac{e^{-\lambda} \lambda^{s-1}}{(s-1)!} > 0 \quad \text{for } s \geq 1.$$

Thus for any positive λ and spare stock of one or more, the number of backorders is a strictly convex function of λ and the assertion is proved. For a spare stock of zero the second derivative is zero, which means that backorders computed from a point estimate are identical to backorders computed from a range of estimates.

The fact just stated is also true for probability distributions like the negative binomial, though difficult to prove analytically. The discrepancy in the number of backorders computed by the two methods may be very large. For any point estimate λ , the extreme case is obtained by selecting s sufficiently large so that expected backorders are approximately zero. Under the alternate method if there are two possible values of mean demand, 0 and λ' , with respective probabilities $(\lambda' - \lambda)/\lambda'$ and λ/λ' , the overall mean is still λ ; but as λ' increases, the expected backorders with the same spare stock s approach λ .

In summary, unless we are certain of the value of true mean demand -- an unlikely occurrence -- we will understate backorders by using a point estimate. The discrepancy may be large and since it varies by item, the resulting allocation of investment will produce inferior results. From another point of view, an underestimate of demand is more serious than a comparable overestimate when expected

backorders is the objective function. This also implies that a point estimate of demand provides insufficient information. The obvious alternative is a Bayesian procedure.

Characteristics of Initial Estimates

As noted above there is some ambiguity as to what an initial estimate is trying to estimate. Is it the most likely value of true mean demand or mode, the median value, the mean, or something else? To answer this question we consider in Table 1 some data reproduced from McGlothlin (7). Now, the distribution of true mean demand is probably skewed to the right, in which case the mean exceeds both the mode and the median. Since the McGlothlin data show that initial estimates are generally much higher than the mean, we can immediately discard the mode and median hypotheses. Table 1 displays twelve groupings of recoverable items from seven weapon systems. Demands were underestimated for only one group; it comprises items on the F-106, whose considerable similarity to the earlier F-102 may account for the relative closeness of the group of estimates to the observed demand -- a factor of $.78=182.4/234.5$. Excluding the Atlas missile system -- for which initial estimates were seven times too high -- and airframe parts which are primarily insurance-type items, we find that for the four groups of nonairframe parts the initial estimates taken as an aggregate were still twice the observed demand. This tends to confirm the belief that the analyst is not really estimating the mean demand, but rather some larger number that has a small probability of being exceeded. Since underestimates are apt to be more serious than overestimates, particularly in the early stages of weapon acquisition, such a bias is altogether reasonable when a point estimate of demand is required.

Estimation of the True Mean Demand Distribution

The problem with Bayesian estimation is that it requires a prior probability distribution for the values of true mean demand. In principle, there are at least four ways to obtain this distribution:

Table 1

RELATION OF INITIAL ESTIMATES TO OBSERVED DEMAND

Weapon System	No. of Items	Avg No. Flying Hours (000)	Demand for All Parts per 1000 Hrs		Corr. of Estimated & Observed Demand
			Estimated	Observed	
F-105:					
Airframe	37	35	73.6	9.3	.09
Non-airframe (A)	25	37	79.5	57.9	.09
Non-airframe (B)	32	32	129.1	44.8	-.09
F-106	51	36	182.4	234.5	.01
Atlas:					
Airframe	56	641 ^a	4.6 ^a	2.0 ^a	.16
Guidance	142	114 ^b	60.0 ^b	7.9 ^b	.13
B-66	33	52	72.6	57.6	.67
C-133:					
Airframe	31	48	22.0	3.0	.35
Non-airframe	29	50	48.8	21.6	.69
B-52:					
Airframe	28	69	30.4	15.8	.21
Non-airframe	33	69	50.4	30.7	.14
F-101/102	23	116	63.1	47.7	.58

^a Atlas airframe operational experience and demand rate are expressed in missile months.

^b Atlas guidance operational experience and demand rate are expressed in equipment months.

1) The estimator can sketch it freehand; 2) he can select a best distribution from a prescribed set; 3) he can estimate a couple of percentage points of a specified probability function; or 4) he can estimate the probability distribution of demand, rather than mean demand, over a fixed time interval. Note that Method 4 is quite different from the other three methods. In the first three methods we were concerned with continuous distributions for true mean value, whereas in Method 4 we are concerned with a discrete distribution of demand realizations. In Method 4 we assume that the compound Poisson parameters are known as a function of true mean demand, and that the distribution of true mean demand is a certain function. Then the parameters of the prior distribution can be estimated.

The author believes that one of the above methods should be adopted by the Air Force where several estimators or groups of estimators make independent assessments for each item. The author also believes that in the interim a procedure is required that incorporates the initial estimate data currently available. Let us now consider such an interim procedure.

Interim Procedure for Estimating True Mean Demand

We shall assume that the initial estimate is the mean of the true mean demand distribution and the distribution itself is gamma. The gamma distribution is a two-parameter, continuous, unimodal distribution defined for nonnegative values and skewed to the right. Figure 1 shows a few gamma distributions to indicate the wide variety of possible shapes. The gamma distribution is given by

$$(11) \quad g(\theta) = \frac{1}{w^v \Gamma(v)} \theta^{v-1} e^{-\theta/w} \quad \begin{array}{l} 0 \leq \theta < \infty \\ v, w > 0 \end{array}$$

where the mean is vw and the variance is w^2 . We have assumed that true mean demand over a unit time period has a gamma distribution. Now we demonstrate that this implies true mean demand over any time period has a gamma distribution.* This property is desirable computationally because the period of time over which a distribution of true mean demand is needed will vary from item to item depending on the time period over which item data has been observed. We shall return to this point later. Furthermore, without this property our prior distribution would have three parameters, rather than two, where the additional parameter would be the one (and only) time period over which true mean demand was gamma distributed.

The property is easily established. If we are interested in the distribution of true mean demand over a time period n times as long,

*This reproductive property is not peculiar to the gamma distribution. All continuous probability distributions behave similarly.

$$g(\theta) = \frac{1}{w \Gamma(v)} \theta^{v-1} e^{-\theta/w}$$

Mean = vw

Variance/mean = w




Fig. 1 -- Gamma prior distribution of true mean demand:
two examples with mean demand of three

then we require the distribution of $n\theta=X$. Since the Jacobian is $1/n$, the new probability distribution $G(X)$ is

$$G(X) = \frac{1}{w^v \Gamma(v)} \left(\frac{X}{n}\right)^{v-1} e^{-X/nw} \frac{1}{n} = \frac{1}{(nw)^v \Gamma(v)} X^{v-1} e^{-X/nw} .$$

Note that we did not compute the probability distribution for the sum of n random observations from a gamma distribution. Rather, we computed the probability distribution for the n -period mean demand on an item whose 1-period mean demand uncertainty is described by a gamma distribution. The true mean demand per unit time is constant though unknown.

Estimation of Parameters

Since the gamma distribution has two parameters, we need two estimation equations. We define the parameters v, w for some arbitrary time period, perhaps six months. Denoting the initial estimate of mean demand as \hat{M}_1 and recalling that the gamma distribution has mean vw , one equation is:

$$(12) \quad \hat{M}_1 = vw .$$

Unfortunately, we have no other data for the second estimation equation. We resort to a hypothesis concerning the variance to mean ratio of the true mean demand distribution, w . We shall refer to this value as the prior distribution uncertainty ratio, to distinguish it from the variance to mean ratio of the compound Poisson demand process.* As a first approximation we shall assume that the uncertainty

*The lognormal distribution is quite similar to the gamma distribution. Because of the simple transformation from a lognormal distribution to a normal distribution facilitating integration, the lognormal was used in our earlier work (3). However, we prefer the gamma distribution in this application because the mean value and uncertainty of the lognormal are complicated functions of the normal distribution parameters.

ratio is constant across all items. This seems to be a reasonable assumption, and empirically the resulting prior distributions look sensible.

Of course, we have not yet specified the constant. In the METRIC computer program the approach is to calculate the stockage policy for several values of the uncertainty ratio. This demonstrates the implications of the assumption on item prior distributions as well as the system stockage policy. This seems to be a meaningful way of capturing the subjective judgment of supply personnel. Moreover, from a statistical point of view we can then examine the robustness of our assumptions. A stockage policy computed under one assumption can be evaluated under another.

Our assumption of a constant prior distribution uncertainty ratio across items is only a first approximation. On the other hand it is far more reasonable than a maximum likelihood point estimate, which surely understates backorders. Our procedure has the advantage that it can adapt to the increasing sophistication of supply managers. METRIC will accept a different uncertainty ratio for each item if it can be provided. Or at a later date METRIC can be modified to accept subjective probability estimates and compute the gamma distribution parameters.

One technical comment should be appended. It was shown in the previous subsection that a gamma distribution with parameters v and w becomes a gamma distribution with parameters v and nw over n time periods (n not necessarily integral). Thus the new mean is vnw and the new variance is $v(nw)^2$ so that the variance to mean ratio or uncertainty is nw over this time period. Therefore, if we estimate the uncertainty ratio over a six-month period as three, then to be consistent we should be prepared to estimate the uncertainty ratio over a year as six.

Combination of Demand Data and Initial Estimates

Suppose we have observed u demands for an item over some period of time, L . We can construct the prior distribution of true mean demand over that period, $g(\theta)$, by the procedure described above. We combine the two types of information with Bayes theorem:

$$(13) \quad \phi(\theta|u) = \frac{p(u|[\theta/\bar{f}])g(\theta)}{\int_{\xi} p(u|[\xi/\bar{f}])g(\xi)d\xi},$$

remembering that $p(u|[\theta/\bar{f}])$ is a compound Poisson density for u demands when the mean customer rate is θ/\bar{f} or mean demand rate is θ . $\phi(\theta|u)$ is the posterior probability that true mean demand is θ , having observed u demands over the time period L .

The computer approximates the integral in Eq. (13) by a sum. Five to ten points, θ_k , appropriately chosen, are usually sufficient. Now $\phi(\theta_k|u)$ was computed for a specific time period, L , corresponding to the item data period, but for the computation of expected backorders in Eq. (2) we need probabilities defined over a period of length T , where T is the average response time. But over the new period the same ϕ 's are appropriate for $\theta_k T/L$. If in addition we expect a change in the level of operations (measured by some program element like flying hours) by a factor a from the data period to the prediction period, the same ϕ 's are appropriate for $a\theta_k T/L$. Since the mean number of customers is the mean demand divided by f , the expected backorders during a period T when u demands have been observed during a period L and with spare stock s is, from Eq. (2):

$$(14) \quad B^*(s|u) = \int_{\theta} \phi(\theta|u) \sum_{x=s+1}^{\infty} (x-s)p(x|[a\theta T/L\bar{f}])d\theta$$

$$\approx \sum_k \phi(\theta_k|u) \sum_{x=s+1}^{\infty} (x-s)p(x|[a\theta_k T/L\bar{f}])$$

The asterisk denotes the Bayesian function. $\phi(\theta_k|u)$ is given by the discrete version of Eq. (13), reducing to $g(\theta_k)$ when there are no demand data to combine with the initial estimates. In METRIC $p(x|\cdot)$ is given by Eq. (1).

For the special case where $p(x|\cdot)$ is Poisson, Eq. (14) can be integrated to yield

$$B^*(s|u) = \sum_{x=s}^{\infty} (x-s)h(x),$$

where $h(x)$ is a negative binomial distribution. This result is well known.

Debiasing the Initial Estimates

As noted above, the initial estimates taken as a group are usually too high. The amount of bias for the group of items can be estimated when demand data are also available. METRIC has the facility of computing this bias and multiplying each initial estimate by this bias factor. Alternatively, a bias factor can be input to METRIC.

Estimation of Compound Poisson Parameters

Throughout this section we have assumed that the compound Poisson is a function of only one unknown parameter. Specifically, METRIC assumes that demand is logarithmic Poisson, which has two parameters defined over a specific period, the mean θ , and variance to mean ratio q . Further, it is assumed that $q = \alpha + \beta\theta$ where α and β are known. It should be pointed out that β is inversely proportional to the length of the period, since the variance to mean ratio q is independent of time. In other words if β is estimated as .5 for a six-month time period, it should be estimated as .25 over a year.

In METRIC we assume that one pair of α, β applies to all items. At first glance this may seem unnecessarily restrictive, but the reader should recall that it is still a generalization of the Poisson assumption, which always seems to understate the variability we observe. Assuming a different q for each item is not a mathematical problem, but one of statistical estimation.

We have employed maximum likelihood procedures to estimate α and β from several sets of data. However, the bivariate estimation appears to be quite unstable, by which we mean that the likelihood function has nearly the same value for a wide range of α, β combinations. For this

reason we simplified our model by assuming $\beta=0$ and applied maximum likelihood procedures again.

Our data consisted of N equal-length periods of data for a large group of items. Since each item is assumed to have logarithmic Poisson demand, the probability distribution of demand during each period is negative binomial, where the parameters for an item are constant over the N periods. Letting u_{in} be observed demand on the i-th item during the n-th time period, the likelihood function is obtained from Eq. (1):

$$\prod_{i,n} \frac{(k_i + u_{in} - 1)!}{(k_i - 1)! u_{in}!} \frac{(\alpha - 1)^{u_{in}}}{\alpha^{k_i + u_{in}}},$$

where α is the variance to mean ratio we seek and $k_i(\alpha-1)$ is the period mean for the i-th item. The simplest method of evaluation is to take logarithms and then, for each assumed value of α , compute the maximum likelihood estimates of each k_i . Since the k_i 's do not interact, this computation is straightforward. The value of α that maximizes the likelihood function is selected.

Kendall (5) shows that for a single item the maximum likelihood estimate of k is a root of a transcendental equation. Thus a numerical solution is required in any case.

In the computations we have performed, the estimated variance to mean ratio has ranged from 1.5 to 2. Our earlier investigations in the single-echelon case (4) indicated that optimal stockage for a variance to mean ratio of two is substantially different from what it is under a Poisson assumption. Section V looks at this sensitivity for METRIC.

Point Estimates of Demand

In the event that a substantial amount of data is available by item, METRIC can utilize point estimates of demand where the uncertainty is then zero. Even when the uncertainty is not close to zero it may be desirable to compute the stockage policy under this assumption for purposes of comparison.

IV. APPLICATION

This section considers how METRIC should be applied.

COST-EFFECTIVENESS DECISIONS

When a weapon system is first phasing in to the Air Force inventory, optimization of stock levels is of interest. As time passes, however, our demand estimates and program elements change, so that any METRIC computer run should include evaluation and redistribution as well as optimization. We remind the reader that by evaluation we mean the computation of expected backorders for any fixed distribution of item stocks between base and depot; by redistribution we mean the allocation to bases and depot of a fixed total stock by item so that expected backorders are minimized.

Effective decision-making depends on a comparison of these three alternatives. For example, the decision to redistribute depends on the increased effectiveness over the evaluated levels and on the cost of redistribution compared with the effectiveness and cost of new procurement. These are the management alternatives. Of course, they may be constrained as in the case where no new procurement money is available, or in the case where increased effectiveness is needed immediately rather than a procurement lead-time from now.

The point to be emphasized here is that this display of management alternatives is more appropriate for decision-making than is a model that purports to produce optimal decisions continuously. There can be only academic value in an "optimal" policy predicated on the assumption that new procurement is possible, when in fact procurement is constrained.

MINIMUM STOCK LEVELS

METRIC allows minimum stock levels to be specified by item for each base, the depot, and the system in optimization or redistribution. Thus if a policy decision is made to stock at least one unit of a particular item at each base and depot, METRIC can optimize and redistribute

subject to that constraint. By setting the minimum stock levels equal to the current stock levels at each base and depot, METRIC can optimize new procurement subject to the constraint that no redistribution of existing assets takes place. Or METRIC can optimize without constraints on minimum base and depot stocks, but with a constraint that minimum system stock on each item equals the current stock. This yields the optimal procurement policy subject to the fact that certain stock is in the system and this stock may be redistributed.

Another application of the minimum stock level option is in redistribution, when a minimum level of depot stock may be desired to provide flexibility for contingencies. We expect that the usual option would be to set minimum stock levels of zero.

MAXIMUM STOCK LEVELS

METRIC allows maximum stock levels to be specified by item for each base, the depot and the system in optimization or redistribution. In addition METRIC recognizes a maximum base stock level, depot stock level, and system stock level which is constant across items.

A particular maximum value on a specific item is then computed as the minimum of two numbers. This is done to provide item flexibility while preserving a simple mechanism for adjusting all maxima simultaneously to the smallest reasonable value. This is desirable because the computer time is roughly proportional to the maximum stock levels.

CHOICE OF COST-EFFECTIVENESS TARGET

METRIC will accept system optimization targets expressed as either total dollars of investment or as expected backorders per item. The system backorders are the sum of backorders at all bases, whereas the system investment is the sum of investment at all bases and the depot. The reader should keep in mind that backorders at a base are simply the sum of all units backordered on all items at a random point in time. Backordered units include those in base repair as well as those being resupplied from depot.

Selection of an appropriate target value is facilitated by knowing investment and backorders under current policy. Since METRIC can reduce the backorders for the same investment or reduce the investment for the same performance, it seems reasonable to choose an intermediate target that should reduce both backorders and investment. The specific target should depend on the availability of funds, the nature of the program, and the degree to which current supply support is deemed satisfactory.

Some additional insight into an appropriate target can be gained by remembering that the absolute value of the Lagrange multiplier is the backorder cost divided by the holding cost rate. As noted earlier, a Lagrange multiplier can be specified as an alternative target.

THE CASE OF DEMAND DATA AND NO INITIAL ESTIMATES

We have concentrated on the case of initial estimates with or without some demand data. Now let us consider the case in which there are no initial estimates or the initial estimates are deemed irrelevant.

The RAND Memorandum on the Base Stockage Model (4) dealt only with the case in which demand data were available, devising a Bayesian prior distribution to be used on all items. The parameters of the prior distribution were estimated from the observed data on all items. The assumption was that before data are collected on an item, our state of knowledge is the same for all items. Now, of course, this is not strictly true because to a supply manager the item cost or even the item name conveys information. For example, we almost always observe a significant negative correlation between cost and demand, but this correlation was ignored in our original technique. Another drawback to that technique was that a data period of a fixed length was needed for all items.

For these reasons we prefer to treat early demand data as an initial estimate of mean demand. The prior distribution uncertainty ratio, i.e., the variance to mean ratio of the gamma prior distribution, should be smaller for items with longer demand history.

COMMON ITEMS

METRIC will most likely be applied by weapon system, since appropriate cost-effectiveness targets and program elements vary from one weapon system to another. This leads to a problem in the determination of stock levels for items common to more than one weapon system. In the event that these common items appear on other weapon systems located at the same set of bases and at no additional bases, we advise considering all demands for these common items in the computation for that weapon system with the lowest backorder/item target.

When many additional bases are involved, we advise estimating depot stock based on all demands using any reasonable methods. This then determines the average response time to each base. For these common items METRIC should be operated in a single-echelon mode, while for all other items it is operated in the multi-echelon mode. To operate METRIC in a single-echelon mode for any number of bases and a particular item, the fraction of units base reparable is set to one. The average base repair time is set equal to the true fraction base reparable multiplied by the average base repair time, plus the true fraction depot reparable multiplied by the average response time determined above.

ESTIMATION OF THE FRACTION OF UNITS BASE REPARABLE

We would like to point out that if the fraction of units base reparable is estimated as one, there will be no depot stock because it is assumed that there is no depot demand. In some cases this will be appropriate, but in the early uncertain stages of weapon acquisition the fraction of units base reparable should probably be estimated to be some value less than one.

ESTIMATION OF THE AVERAGE BASE REPAIR TIME

Suppose that we observe the repair times on a large group of base reparable items and then use the average as our estimate of the base repair time. The resulting number will probably be much too large and dependent on the stockage policy. For example with a large

inventory on each item there is no particular hurry about returning a reparable carcass to a serviceable condition.

We are really interested in a standard base repair time for an item (independent of stockage policy) which reflects a normal repair time. Our mathematical model assumes that the repair time for each demand on a particular item is independently and identically distributed. However, in practice, maintenance can schedule repair opportunistically to improve support performance. As a result, a seven-day resupply time which fluctuates only slightly from the average can be expected to provide lower support performance than a seven-day average base repair time. As an approximation, the base repair time should be estimated at some lower value.

V. DETAILED DESCRIPTION OF METRIC COMPUTER PROGRAM, WITH AN EXAMPLE

ORGANIZATION OF THE COMPUTER PROGRAM

This section describes the operation of METRIC, illustrated by computer output obtained for the F-111. For this weapon system we had data on unit cost and initial estimates of mean demand per flying hour for 652 recoverable items. Since we ourselves fabricated the rest of the METRIC input data on these items for this example, the specific results are not applicable to the provisioning of the F-111.

Before turning to the example, some general comments are in order on the organization of the computer program that will be particularly relevant to those familiar with the base stockage program (6). In METRIC each item is processed individually. Items are no longer aggregated into cost-demand categories because there are many more item variables in the multi-echelon case. Consequently, the computer program is substantially slower. On the other hand, there are major advantages to this redesign. The following is a comparison of the parameters for METRIC with those for the base stockage program on the IBM 7040-7044 with a 32,000-word memory.

Maximum Value of Parameter	METRIC	Base Stockage
<u>System</u>		
Number of line items	10,000	2000
Number of bases	20	1
<u>Item</u>		
Number of depot repair times	10	2
Number of uncertainty codes	10	1
Number of observed demands	unlimited	49
<u>Item and Base*</u>		
Number of base repair times	200	2
Number of base repair fractions	200	2
Number of order and shipping times	200	2
Number of essentialities	200	1
Number of program elements	200	1
<u>Item Stock</u>		
Each base	40	20
Depot	49	--
System	100	--

* In METRIC each of the five variables by item and base is specified by a code from 0 to 9 which references a set of 20 values, one for each base, producing a total of 200 possibilities.

In the base stockage program most of the computer time is used in allocating investment to the items. For example, allocating an average of k units to each of n items requires an amount of time proportional to kn^2 . Each allocation is preceded by searching a list of length n , and there are kn such searches. Of course, aggregation helps to decrease the size of n .

In a typical METRIC computation, allocation consumes only about one-third of the total computer time. Furthermore the allocation time is approximately proportional to n , not n^2 , because the Lagrange multipliers, corresponding to the solutions, are estimated from a sample of less than one hundred items. Using this estimate, the final adjustment to the exact optimum for the entire set of items is obtained by passing a scratch tape with information on all items a small number of times. The time to estimate the Lagrange multiplier is nearly constant, whereas the tape time is proportional to n . As noted under "Multi-Echelon Theory," allocation begins with zero stock on all items. An alternative would be to guess the value of the Lagrange multiplier and start the allocation at that point. However, for the typical example described in this section, the alternative procedure would reduce computer time by only two percent.

METRIC is composed of about 900 FORTRAN IV instructions. The computer time is approximately proportional to the product of (number of items) x (number of Bayes states) x (number of bases) x (maximum system stock). In the example below, for the standard case with two targets and with or without redistribution and evaluation, the computer required six seconds per item.

EXAMPLE

As noted above, we had data on unit cost and initial estimates of mean demand per flying hour for 652 recoverable items. We selected arbitrary values of observed demand over time periods ranging from six months to a year, by item, for a base flying an average of 1000 hours per month. For the prediction period, our system consists of a depot and three bases, the bases having average monthly flying programs of 500, 1000, and 2000 hours, respectively. For simplicity we assumed

that all base resupply times were .25 months, and that for each item there was a base repair time of .25 months, a depot repair time of 1.0 months, and a fraction of units base repairable of .5. The minimum stock levels were zero and maximum stock levels were 10 for each base, 20 for the depot, and 40 for the system. Redistribution and evaluation were not performed.

Many kinds of sensitivity tests would be interesting, and should be the subject of a detailed investigation. In this Memorandum, however, we shall consider only two types: 1) sensitivity to the uncertainty ratio of the initial estimate, i.e., the variance to mean ratio of the gamma distribution; and 2) sensitivity to the variance to mean ratio of the logarithmic Poisson demand process. Each type of sensitivity is examined by comparison with our standard case, in which the uncertainty ratio over a six-month period was taken as three and the logarithmic Poisson variance to mean ratio was two. The targets for the standard case were .10 and .02 backorders/item. With 652 items, these targets imply that at a random point in time the expected numbers of backordered units in the system are 65.2 and 13.0.

UNCERTAINTY RATIO OF THE INITIAL ESTIMATE

To indicate the effect of the uncertainty ratio we shall examine two sample items. For the first item in Table 2 the initial estimate of demand during a six-month display period, with an average of 1000 flying hours per month, was 9.12. Assuming an uncertainty ratio of 3, the other parameter of the gamma distribution from Eq. (12) is then 3.04. A set of seven possible values of true mean demand, θ_i , was selected ranging from 0.81 to 34.40, and the prior distribution was computed where

$$g(\theta_i) = \text{prob} \left\{ \frac{\theta_{i-1} + \theta_i}{2} \leq \theta \leq \frac{\theta_i + \theta_{i+1}}{2} \right\} .$$

The computed mean of the gamma distribution of 9.57 differs from the true mean of 9.12 because of the discrete approximation to the density function. One demand was assumed observed during a six-month period

Table 2
FIRST SAMPLE ITEM, UNIT COST \$922

Gamma Distribution Uncertainty Ratio, w	Gamma Distribution Parameter, v	Display Period (months)	Months of Demand	Observed Demand	Computed Gamma Mean	i						
						1	2	3	4	5	6	7
3	3.04	6	6	1.00	9.57	.017	.075	.124	.308	.334	.131	.010
					3.88	.082	.364	.323	.217	.014	.000	.000
						.81	2.44	4.07	6.66	11.52	19.90	34.40

$u = c$ 1 2 3 4 5 6 7 8 9 10+
 $p(u)$.135 .147 .142 .126 .106 .086 .067 .052 .039 .029 .072

Allocation (total 15) $\frac{D}{3}$ $\frac{1}{3}$ $\frac{2}{3}$ $\frac{3}{3}$ $\frac{4}{3}$ $\frac{5}{3}$

1	9.12	6	6	1.00	9.18	.000	.014	.099	.243	.357	.233	.053
					5.84	.001	.121	.391	.346	.128	.011	.000
						.95	2.84	4.73	6.71	9.14	12.46	16.98

$u = c$ 0 1 2 3 4 5 6 7 8 9 10+
 $p(u)$.036 .071 .097 .111 .113 .107 .096 .082 .068 .054 .164

Allocation (total 17) $\frac{D}{4}$ $\frac{1}{4}$ $\frac{2}{4}$ $\frac{3}{4}$ $\frac{4}{4}$ $\frac{5}{4}$ $\frac{6}{4}$

Table 3
SECOND SAMPLE ITEM, UNIT COST \$1595

Gamma Distribution Uncertainty Ratio, w	Gamma Distribution Parameter, v	Display Period (months)	Months of Demand	Observed Demand	Computed Gamma Mean	i												
						1	2	3	4	5	6	7						
3	.55	6	12	40.00	2.57 15.42	$g(\theta_1)$ $\phi(\theta_1)$.880 .000	.094 .006	.020 .145	.005 .593	.001 .253	.000 .002	.000 .000					
						θ_1	1.97	5.90	9.83	14.46	20.97	30.42	44.12					
							u^a	0	1	2	3	4	5	6	7	8	9	10+
							$P(u)$.000	.001	.003	.007	.012	.018	.026	.034	.041	.049	.809
Allocation (total 24) $\frac{D}{6} \frac{1}{6} \frac{2}{6} \frac{3}{6} \frac{4}{6} \frac{6}{6} \frac{8}{6}$																		
1	1.65	6	12	40.00	2.31 9.80	$g(\theta_1)$ $\phi(\theta_1)$.952 .000	.047 .265	.001 .631	.000 .088	.000 .015	.000 .001	.000 .000					
						θ_1	2.10	6.30	10.49	14.00	17.13	20.97	25.66					
							u^b	0	1	2	3	4	5	6	7	8	9	10+
							$P(u)$.004	.013	.026	.041	.055	.067	.076	.081	.082	.081	.474
Allocation (total 20) $\frac{D}{4} \frac{1}{4} \frac{2}{4} \frac{3}{4} \frac{5}{4} \frac{7}{4}$																		

and the corresponding Bayesian posterior distribution, $\mathcal{D}(\theta_1)$, was computed. Note that this value of observed demand is much lower than expected from the initial estimate, so that the posterior distribution with a mean of 3.88 is quite different from the prior.

The above information for an uncertainty ratio of three can be compared with the corresponding information for an uncertainty ratio of one. In the latter case with $w=1$ the value of v must be 9.12. The posterior mean is 5.84, indicating a greater reliance on the initial estimate or less uncertainty. However, it is difficult to compare the two posterior distributions since the θ_1 values are different. For this reason we have computed $p(u)$, the probability distribution for the number of demands to be observed in a six-month period when demand is logarithmic Poisson with a known variance to mean ratio, q . For instance, there is a probability of .135 that no demands will be observed in a six-month period at a base with 1000 flying hours per month when the uncertainty is three, initial estimate is 9.12, variance to mean ratio of the logarithmic Poisson is two, and past demand in six months is one. The probabilities of zero, one, two, ..., nine, and more than nine demands are shown. A comparison of these two sets of probabilities corresponding to different uncertainty indicates a substantial difference. Yet the remarkable fact is that when \$18.28 million is allocated among the 652 items under the two different uncertainty ratios, the allocations are very similar, differing by only one unit at depot and one unit at base three. The amount of \$18.28 million is used for comparison because it yields the target of .10 expected backorders per item in the standard case.

Similar observations apply to the second sample item in Table 3 for which, by contrast, observed demand is much higher than the initial estimate. We have purposely chosen two items whose initial estimates are very different from the observed data. Most items will show even less difference in allocation under the two uncertainty ratios.

From the system point of view we can summarize sensitivity to the uncertainty ratio by Table 4. For example, in the first column

Table 4

SENSITIVITY TO UNCERTAINTY RATIO WHEN
LOGARITHMIC POISSON VARIANCE/MEAN = 2.0

		Assumed Uncertainty		
		3	1	
(A)	Expected Backorders per Item	True Uncertainty	Bases	Bases
		3	.100 .024 .031 .045	.109 .025 .034 .050
	1	.112 .027 .035 .050	.103 .025 .032 .046	
	Investment (in \$ million)	Total Depot	\$18.285	\$18.287
	Base 1	5.781	5.715	
	Base 2	2.220	2.275	
	Base 3	3.921	3.953	
		6.364	6.345	

		Assumed Uncertainty		
		3	1	
(B)	Expected Backorders per Item	True Uncertainty	Bases	Bases
		3	.020 .006 .006 .008	.023 .007 .007 .010
	1	.023 .007 .007 .009	.020 .006 .006 .008	
	Investment (in \$ million)	Total Depot	\$34.031	\$34.070
	Base 1	9.056	8.760	
	Base 2	5.272	5.341	
	Base 3	7.991	8.067	
		11.712	11.902	

NOTE: Uncertainty ratios are defined relative to a six-month time period.

of the upper matrix we see the result of allocating stock under the assumption that the uncertainty ratio is 3. If the true ratio is also 3 the expected number of backorders is .100 distributed as .024 at Base 1, .031 at Base 2, and .045 at Base 3. Looking at either row in the upper matrix we find that an incorrect uncertainty assumption leads to an increase in expected backorders/item of about 9 percent* (.109/.100 or .112/.103). In the lower matrix, with almost twice as much investment, the degradation is on the order of 15 percent (.023/.020). As we have seen earlier, the two uncertainty ratios can lead to substantially different probabilities, so the small degradation evidenced here is somewhat surprising and very reassuring. It appears that correct estimation of the uncertainty ratio becomes more important with increasing investment. This tends to support our belief that in the early stages of weapon acquisition, when the uncertainty ratio is hardest to estimate, the investment targets should be kept as low as possible to prevent misallocation. This does not necessarily result in a small range of items on which positive stock is maintained. In the standard case with a .02 backorder/item target, all items have positive stock levels; and with a .10 backorder/item target, only nine items have zero stock levels.

We should remind the reader that system investment is the sum of investments at the bases and depot, whereas system backorders/item is the sum of backorders/item at the bases. Note that the fraction of system investment allocated to the depot decreases as investment increases.

VARIANCE TO MEAN RATIO OF THE LOGARITHMIC POISSON PROCESS

Table 5 is a display of the system sensitivity to the variance to mean ratio of the logarithmic Poisson process. The upper left-hand corner of each matrix contains the standard case, as in Table 4,

*We believe the relative error is more meaningful than the absolute error because expected backorders decreases approximately exponentially with increasing investment.

Table 5

SENSITIVITY TO LOGARITHMIC POISSON
VARIANCE/MEAN WHEN UNCERTAINTY RATIO = 3.0

		Assumed Variance/Mean	
		2	1
Expected Backorders per Item	True Variance/Mean 2	Bases .024	Bases .027
		.100 .031 .129 .041	.045 .061
	1	Bases .009	Bases .004
		.031 .009 .017 .005	.013 .008
Investment (in \$ millions)	Total	\$18.285	\$18.287
	Depot	5.781	6.062
	Base 1	2.220	2.509
	Base 2	3.921	3.700
	Base 3	6.364	6.016

		Assumed Variance/Mean	
		2	1
Expected Backorders per Item	2	Bases .006	Bases .009
		.020 .006 .036 .012	.008 .015
	1	Bases .001	Bases .000
		.002 .001 .000 .000	.000 .000
Investment (in \$ millions)	Total	\$34.031	\$34.034
	Depot	9.056	8.501
	Base 1	5.272	5.795
	Base 2	7.991	8.118
	Base 3	11.712	11.620

with a variance to mean ratio of two. Based on our maximum likelihood calculations mentioned earlier, this value of two is about as large as we would expect to observe. The smallest possible value is one corresponding to the Poisson.

In the upper matrix we see that an incorrect assumption of variance to mean ratio equal to one produces a 29-percent increase in expected backorders (.129/.100). By contrast, an incorrect assumption of variance to mean ratio equal to two yields an 82-percent increase in expected backorders (.031/.017). In the lower matrix, with almost twice the investment, an incorrect assumption of variance to mean ratio equal to one results in an 80-percent increase in expected backorders (.036/.020).

Clearly, the logarithmic Poisson variance to mean ratio is an important parameter. As in the sensitivity tests of the uncertainty ratio, it appears that correct estimation becomes more important with increasing investment. By looking at the columns in Table 5 we see that if backorder targets are specified, rather than investment targets, the resulting investment is extremely sensitive to variance to mean ratio.

ITEM DETAIL FOR THE STANDARD CASE

Table 6 displays the allocation for a sample of items in the standard case, where the uncertainty ratio is three, the logarithmic Poisson variance to mean ratio is two, and the target is .10 backorders per item. The purpose of Table 6 is to show allocation as a function of unit cost and demand, where the demand value is the mean of the posterior distribution for a base with a level of operations one. The first two items were described in detail in Tables 2 and 3. We also display elapsed depot time, which is the average time delay at depot due to the fact that the depot does not always have a serviceable unit on the shelf when a demand occurs. Since the average depot repair time was taken as one month on all items, the elapsed depot time must be between zero and one.

Table 6

SAMPLE OF ITEM DETAIL FOR THE STANDARD CASE
WITH TARGET OF .10 BACKORDERS/ITEM

Item No.	Cost (\$)	Demand	Stockage					Elapsed Depot Time
			Sum	Depot	Bases			
					1	2	3	
1	922	3.88	15	3	3	4	5	.14
2	1,595	15.42	24	6	4	6	8	.15
3	17,329	1.06	1	1	0	0	0	.37
4	35,370	14.73	10	4	1	2	3	.29
5	159	21.76	40	12	8	10	10	.03
6	444	6.05	21	4	4	6	7	.11
7	84,498	19.09	7	4	0	1	2	.38
8	110,500	3.29	0	0	0	0	0	1.00
9	5,687	.16	0	0	0	0	0	1.00
10	393	.18	6	1	1	2	2	.32
11	3,460	.53	3	1	0	1	1	.34
12	3,200	3.77	11	2	2	3	4	.26
13	29	.91	17	1	4	5	7	.36
14	10,925	6.18	9	3	1	2	3	.20
15	884	7.42	20	4	4	5	7	.14
16	204	5.30	25	3	6	7	9	.18
17	1,402	5.29	16	3	3	4	6	.18
18	48,451	19.09	11	4	1	2	4	.38
19	369	15.35	30	6	6	8	10	.15
20	220	.08	5	1	1	1	2	.31

SUMMARY STATISTICS FOR THE EXAMPLE

Over the entire group of 652 items, the average unit cost per item was \$3778 and the average predicted demand per item for the system was 16.10 demands during a six-month period. We remind the reader that the system consists of three bases with levels of operation of .5, 1, and 2. Due to the positive correlation between demand and unit cost in our data, the expected dollar value of daily demand for the system was \$439,000, which is substantially larger than the uncorrelated product $3778 \times 16.10 \times 652 / 180 = \$220,000$. Consequently, the targets of .10 and .02 backorders/item are obtained with investments of $18.285 / .439 = 41.68$ and $34.031 / .439 = 77.58$ days of supply.

As noted earlier the absolute value of z_j , the Lagrange multiplier, is the imputed ratio of the backorder cost divided by holding cost rate at Base j , where in our example this ratio is constant for the several bases. The two Lagrange multipliers corresponding to targets of .10 and .02 backorders/item are -1.577×10^5 and -7.194×10^5 , respectively. Thus, if we estimate the holding cost rate per year as .20 (expressed as a fraction of unit cost), the imputed costs of a backorder per day are $(1.577 \times 10^5 \times .2) / 365 = \86.54 and $(7.194 \times 10^5 \times .2) / 365 = \394.19 , respectively.

VI. GENERALIZATIONS OF THE RECOVERABLE ITEM MODEL

In this section we discuss two theoretical extensions of the recoverable item model, though neither is incorporated in METRIC.

BAYESIAN ESTIMATION FOR THE FRACTION OF UNITS BASE REPARABLE

In METRIC, the fraction of units reparable at Base j on a particular item, r_j , was assumed to be a known or estimable constant. By contrast, true mean demand, base repair time, depot repair time, and order and shipping time were variables. Let us consider how a Bayesian procedure could have been used to represent variability in the fraction of units reparable. As a simple example, assume that there are two possible values for the fraction of units reparable at Base j for a specific item, $r_j^{(1)}$ and $r_j^{(2)}$, with associated probabilities $z_j^{(1)}$ and $z_j^{(2)}$. Then the compound Poisson demand process at depot has a mean

$$\sum_{j=1}^J \lambda_j (1 - r_j^{\xi_j}) ,$$

with probability

$$\prod_{j=1}^J z_j^{\xi_j} ,$$

where each $\xi_j = 1$ or 2 . Since ξ_j may have two values and a particular depot mean is composed of J such ξ_j , there are 2^J possible values of mean demand to consider. For each particular value of mean depot demand, Stages 1 and 2 under "Multi-Echelon Theory" would have to be computed and multiplied by the appropriate probability

$$\prod_{j=1}^J z_j^{\xi_j} .$$

Of course, the λ_j are computed by a Bayesian procedure as well. Thus, if there are 10 Bayes states in the approximation for demand and 2^J

possibilities arising from r_j , there are (10×2^J) required computations. For a reasonable value of J such as 10 bases, the computation is not feasible.

A simpler Bayesian computation results when we consider the fraction base reparable to be the same at all bases, though unknown. Suppose there are K Bayes states $r^{(1)}, r^{(2)}, \dots, r^{(K)}$. Instead of 2^J calculations of Stages 1 and 2, this procedure entails only K computations. Furthermore, this model is probably more reasonable than the more complicated formulation in which the probability of a large fraction reparable at Base j is independent of the fraction reparable at other bases. However, we do not believe that the complications of estimating the Bayesian probabilities for fractions base reparable and additional computations warrant inclusion in METRIC.

INCLUSION OF CONDEMNATIONS

One of the assumptions for METRIC is that the system is conservative, which means that there are no condemnations. As noted above, this seems to be a reasonable approximation for Air Force recoverable spare parts. However, we shall present the mathematics for generalizing that assumption.

Consider a single item whose condemnation rate is γ . Let r_j be the fraction of units at base j which are base reparable as before, but let $(1-r_j)$ be the fraction of units either depot reparable or condemned. We assume equally spaced procurement intervals of τ . Whenever a condemnation occurs, depot stock is reduced, and we suppose that depot stock is sufficiently large so that the probability of exhausting depot stock during a procurement cycle is negligible.

The average response time for base j is now a function of time t because the depot spare stock is a function of time. Generalizing the expression obtained in Stage 2 of Sec. III yields

$$T_j(t) = r_j A_j + (1-r_j)(0_j + \Delta(s_0, t)D) \quad 0 \leq t \leq \tau,$$

where $\Delta(s_0, t)$ is the generalization of $\delta(s_0)$. Immediately following a procurement delivery when $t=0$, the depot stock level is s_0 so that

$\Delta(s_0, 0) = \delta(s_0)$. After a positive time t has elapsed, the depot stock level will be $(s_0 - k)$, where $0 \leq k \leq s_0$, provided that k condemnations have taken place. Under the assumption that all demands by a particular customer are either base reparable, depot reparable, or condemned, the probability distribution of condemned items is compound Poisson with a customer rate $\lambda\tau$. As noted earlier we assume that the probability of more than s_0 condemnations is negligible, but we obtain a better approximation to the delay if we assume that in such cases the depot spare stock is zero. Generalizing Eq. (5) yields

$$\begin{aligned} \Delta(s_0, t) &= \frac{1}{B(0|\lambda D)} \left\{ \sum_{k=0}^{s_0-1} B(s_0 - k|\lambda D) p(k|\lambda\tau\gamma) + B(0|\lambda D) \sum_{k=s_0}^{\infty} p(k|\lambda\tau\gamma) \right\} \quad 0 \leq t \leq \tau \\ &= \frac{1}{B(0|\lambda D)} \left\{ \sum_{k=0}^{s_0-1} B(s_0 - k|\lambda D) p(k|\lambda\tau\gamma) \right\} + \sum_{k=s_0}^{\infty} p(k|\lambda\tau\gamma) \quad 0 \leq t \leq \tau. \end{aligned}$$

The expected backorders at base j with a base stock level s_j is then a function of t :

$$B(s_j|\lambda_j T_j(t)) \quad 0 \leq t \leq \tau.$$

We are interested in the time average of this statistic which can be approximated by evaluating the function at several points t_1, t_2, \dots, t_n in the interval $0, \tau$ and averaging the result. This function generalizes Stage 2 of the multi-echelon procedure described in Sec. III.

APPENDIX

Theorem Let $\Xi_i(m_i)$ be a monotonic decreasing function with a finite lower bound defined for $m_i=0,1,2,\dots$. Define $\Xi'_i(m_i)$ on the boundary of the convex hull of $\Xi_i(m_i)$ such that $\Xi'_i(m_i) \leq \Xi_i(m_i)$. Suppose the objective is to find the $\{\bar{m}_i\}$ which

$$(A.1) \quad \min_{\{m_i\}} \sum_i [c_i m_i + \Xi_i(m_i)] \quad m_i=0,1,2,\dots,$$

where the $\{c_i\}$ are nonnegative and not all zero. Then unique optimizing $\{\bar{m}_i\}$ are determined by

$$(A.2) \quad \begin{aligned} c_i + \Xi'_i(\bar{m}_i+1) - \Xi'_i(\bar{m}_i) &\geq 0 \\ c_i + \Xi'_i(\bar{m}_i) - \Xi'_i(\bar{m}_i-1) &< 0. \end{aligned}$$

Proof Since the objective function in Eq. (A.1) is a sum of independent functions for each item i , we can disregard the summation sign and deal with a single item. We shall begin by showing that when Ξ_i in Eq. (A.1) is replaced by Ξ'_i , Eq. (A.2) has a unique optimal solution. Then we shall demonstrate that this solution is also appropriate for Ξ_i .

It is clear that if $\{\bar{m}_i\}$ exist which satisfy Eq. (A.2), then Eq. (A.1) is solved for Ξ'_i . Define $\Xi'_i(-1)=\infty$. We use the fact that a monotonically decreasing sequence bounded below has a limit. This shows that Eq. (A.2) has at least one solution for any nonnegative value of c_i . But Eq. (A.2) implies that the set of points $\Xi'_i(\bar{m}_i-1)$, $\Xi'_i(\bar{m}_i)$, and $\Xi'_i(\bar{m}_i+1)$ is strictly convex. Since the function Ξ'_i is convex, \bar{m}_i is unique.

Consider any point m_i at which Ξ_i was decreased to form Ξ'_i . The value of $\Xi'_i(m_i)$ lies on a secant line so that $\Xi'_i(m_i+1) - \Xi'_i(m_i) = \Xi'_i(m_i) - \Xi'_i(m_i-1)$. And thus, the artificial point m_i cannot satisfy Eq. (A.2). Replacing $\Xi_i(m_i)$ by $\Xi'_i(m_i)$ can only decrease the minimum in Eq. (A.1), but as we have just seen any \bar{m}_i satisfying Eq. (A.2) implies that at point $\Xi_i(\bar{m}_i) = \Xi'_i(\bar{m}_i)$, solving the original problem.

In other words if Eq. (A.1) is minimized for $\{\Xi'_i\}$ and then each Ξ'_i is artificially increased at some values $m_{i1}, m_{i2} \dots$, not coinciding with the solution \bar{m}_i , the original solution is obviously unaffected.

REFERENCES

1. Brooks, R., and A. Geoffrion, Finding Everett's Lagrange Multipliers by Linear Programming, The RAND Corporation, P-3315, February 25, 1966; also to be published by Operations Research.
2. Everett, H., "Generalized Lagrange Multiplier Method for Solving Problems of Optimal Allocation of Resources," Operations Research, Vol. 11, No. 3, May-June 1965, pp. 399-417.
3. Feeney, G. J., and C. C. Sherbrooke, "The (s-1,s) Inventory Policy Under Compound Poisson Demand," Management Science, Vol. 12, No. 5, January 1966, pp. 391-411; also published by The RAND Corporation, RM-4176-1-PR, March 1966.
4. Feeney, G. J., and C. C. Sherbrooke, A System Approach to Base Stockage of Recoverable Items, The RAND Corporation, RM-4720-PR (DDC No. AD 627644), December 1965.
5. Kendall, M. G., and A. Stuart, The Advanced Theory of Statistics, Vol. 2, Hafner Publishing Company, New York, 1961, p. 72.
6. 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 (DDC No. AD 615261), April 1965.
7. McGlothlin, W. H., Levelopment of Bayesian Parameters for Spare-Parts Demand Prediction, The RAND Corporation, EM-3699-PR (DDC No. AD 411356), July 1963, p. 7.
8. Petersen, J. W., R. M. Paulson and W. A. Steger, Supply and Depot-Repair Interactions: A Case Study of Electronics Support, The RAND Corporation, RM-2365, April 1959.
9. Sherbrooke, C. C., Discrete Compound Poisson Processes and Tables of the Geometric Poisson Distribution, The RAND Corporation, RM-4831-PR, July 1966.
10. Weifenbach, Annette, Base Maintenance Activity and Repair Cycle Times, The RAND Corporation, RM-5027-PR, September 1966.

DOCUMENT CONTROL DATA

1. ORIGINATING ACTIVITY THE RAND CORPORATION		2a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	
		2b. GROUP	
3. REPORT TITLE METRIC: A MULTI-ECHELON TECHNIQUE FOR RECOVERABLE ITEM CONTROL			
4. AUTHOR(S) (Last name, first name, initial) Sherbrooke, Craig C.			
5. REPORT DATE November 1966		6a. TOTAL No. OF PAGES 63	6b. No. OF REFS. 10
7. CONTRACT OR GRANT No. AF 49(638)-1700		8. ORIGINATOR'S REPORT No. RM-5078-PR	
9a. AVAILABILITY/LIMITATION NOTICES DDC-1		9b. SPONSORING AGENCY United States Air Force Project RAND	
10. ABSTRACT A description of METRIC, a mathematical model of a base-depot supply system translated into a computer program. METRIC is capable of determining base and depot stock levels for a group of recoverable items of the type that accounts for 78 percent of the Air Force's investment in spares. Designed for application at the weapon-system level, METRIC's primary aim is to optimize system performance for specified levels of system investment. This is probably the first multi-echelon, multi-item model ever proposed for Air Force implementation.		11. KEY WORDS Logistics Inventory control Models Maintenance Computer programs	