

MEMORANDUM

RM-4054-PR

APRIL 1964

SYSTEMS ANALYSIS  
AND SUPPLY MANAGEMENT

G. J. Feeney and C. C. Sherbrooke


PREPARED FOR:

UNITED STATES AIR FORCE PROJECT RAND

---

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

---



MEMORANDUM

RM-4054-PR

APRIL 1964

SYSTEMS ANALYSIS  
AND SUPPLY MANAGEMENT

G. J. Feeney and C. C. Sherbrooke

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.

DDC AVAILABILITY NOTICE

Qualified requesters may obtain copies of this report from the Defense Documentation Center (DDC).

---

The RAND Corporation

1700 MAIN ST. • SANTA MONICA • CALIFORNIA • 90406



PREFACE

This Memorandum describes some new ideas in the application of systems analysis to supply management. The language is nontechnical, following the text of a briefing presented to the Defense Supply Agency on January 20, 1964 and to Headquarters USAF on January 21, 1964. Technical memoranda with complete mathematical derivations are in preparation.

The study was initiated in the summer of 1962 in response to a request from Logistics Plans, Headquarters USAF. At that time the focus of attention was base stockage policy for recoverable spare parts. A new approach was developed for that problem, tested with Andrews AFB data, and documented in RM-3644-PR.\* Representatives of Logistics Plans, Headquarters Air Defense Command, and Air Force Logistics Command are working in close cooperation with RAND to plan limited implementation of the new technique.

The salient feature of the new system approach, however, is its applicability to a wide variety of supply management problems -- not merely base stockage of recoverable spare parts. The objective in this Memorandum is to show that systems analysis can become a powerful, new approach to a broad spectrum of problems.

---

\* G. J. Feeney, J. W. Petersen, and C. C. Sherbrooke, An Aggregate Base Stockage Policy for Recoverable Spare Parts, The RAND Corporation, RM-3644-PR, May 1963.



## SUMMARY

This study maintains that supply management problems should be approached from the viewpoint of system analysis. By this we mean that individual item decisions should be based on analysis of system behavior. Two broad areas of research are discussed: stock policy and demand analysis.

### STOCK POLICY

The traditional approach to stock policy determines the most economical item decisions under the assumption that stockout cost and holding cost are known. This approach has two serious flaws. First, stockout cost and holding cost are extremely difficult to measure. Supply policy is governed by arbitrary estimates which are not easily defended. Second, these methods tell us nothing about what will happen to the system as a whole; particularly, how much investment will be required and what level of supply performance will be realized.

This Memorandum describes a new, system-oriented approach to stock policy. Instead of attempting to "measure" the ratio of stockout cost to holding cost, we treat this ratio as a management control variable. A special technique has been developed to display the system performance and system investment which result from alternative settings of the control ratio. It can be shown that these alternatives are "efficient." Each alternative provides maximum system performance for a given level of system investment.

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

## DEMAND ANALYSIS

The traditional approach to demand analysis calculates an item's issue rate (demand observed over some past period divided by the length of the period) and assumes that future demand will be some random variation around this level. Such an approach is adequate if the item has relatively high demand, if a relatively long history is available, and if the history is relevant. But many of the most important items, particularly high-cost, low-demand spare parts, fail to meet one or more of these three essential requirements.

A new approach to demand analysis, based on a mathematical technique called Bayesian inference, is described in this Memorandum. This approach exploits the surprising fact that we can increase our knowledge of an item by analyzing the behavior of the other items in the system. Instead of trying to estimate the item's true average demand, this approach estimates the probability that the item's average demand is at one of several levels. We think this is a much more realistic way of characterizing what we can learn from an item's demand history. These probabilities can now be used to appraise the true risks and potential payoff for various stock levels. Data-processing procedures can be designed to use information of this kind with little or no increase in processing complexity.

Bayesian demand analysis has two immediate implications for improved supply management. First, because it wrings maximum relevant information from available demand data, this approach promises substantial improvement in the efficiency of stockage decisions. Specifically, by using this kind of analysis we should be less liable to overestimate demand and buy too much because of a random surge in demand, and less liable to underestimate demand and buy too little because of a random decline in demand. Second, because the approach can be applied to any time period, it is extremely flexible. It should be possible to eliminate many of the policy problems now created by items that do not meet the requirements of daily-issue-rate computation: items with low demand, and items with erratic demand patterns. In the framework of Bayesian demand analysis, policies can be developed that prescribe action unambiguously.



ACKNOWLEDGMENTS

Jim Petersen's contagious enthusiasm and exasperating skepticism were an essential stimulus to this study. Harry Campbell's frequent counsel was of great value.

General Paul L. Barton's interest and encouragement, and the suggestions of Colonels Vernon J. Taylor, Herbert J. Gavin, Wendell G. McKoy, Ira M. Kay, and of Major Jack R. Smith, provided invaluable support.



CONTENTS

PREFACE .....	iii
SUMMARY .....	v
ACKNOWLEDGMENTS .....	vii
Section	
I. INTRODUCTION .....	1
II. STOCK POLICY .....	2
Item Approach .....	2
System Approach .....	5
Management Implications .....	7
Extensions .....	9
III. DEMAND ANALYSIS .....	10
System Approach .....	10
Management Implications .....	14
Extensions .....	15
IV. CONCLUSIONS .....	16



## I. INTRODUCTION

Why should supply management be concerned about system analysis? Because supply operations are complex systems involving interrelated decisions on thousands of line items. In fact, supply management is system management in the fullest sense. System analysis, the study of the relationship between the behavior of a system and the behavior of its parts, presents significant new opportunities for profitable innovation in supply management.

For several years, The RAND Corporation has investigated advanced applications of systems analysis to supply management, working closely with several interested groups in the Air Force. This Memorandum will summarize our major research findings in two broad areas, stock policy and demand analysis. This discussion has two objectives: to present some new ideas that have come out of the research, and to illustrate, through these ideas, the general thesis that systems analysis can be useful in many areas of supply management.

The emphasis here is on basic concepts and general implications for supply management. Theoretical details and specific applications will be described by the authors in subsequent RAND Memoranda.

## II. STOCK POLICY

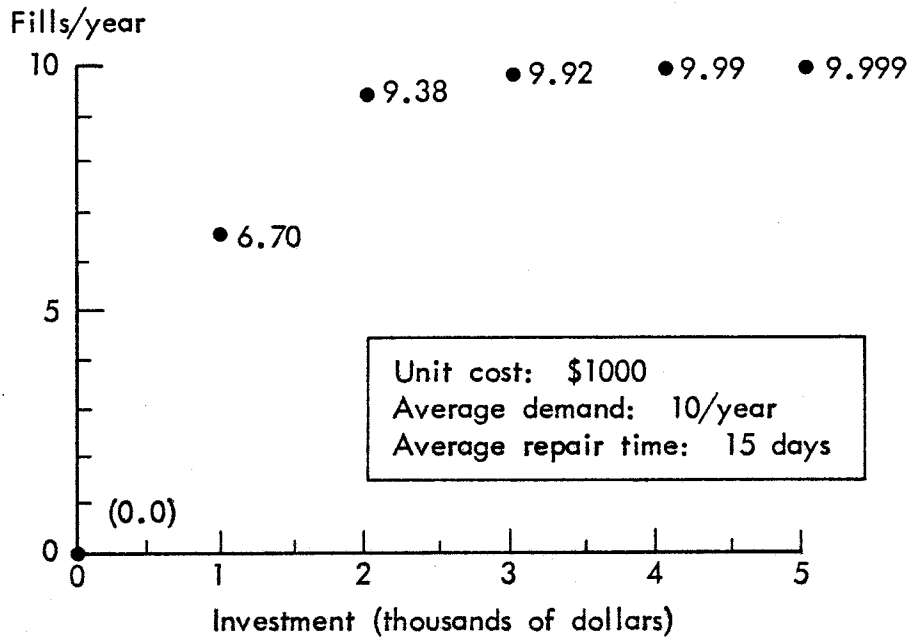
Supply Management is confounded by a paradox. On the one hand, its problems are system problems. Determining realistic performance standards, maintaining budget compliance, developing objective requirements justifications, and attaining balanced adjustment to unexpected changes in customer demand or resource availability are all problems that require analysis and control of a supply system as a whole. On the other hand, stockage policies must be stated and stockage decisions must ultimately be made at the item level. In order to get at the system we have to work through the items. This is like trying to move a mountain with tweezers, one grain at a time.

It is tempting to dismiss the problem by assuming that if decisions are diligently made on each item the system will take care of itself. This might be called the item approach to stock policy. We believe its basic assumption is incorrect, and that this approach cannot cope with the fundamental problems of supply management. On the contrary, the system should control the items, not the items the system. This section discusses how a system-oriented stock policy might be achieved.

### ITEM APPROACH

It is useful to begin with a careful review of the item approach to stock policy. For clarity we will use a specific example: a recoverable aircraft spare part with unit cost of \$1000, average demand of 10 units a year, and average repair time of 15 days. The item's supply performance can be measured in terms of successful fills -- the number of times a demand can be met immediately, with no repair delay. How many units of this item should be stocked at a base to back up the repair process?

Under the common assumption that demand uncertainties are governed by a Poisson probability distribution, straightforward application of inventory control theory will tell us the average number of successful fills that can be expected for various levels of investment in backup stock. This relationship is summarized in Fig. 1. Thus, if no spares are provided, the investment is 0, and there will be no successful fills,



Value	0	670	938	992	999	999.9	(Fill worth: \$100)
Cost	0	200	400	600	800	1000.0	(Inv. cost: .20)
Net	0	470	538	392	199	-0.1	

Value	0	3350	4690	4960	4995	4999.5	(Fill worth: \$500)
Cost	0	100	200	300	400	500.0	(Inv. cost: .10)
Net	0	3250	4490	4660	4595	4499.5	

Fig. 1—Item performance/item investment

since every demand will have to wait through the repair delay. A \$1000 investment -- one unit of stock -- buys an average of 6.7 fills per year; a \$2000 investment, 9.38 fills; and a \$3000 investment, 9.92 fills. A \$5000 investment provides virtually complete protection: an average of approximately 9.999 fills per year out of an average of 10 demands.

This analysis determines what stock to carry if we also specify two critical parameters: what a fill is worth, and how much to charge per dollar invested in stock (cost of holding inventory). The sample calculations at the bottom of Fig. 1 were made under two different assumptions: first, that each fill is worth \$100 and investment costs 20 per cent a year; second, that each fill is worth \$500 and investment costs 10 per cent a year. In the first case, two units of stock are best because they provide the largest net difference between the value of fills and the cost of investment. In the second case, because the assumed value of a fill is higher and the assumed cost of investment is lower, a stock of three units is best.

Thus the item approach measures the relationship between item investment and item performance and provides a basis for selecting what seems to be the most profitable stock level. Why isn't this the complete answer to the stockage problem?

Serious objections to this approach can be raised at two levels. First, it has often been observed that it is extremely difficult to determine the critical parameters -- fill value and investment cost. How much is it worth to prevent a shortage? How much should we charge for money tied up in stock? In most real situations we can do little better than make an arbitrary guess. A second and much more fundamental objection is that the item approach tells us nothing about what is going to happen to the system. It is all very well to say that the "proper" charge for invested funds is 15 per cent per year, but we are in trouble if the resulting item decisions lead to system requirements that exceed available budgets. Similarly, a special study that decrees a shortage costs \$357 is of little comfort if the resulting system performance causes widespread customer dissatisfaction.



### SYSTEM APPROACH

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

In order to show what this technique is and how it works, let us return to the sample item discussed earlier. Notice that the relative profitability of each alternative depends only on the ratio of fill value to investment cost. Specifically, if we determine how much investment each fill is worth we can immediately compare each of the stockage alternatives and find which is most profitable. In the first case, for example, instead of assuming that a fill is worth \$100 and investment costs 20 per cent per year, it would have been sufficient to assume that one fill a year is worth a \$500 investment. A simple comparison shows that the same alternative, two units of stock, provides the greatest net gain. Similarly, in the second case it would have been sufficient to assume that each fill is worth \$5000 and thus that three units of stock provide the most economical alternative.

Now suppose, instead of attempting to measure the "true" value of the fill/investment ratio, we treat the ratio as a management-control variable -- a "control ratio." We can find out what would happen to the supply system as a whole if a particular ratio is applied to the stock decisions of every item in the system. Figure 2 shows three examples of such an analysis for a system composed of 728 recoverable spare parts from the F-101 interceptor at Oxnard Air Force Base. Thus, if we assume that each fill is worth a \$1000 investment and apply this control ratio to stock decisions on each item, the system result is a fill rate of 66 per cent (that is, on the average 66 per cent of the demands on the system will be met on time) and a system investment of \$270,000. If the ratio is increased to \$3000, the system fill rate becomes 85 per cent and system investment is \$800,000. A \$10,000 ratio produces a system fill rate of 94 per cent and requires a system investment of \$1,580,000.




Fig. 2 — System performance/system investment

It is possible to show that each of these system alternatives has an extremely important property. In economic terminology, they are efficient alternatives. This means that each alternative provides maximum system performance for a given level of system investment or, equivalently, each provides minimum investment for a given level of performance.

In this case, for example, if \$270,000 is all you have to invest, there is no way to obtain a fill rate higher than 66 per cent. With an \$800,000 investment, the highest attainable fill rate is 85 per cent; with \$1,580,000, the maximum is 94 per cent. Similarly, a 66-per-cent fill rate is unattainable with less than \$270,000, an 85-per-cent rate with less than \$800,000, and a 94-per-cent rate with less than \$1,580,000.

Figure 3 shows the final step: a complete display of all the system alternatives and the associated values of the control ratio. Thus, a system fill rate of 80 per cent requires a system investment of \$600,000 and is associated with a control ratio of \$2,500. That is, if item decisions were made under the assumption that a fill is worth a \$2,500 investment, the system result would be a fill rate of 80 per cent and a system investment of \$600,000. Similarly, a system investment of \$1,000,000 will provide a system fill rate of 88 per cent and will be associated with a control ratio of approximately \$4,000.

#### MANAGEMENT IMPLICATIONS

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

This study is still in the research phase; a good deal of development work must be done before large-scale applications will be completed. But there is already reason to believe that this approach has important implications for supply management. Recall the key problem discussed earlier. First, in the evaluation of system-performance potential, this

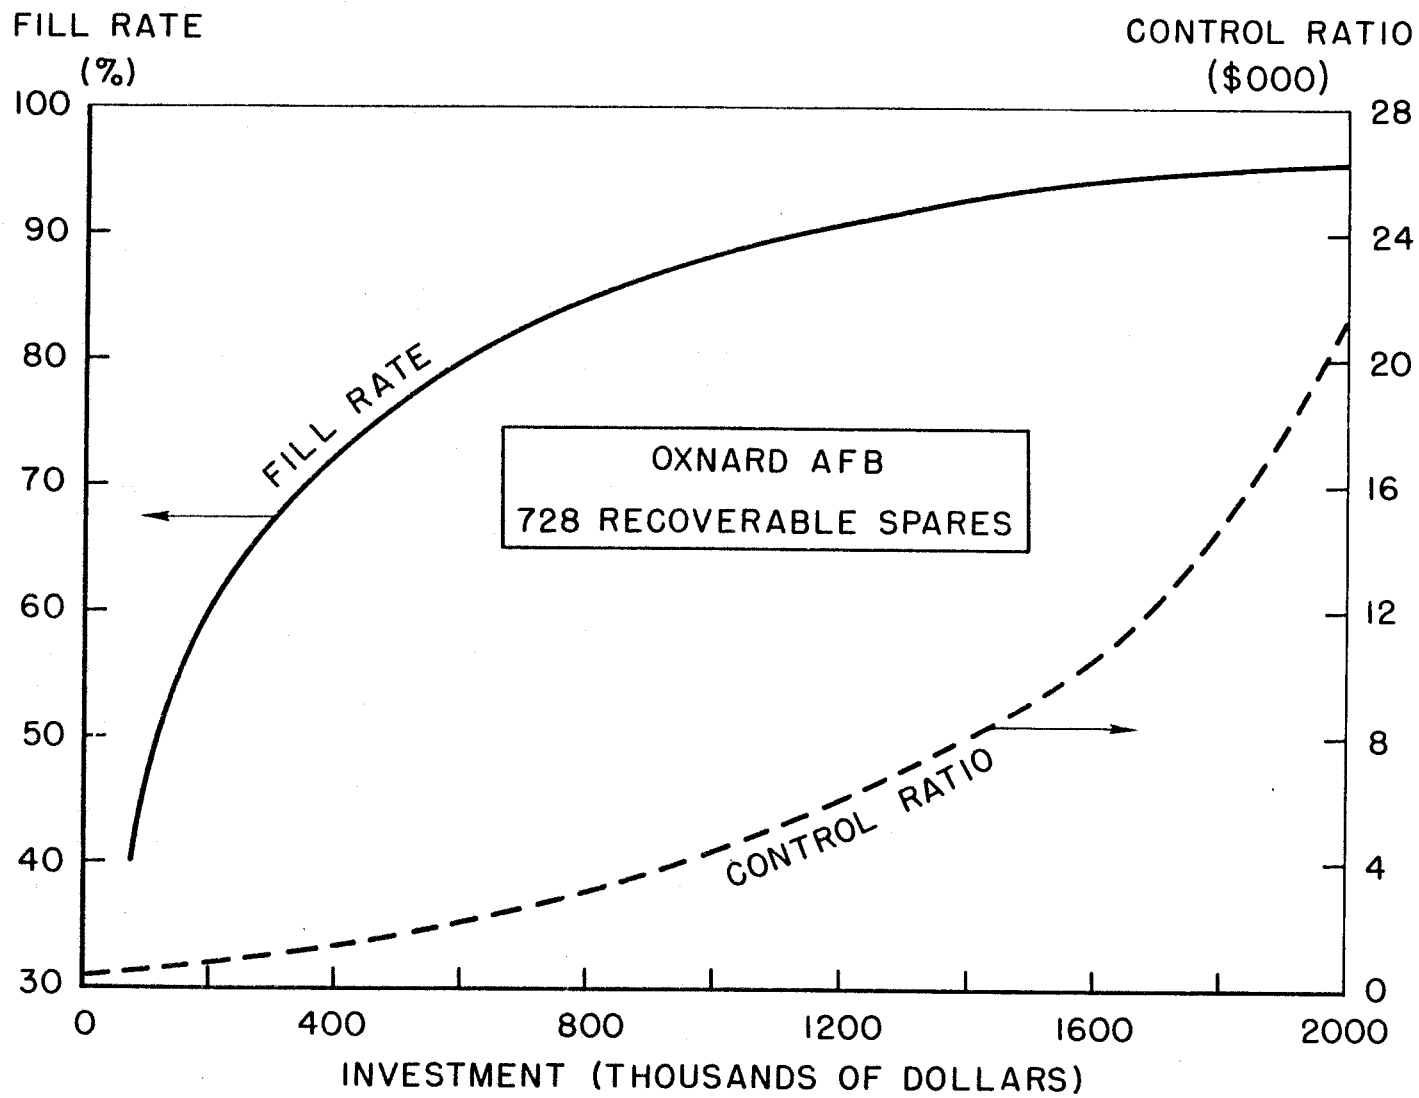


Fig. 3 — System performance/system investment

kind of analysis shows what level of performance is realistically attainable for a given budget, and furnishes a benchmark against which actual performance can be measured. Second, it provides a mechanism for keeping a system within budget constraints. Third, it permits comprehensive justification for budget requirements. In addition to total investment, we can graph system performance against new purchases and thus provide a quantitative display of the performance degradation that might result from inadequate funding or, alternatively, a display of the funds that might be made available through more austere service levels. Fourth, it can be applied continuously, using the control ratio as a management tool to assure balanced response to unexpected changes in customer demand and available funds.

#### EXTENSIONS

This discussion has examined only one case: a system with a relatively small number of high-cost, recoverable aircraft spares. It should be emphasized that this example, while important in itself, was selected because it simplified the discussion. This same approach is equally applicable to large systems with tens or hundreds of thousands of items. Indeed, for technical reasons, the larger the system, the more accurate the results will be. The Performance/Investment curve was generated through a mathematical procedure which is extremely general. The process has been used successfully in a variety of different systems and problems, including base stockage policy and war reserve requirements.

We are presently investigating the extension of this general approach to several more advanced problems, including initial provisioning, redistribution, allocation among subsystems, tradeoff between maintenance capability and aircraft spares, differential support based on mission essentiality, system sensitivity to change in data-analysis techniques, repricing of inventories in relation to changes in demand, and disposal criteria. It appears that this approach can be applied in each of these problem areas to yield significant improvements in supply operations.

### III. DEMAND ANALYSIS

In contrast to stock policy, which obviously involves consideration of the supply system as a whole, our second area, demand analysis, is a detailed, technical process that would seem to be confined entirely to individual items. Here, too, for more subtle reasons, we will find that system analysis provides new insights that may significantly improve the effectiveness of supply management.

Demand analysis is essentially a learning process, by which an item's specific demand experience is transformed into a reasonable estimate of future demand possibilities. The commonest procedure is to calculate the item's issue rate by adding up all demands experienced during some past period and dividing by the length of the period. Future demand is then characterized as some kind of random variation around the issue rate.

Such an approach is entirely adequate if the item has relatively high demand, if it has a relatively long and accurate demand history, and if the history is relevant. But it is well known that many items, very often the most important ones, fail to satisfy one or more of these three essential requirements. How do we estimate demand for these items? What, for example, is a "thirty-day supply" of an item that has had three demands in the last two years, or no demands at all in the past six months? What is a "thirty-day supply" of an item that was introduced into the system only four months ago, or an item sensitive to operations that expand or contract to unforeseeable levels? In cases like these, which are common in almost all supply systems, the issue-rate technique breaks down.

#### SYSTEM APPROACH

What remedies can be found? One approach is to develop more sophisticated analytical techniques. There is a growing literature about new forecasting and prediction procedures such as exponential smoothing, time-series analysis, and so on. Unquestionably, these developments can be valuable to supply management, particularly in dealing with high-cost, high-demand items with life-cycle or other

systematic fluctuations in average demand levels. Our purpose here, however, is to describe a new and different technique that might be called the system approach to demand analysis. This approach is based on the surprising fact that we can increase our knowledge of an item by studying the behavior of the other items in the system.

As before, we will use a simple system to illustrate the basic ideas. Imagine a system with many items but with the special property that only three demand distributions are present. Each item's demand is governed by one of these three distributions. We will call these three demand distributions the components of the system.

Table 1A shows the six-month demand probabilities of each component. Thus, each item in Component A has a mean demand of 0.25 units in six months or an average of 1 demand every two years. On the average, 77.88 per cent of the items in this component will show no demand in a six-month period, 19.47 per cent will show one demand, 2.43 per cent 2 demands, etc. Each item in Component B has a six-month mean of 1. On the average 36.79 per cent of these items will show no demand over six months, 36.79 per cent 1 demand, 18.39 per cent 2 demands, etc. Each item in Component C, finally, has a six-month mean of 3. Only 4.98 per cent of these items will show no demand in six months; on the average, 14.94 will show 1 demand, etc.

Table 1B shows what we call the structure of the system, specifying what percentages of the items are in each of the three components. In this case, 60 per cent of the items are in Component A, 30 per cent in Component B, and 10 per cent in Component C. This table also shows the distribution of items by component and six-month demand level. Thus, on the average, 46.73 per cent of the items will show no demand and be in Component A, 1.84 per cent of the items will show 3 demands and be in Component B, 1.01 per cent of the items will show 5 demands and be in Component C, etc. The last line of Table 1B gives the distribution of system demands resulting from this structure. On the average, 58.27 per cent of the items in the system will show no demands in a six-month period, 24.21 per cent will show 1 demand, 9.22 per cent 2 demands, etc.





Now we come to the crucial step. The structural analysis shows what percentages of the items at a given demand level are in each of the three components. This means that even if we do not know which component a given item is in, we can observe the item's demand during a six-month period and estimate the probability that it is in any one of the three categories.

For example, since 58.27 per cent of all items are expected to have no demand and 46.73 per cent are expected to have no demand and be in Category A, it follows that  $46.73/58.27$ , or 80.20 per cent of the items observed to have no demand must have been in Category A. This process of estimating the probability that an item's true demand is in a particular category, based on the observed demand during some period and on the over-all structure of the system, is a particular application of a general technique called Bayesian inference, after Thomas Bayes, who first described the technique some two hundred years ago.

Table 1C summarizes these probabilities. Thus, in this particular system, we know that if we observed 1 demand for an item during the last six months there is a 48.24 per cent chance that the item is in Category A, a 45.60 per cent chance it is in Category B, and a 6.15 per cent chance it is in Category C. Similarly, if we observed 4 demands for an item during the last six months there is still a small chance -- 0.28 per cent -- that its true average is only 1 demand every two years, and a fairly good chance -- 21.44 per cent -- that its true average is only 2 demands a year.

In this simple example, the system structure was assumed to be known in advance. Of course, in a real system the structure would not be known. In our original work in this area, we developed an estimate of the structure intuitively; that is, we took a representative set of demand components and used judgment to guess at the system mix. Such an approach was justified, in part at least, because the results did not appear overly sensitive to errors in the specification of the underlying system structure. More recently we have discovered a procedure that makes such guesswork unnecessary. This procedure uses the system's demand profile -- what percentage of the items had

no demands, 1 demand, 2 demands, etc., -- to produce a direct estimate of the underlying demand components and the system mix.

To simplify the discussion, we used an example with only three demand components with means of 0.25, 1.0, and 3.0. In actual applications, of course, a much broader set of components would be used, spanning the entire range from very low to very high demand levels. This poses no technical problems. Indeed, in most cases we treat the underlying components as a continuum and estimate the structure as a continuous distribution over the whole range of possible demand levels.

#### MANAGEMENT IMPLICATIONS

Notice that in this approach we do not conclude with an estimate of the item's true average demand, but with estimates of the probability that the item's average demand is at one of various levels. We think this is a much more realistic way of characterizing what we can learn from an item's demand history. These probabilities can now be used to appraise the true risks and potential payoffs for various stock levels. Data-processing procedures can be designed to use information of this kind with little or no increase in processing complexity.

Bayesian demand analysis has two immediate implications for improved supply management. First, because it wrings maximum relevant information from available demand data, this approach promises substantial improvement in the efficiency of stockage decisions. Specifically, by using this kind of analysis, we should be less liable to overestimate demand and buy too much because of a random surge in demand, and less liable to underestimate demand and buy too little because of a random decline in demand. Second, the approach is extremely flexible because it can be applied to any time period. It should be possible to eliminate many of the policy problems now created by items that do not meet the requirements of daily-issue-rate computation -- items with low demand, short histories, and erratic demand patterns. In the framework of Bayesian demand analysis, policies can be developed that prescribe action unambiguously.

EXTENSIONS

A number of more advanced applications of these general concepts are being studied. For example, we have noticed that supply systems that appear on the surface to be quite different turn out to have very similar underlying demand structures. We think the cross-sectional mix of demand levels and item costs are critical system parameters through which it may eventually be possible to analyze and compare aircraft systems in terms of such parameters as speed, range, payload, and the like.

We believe that these methods could be employed profitably throughout the life of a weapon system. In initial provisioning, estimates of demand would be used. These item estimates can be adjusted by a Bayesian procedure if data is available that suggests systematic over- or underestimation or central tendency errors. During the phase-in of a new weapon, Bayesian analysis can be used to combine initial estimates with experience data to obtain revised estimates.\* After a reasonable period of experience, such as six months or a year, the past item demand would be used exclusively.

---

\* See W. H. McGlothlin, Development of Bayesian Parameters for Spare-Parts Demand Prediction, The RAND Corporation, RM-3699-PR, July 1963.

IV. CONCLUSIONS

Supply management is system management. Procedures that look no further than the individual line item are not adequate. It is essential that new management techniques be developed that take account of the complex relationship between the behavior of line items and the behavior of the supply system as a whole. We have examined two such techniques, one concerned with stock policy, the other concerned with demand analysis. We believe each provides persuasive evidence of the basic theme of this discussion: that system analysis presents significant new opportunities for profitable innovation in supply management.