MEMORANDUM
RM-5725-PR
SEPTEMBER 1968

WAR RESERVE SPARES KITS SUPPLEMENTED BY NORMAL OPERATING ASSETS

Robin B. S. Brooks and John Y. Lu

PREPARED FOR:
UNITED STATES AIR FORCE PROJECT RAND

The RAND Corporation
SANTA MONICA • CALIFORNIA

MEMORANDUM
RM-5725-PR
SEPTEMBER 1968

WAR RESERVE SPARES KITS SUPPLEMENTED BY NORMAL OPERATING ASSETS

Robin B. S. Brooks and John Y. Lu

This research is supported by the United States Air Force under Project RAND—Contract No. F44620-67-C-0045—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 study should not be interpreted as representing the official opinion or policy of the United States Air Force.

The RAND Corporation
1700 MAIN ST • SANTA MONICA • CALIFORNIA • 90406

Published by The RAND Corporation

PREFACE

This Memorandum describes a model used to assist the Air National Guard in designing their war reserve spares (WRS) kits. It extends earlier classified work that RAND did for the Aerospace Defense Command. The primary difference between this model and the earlier one is that in this model we take explicit account of the role that normal operating assets might play in supplementing WRS. The possible application of this technique to the problem of designing Tactical Air Command mobility kits is also pointed out.

The Memorandum is intended for operations analysts who are concerned with Air Force logistics planning.

SUMMARY

In peacetime, base stock levels of spares are determined on the assumption of normal resupply from the depot. In the event of war, however, a unit must be prepared to operate from stock on hand for a period of time without being resupplied from the depot. This Memorandom describes a mathematical model for determining such war reserve stockage (WRS) requirements. Specifically, the model solves the following kind of optimization problem: find the least-cost WRS kits that will keep the probability of a stockout after K cannibalizations less than or equal to some target objective α. The user of the model specifies the number of allowable cannibalizations, and the level of protection that the kit is supposed to provide.

The key assumptions are the following: (1) the quantity demanded for an item is distributed by compound Poisson; (2) the quantity demanded for an item is independent of the demand for other items; (3) the wartime demand does not depend on the number of serviceables that might be available from normal operating stocks; (4) there are no holes in the aircraft at the outset of a contingency; and (5) items to be considered for inclusion in the kit are equally essential for all missions and all such items can be identified.

We call the probability of meeting all spares demands the operational rate. One interesting feature of this model is that in the probability computation it takes into account the possibility of utilizing normal base operating assets. It should be noted, however, that just because the base supply has, say, 5 units of an item, it is not assumed that there will necessarily be 5 units available at all

times. In fact, we compute the probability that 1, 2, 3, 4, or all 5 units will either be in maintenance or in the pipeline from the depot. These probabilities are then used to compute the expected number of serviceable units on hand at a random point in time.

To maximize the operational rate, the well-known technique of marginal analysis is used. The technique assigns stock to an item in the order indicated by the increase in operational rate per dollar invested. In other words, it finds the item with the largest ratio of increment in operational rate to increment in cost (which is the item's unit cost) and assigns the first unit, then looks for the next largest ratio, and so on until operational rate is increased to the limit of a given budget constraint.

A sensitivity analysis was performed, based on the data from an Air National Guard unit. If peacetime levels were explicitly taken into account when designing a WRS kit, results indicated that a cost saving of nearly 40 percent would be effected without degrading base supply performance in wartime.

CONTENTS

I. INTRODUCTION

SCOPE

This Memorandum describes a mathematical model that was used to assist the Air National Guard in designing war reserve spares (WRS) kits. Although the study was done explicitly for the Air National Guard, we feel that it also has some relevance to the problem of designing WRS kits under the current mobility concept of the Tactical Air Command (TAC). (This point will be brought up again in Sec. IV.)

BACKGROUND

The Air National Guard study extends an earlier effort for the Aerospace Defense Command (ADC) in which RAND developed a model to compute WRS kits that ADC would use at their dispersal sites.
In the summer of 1966, ADC, together with the Air National Guard, asked RAND to help them construct WRS kits to be used by Air National Guard units assigned to an air defense mission. The Air National Guard operates under a fight-in-place concept, rather than a dispersal concept that ADC uses. This difference should be reflected in the design of WRS kits, because the Air National Guard can supplement their war reserve spares in the event of war with assets that they use for normal day-to-day operations. It seemed, therefore, that WRS kits for the Air National Guard could be designed at considerably smaller cost than kits used by ADC.

PRINCIPAL FEATURES OF THE MODEL

A mathematical model was developed and programmed in FORTRAN IV for the IBM 7044 computer at RAND. The model solves the following kind of optimization problem: it finds the least-cost WRS kits that will keep the probability of a stockout after K cannibalizations less than or equal to some target objective α. The user specifies the number of allowable cannibalizations, K, and the level of protection that the kit is supposed to provide, α. For instance, if we set $K = 1$ and $\alpha = 10$ percent, the model will compute a WRS kit that will provide enough parts support to keep the number of aircraft grounded due to lack of spares at one or less, with a probability of 90 percent.

We call the probability of meeting all spares demand the kit's operational rate. This is used as a measure of kit performance. The probability computation considers a number of factors. The first feature of the model is that it takes into account the possibility of utilizing normal base operating assets. It should be pointed out that just because the base has, say, 5 units of an item, we do not assume it will necessarily have 5 units on hand at all times. In fact, we compute the probability that 1, 2, 3, 4 or all of these units will either be in maintenance or in the pipeline from the depot. The kit computation then takes these probabilities into account to assess the availability of the normal operating assets of this item in wartime. In order to satisfy subsequent demands, the second feature the model takes into account is the possibility of cannibalizing an aircraft that is already grounded. The third feature is that the model considers the possibility that some parts can be repaired within an aircraft's turnaround time. The fourth feature is related to demand forecast. Just as in the ADC study, past demand data are treated in a Bayesian fashion for predicting future demand. Finally, percentage base repair is treated as an unknown factor to be predicted for kit optimization. Its past data are also treated in a Bayesian fashion similar to the demand prediction.

ORGANIZATION OF THE MEMORANDUM

The outline of the Memorandum is as follows: Section II contains statements of the assumptions on which the model is based, a discussion of the measure of performance and its computation, and a description of the optimization method. In Sec. III, sample data are used to illustrate the effect of considering ordinary base stock levels when computing the costs of war reserve spares requirements for a WRS kit. And in Sec. IV, several possible extensions of the model are discussed--in particular, its applicability for computing TAC mobility kits.

II. THE MODEL

ASSUMPTIONS

1. The quantity demanded for an item in a fixed period is assumed to follow a stationary probability distribution. In particular, the form of assumed distribution is Geometric (or stuttering) Poisson. This provides a flexibility for modeling the situation in which item demands have a variance-to-mean ratio greater than unity.

2. The quantity demanded for an item is independent of other items. This assumption is appropriate as applied to a set of recoverable assemblies as long as there are not many complex LRU (Line Replaceable Unit) module relationships.

3. The wartime demand does not depend on the number of serviceables that might be available from day-to-day operating stocks. In other words, the distributions of quantities demanded in peacetime and in wartime are independent of each other.

4. We assume that there are no holes in the aircraft at the outset of hostilities, i.e., all the aircraft are operationally ready. This probably is not a true picture of the real world, but a crude approximation of an expediting process that takes place whenever an aircraft is grounded for lack of spares.

5. Items to be considered for inclusion in the kit are equally essential for all missions and all such items can be identified.

OPERATIONAL RATE, CANNIBALIZATIONS, AND NORS

"Operational rate" is the yardstick used to measure the effectiveness of a WRS kit. The operational rate is defined as the probability that the kit, together with those serviceables available from normal base stocks, can meet spare parts demand during the emergency. What does an operational rate of, say, 90 percent mean? One statistical interpretation is that if a situation arises requiring use of the kit, then 90 percent of the Air Force units using the kit will not be hampered from performing their mission because they lack spares.

Because of assumption 2, operational rate may be expressed as the product across all items of the individual item operational rates. (The operational rate for a single item is the probability of meeting all demands for that item.) Thus if we let $g_j(k_j)$ be the operational rate for item j when there are k_j units of the item in the kit and base supply, then

$$(1) \qquad \prod_j g_j(k_j)$$

gives the overall operational rate. Note that $g_j(k_j)$ is simply the probability that during the war there will be no more than k_j demands for item j in excess of serviceables available from the combined assets of the kit and base supply. Now suppose that there is one aircraft that can be cannibalized for parts, and let a_j be the number of units of item j that are used on one aircraft. Then, in effect, for item j we have a_j units that may be added to the k_j units in the system, so that the probability of having a stockout on item j has been decreased from $g_j(k_j)$ to $g_j(k_j + a_j)$, and the overall operational rate has been increased from expression (1) to

$$\prod_j g_j(k_j + a_j).$$

In fact, in general, if there are c aircraft available for cannibalization, then the effective stock level for item j becomes $k_j + ca_j$, so the operational rate increases from (1) to

$$\prod_j g_j(k_j + ca_j). \tag{2}$$

As it turns out, it is just as easy to optimize the generalized operational rate given by (2) as it is to optimize the special case given by (1).

In the foregoing, we have interpreted (2) as the probability of meeting all demands for spares during the war, given that there are c aircraft available for cannibalization. This interpretation is also the one given in the ADC study. There is, however, a slightly different interpretation of (2) that may be more meaningful. We may rephrase our interpretation of (2) as the probability of meeting all demands for spares without having to cannibalize more than c aircraft. Thus, if we consolidate parts shortages on as few aircraft as possible, we may interpret (2) as the probability that there will be no more than c NORS (Not Operationally Ready--Supply) aircraft.

To compute (2), it is obvious that we need only be able to compute $g_j(k)$ for any j and any k, i.e., the operational rate for an individual item j when k units of the item are in the kit and there are no aircraft available for cannibalization.

COMPUTATION OF OPERATIONAL RATE FOR A SINGLE ITEM

In this subsection we describe the computation of the operational rate for a single item, i.e., the probability that the quantity on hand at the beginning of the emergency (WRS plus serviceables from normal base stockage) is sufficient to cover wartime demands. This probability, of course, takes into account the uncertainty associated with the number of wartime demands. It also takes into account the uncertainty associated with other relevant quantities listed below.

The Amount of On-Hand Serviceables from Normal Base Operations (Exclusive of WRS). It is obvious that if this quantity is high, then there is a good chance of meeting all demands for the item. It is also obvious that the amount of on-hand serviceables is uncertain--since it depends, in part, on the demands for the item just prior to the emergency.

Wartime and Peacetime Demand Rates. To compute the probability distribution for wartime demands, we need to know the item's expected wartime demand. Unfortunately, this is not known exactly; all we have are data on the number of demands experienced during some past period. It is perfectly possible, for example, that an item with a low, but positive, demand rate could go for six months or a year without experiencing any demands, and yet be one of the first items demanded during the war. As pointed out above, the amount of on-hand serviceables at the beginning of the war has an important effect on the chances of meeting all wartime demands for the item. This amount is determined, in part, by the peacetime demand rate. Hence the item's peacetime demand rate influences the chances of meeting all the wartime demands. But just like the wartime rate, the peacetime rate is not known with certainty either.

Percentage Base Repair. On-hand inventory of serviceables is influenced not only by peacetime demand rates, but also by the peacetime response time for an item, i.e., the time between the turn-in of a carcass and the receipt of a serviceable replacement for it at base supply. (The replacement will either be the original carcass repaired, or, if the repair cannot be accomplished at base level, a receipt from the depot.) Since base repair time is generally shorter than depot order and shipping time, then the higher the percentage of repair accomplished at base, the higher, on the average, the on-hand serviceables will be. Base repair is an even more important influence during the emergency. It may be that if an item can be repaired at all, then it can be repaired within the turnaround time of the aircraft. If this is the case, then the item's wartime demand rate should be reduced by an amount proportionate to the percentage of base repair. Unfortunately, we cannot be certain what the expected percentage base repair on an item should be. For most items, the amount of relevant experience is so small that percentage base repair--just like demand rates--should be regarded as an uncertain quantity.

For the sake of definiteness in what follows, let us assume that we are talking about an item that can be repaired during the war, within the turnaround time of the aircraft--in those cases where the repair can actually be done at base level. We assume, for now, that the wartime demand has a Poisson distribution. We begin by assuming that expected wartime demand, percentage base repair, and the amount of on-hand serviceables from normal base stocks are, contrary to what was said above, known quantities denoted by λ, ρ, and X, respectively. Let k denote the quantity of the item in the WRS kit. Then the probability that all wartime demands for the item can be met via either the kit, repair, or on-hand peacetime serviceables is simply

$$F[k + X;\ \lambda(1 - \rho)], \tag{3}$$

where $F[\cdot;\nu]$ denotes the cumulative Poisson distribution function with mean ν.

$$F[y;\nu] = \sum_{j=0}^{y} e^{-\nu}\nu^j/j! \quad .$$

Unfortunately, as we pointed out before, we do not know X. Under suitable assumptions, in particular that peacetime demand has a Poisson distribution, however, we do know its approximate probability distribution [4]. In fact, if we denote average repair time and average resupply time by R and S, respectively, so that average response time is $(1 - \rho)S + \rho R$, then the probability that X takes on the value x is given by $f[q - x; \mu((1 - \rho)S + \rho R)]$ for $x > 0$, and $1 - F[q - 1; \mu((1 - \rho)S + \rho R)]$ for $x = 0$, where q is the normal peacetime stock level, μ is the item's peacetime demand rate, and $f[\cdot;\nu]$ is the Poisson density function with mean ν. Thus, if we form the expected value of (3) with respect to X, we obtain

$$(4) \qquad F[k; \lambda(1 - \rho)][1 - F(q - 1; \mu((1 - \rho)S + \rho R)]$$

$$+ \sum_{x=1}^{q} F[k + x; \lambda(1 - \rho)]f[q - x; \mu((1 - \rho)S + \rho R)]$$

for the probability of meeting all wartime demands.

In passing from (3) to (4), we have removed the assumption that X, the on-hand serviceables from normal stocks, is known. In much the same way, we remove the assumptions that λ, ρ, and μ are known with certainty. Let us assume instead that our uncertainty--and knowledge--

about ρ may be reflected by m numbers $\rho_1,\ldots,\rho_m$ and m probabilities $r_1,\ldots,r_m$ where r_i represents the probability that $\rho = \rho_i$. Similarly, we assume that our knowledge about λ and μ may be reflected in n pairs of numbers $(\lambda_1,\mu_1),(\lambda_2,\mu_2),\ldots,(\lambda_n,\mu_n)$ and n probabilities $d_1,\ldots,d_n$, where d_j is the probability that λ takes on the value λ_j and μ takes on the value μ_j. (The derivation of these numbers and probabilities will be explained below.) Then, if we sum (4) over the different values that λ, ρ, μ can take on--in each case multiplying it through by the probability that these values will occur--we obtain

$$(5)\quad g(k) = \sum_i r_i \sum_j d_j \Big\{F[k;\ \lambda_j(1-\rho_i)][1 - F(q-1;\ \mu_j((1-\rho_i)S + \rho_i R)]$$
$$+ \sum_{x=1}^{q} F[k+x;\ \lambda_j(1-\rho_i)]f[q-x;\ \mu_j((1-\rho_i)S + \rho_i R)]\Big\}$$

as the operational rate for a single item.

It remains to determine the numbers ρ_i, r_i, λ_j, μ_j, d_j, the probability distributions for percentage base repair and wartime and peacetime demand rates. These are derived by applying the "objective Bayes" technique [2]. We first describe how the numbers $\rho_1,\ldots,\rho_m$ and $r_1,\ldots,r_m$ are derived (these give the probability distribution for the percentage base repair).

We assume that ρ, the percentage base reparable, has an _a priori_ probability distribution that we then approximate by a discrete distribution. This approximation is obtained by dividing the interval between 0 and 1 into m equal intervals. We let $\rho_0 = 0$, and let ρ_i be the midpoint of the i^{th} interval for $i > 0$. We wish to associate with ρ_i

the probability r_i^0 that was originally associated with the i^{th} interval. (In the case where $i = 0$, r_i^0 will be the probability originally associated with 0.) The numbers r_i^0 are estimated by simply letting r_i^0 be the proportion of items that, within the data period, have had a percentage base repair that falls in the i^{th} interval.*

For an individual item, the numbers r_i (a posteriori probabilities) are obtained by applying Bayes' rule to the item's repair data and the a priori probabilities r_i^0. Specifically, suppose that during the data period N units of the item have been turned in to base maintenance, but of these N only K units could be repaired on the base. The probability of this event, given the number N of turn-ins and given that the base repair percentage is actually ρ_i, is simply

$$\binom{N}{K}\rho_i^K(1 - \rho_i)^{N-K}.$$

Hence, by Bayes' rule, the a posteriori probability that $\rho = \rho_i$ is given by

$$(6) \qquad r_i = \frac{r_i^0\binom{N}{K}\rho_i^K(1 - \rho_i)^{N-K}}{\sum_j r_j^0\binom{N}{K}\rho_j^K(1 - \rho_j)^{N-K}} = \frac{r_i^0\rho_i^K(1 - \rho_i)^{N-K}}{\sum_j r_j^0\rho_j^K(1 - \rho_j)^{N-K}} .$$

We turn now to the determination of the demand rates λ_j, μ_j, and their associated probabilities, d_j. We assume that λ and μ stand in the relation

---

*This procedure for determining the number r_i^0 is entirely analogous to the use of objective Bayes technique in [2] and [3] except that there it was applied to demand rates rather than percent base repair.

$$\lambda = C\mu ,$$

where C, a constant across all items, reflects the relative level of activity during war and peacetime. (We have computed C as the ratio of the number of sorties to be flown during the war to the number of sorties per unit time flown during peace.*) Thus, if we know the μ_j, we may compute the λ_j by the formula

$$\lambda_j = C\mu_j . \qquad (7)$$

Finally, the numbers μ_j and d_j are determined exactly as in the RAND base stockage model [3] (another application of the objective Bayes approach).

The assumption that peacetime and wartime demand have Poisson distributions may be relaxed by assuming some compound Poisson distribution instead. In the work for the Air National Guard, we assumed a geometrically compounded Poisson distribution. This made it necessary to specify an additional piece of data, the variance-to-mean ratio, which we assumed to be the same for all the items.

In the foregoing, we assumed that the item whose operational rate we were computing could be repaired, in those cases where repair could be accomplished at all, within the aircraft's turnaround time. For items that cannot be repaired within the turnaround time, the term $\lambda(1 - \rho)$ in (3) and subsequent expressions should be replaced by λ.

---

*There is an assumption implicit in this statement, viz., demands are correlated to the number of sorties. If this assumption is not appropriate one may have to consider the ratio of total flying hours per unit time or some other appropriate index for measuring the differences in the activity levels of the two situations.

OPTIMIZATION

Let b_j be the unit cost of item j. Then the cost of the WRS kit will be $\sum_j b_j k_j$, where k_j is the quantity of item j in the kit. The problem of finding the least-cost WRS kit that will have a prespecified operational rate s may then be phrased as follows:

$$\text{(7)} \qquad \begin{aligned} &\text{Minimize } \sum_j b_j k_j \\ &\text{Subject to } \prod_j g_j(k_j + ca_j) \geqq s. \end{aligned}$$

It turns out to be more convenient to replace operational rate by its logarithm. When we do this, (7) becomes

$$\text{(8)} \qquad \begin{aligned} &\text{Minimize } \sum_j b_j k_j \\ &\text{Subject to } \sum_j \log g_j(k_j + ca_j) \geqq \log s. \end{aligned}$$

This problem may be solved by marginal analysis. We start by setting $k_j = 0$ for each item j. We then find an item j for which the ratio

$$\frac{\log g_j(k_j + ca_j + 1) - \log g_j(k_j + ca_j)}{b_j}$$

is a maximum, and increase the corresponding stock level k_j by one unit.

We continue in this way until the constraint in (8) (and therefore the constraint in (7)) is just satisfied.*

---

*Strictly speaking, problem (8) is only approximately solved by this procedure. However, the solution obtained by marginal analysis is "efficient" in the sense that no other policy is better with respect to one criterion (cost or performance) without being worse with respect to the other [5]. It should also be pointed out that when demand rates are large, the functions g_j may not be concave. In this event they should be replaced by their concave majorants as in [3].

III. SENSITIVITY ANALYSIS

One of the interesting features of the model described in Sec. II is that it takes account of ordinary base stock levels when computing the war reserve spares requirements. In this section, we use some numerical results to illustrate how such a consideration affects the cost of a WRS kit. For the sensitivity analysis pertaining to other features of the model, such as the ability to incorporate cannibalization into a spares requirement computation and the use of a Bayesian technique for demand prediction, the reader is referred to the classified work we did for ADC and [5].

To examine the impact of base stock levels on the cost of a WRS kit, we have computed some WRS kits by four different methods. The computation is based on data from an F-102 Air National Guard Squadron at Boise Airport, Idaho. The data came from 16 aircraft over a 6-month period. During that time the squadron flew a total of 2218 hours and 1299 sorties. In all, we considered 245 items as candidates for the kit. The results of this computation are presented below.

	Method	Cost (In $ thousand)
I.	Peacetime levels not used during war	332
II.	Peacetime levels used during war, but not in the optimization	222
III.	Peacetime levels used during war and in optimization	137
IV.	Peacetime levels computed by BSM[a] and used in both war and optimization	116

NOTE: 90-percent performance.

[a]The RAND Base Stockage Model.

Method I assumes that the only supply support during the war would come from the WRS kit. This assumption would be valid were the Air National Guard to operate under a dispersal concept similar to that used by units of the Aerospace Defense Command. In the second computation, it was still assumed that there were no normal operating assets for the purpose of kit optimization. We then came up with a kit that, when evaluated jointly with the normal peacetime assets, resulted in 90-percent performance. This kit cost $222,000. The third method explicitly takes into account the levels at Boise Airport. These base stock levels are a mixture of Chapter 11 stock levels and negotiated levels. A target for performance was again set at 90 percent. The resulting kit cost $137,000.

The meaningful comparison for the Air National Guard is between the second and third kits. By taking peacetime levels into account explicitly when designing a WRS kit, we can effect a cost saving of nearly 40 percent without degrading base supply performance in wartime.

The computation method for the last kit was the same as that for the third kit, except that peacetime levels were computed using the RAND Base Stockage Model. We were interested to see what effect the "efficient" base stock levels in peacetime would have on the cost of a WRS kit. We computed the dollar investment at Boise Airport at its peacetime levels* and used the same dollar figure as an investment constraint in maximizing peacetime fill rate. The results indicate that the peacetime levels computed by the Base Stockage Model seem

---

*This figure was estimated to be about $538,000.

to have provided a somewhat better set of levels with which to start the war reserve computation.

These numerical results indicate that if we could explicitly take account of the availability of peacetime assets to support wartime operations when we design a WRS kit, we could either substantially reduce costs or substantially increase supply performance.

IV. EXTENSIONS OF THE MODEL

In this section we briefly discuss three possible extensions of the model and its application.

THE DESIGN OF MOBILITY KITS

TAC mobility kits represent another type of WRS kit. The main difference between the TAC mobility kits and the ADC-ANG kits is that a mobility kit is taken with a unit when it deploys, and is therefore used at a site remote from the one at which the unit normally operates. This means that there is no possibility of supplementing WRS with peacetime stocks, unless some peacetime stocks may be deployed as well. In this event, we may well want to design mobility kits taking explicit account of the possibility of augmenting them from peacetime stocks at the time of deployment.

JOINT OPTIMIZATION OF PEACETIME AND WRS STOCK LEVELS

The model we have described here sets only the WRS levels; it treats peacetime stock levels as externally determined data. In other RAND stockage work, models have been developed for determining the peacetime levels. These models (as well as current Air Force policy) ignore the possible role of peacetime stocks in augmenting WRS. It might be possible to derive more cost-effective stockage policies by optimizing both the WRS and the wartime levels at the same time in order to meet constraints on both the peacetime and the wartime performance of the combined stockage policy. The theory for such a joint optimization has already been developed [1].

USE OF NORS AS A PERFORMANCE CRITERION

In Sec. III we pointed out a connection between operational rates for various levels of cannibalization, and the probability distribution for the number of NORS aircraft. Since expected number of NORS is perhaps a somewhat more meaningful performance measure than operational rate, one might well ask why we do not optimize expected NORS rather than operational rate. Unfortunately, the expression for expected NORS is not nearly as tractable mathematically as is operational rate. In fact, the marginal analysis technique of the preceding section is inapplicable, at least directly, to the problem of optimizing expected NORS. Thus, how to directly optimize expected NORS is a topic of current research. Fortunately, however, kits that give high operational rates tend to give low expected NORS, so that when we optimize operational rate, we are, to a large extent, optimizing expected NORS.

REFERENCES

1. Brooks, R.B.S., Some Linear Programming Applications to Stockage Problems, The RAND Corporation, RM-5422-PR, September 1967.

2. Feeney, G.J., and C.C. Sherbrooke, An Objective Bayes Approach for Inventory Systems, The RAND Corporation, RM-4362-PR, March 1965.

3. -----, A System Approach to Base Stockage of Recoverable Items, The RAND Corporation, RM-4720-PR, December 1965.

4. -----, The (s-1,s) Inventory Policy Under Compound Poisson Demand: A Theory of Recoverable Item Stockage, The RAND Corporation, RM-4176-PR, March 1966.

5. Fox, B., "Discrete Optimization via Marginal Analysis," Management Science, November 1966.