AD 408 943

c 3

408 943

JUN 25 4 04 PM '63
OFFICE OF THE SECRETARY OF DEFENSE

MEMORANDUM
RM-3644-PR
JUNE 1963

AD No. 408943

DDC FILE COPY

AN AGGREGATE BASE STOCKAGE POLICY FOR RECOVERABLE SPARE PARTS

G. J. Feeney, J. W. Petersen and C. C. Sherbrooke

PREPARED FOR:

UNITED STATES AIR FORCE PROJECT RAND

The RAND Corporation
SANTA MONICA • CALIFORNIA

$4.60

(5) 739 200

(4) $4.60

AFRDP
Received
JUN 24 1963

(14) memo. no.

MEMORANDUM
RM-3644-PR
JUNE 1963

(6) AN AGGREGATE BASE STOCKAGE POLICY FOR RECOVERABLE SPARE PARTS

(10) G. J. Feeney, J. W. Petersen and C. C. Sherbrooke

(7) NA
(8) NA
(9) NA
(11) Jun 63.
(12) 40p.
(13) NA

(15) This research is sponsored by the United States Air Force under Project RAND—contract No. AF 49(638)-700 monitored by the Directorate of Development Planning, 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.

(16-17) NA
(18-19) NA
(20) U.

The RAND Corporation
1700 MAIN ST • SANTA MONICA • CALIFORNIA

(21) Report on Proj RAND.

AFRDP Received JUN 24 1963

RAND PUBLICATION

RM-3644-PR

AN AGGREGATE BASE STOCKAGE POLICY FOR RECOVERABLE SPARE PARTS

G. J. Feeney, J. W. Petersen and C. C. Sherbrooke

June 1963

In brief

This Memorandum describes the initial results of a digital computer program that considers repair characteristics and unit cost by item in the computation of item stock levels. These levels will achieve a specified aggregate base fill rate across all recoverable items with the least dollar investment in base recoverable inventory. ~~Preliminary tests of the stockage model, while not conclusive, are very encouraging.~~ The test consisted of taking demand data for a sample of 2802 recoverable items at Andrews Air Force Base and using the first six months of data as model input to compute item stock levels required to achieve a range of aggregate base fill rates. Demands for these items for the next six months were then compared with these stock levels, ~~in order to estimate the support performance that would have resulted if this method of setting stock levels had been used at the base.~~ It was found that the "actual" fill rates differed by less than 5 per cent from the target fill rates that had been used in setting the stock levels. The most important result of the present study is the method which uses both repair characteristics and unit cost to analyze stock requirements across a large group of items. ~~The method looks promising not only as a means of setting base stock levels but also as a management control device for monitoring base supply effectiveness. 40 pp. incl. Illus.~~

(A)

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

PREFACE

This Memorandum describes and presents the initial results of a new technique for computing base stock levels for reparable spare parts. RAND initially started this effort in response to a request by Logistics Plans, Headquarters USAF. The study that resulted represents close cooperation and participation between representatives of RAND and the Air Force.

Essentially the information to follow is the text of a briefing presented at Headquarters, Air Force Logistics Command in Dayton on 24 January 1963. Although charts and data were prepared, the briefing time did not permit a technical presentation of all aspects of the stockage policy; therefore a brief review of some of the technical information appears in the Appendix of this Memorandum.

SUMMARY

This study is concerned with the problem of base level stockage policy for recoverable items. Resupply for a recoverable item can come from base parts repair, a requisition on depot stock, or some combination of the two depending on the complexity of the item, its malfunction, and the base repair capability. Base repair cycle time may differ significantly from depot resupply time so that the amount of stock required to achieve a given level of support protection at a base will be a function of item repair characteristics.

This report presents initial results of a method that considers repair characteristics and unit cost in the computation of a set of item stock levels that can achieve a given aggregate base fill rate* across all recoverable items with the least dollar investment in base recoverable inventory. Alternatively the model can maximize aggregate base fill rate for a specified investment in recoverable inventory.

Preliminary tests of the stockage model, while not conclusive, are very encouraging. The test consisted of taking demand data for a sample of 2802 recoverable items at Andrews Air Force Base and using the first six months of data as model input to compute item stock levels required to achieve a range of aggregate base fill rates. Demands for these items for the next six months were then compared with these stock levels in order to estimate the support performance that would have resulted if this method of setting stock levels had been used at the base. It was found that the "actual" fill rates

*Base fill rate is defined as the portion of total demands for support that can be met without delay from inventory on hand.

differed by less than 5 per cent from the target fill rates that had been used in setting the stock levels.

The most important result of the present study is the method which uses both repair characteristics and unit cost to analyze stock requirements across a large group of items. The method looks promising not only as a means of setting base stock levels but also as a management control device for monitoring base supply effectiveness.

CONTENTS

I. INTRODUCTION

This Memorandum is concerned with the problem of base stockage for recoverable spare parts. The function of base stock is to provide an acceptable level of support protection against the variability of demand and resupply cycle time. For recoverable items, base resupply can come from: base parts repair, a requisition on depot stocks, or some combination of the two, depending on such factors as base repair capability and the technological complexity of the line item. This Memorandum describes a method of setting item stock levels at a base in order to achieve a given aggregate fill rate for all recoverable items with the least dollar investment in base recoverable inventory. Base fill rate is defined as the portion of total demands for supply support that can be met without delay from inventory on hand.

Section II considers the base recoverable stockage problem and describes the general method used to analyze it. Section III contains a more detailed description of the method, and illustrates each step of its operation. Section IV examines the sensitivity of the stockage results to variations in key input assumptions. The Memorandum closes in Sec. V with a discussion of the implications of the study for Air Force stockage policy.

II. THE PROBLEM AND GENERAL METHOD OF STUDY

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

Present Air Force policy authorizes a base to establish a 30-day stockage objective for Cost Category I and II-R items. That is, a base is authorized an objective equal to average monthly issues.* We may note that this policy does not distinguish whether an item is base or depot reparable and yet, very clearly, the amount of support protection provided by a given amount of stock would be quite different if an item were base or depot reparable, and if base repair and base resupply cycle times differed significantly. For example, 30 days stock for an item that is 100 per cent base reparable with an average

*Air Force Manual 67-1, Vol. 2, Chapter II. The distribution system operates on the basis of stock control levels which includes the base stockage objective plus the number of days of stock required for normal resupply action.

Thirty days is the peacetime objective. If war reserve materiel is authorized it is treated as an additive and not normally considered available for use except in an emergency.

repair time of 3 days represents one level of protection; another level, offering considerably less support protection, is 30 days stock for an item that is 100 per cent depot reparable with a depot order and shipping time of 3 weeks.

Ideally, what we would like to have is a stockage policy that balanced the incremental cost of placing an additional unit of stock at base level with the value of increased operational effectiveness resulting from the change in stock level. Granting practical problems of implementing such an ideal system, it is significant to note that present Air Force policy does not even attempt to consider this trade-off when establishing base stock levels for Category I and II-R items. Current policy calls for a 30-day item stock level regardless of whether the unit cost is $50 or $50,000, or if the item is base or depot reparable.

Base stockage is too important an element of support to allow one to feel at ease with the above incongruities. Not only do base stocks represent a very large dollar investment, but also base supply effectiveness has a direct impact on the operational capabilities of the weapons being supported. Under such circumstances, it would seem desirable to have a policy that considers support protection and stockage cost in a more systematic manner.

The base stockage problem described above has been the subject of many previous studies. The innovation of the present method of study is a model developed to consider both item repair characteristics and cost in the computation of a particular set of item stock levels necessary to achieve a predicted aggregate base fill rate with the least dollar investment in recoverable item inventory (or alternatively, the model can maximize base fill rate for a specified investment). Figure 1 schematically depicts the operation of this model, which is run on an electronic computer. Given various inputs such as demand rates, base repair time distributions, base resupply time distributions and percentage issues base repaired, the model begins by simulating support operations in order to determine the relationship between item stock levels and item fill rate (i.e., the portion of total demands that could be filled from stock on hand.) Of course, when stock levels are actually established, we don't know item demand rates with certainty. Usually all we know is the number of demands that were experienced over some finite observation period and, therefore, given these observed demands we must predict the relation between stock levels and item fill rates. The second part of the model accomplishes this step and will be explained later. For the moment, consider the output from this second step, i.e., the relationship between item stock levels and predicted fill rate. This output is merged with additional information concerning the total number and unit costs of recoverable items at the base. The third part of the model analyzes the stockage requirements to determine the set of item stock levels that will achieve specified aggregate base fill rates with the least dollar investment in stock.

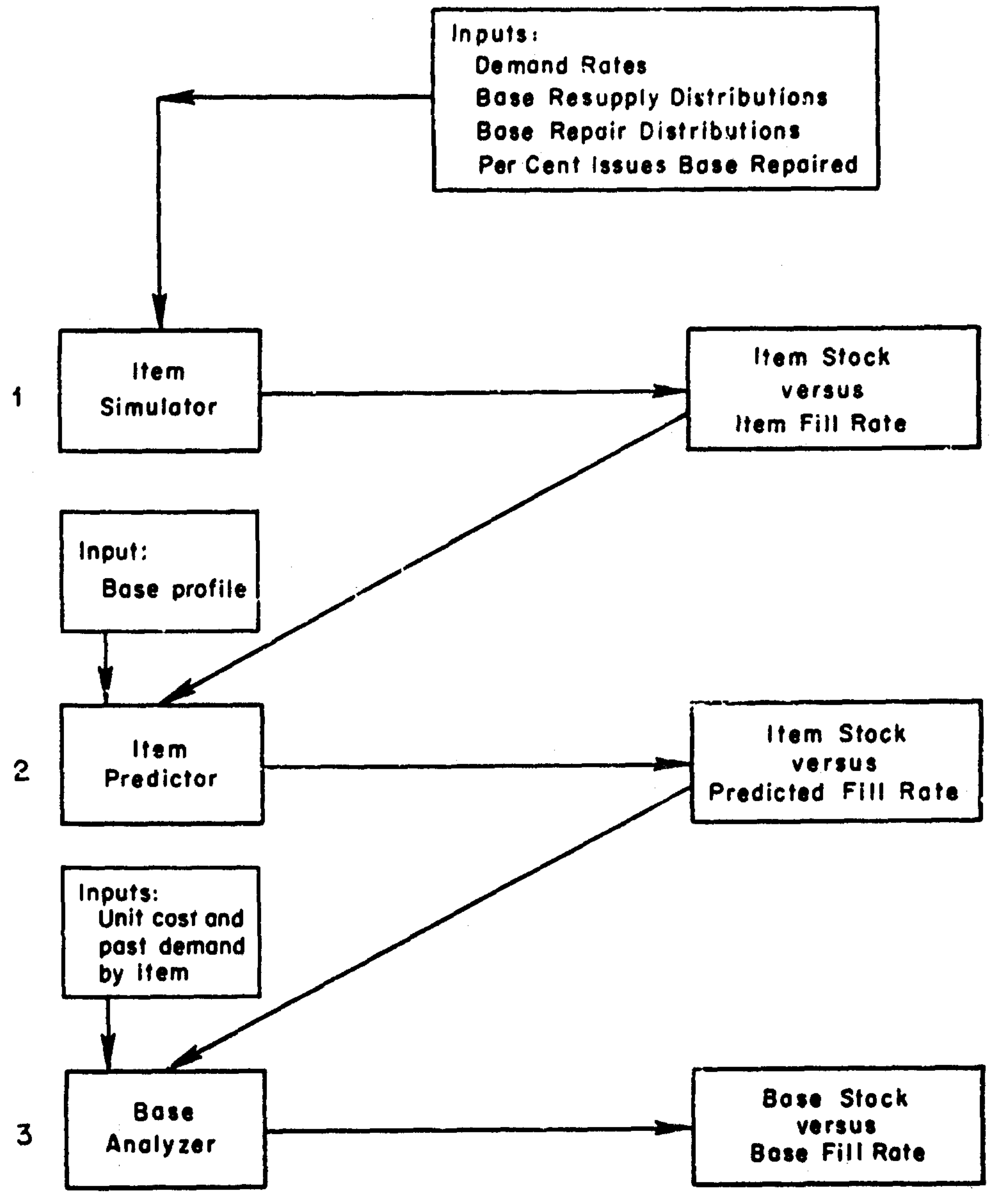

Fig. 1— Schematic Representation of Base Stockage Model

III. ILLUSTRATIONS OF MODEL OPERATIONS

So much for an overview of the stockage model; we will now review each step of its operations, and give illustrations of inputs and outputs. First, we will consider representative inputs run to date.

Figure 2 presents an approximation of the relationship between demand variability and mean demand rate at an Air Force base.* The standard deviation is a measure of variability such that about two-thirds of the months you would expect the demand for a part to fall within a range of the mean, plus or minus one standard deviation; about 95 per cent of the months you would expect the demand for a part to fall within a range of the mean, plus or minus 2 standard deviations, etc.** There are two points to make with regard to Fig. 2. First, since very few recoverable items have demand rates higher than five per month, the figure depicts the relevant range of demand rates. Second, the demand variability shown is very high, and particularly for lower demand rates could only be explained by long periods of zero demand followed by occasional high or peak demands. We will return to this problem of demand variability later on.

Next, we need to know something about resupply or base repair cycle times. Figure 3 summarizes two of the cases we have examined

---

*See W. H. McGlothlin and Eloise E. Bean, Application of the Bayes Technique to Spare-Parts Demand Prediction, RM-2701 (ASTIA No. AD 255161), The RAND Corporation, January 3, 1961, and Bernard Okun, Experimental Design, Test, and Evaluation of an F-100D Flyaway Kit, RM-2233 (ASTIA No. AD 210498), The RAND Corporation, October 31, 1958.

**Strictly speaking, this statement is true only if the variable is normally distributed.

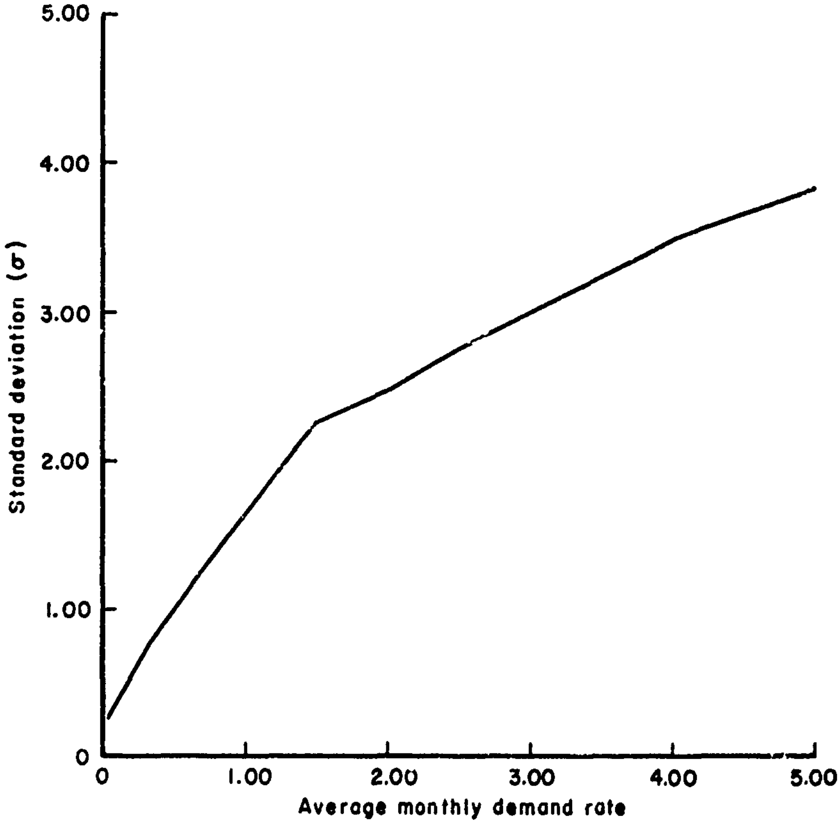

Fig. 2 — Pattern of standard demand variability (relation between mean demand and standard deviation)

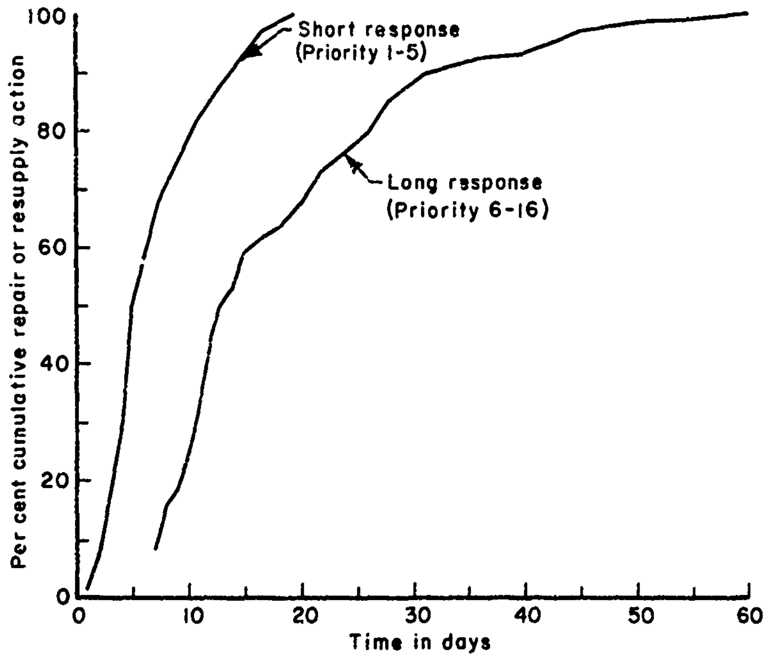

Fig. 3 — Base repair times and depot resupply times used in simulation experiments

to date. As we have used it here, the reader can view response as base repair cycle time, depot resupply cycle time or some combination of both. As a matter of interest, the two cases of response shown represent actual base resupply conditions that existed at sample Air Force bases in late 1959; short response represents resupply times for priority 1-5 requisitions; long response corresponds to priority 6-16 requisitions.* Although the Air Force has no formal system to measure base repair cycle times, sample studies indicate that typical routine base repair cycles are on the order of one week or less.**

A given set of the above inputs consisting of a demand rate and measure of variability, a base resupply distribution, a base repair cycle distribution, and the percentage of issues that could be repaired at base, were fed into the first part of the model and simulated over a period of 20,000 individual demands. This operation determined the relationship between item stock levels and item fill rates for each specific set of inputs being considered.***

Table 1 presents a sample output from the first portion of the model. The example shown is for an item with a mean monthly demand of 0.19, or an average of one demand in about five months. The item

---

*Annette Heuston, R. M. Paulson, A. H. Rosenthal, Base-Depot Requisitioning Pipeline Times, RM-2656, The RAND Corporation, November 1, 1960.

**Annette Weifenbach, The Base Repair Cycle for the F-102 Fire Control System, RM-2418 (ASTIA No. AD 231548), The RAND Corporation, July 21, 1959.

***A description of the operation of the simulation model appears in the Appendix.

is assumed to be 100 per cent depot recoverable with an average resupply time distribution from depot of 6.735 days and a standard deviation of 4.430 days.* The simulation breaks time into six-month intervals and computes that 42.1 per cent of the six-month intervals would pass without a single demand, 27.6 per cent would experience a demand of one, etc. The remaining part of the table shows the item fill rate that could be expected given the six-month demands and stock levels indicated.** Notice that one demand during a six-month period with a stock level of one provides a 0.999 fill rate --- not 1.000. This happens because of the small probability that the one unit of stock is in the repair pipeline when the demand occurs.

Table 1 depicts the relation between stock levels and fill rates for a given demand rate. Part of the problem is that we do not normally know the demand rate of an item but must estimate it on the basis of past experience. Typically, the amount of past experience will be limited, which means that a prediction based on such experience could provide a very misleading picture of the underlying failure rate. For example, in Table 1 we may note that in a specific six-month period it

---

*We simulated resupply variability by taking random draws from the empirical resupply distribution called "short response" in Fig. 3. Demand variability was simulated by taking random draws from a stuttering Poisson distribution. For the characteristics of this distribution see W. S. Jewell, "The Properties of Recurrent-Event Processes," Journal of the Operations Research Society of America, Vol. 8, No. 4, July-Aug. 1960, pp. 446-472.

**The expected fill rate for each stock level is the ratio of the expected number of fills divided by the expected number of demands. The expected number of demands is obtained by multiplying the frequency column with the issue column, i.e., (0.421 x 0) + (0.276 x 1) + . . . (0.001 x 8) = 1.130. The expected number of fills is obtained by weighting each of these entries by the appropriate fill rate, i.e., if the stock level is 1, the computation is (0.421 x 0 x 1.0) + (0.276 x 1 x 1.0) + (0.156 x 2 x 0.742) . . + (0.001 x 8 x 0.400) = 0.818. The expected fill rate for a stock level of 1 is 0.818/1.130 = 0.724.

would be possible to experience six demands for an item that had a true failure rate of only 0.19 per month. While there are only six chances in a thousand of this occurring, the point is it could happen, and if we were to base our estimate of demand on one particular period we could make a sizeable error.

Table 1

ITEM STOCK LEVELS AND CONDITIONAL ITEM FILL RATES FOR AN ITEM WITH A MEAN MONTHLY DEMAND OF 0.19

6-Month Demands	Frequency	Conditional Fill Rate Given Stock Level of					
		1	2	3	4	5	6
0	0.421	1.000	1.000	1.000	1.000	1.000	1.000
1	0.276	0.999	1.000	1.000	1.000	1.000	1.000
2	0.156	0.742	1.000	1.000	1.000	1.000	1.000
3	0.080	0.633	0.912	1.000	1.000	1.000	1.000
4	0.038	0.556	0.833	0.965	1.000	1.000	1.000
5	0.017	0.534	0.806	0.946	0.989	1.000	1.000
6	0.006	0.471	0.749	0.883	0.950	0.991	1.000
7	0.003	0.423	0.720	0.891	0.949	0.983	0.994
8	0.001	0.400	0.687	0.825	0.925	0.975	0.987

NOTE: The table's computations are based on the following simulated conditions: the item is 100 per cent depot recoverable, has a mean monthly demand of 0.19, and a mean resupply time of 6.735 days.

Besides knowing the previous number of demands for an item over a given period, we also know, from the nature of spare parts demands, that most items have low demand rates. As an example, consider Table 2, which shows the distribution of recoverable items by number of issues at Andrews Air Force Base for a six-month period ending in April, 1962.*

*Data from Andrews Air Force Base, Stock Balance and Consumption Report, October, 1961 to April, 1962.

While the actual distribution will of course vary with the length of the observation period, the point here is that every demand study we have seen shows a similar distribution, in that most items experience zero demands or one demand over the observation period and a very few have large demands. Further, the items that experienced zero demands over a period of six months cannot be assumed to have zero demand rates. For example, 23 per cent of the recoverable items that experienced zero demands at Andrews during the given 6-month period had one or more demands during the next 5.5-month period.*

Table 2

REPARABLE ITEMS BY NUMBER OF ISSUES[a]

No. of Issues	% of Reparable Items
0	73.7
1	11.3
2	4.8
3	2.8
4	2.0
5	1.2
6	0.7
7	0.6
8	0.5
9	0.3
10 and over	2.1

[a]Andrews Air Force Base, October, 1961 to April, 1962.

Since the above distribution has appeared with minor variations in all our studies of item demand,** it seems reasonable to postulate

*Data from Andrews Air Force Base, Stock Balance and Consumption Report, April 1962 to September 15, 1962.

**See M. Astrachan, B. Brown, and J. Houghten, A Comparative Study of Prediction Techniques, RM-2811, The RAND Corporation, December 1961, pp 33-35.

that real demand rates for spare parts must be distributed in a similar fashion. In the second part of our model we use this assumption in predicting the relation between stock levels and fill rates for future periods. Essentially what we do is compute fill rate as a conditional probability, taking into account both the number of demands experienced in the previous period by each line item and the distribution of demand rates across all items. That is, given the number of demands experienced with a given line item we first determine the probability that it came from a given demand rate. Then, given this demand rate we predict the relation between stock levels and fill rate for some future period.* An example of the output from this portion of the model is shown in Table 3.

In Table 3, each row shows the relation between stock level and forecasted fill rate for items with a given number of demands in a previous six-month period. It is interesting to note that a stock level of one for items with no demand during a previous six-month period would be expected to provide only a fill rate of 76 per cent. Given the short response conditions under which Table 2 was computed, this not only means that many of these items would have several demands during the next six-month period, but that many of these demands would be in multiples.

Given the information in Table 3, and knowing the issues and unit cost of each reparable item at a base, the model then goes into its third phase, calculating the particular set of item stock levels that

---

*This is the so-called Bayesian estimating procedure. A technical discussion of this estimating method appears in R. Schlaifer Probability and Statistics for Business Decisions, McGraw-Hill, New York, 1959, pp. 330-339.

will achieve a given base fill rate at the lowest stock investment. The third phase operates essentially as follows: the model assumes a zero stock level for every item at the base and then poses the question, where could one get the greatest amount of fill protection per dollar unit of stock? The model would assign a unit of stock to this item and ask where it could get the next largest amount of

Table 3

STOCK LEVEL AND PREDICTED FILL RATE
(Short Response)

Previous 6-Month Demand	Predicted Fill Rate with Stock Level of						
	1	2	3	4	5	6	7
0	0.757	0.941	0.985	0.996	0.999	1.000	1.000
1	0.671	0.892	0.964	0.987	0.995	0.998	1.000
2	0.615	0.852	0.942	0.976	0.990	0.995	0.998
3	0.557	0.806	0.914	0.960	0.982	0.991	
4	0.511	0.767	0.890	0.947	0.975	0.983	
5	0.479	0.740	0.873	0.938			
6	0.448	0.712	0.854				
7	0.429	0.693					
8	0.391						
9	0.384						

protection per dollar unit of stock, and so on.* Naturally, the method begins by stocking low-cost, high-demand reparables, and

*This method of allocating budget between items is similar to that described in earlier studies of flyaway kit design which assumed the constraint to be the total weight of the kit rather than investment. The principal difference is that in a flyaway kit, one is attempting to achieve maximum fill rate over a fixed period of time, e.g., a month, without resupply. The present model seeks to maximize the "steady state" fill rate over an indefinite period of time with resupply that can vary as an item characteristic. For a discussion of flying kit design see Okun, op. cit.

works down to low-demand, high-cost items. The allocation process continues until the target aggregate base fill rate has been obtained, at which time the computer prints out the item and summary results.

Table 4 is a sample of the item detail that the model produces. Such information is printed out for every reparable item at the base, with each column showing the particular set of item stock levels that will achieve the given aggregate target base fill rate at the least

Table 4

SAMPLE OF ITEM DETAIL
Part 3 of the Stockage Model

Stock No.	Unit Cost	Previous 6-Month Demand	Stock Level for an Aggregate Target Base Fill Rate of									
			0.700	0.750	0.800	0.850	0.900	0.925	0.950	0.975	0.990	0.995
12706057962	$14,627	11	0	0	0	0	1	2	3	4	6	7
12706059977	15,368	2	0	0	0	0	0	0	0	0	1	2
12706088004	420	1	0	1	1	2	2	2	3	3	4	5
12706088005	50	0	2	2	2	2	3	3	3	4	5	5
12706098444	63	4	5	5	6	6	7	7	8	9	10	10
12706098513	318	0	0	1	1	1	1	2	2	3	3	4
12706098519	207	0	1	1	1	1	2	2	3	3	4	4
12706100140	500	1	0	0	0	1	1	1	2	2	3	3

dollar investment in stock. At the end of the item detail the model produces a summary of results across all reparable items. Table 5 is an example of such a stockage summary.

The particular example shown in Table 5 happens to be for some 2802 XD-I, XD-II, and XB-I items at Andrews Air Force Base, and was computed on the basis of the April 1962 Stock Balance and Consumption Report. Table 5 not only shows the dollar investment required to achieve various target fill rates, but also shows the range of items that would have to be stocked in order to achieve them. It is interesting to note how rapidly investment and range increase as we attempt to achieve higher fill rates.

Table 5

SUMMARY OF RANGE AND COST OF BASE STOCK LEVELS FOR VARIOUS TARGET FILL RATES[a]
(2802 Reparable Items)

Target Base Fill Rate	Stock Investment ($ Million)	Range of Items Stocked (%)
0.700	0.631	57.1
0.750	0.833	65.3
0.800	1.159	74.2
0.850	1.586	81.0
0.900	2.276	88.2
0.925	2.972	92.7
0.950	3.820	96.0
0.975	5.009	98.7
0.990	6.994	99.7

[a]Andrews Air Force Base, April 1962.

IV. RESULTS WITH ALTERNATIVE INPUT ASSUMPTIONS

In this section we will examine three alternative input assumptions and how each affects the results of the model. These assumptions concern resupply times, demand variability, and the distribution of demand rates across items. We will study each of these aspects of the problem in turn.

RESUPPLY TIMES

In Fig. 3, we depicted two sets of resupply response which cover the interesting range of response in the ZI. Holding all other inputs constant, these two levels of response were run in the model to determine what impact they would have on the stock required to achieve a given fill rate. Figure 4 shows the summary results from these runs.

Several points can be made regarding Fig. 4. First, as we would expect, differences in resupply response have an impact on stockage requirements. In this case, cutting the average resupply time from 17 to 7 days made it possible to maintain a given base fill rate with approximately 80 per cent of the investment in stock.* Consequently, it follows that where significant differences exist between base repair cycle length and depot resupply time, it would pay to investigate the possibility of setting item stock levels as a function of per cent issues base repaired.

---

*Responsive resupply has further value in that when a stockout does occur it will have a shorter duration. A more complete discussion of this aspect of the support problem appears in J. W. Petersen, H. W. Nelson, R. M. Paulson, The Costs and Benefits of Responsive Support Operations, The RAND Corporation, RM-2871, October 1962.

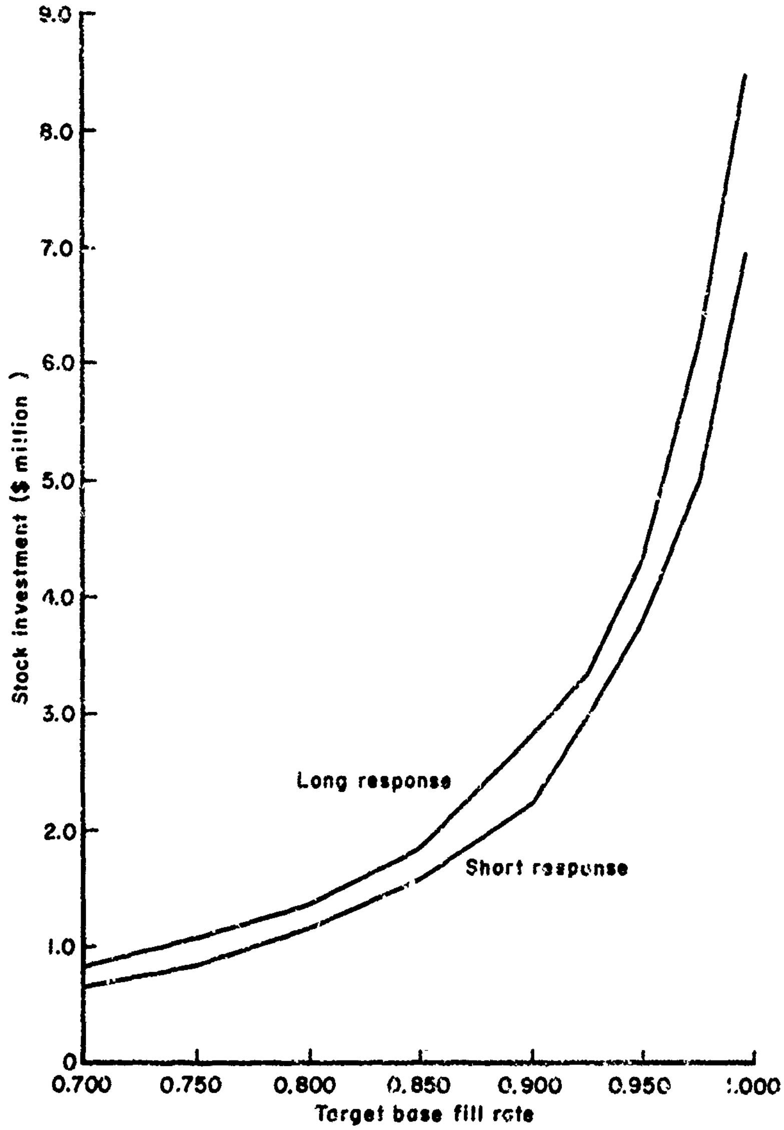

Fig. 4 — Relation between stock investment and base fill rate under conditions of short and long response

DEMAND VARIABILITY

So much for sensitivity to resupply response We next consider an alternative assumption of demand variability. The variability depicted in Fig. 2, and assumed in all of the runs discussed to this point, was based on prior studies of Air Force parts issue experience. The variability is much higher than one would expect from a random failure process, and, therefore, we have the problem of how to explain such variability. Several possible reasons come to mind. For example, the issue experience used in the computations might have been contaminated with recording errors of various sorts. Possibly, what we are treating as an issue may, in fact, represent nothing more than the transfer of stock from base supply to pre-issue or to Flyaway Kits. Or probably of greater importance, what is being treated as a normal issue may represent some non-recurring maintenance action, such as a modification program.

In order to examine the sensitivity of base stock requirements to assumptions of demand variability, we made an additional run of the model, using all the same inputs but reducing demand variability. Figure 5 depicts the relation between mean and standard deviation for the low variability case. For comparison purposes, we also show the demand variability relationship used in previous runs. The low variability case is essentially that of a Poisson distribution, whose standard deviation is equal to the square root of the mean. As a matter of interest, this is the sort of variability we would expect if we were dealing with a random failure process.

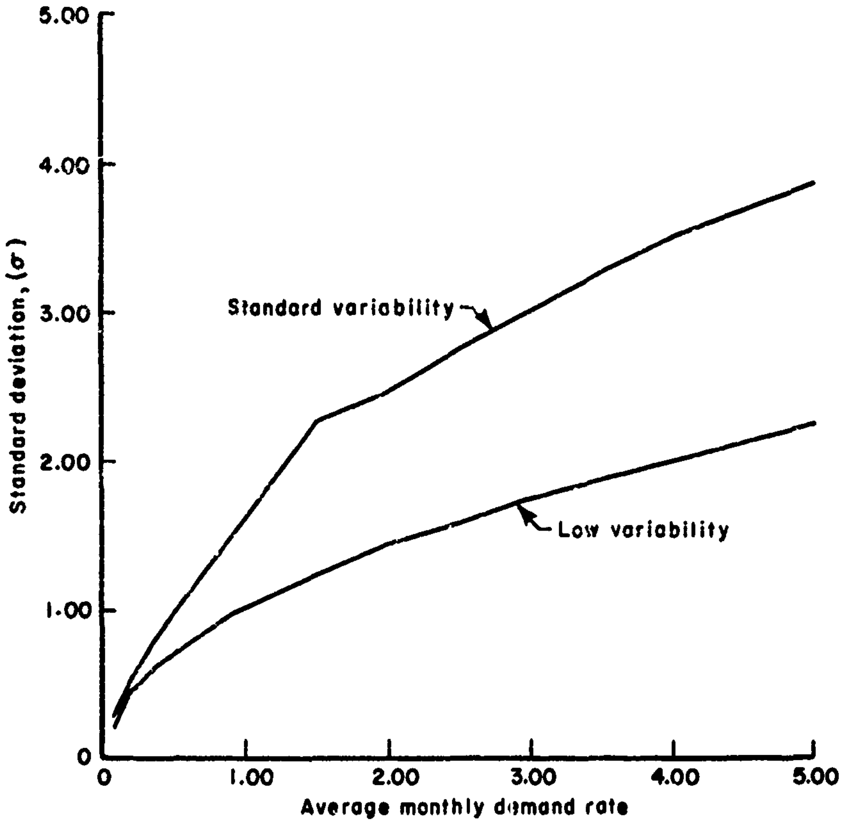

Fig. 5 — Pattern of standard and low demand variability (relation between mean demand and standard deviation)

Figure 6 shows the summary results of operating the model with the two estimates of demand variability. As expected, stockage requirements are very sensitive to assumptions of demand variability, the low variability run achieving a given fill rate with somewhat less than half the stock investment required by the previous run.

Since stockage requirements are so sensitive to demand variability, it is important to be able to test the reasonableness of any variability assumption. Fortunately, there are at least two ways of doing this. In the first place, our "standard variability" case is based on empirical estimates of variability found to exist for a sample of B-52 and Falcon missile spares.* A second check can be obtained by examining the transition probabilities computed analytically by the model. These transition probabilities are the probability that an item having μ demands during a six-month period will experience μ' demands during the next six-month period. By comparing these analytic results, which are very sensitive to variability assumptions, with the empirical data from Andrews, we could check the reasonableness of our variability assumption. On this basis we could conclude that the standard variability assumption was reasonable.

BASE PROFILE

Our third assumption concerns base profile. In the second section of the model it is necessary to input a set of weights (base profile) relevant to the proportion of total line items that fall in given demand classes. The initial runs were made with a profile developed on the basis of Andrews Air Force Base experience. This profile

*McGlothlin and Bean, *op. cit.*

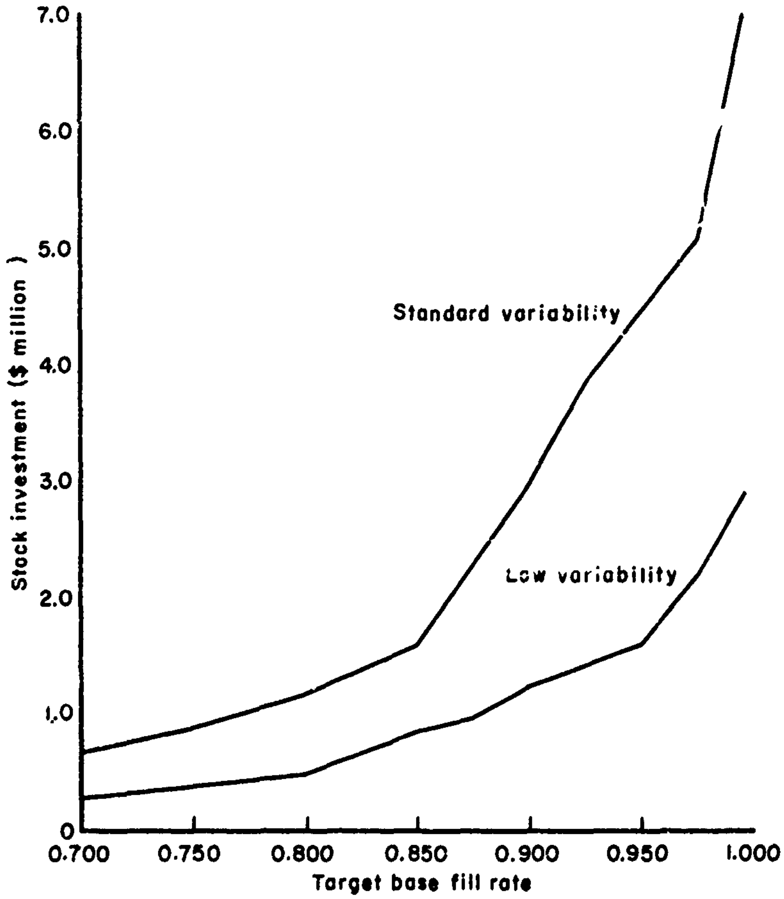

Fig. 6 — Base fill rate versus stock investment under assumptions of standard and low demand variability

assumed there were 11 underlying demand classes and that all reparables fell into one of these classes.* Profile 1 of Table 6 shows the base demand classes and the set of weights used to represent Andrews Air Force Base. By cumulative multiplication of demand rate times proportion of items in that class (weight) we get the average number of demands per item per month which was 0.26 for Andrews during the six-month period from October 1961 to April 1962. The depicted weights produce the same average.**

The interesting point is that the above weights influence the base stockage requirement, and, therefore, we were interested in determining how sensitive the results were to the specific weights used. In order to examine this question, we made an additional model run using the set of weights shown in Table 6 as Profile 2. This set of weights also results in an average demand per item per month of 0.26; however, except for the first demand class the distribution is completely rectangular. Profile 2 is not intended to be a reasonable representation of Andrews reparable items. This unlikely distribution was deliberately chosen in order to assess the impact of incorrect weights on stockage requirements.

---

*This is not strictly true since there were a few extreme items which would not reasonably fit into one of our 11 classes. As a matter of fact, any item with more than 30 issues during a six-month period was excluded from our computation, and it was assumed that stock levels for such items would be determined on an exception basis. There are few reparable items with demands above 30. In a six-month period only 26 items out of a sample of 2828 were excluded for this reason.

**The actual assignment of weights was done on the basis of a priori estimates. This is not as crucial as it might at first appear. Since the total is constrained at the low level of 0.26 demands per item per month, it is clear the bulk of the items have low demand rates which, of course, is in conformance with experience (see Table 2).

The change to Profile 2 did have an impact on stockage requirements, although the impact turned out to be less than we intuitively expected. Profile 2 results in the stockage of a smaller range of items (Fig. 7), and the dollar investment required to achieve a given target fill rate was less than the investment required under Profile 1.

Table 6

DEMAND CLASSES AND BASE PROFILES USED TO REPRESENT REPARABLE DEMANDS AT ANDREWS AIR FORCE BASE

Average Monthly Demand	% of Reparable Line Items	
	Profile 1	Profile 2
0.04	0.721	0.900
0.19	0.138	0.010
0.35	0.052	0.010
0.86	0.025	0.010
1.50	0.018	0.010
2.00	0.011	0.010
2.50	0.009	0.010
3.00	0.008	0.010
3.50	0.007	0.010
4.00	0.006	0.010
5.00	0.005	0.010
Average monthly demand/item	0.26	0.26

The interesting question, of course, is the impact incorrect weights would have on the resulting support performance. In order to examine this, we took the next six months of data from Andrews Air Force Base and played the demands against the item stock levels computed, using each of the two profiles. Table 7 summarizes these results.

There is some degradation in support performance in that "actual" or evaluated fill rates fell below target rates in every instance,

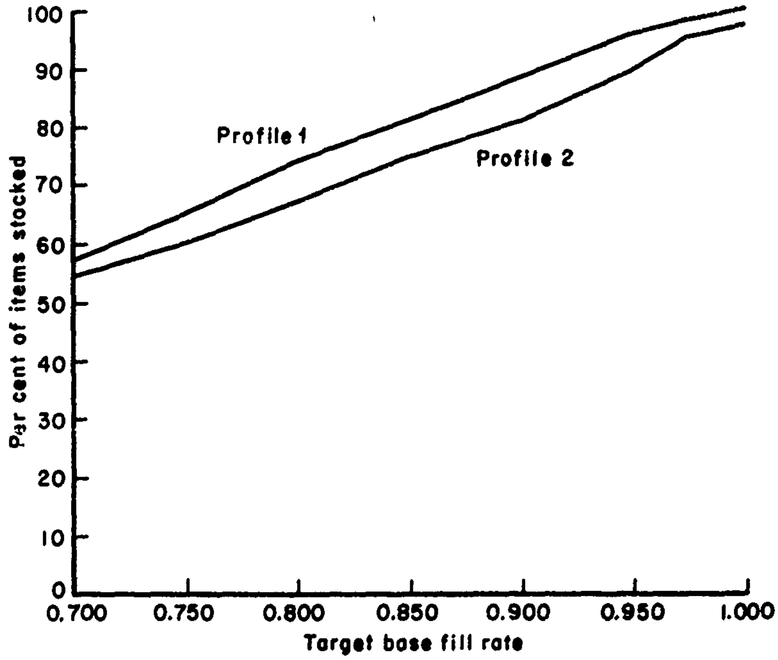

Fig. 7—Range of items versus base fill rate: Profiles 1 and 2

although the amount of degradation was smaller using Profile 1. The difference in support performance decreases as we move to the higher target fill rates, suggesting that the sensitivity to base profile is a function of the target fill rate. In general, however, bearing in mind the extreme or unlikely nature of Profile 2, we may conclude that the model is not overly sensitive to the range of reasonable assumptions one might make concerning base profile.

Table 7

ANDREWS AIR FORCE BASE SUPPORT PERFORMANCE
AS A FUNCTION OF BASE PROFILE
(Figures in Per Cent)

Target Fill Rate	Evaluated Fill Rate	
	Profile 1	Profile 2
70.0	65.0	57.0
75.0	71.1	61.0
80.0	76.6	70.0
85.0	82.9	75.0
90.0	87.6	82.0
92.5	91.4	86.5
95.0	94.1	90.0
97.5	96.3	93.2
99.0	98.3	96.5

V. IMPLICATIONS OF THE STUDY

By far the most important result of the present study is the method developed to analyze stock requirements across a large group of items. The method looks interesting not only as a means of setting base stock levels but also as a management control device for monitoring base supply effectiveness. The real question is how well would the technique work under conditions of actual support operations? We know that the internal logic of the model is all right, that is, that the model is doing what we say it is. The real question is whether or not the input assumptions are reasonable. Undoubtedly in the real world, there are frictions and problems such that we would not expect model results and actual base results to exactly coincide. The question is how close can the model come to predicting aggregate base fill rates for reparable items? If the technique can do even reasonably well, then it appears to have considerable potential for Air Force support operations.

The preliminary tests of the model, while not conclusive, are certainly encouraging. As mentioned, we had one year's data from Andrews Air Force Base covering issues of reparable items from October 1961 to October 1962. Using the first six months of demand experience to compute stock levels for a range of target fill rates in the manner described in the previous section, we then played the second six months' demands against these levels to estimate what sort of support performance we would have achieved.* The results are shown in Table 8,

*The method of estimation is covered in the discussion of evaluated base fill rate, p.40.

along with estimates of the fill rate that would have been obtained at the base if all items had been at authorized levels. It is rather interesting to compare the model results with performance and dollar investment of the authorized levels at Andrews.

While these results are encouraging, it would be highly desirable to make a more rigorous test of the stockage method. Air Staff at

Table 8

ANALYSIS OF ANDREWS AIR FORCE BASE SUPPORT PERFORMANCE UNDER ASSUMED STOCKAGE CONDITIONS (2802 Reparable Items)

% Target Fill Rate	% "Actual" Fill Rate	Stock Investment ($ Million)
70.0	65.0	0.631
75.0	71.1	0.833
80.0	76.6	1.159
85.0	82.9	1.586
90.0	87.6	2.276
92.5	91.4	2.972
95.0	94.1	3.820
97.5	96.3	5.009
99.0	98.3	6.994
Authorized Base Levels	61.3	2.521

Headquarters USAF has concurred in this desire, and has ordered the establishment of a task group to service test the technique at one or more sample bases. The data collection program for this test has been initiated, and it is the intent of the task group to make the stock level computations for the test base in summer or early fall of this year.

A more complete evaluation of the reparable item stockage technique just described must await service testing results. In the meantime, however, we can note several additional implications of the present stockage study.

Because any stockage policy must operate implicitly or explicitly with some assumption of demand variability, it is important to note how sensitive stock requirements are in this area. Unfortunately, there is a lack of data with which to estimate base demand variability, and more important, we have little understanding of what causes such extreme fluctuations in demand. Perhaps a large part of variability is simply erroneous reporting. If so, improved reporting quality could produce large reductions in stock investment. Clearly, we need to know more about the nature and causes of fluctuations in demand. Without such information, we are not in a position to decide what portion of demand variability should legitimately be covered by base stockage policy. It is hardly necessary to point out that this is a basic question, which from a management point of view is as important as choosing a specific target base fill rate.

As noted in previous studies, more responsive resupply does reduce the amount of stock required to achieve a given fill rate. From this it follows that in instances where there is a significant difference between base repair cycle length and depot resupply time, it would be beneficial to establish stock levels for an item as a function of percentage issues base repaired.

No stockage policy can eliminate stockouts; expedited deliveries will still be required from time to time due to the vagaries of demand.

But the achievement of a specified target base fill rate with minimum stock investment will cause the costs of expediting to be incurred in resupplying the high-unit-cost items, which represents a more efficient use of support resources.

APPENDIX

Three distinct parts compose the base stockage model depicted in Fig. 1; these are an Item Simulator, Item Predictor and Base Analyzer.

ITEM SIMULATOR

We assume that each stock item belongs to a process k where $k = 1, 2, 3, \ldots, K$. Process k is a stuttering Poisson demand distribution with mean μ_k, variance σ_k^2, base repair distribution b_k, base resupply distribution d_k, and per cent of base repair r_k.

The item simulator computes:

$p_i(u/k)$ = pr {u demands during period i/item is in process k}.

$F_i(k, u, s)$ = item fill rate during period i if item is in process k, has u demands and a stock level s.

Table 1 provides an example of the Item Simulator output. The process k in this table has a mean demand per month of 0.19, a standard deviation of 0.565, a base resupply distribution b_k (not shown in Table 1) with a mean of 6.735 days and a standard deviation of 4.430 days, and issues that are 100 per cent depot recoverable. The frequency column shows $p_i(u/k)$, with i denoting a six-month period. The main body of the table gives $F_i(k, u, s)$. In representing Andrews Air Force Base, we chose 11 processes with the mean monthly demands shown in Table 6. Each process results in a page of output like Table 1.

Let us briefly describe the logic of the simulator. For any process k, the simulator computes a system response distribution by multiplying the base repair distribution b_k by the per cent base repair r_k and adding to the base resupply distribution d_k multiplied by $(1 - r_k)$. The simulator samples the time between demands for some 10,000 or 20,000 demands from a stuttering Poisson demand distribution with mean μ_k and variance σ_k^2. Whenever a demand occurs, a reparable is turned in and the response distribution is sampled to determine when the reparable becomes serviceable (repair or resupply). We assume stock levels of 1, 2, ..., n and compute the number of fills/number of demands for each six-month period. The percentage of six-month periods in which u demands occur is $p_1(u/k)$. And $F_1(k, u, s)$ is the number of demands that could be filled from available stock, divided by the number of demands in all six-month periods during which u demands occur, assuming a stock level s. In computing the number of fills,we observe the rule that a unit of stock being returned to the shelf in a serviceable condition will always be used to satisfy a back order if one exists, i.e., no demand is lost.

ITEM PREDICTOR

The Item Predictor is an analytic calculation used to forecast the fill rate for a period of time during which there is a stock level s based on the knowledge that u demands occurred during a previous period. In addition to the output from the Item Simulator we require "a priori" estimates, $w(k)$, that an item is in process k. This profile, examples of which appear in Table 6, does not assign an item to a specific process but rather gives the percentage of all items that belong to each of the k processes.

Thus

$$w(k) = pr \quad \{\text{item is in process } k\}.$$

We want to estimate the probability that an item experiencing u demands during a fixed time period i is in process k. A simple application of Bayes Theorem will provide this estimate.

$$q_i(k/u) = \text{pr}\ \{\text{item is in process } k/u \text{ demands during period } i\},$$

(1)
$$= \frac{p_i\ (u/k)\ w(k)}{\sum_{k=1}^{K} p_i(u/k)\ w(k)}\ .$$

We can now calculate the "actual" or evaluated fill rate that would have been achieved during period i if the item had a stock level s and u demands. As noted earlier, if empirical fill rates are available, they can be compared with or substituted for the following calculation.

$$EFR_i(u,s) = \text{evaluated item fill rate during period } i \text{ if item has } u \text{ demands and a stock level } s,$$

(2)
$$= \sum_{k=1}^{K} F_i(k,\ u,\ s)\ q_i(k/u).$$

In order to predict fill rate for some future period of time, we provide the following analytic calculation of the transition probabilities. Empirical data are also useful at this step.

$$T(u'/u) = \text{pr}\ \{\text{item has } u' \text{ demands in next period/ } u \text{ demands in previous period}\},$$

(3)
$$\sum_{k=1}^{K} p_2(u'|k) \quad q_1(k|u).$$

The subscripts show that the q's are calculated from the first time period and the p's from the second. If the two time periods have different lengths, the subscripts are required.

Finally, we calculate the predicted fill rate during a future time period.

PFR(u, s) = predicted item fill rate for next period if item had u demands during the previous period and a stock level s,

$$= \frac{\sum_{u'=1}^{\infty} u' \ EFR_2(s, u') \ T(u'/u)}{\sum_{u'=1}^{\infty} u' \ T(u'/u)} . \tag{4}$$

The numerator gives the expected number of fills and the denominator expresses the expected number of demands during the future period. Table 3 shows output from the Item Predictor. In this table we have consolidated K pages of output from the Item Simulator into one page of output from the Item Predictor, in which the k processes are now implicit rather than explicit.

BASE ANALYZER

The base analyzer employs the item characteristics of cost and demand during a time period to make item stockage decisions leading to a target base fill rate with minimum investment (or equivalently, maximum base fill rate for a specified investment).

Let

c = item cost;

$\bar{u}(u)$ = expected demand for an item during the next period if item had u demands during the previous period,

$$= \sum_{u'=1}^{\infty} u' T(u'/u);$$

β_F = dollar value of a fill ;

β_I = per cent holding cost (interest, warehousing, obsolescence).

β_F is simply the stockout cost, usually unknown. Whenever we assign positive stock, we impute a positive value of β_F. Similarly β_I has some finite positive value. We shall assume β_F and β_I are constant over all items. Then

$\pi(s)$ = expected gain (loss) when s units are stocked for an item which had u demands during previous period,

$$= \beta_F \, \bar{u}(u) \, \text{PFR}(u, s) - \beta_I \, sc. \tag{5}$$

Maximizing $\pi(s)$ we obtain

$$\frac{\Delta\pi}{\Delta s} = 0 = \beta_F \, \bar{u}(u) \left[\text{PFR}(u, s+1) - \text{PFR}(u, s) \right] - \beta_I c;$$

$$\frac{\beta_I}{\beta_F} = \lambda = \frac{\bar{u}(u) \left[\text{PFR}(u, s+1) - \text{PFR}(u, s) \right]}{c} . \tag{6}$$

Supplied with an estimate of $\beta_I/\beta_F = \lambda$, we could easily determine the optimum stockage policy for each item. But λ is not specified when we seek to obtain some target value for the predicted base fill rate at minimum investment. The predicted base fill rate is defined as:

$PBFR(s_1, s_2, \ldots s_N)$ = Predicted Base Fill Rate when item stock levels are $s_1, s_2 \ldots s_N$,

(7) $$= \frac{\sum_{t=1}^{N} \bar{u}(u_t)\; PFR(u_t, s_t)}{\sum_{t=1}^{N} \bar{u}(u_t)} .$$

Let us examine Eq. (6). The quantity λ expresses marginal fills per dollar investment in stock. By item, $\bar{u}(u)$ and c are constants, and under the assumption that $\left[PFR(u, s+1) - PFR(u, s)\right]$ is monotone decreasing, it can be seen that for each value of s there is a unique value of λ. As s increases, λ decreases. Of course, the right hand side of Eq. (6) is discrete valued so that λ is also discrete.

The procedure for allocating stock is now obvious. Initially we set all stock levels $s_1, s_2, \ldots s_N$ equal to zero. Then by employing Eq. (6) for each item, the computer allocates a unit of stock to that item giving the largest marginal fill per dollar invested (largest value of λ). Then the predicted base fill rate is calculated from Eq. (7) to determine whether the target fill rate has been achieved. The process is repeated, with a unit of stock being allocated to the item giving the next largest marginal fill per dollar invested and so on until the target fill rate is exceeded. At this point a corresponding value λ^* has been defined. Since

$$\lambda = \frac{\beta_I}{\beta_F} ,$$

we know the ratio of holding cost to the value of a fill implied by the specification of our target fill rate. The investment in stock is simply

$$\sum_{t=1}^{N} c_t s_t ,$$

and by the allocation procedure we know that this is a minimum.

Several technical comments are in order. The problem is actually a non-linear integer program. Due to the discreteness of the problem, the solution described above will overshoot the target fill rate by a small amount. Often trial and error methods will make it possible to achieve a fill rate closer to the target at a slightly lower cost. In the usual case of several hundred items, however, the overshoot is extremely small, making the further refinement unwarranted.

An alternative method of deriving Eq. (6) is to take the first differences in Eq. (7), subject to the constraint on investment. In this formulation it is easy to see that maximum fill rate for a given investment is equivalent to minimum investment for a specified fill rate.

We have assumed that the value of a fill is the same for all reparables. As a first approximation over the class of reparables this seems reasonable, but it is easy to modify this assumption. Suppose the value of a fill on one item is taken to be ten times the value of a fill on any other. This means the β_F for that item is ten times larger, or equivalently that λ in Eq. (6) is one-tenth of its previous value. But this same result can be achieved by dividing the item cost c in the right hand side of the equation by 10. On a system basis we thus obtain "shadow prices" by dividing unit cost by the value of a fill.

In the cases run to date, the marginal fill rate

$$\left[PFR\ (u,\ s+1)\ -PFR\ (u,\ s)\right]$$

has always been a monotone decreasing function of s. However, it is possible to construct examples violating this condition for the smallest values of s so that Eq. (6) does not have a unique solution. For s sufficiently large, the monotonicity requirement must be met because the marginal fill rate approaches zero. As an illustration consider Fig. 8. Note that the marginal fill rate $\left[PFR\ (u,\ 2)\ -PFR\ (u,\ 1)\right]$ is greater than $\left[PFR\ (u,\ 1)\ -PFR\ (u,\ 0)\right]$. Consider the quantity PFR (u, s)/s representing average fill rate per unit of stock. The slope of the secant through the origin gives the value of PFR (u, s)/s which is maximized by the dotted line at s = 3. Since the average fill rate per unit of stock is less at s = 1, 2, these points are inferior. Subject to the target base fill rate that is specified, our item decision will be to stock 0, 3, or more units.

We can accomplish the desired result by altering the predicted fill rates. Suppose that s^* gives the largest average fill rate per unit of stock, PFR (u, s^*)/s^*. (In Fig. 8 the value of s^* was 3.) Replace PFR (u, 1) by PFR (u, s^*)/s^*, PFR (u, 2) by 2 PFR (u, s^*)/s^*, . . . and PFR (u, s^* -1) by (s^* -1) PFR (u, s^*)/s^*. In each case the replacement number, taken from the secant line, is at least as large as the replaced number. We have equated the marginal fill rates for stock levels 0 through s^* -1. Therefore, the allocation procedure of the Base Analyzer will result in a stockage of 0 units or at least s^* (unless the target base fill rate is exceeded between these steps). The non-uniqueness is avoided and the correct item stock decisions result.

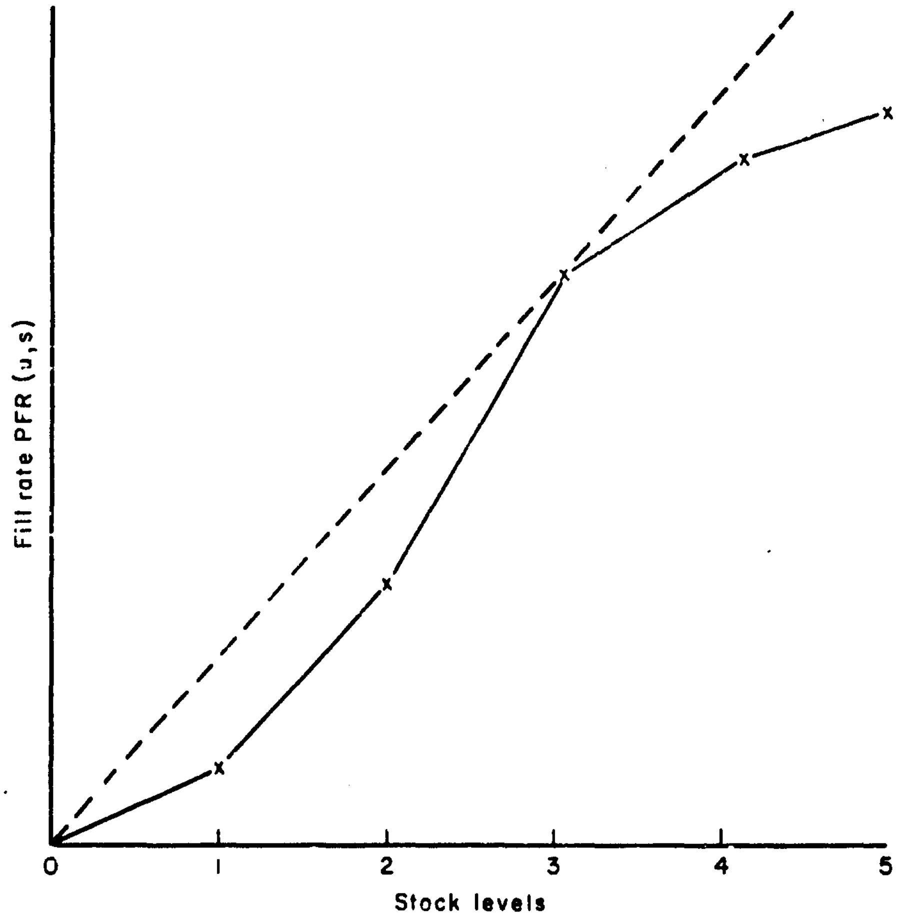

Fig. 8—Example with a non-monotonic marginal fill rate

Finally we define an evaluated base fill rate so that the "actual" fill rate during the second period at Andrews Air Force Base, established on the item stock levels computed by the Base Analyzer and the actual demand during the second period, could be compared with the target or predicted base fill rate. In a manner analogous to Eq. (7) we have

$$\text{EBFR}\ (s_1, s_2 \ldots s_N) = \text{evaluated base fill rate when item stock levels are } s_1, s_2, \ldots s_N,$$

$$= \frac{\sum_{t=1}^{N} u_t \ \text{EFR}\ (u_t, s_t)}{\sum_{t=1}^{N} u_t} .$$