118984

Methodology Transfer Paper 2

Evaluating Inventory Control Decisions

GAO

Institute For Program Evaluation

United States General Accounting Office

June 1982

Methodology Transfer Paper 2

Evaluating Inventory Control Decisions

Institute for Program Evaluation

June 1982

PREFACE

This document is intended to assist government managers in evaluating the efficiency of inventory decisions. Inventory levels are governed by decisions about when and how much to replenish stock. The basic premise of this document is that inventory systems are managed most efficiently when inventory levels are such that the sum of the costs of carrying inventories, replenishing inventories, and running out of stock is minimized.

We are issuing this document as part of the General Accounting Office response to title VII of the Congressional Budget Act of 1974 (Pub. L. No. 93-344), which requires us to develop and recommend methods for reviewing and evaluating government programs and activities. We believe it will help expand congressional activities in investigating alternative ways of achieving program objectives that are "preferable according to cost effectiveness criteria or other explicit standards," as the Act intends.

In recent years, interest in analytic techniques for inventory management has been increasing for several good reasons. One is that inventory investment has had a tendency to become larger than most managers would like. Moreover, new operational requirements, especially in the military, have intensified the pressure for faster and more reliable service from inventory systems. At the same time, there has been a trend toward using budgets more efficiently, while the use of high-speed computers has made it possible to implement more sophisticated control procedures. This is an opportune time to question and evaluate the extent to which analytic techniques can help inventory managers make decisions that will minimize costs.

It is important to note, however, that inventory managers cannot control the demand for inventory other than by setting service level objectives. Therefore, it is not likely that total procurement actions over any period of time will be reduced, although average inventory investment might be. Rather, reductions can be made in the cost of operating a system--that is, what is reduced is the sum of the costs of carrying inventories, being out of stock, and placing orders. Thus, potential cost savings are generally operational and recurring rather than one-time savings achieved by cutting inventory requirements, and they can be significant. Even in relatively small systems, proper decisionmaking can reduce operational costs by hundreds of thousands of dollars each year.

We believe this document will be useful not only to managers who are interested in operating inventory systems efficiently but also to evaluators and auditors. This is partly because we intend it to be useful on the job, and to this end we have presented checklists and other guidance material as aids in insuring that proper evaluative considerations are being made. In addition, we have placed much of the discussion in a perspective that is different from most of the literature on the subject in the expectation that this will be useful for novices but also for knowledgeable practitioners as well.

Many of the examples we present are taken from experience and data obtained during various GAO evaluations of Federal inventory systems. Much of the philosophy about the nature of inventory systems and how they should be managed reflects the influence of Eliezer Naddor of The Johns Hopkins University. He has served GAO in an advisory capacity for many years. For information about the contents of this document, contact Institute for Program Evaluation, Associate Director, Methodology Development and Measurement Assistance Group.

This paper is one of a series of methodology transfer papers developed by the Institute for Program Evaluation. The purpose of a methodology transfer paper is to provide GAO staff with a clear and comprehensive discussion of the basic concepts of an evaluation methodology. Additionally, transfer papers explain both the general and the specific applications and procedures for using the evaluation methodology. The first paper in this series, Causal Analysis: A Method to Identify and Test Cause and Effect Relationships in Program Evaluations, and the present one will be followed by others in preparation.

Eleanor Chelimsky

Eleanor Chelimsky, Director
Institute for Program Evaluation

Contents

ABBREVIATIONS

C	Cost
D	Demand
EOQ	Economical order quantity
EWMA	Exponentially weighted moving average
F	Fill rate
GAO	U.S. General Accounting Office
L	Level, order level
Lt	Lead time
Q	Quantity, order quantity, economical order quantity
R	Reorder point
RL	Reorder point-order level
RQ	Reorder point-fixed order quantity
S	Standard deviation
t	Time
TEC	Total expected cost
TL	Fixed interval-order level
$\overline{X}$	Mean

CHAPTER 1

INTRODUCTION

This document is for managers, evaluators, auditors, and others who conduct studies in government inventory systems. We focus on two related replenishment questions that inventory managers face: (1) when to order from an outside vendor or the internal production system and (2) how much to order each time an order is placed. Answering these questions properly can assure an organization that it holds just enough inventory or stock to meet its demand commitments economically.

Although deciding when and how much to order is only one of several functions of inventory or supply management, it is probably the heart of supply management, because it directly affects the cost of operating all the other supply management functions. One writer, for example, suggests that supply management contains seven major functions, which (listed in chronological order) are managing, cataloging, determining requirements, procurement, distribution, overhaul, and disposing of material. (Kuhlman, 1969, p. 7) 1/ Thus, a decision to place orders frequently will affect the cost of operating the supply system's procurement function. Similarly, a decision to place large orders infrequently will affect the cost of storing the inventory until it is needed.

An objective of inventory decisionmaking is to arrive at a set of optimal rules for deciding when to place an order and how much to order. Optimal rules assign values to controllable variables that minimize the sum of all relevant inventory costs. One problem in arriving at optimal rules is that these costs conflict with each other. For example, buying in large quantities reduces the probability of stockouts, reduces the cost of placing and receiving replenishment orders, and tends to result in lower prices per unit of stock and lower transportation costs, but the larger quantities increase investment and tend to increase the cost of holding the inventory. Thus, inventory managers must seek a balance between the costs of holding inventory, ordering additional quantities, and running out of stock.

Another problem in deciding when and how much to order is that the decision involves risk. To acquire additional stock, it is necessary to forecast future demand and lead time, but forecasts are risky because demand usually varies from month to month and delivery time or procurement lead time varies from order to order. In striving for a reasonable balance between costs, managers must provide for these variations and their ensuing risks.

1/Interlinear bibliographic references are cited in full in appendix II.

If managers are to minimize costs, their decisions about when and how much to order must be made in view of conflicting cost functions, forecasts of demand rates, and forecasts of replenishment lead times. A number of sophisticated techniques, or models, have been developed for doing this, and it is our purpose to show how to evaluate the appropriateness of the techniques that are used in government organizations. In chapters 2 through 5, we discuss the four factors by which inventory systems are characterized--namely, policy or when to order and how much to order; the costs of carrying, shortage, and replenishing; constant and variable demands; and lead time, whether insignificant or significant, constant or variable. In chapter 6, the final chapter, we discuss the techniques that are available for considering these factors together to minimize cost.

The three policies that are used the most are reorder point-fixed order quantity, fixed interval-order level, and reorder point-order level. We give an illustration of each in chapter 2.

In chapter 3, we focus on the cost elements that should be considered in developing rules for deciding when and how much to order. It is not enough to be able to say "These are the rules"; good inventory management can also say ". . . and this is how much it is going to cost." Time and quantity are the variables that are subject to control, and the problem is to find their specific values for minimizing total cost.

Organizations keep inventories to meet demands, fill orders, and satisfy demands, but only rarely do they have sufficient knowledge about what generates demand to predict demand patterns with certainty. In some cases, demand patterns show a regularity that, for purposes of an adequate approximation, can be treated as a certainty. Often, however, managers have to describe demand in probabilistic terms and to assume that a random process generates demand. Accordingly, in chapter 4 we discuss the probability of demand and demand forecast.

In chapter 5, we discuss the nature and measurement of lead time--the time between scheduling a replenishment or placing an order and its actual addition to stock. It is of concern because some of this time cannot be controlled by the decisionmaker. The resulting uncertainty must be considered in trying to optimize decisions.

In the last chapter, we discuss analytic, simulation, and heuristic techniques for assigning values to policy variables. With an analytic approach, managers construct a mathematical model of the system to be studied with which they can determine the set of operating rules that minimizes costs. With simulation, they design and conduct experiments with a model of a real system for the purpose of either evaluating various decision policies for the system's operation or understanding the system's behavior. In heuristic considerations, managers study insights from analytic approaches or experience with and intuitions about a system.

Various professional organizations have issued formal guidance for analyzing and auditing many of the functions of supply management systems. (Among them are Barden, 1973; Institute of Internal Auditors, 1970; McCarthy and Morison, 1975; and National Association of Accountants, 1964.) Guidance for analyzing how well management decisions minimize costs is, however, sketchy, at best. We hope that this document can help fill in that gap in our practice and knowledge of government supply management.

CHAPTER 2

INVENTORY POLICY OPTIONS

Answering systematically the questions of when to place an order and how much to order is inventory policy decisionmaking. The question of when to order is usually answered in one of two ways--inventory should be replenished after a specified, fixed interval of time, t, or inventory should be replenished when it has been reduced to a specified number of units, R, called the "reorder point."

With the first answer, orders are placed each time an inventory count is made if there have been any demands at all since the last count, whether the counts are daily, weekly, quarterly, or at some other interval. Thus, for example, if inventory levels are reviewed monthly, either physically or by computer, and in May, some stock, regardless of the amount, has been depleted since April, then the manager places an order.

With the second answer, orders are placed only when the inventory position--usually defined as the amount on hand and available for issue plus the amount on order and due in from suppliers--is less than or equal to the reorder point. The interval of the inventory counts is not what determines the answer. Thus, if the reorder point is 25 units and an inventory count reveals an inventory position of 20 units, then an order should be placed. If the inventory position is 26 or higher, no order should be placed.

The question of how much to order is also usually answered in one of two ways--the quantity to be ordered is a predetermined, fixed quantity, Q, or the quantity to be ordered will bring the inventory position to a specified level, L, called the "order level."

The first answer means simply that whenever a manager decides to place an order, the quantity is already known. For example, if it has been established that 50 units will be ordered whenever the inventory position--the amount on hand and on order--falls below a certain amount, then 50 units will be ordered.

The second answer means that the amount that is ordered is sufficient to bring the appropriate inventory position up to a maximum order level. For example, if the established order level is 100 units and the inventory position is 25 units, and if it is time to order, then the order quantity will be 75 units. If the inventory position had been only 20 units, the order quantity would be 80 units. Thus, the order quantity varies according to the inventory position at the time an inventory count or review is made.

The four ways of answering the two questions suggest four possible decisionmaking policies, but in practice only three are used. These are (1) the reorder point-fixed order quantity, or RQ, policy; (2) the fixed interval-order level, or TL, policy; and (3) the reorder point-order level, or RL, policy.

REORDER POINT-FIXED ORDER QUANTITY POLICY

Under the reorder point-fixed order quantity policy, the order size is held constant and the time between placing orders is allowed to vary to compensate for variability in usage rates. Orders are placed only if the inventory position is less than or equal to some specified amount, the reorder point, at the time inventory levels are reviewed, whether by physical count or by computer analysis. 1/ The quantity ordered is some fixed amount called the order quantity. If this quantity is not sufficient to bring the new inventory position above the reorder point, then multiples of the order quantity are ordered until the position is more than the reorder point.

Figure 1 on the next page diagrams this policy. To keep the presentation simple, we assume that lead time is insignificant. All orders (all positive Q's) are equal. Orders are placed at the time of review only if the inventory position is less than the reorder point--in the diagram, this is periods 1, 3, and 4. At period 2, the time of the second review, the inventory position was above the reorder point, so no order was placed; the quantity ordered is zero.

FIXED INTERVAL-ORDER LEVEL POLICY

As its name suggests, in the fixed interval-order level policy, orders of varying sizes are placed to replenish stock at fixed intervals of time. For example, a stock count might be taken once a month, at which time the manager places a replenishment order, basing it on the amount used or demanded since the last review and taking into account the forecast for the next period. The decisionmaker must decide what the fixed interval or the review cycle and what the rules governing the order size ought to be. Figure 2 shows how this policy operates.

1/A distinction is usually made between continuous and periodic review. In a continuous review, the state of a system is known at each point in time because each transaction (demand, placement of order, receipt of shipment, and so on) is recorded and reported as it occurs. In a periodic review, the state of the system is examined only at discrete, usually equally spaced, points in time. We do not emphasize this distinction in this document. In reality, most systems are reviewed periodically and for some the period is quite small. In the so-called transactions reporting system, for example, computer printouts are given to managers daily. While this sounds like continuous review, it still entails the lapse of one day.

Figure 1

Reorder Point-Fixed Order Quantity Policy

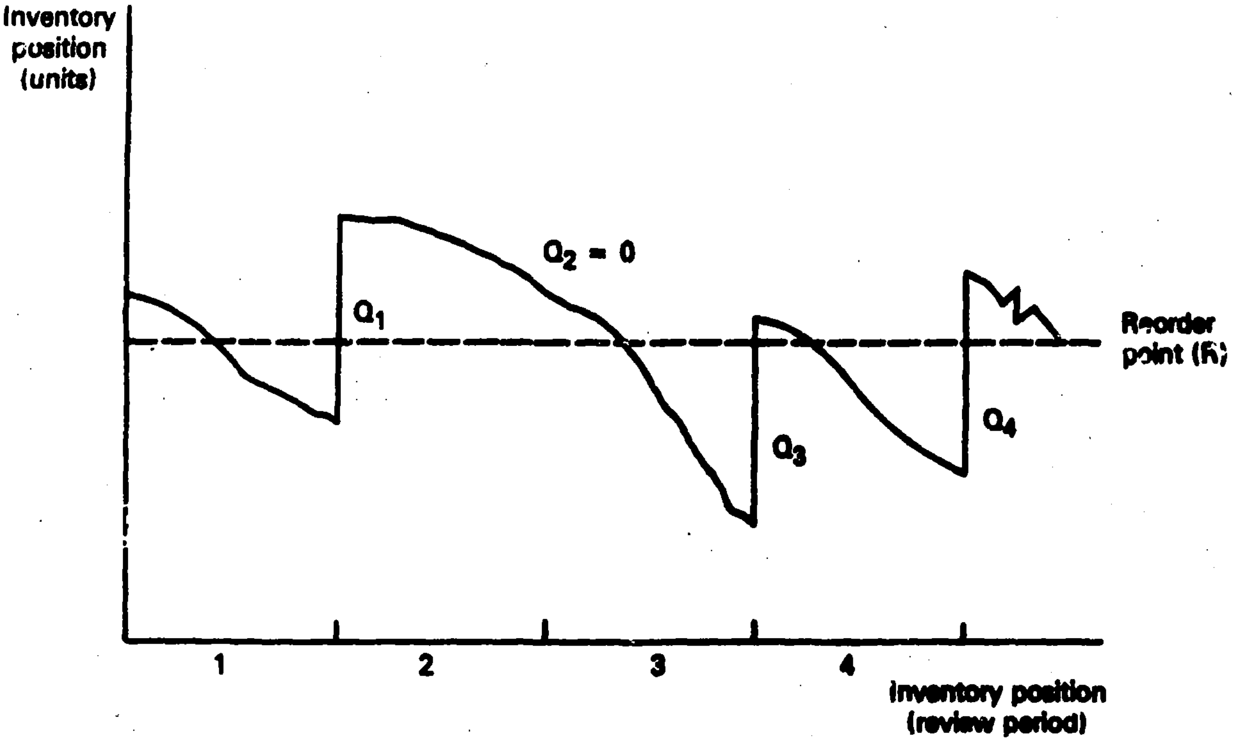

Figure 2

Fixed Interval-Order Level Policy

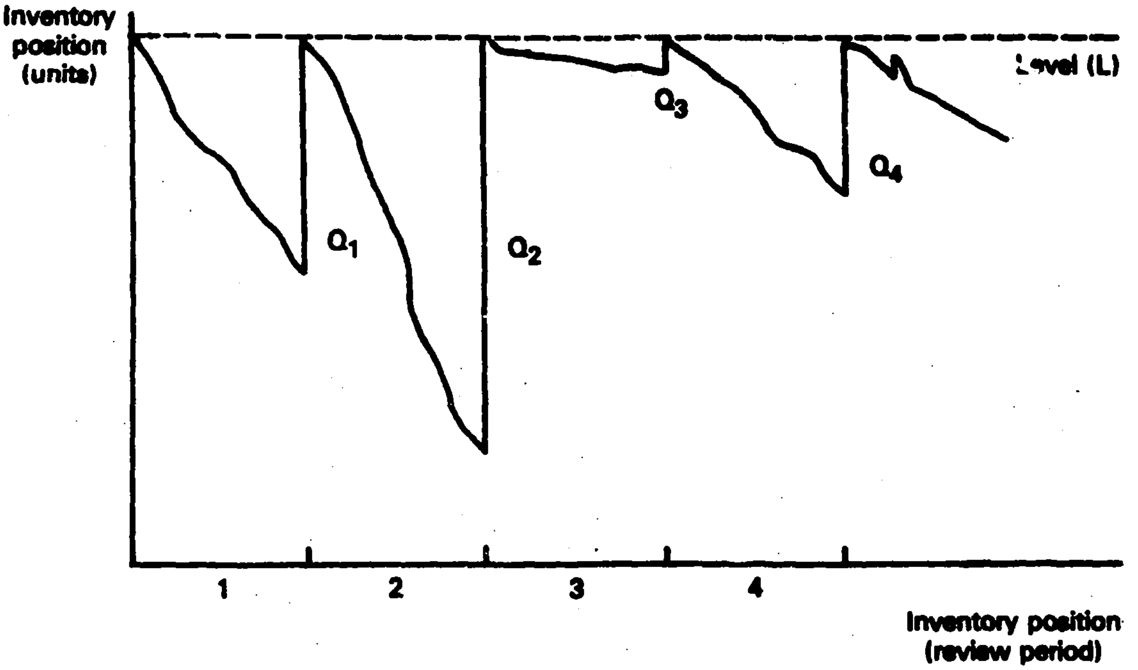

Figure 3

Reorder Point-Order Level Policy

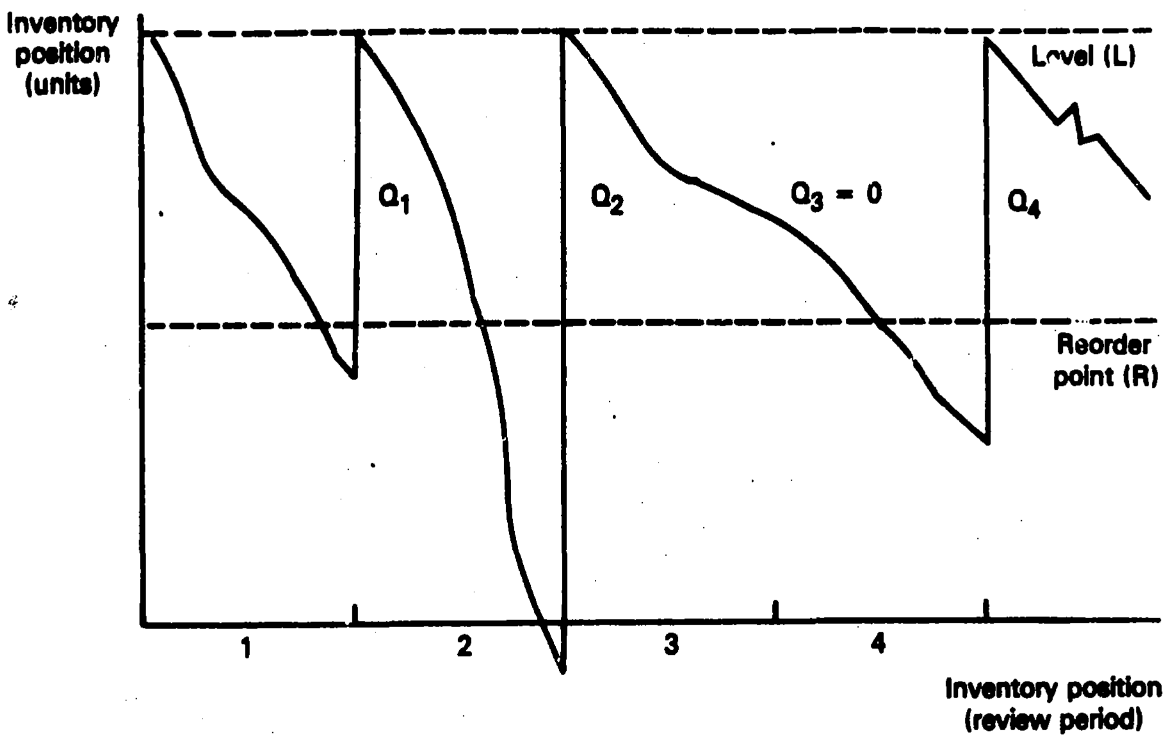

Note that the order size varies, being relatively large in the second period but small in the third period.

REORDER POINT-ORDER LEVEL POLICY

The reorder point-order level policy is a blend of the two others. Inventory stocks are reviewed regularly and periodically, and stocks are replenished only when they have fallen to or below a specified level. The order that is placed will bring the amount of stock on hand and on order up to a specified maximum level.

The ordering rule for the reorder point-order level policy is very simply stated. At the time of the review, if fewer than R units are available on hand and on order, then order enough to bring stock up to level L; otherwise, do not order. (Magee and Boodman, 1967, p. 136) Operation of this rule is portrayed in figure 3. The system shown here operates identically to that shown in figure 2 except for the third review, when the stock is above the minimum level R and no order is placed.

SELECTING A POLICY OPTION

Recall that in deciding when and how much to order, managers should seek to minimize the costs of carrying inventory, running out of stock, and ordering additional quantities. Generally, the reorder point-order level policy is less costly than either of the two other policies. (Hadley and Whitin, 1963, pp. 364-74; Naddor, 1966, pp. 314-19) They may be justifiably selected for other reasons, however.

The reorder point-fixed order quantity policy is employed most usefully where close control is not necessary because of low activity or a low item value. Both characteristics typically cause order sizes to be large, economical, and infrequent. One drawback of the fixed order quantity system, however, is that two items can reach their reorder points at different times, making it difficult to group orders going to the same supplier.

The fixed interval-order level policy makes possible tighter, more frequent control over high-value inventories. Inventories may be reviewed weekly or even daily. One problem with it is that it is possible to place small, uneconomical orders (as in the third period in figure 2). On the other hand, it is also possible to group orders for a number of individual items from one supplier so that they can be shipped in carload lots, thus lowering costs.

CHAPTER 3

COSTS THAT VARY WITH INVENTORY POLICY

Balancing opposing costs is at the heart of all inventory problem-solving. These are the costs that vary with changes in inventory policy. They are characteristically not the costs reported in summary accounting records. Cost information taken from accounting records typically requires reorganizing or restating to arrive at cost definitions suitable to a particular problem. Often, costs must be derived by experimental or statistical methods.

Developing and measuring inventory-related costs require differentiating accounting costs for historical and financial reporting from operational or functional costs. Operational costs are essentially the out-of-pocket expenditures or forgone opportunities for profit that are considered in arriving at policy decisions or day-to-day management decisions, and their magnitudes are affected by changes in inventory decision variables such as replenishment size. The three types of cost that may be important in determining inventory policy are the costs of carrying items in inventory, the costs associated with demands occurring when an item is out of stock (these are called "shortage costs"), and the costs associated with replenishing the units stocked.

CARRYING COSTS

The costs of carrying or holding inventory are roughly proportional to the size and value of the inventory. They have several components. (England and Leenders, 1975, pp. 359-60)

1. The cost of money tied up in carrying the inventory is often referred to as an "opportunity cost," money or capital tied up in carrying inventory that could be put to alternative uses. Thus, an opportunity for its earning a return may be forgone. Although the concept of profitability may not be relevant to government, alternative uses of limited funds are real and frequent. The minimum cost assigned to this component is the return that could have been earned by keeping the money in a bank and includes interest on money invested in the inventoried items, in land and buildings to hold the inventory, and in inventory handling and control equipment.

2. Storage space cost is the cost associated with the physical space required to house the inventory of a given item. It depends on alternative uses for the storage space. If it is available and no alternative uses are prevalent, the space is essentially free. If alternatives do exist, then the storage cost is an opportunity cost.

Storage space costs can include rent on the storage facility, taxes and insurance on the building, depreciation on the building and warehouse installation, the cost of maintenance and repairs, utility charges for heat, light, and water, and the salaries of security and maintenance personnel. Moreover, in many situations, storage costs vary directly with the quantity of inventory in storage. For example, storage space might be rented as needed, electricity might vary with the number of items requiring refrigeration, and taxes might be levied on various inventory values.

3. Shrinkage costs are associated with items that shrink in value during storage. The shrinkage of the inventory value can result from physical deterioration, obsolescence, or pilferage and represents a cost that must be assigned to carrying inventory. If the inventory is insured against this risk, then the cost of insurance should be included.

4. Inventory service costs are the results of the need to provide prompt service to an organization's clients. The costs connected with providing this service include labor costs in handling and maintaining stocks, clerical expenses in keeping records, and employee benefits for warehouse and administrative personnel. Inventories may also be taxed, and the tax assessments should be included as a cost for providing a service.

5. Handling-equipment costs, in addition to the investment cost in money tied up in warehouse equipment, include taxes and insurance on equipment, depreciation on equipment, fuel expense, and maintenance and repair costs.

Carrying cost is customarily computed as a fraction of the cost of items carried in inventory per unit of time. For example, a one-dollar item may cost 25 cents per year to carry--that is, one-fourth, or 25 percent, of its value. The fraction, usually referred to as a percentage, depends on the nature of the cost elements. For many government applications, its numerical value varies from about 20 percent to about 40 percent. Unfortunately, this value cannot be determined directly from present accounting records without a certain amount of research.

Some have deemed this method of assessing the holding interest rate to be somewhat meaningless. (Lewis, 1975, p. 109) Some have proposed that the value that is used should be a control variable that management alters in light of changes in an existing financial situation. High values of the rate would be used during recession to justify reducing inventories, and low values would be used during expansion to allow for increased stockholding. We do not advocate this proposal because it defeats the purpose of making decisions that minimize costs. A better way to reduce or raise inventories is by using service level or fill rate goals, which we discuss in the next section.

SHORTAGE COSTS

The cost of not having inventory available when it is needed is called the "shortage" or "stockout" cost. In this definition, it is important to take note of the phrase "when needed": if the system is out of stock but stock is not needed, no shortage cost is incurred. (See U.S. General Accounting Office, 1981, pp. 41-45, for an application of this concept.) The shortage cost can be one of two types, depending on the reaction of the prospective customer to the item's being out of stock. If the customer is willing that delivery be delayed, the system can typically institute an emergency expediting procedure to get stock. This is called a "back order." The sale is thus not lost, but additional costs such as for special telephoning, shipping, and handling are incurred.

The other kind of shortage or stockout cost is the sale that is lost because the customer is not willing to wait but seeks an alternative source for the item. In government or other organizations with "captive" customers, the cost of a stockout might have to be measured in terms of its consequences to the customer, especially if the customer is forced to choose alternatives that are less than economical. For example, the Veterans Administration stocks goods only if they can be sold to customers (generally hospitals) at a certain percentage lower than in the local market (currently 15 percent). If a hospital requests an item that is out of stock but cannot wait for it--items such as medicine and food may be needed immediately--then there is an additional cost to the government (of at least 15 percent) because the customer must pay a higher price for the commodity.

Many managers state that they do not allow shortages to occur. This usually means that they assume that the unit shortage cost is infinite. This may sound unrealistic but it is not. What the managers are really saying is that the unit shortage cost is relatively high and, therefore, they intend to have very few shortages.

Many other managers state categorically that the unit shortage cost cannot be measured, and it is true that often it cannot be measured precisely. Nevertheless, managers who may not be able to state the numerical value of a stockout may make decisions affecting surpluses and shortages that imply that such a value exists. For example, many organizations strive to achieve a given fill rate or service level--that is, a set percentage of orders filled from stock. Opinions about the definition of service level vary from "the probability of not running out of stock" to "the proportion of annual demand not satisfied." The former has the advantage of being fairly easy to calculate but the disadvantage of being easily misinterpreted. From a customer's point of view, the latter is preferred. The former is actually a "vendor service level" whereas the latter is a "customer service level." (See Lewis, 1975, pp. 112-13, 155-69, for a good discussion of this.) For estimating shortage costs with the approach we describe, these differences of definition are not significant.

Establishing a fill rate or service level implies that the organization can tolerate stockouts. A 90 percent fill rate, for example, means that it is acceptable to be out of stock 10 percent of the time. By permitting stockouts, managers can reduce the amount of inventory on hand at any one time and, thus, reduce carrying costs. But because stockouts also cost money, managers must balance these opposing costs when deciding when and how much to order.

If shortage and carrying costs are known, selecting optimal decision rules will balance the expected cost of carrying an incremental unit in inventory against the expected cost of not carrying that unit. (We discuss this in chapter 6.) When shortage costs are not known, the fill rate may be used to assign an implicit value to them, as in the following formula:

$$C_2 = C_1[F/(1 - F)]$$

F is the desired fill rate, C_1 represents unit carrying costs, and C_2 represents unit shortage costs. A 90 percent fill rate, therefore, implicitly values the cost of maintaining one unit on back order at nine times the cost of carrying one unit in inventory for the same length of time. (See National Association of Accountants, 1964, pp. 47-58, 107-12, for a relatively nontechnical discussion and Naddor, 1975a, p. 1240, for a mathematical derivation.)

REPLENISHMENT COSTS

Replenishment costs, sometimes referred to as "procurement" or "order" costs, arise in many different ways and can vary considerably among systems. For example, processing an order through accounting and purchasing may include costs for paper, postage, and labor, for telephoning vendors, and for using computer time in making computations and updating records. Transportation may also be borne directly by the inventory system, as may receiving costs in uncrating, inspecting, and testing the goods.

Some replenishment costs depend on the quantity ordered; others do not. Transportation costs, for example, and part of the receiving and part of the inspection costs depend on the quantity ordered and could be included as part of the cost of the item procured. The cost of placing an order also depends sometimes on the size of a contract to be processed, with larger contracts requiring more administration and review while orders in smaller dollar amounts may require very little review. For example, orders of $2,500 or less may cost as little as $10 while costs may rise to $45 for orders involving contracts between $2,500 and $10,000 and to $75 for orders in excess of $10,000.

Costs that are independent of the quantity ordered include costs for paper, postage, and telephones as well as the labor in processing the order. Some parts of receiving and inspection costs are also independent of the order size, as are the costs of setting up for a production run if the inventory system controls the plant in which the item is made. Every order incurs costs that do not depend on the quantity ordered. They are often referred to as "fixed costs" even though they may vary somewhat from one order to another.

Replenishment costs may be difficult to measure from standard accounting records. One approach might be to divide the total period cost of the procurement department by the number of orders it processes in that period. This would assume, erroneously, that all procurement costs are borne by the procurement department and that all its costs are variable, depending on the number of orders processed. A better approach to measuring the various components of replenishment cost is to analyze a representative sample of orders for the cost of placing orders over time. This would help determine both the fixed and the variable components of procurement costs.

The number of different items pertaining to one order is an important consideration in determining the unit replenishment cost. When two or more different items may be ordered simultaneously, only one replenishment cost exists. The items are in a sense sharing the cost of such things as transportation and paperwork. If the items were to be ordered individually, these costs would likely be higher and, correspondingly, the replenishment cost would be higher. Thus, decisions should not be made by considering each item individually.

OTHER COSTS

Operating the information-processing system

In operating an inventory system daily, managers will find that the cost of obtaining and processing the information necessary for decisionmaking clearly depends on the type of system it is and on the policies they are guided by. The cost of a computer's continuously or periodically updating inventory records, the cost of physically counting the inventory, and the cost of making demand forecasts should all be recognized as a part of the cost of carrying inventory--that is, as inventory service costs. These costs are often sizable, however, and may be worthy of special consideration when managers evaluate the inventory management system.

Quantity discounts

The amount paid to vendors, representing the cost of units procured, is relevant for analysis only if quantity discounts are allowed. That is, when the unit price of an item is adjusted for the quantity purchased, a decisionmaker can vary the item's unit price by varying the replenishment size.

Suppose, for example, that an organization expects to use 1,200 units of an item a year and the unit price is $10. The total amount to be paid annually to the vendor is $12,000 regardless of whether replenishment is scheduled at 1,200 units once a year or 100 units 12 times a year. The decisionmaker cannot influence the total amount to be paid to the vendor by altering inventory procurement policy. The vendor, however, might offer quantity discounts at some rate such as

Order quantity	Unit price
1-100	$10.00
101-300	9.50
301-600	9.00
601-1,200	8.50

If 100 units are procured 12 times a year, the total amount paid the vendor is $12,000, but if 600 units are procured twice a year, the vendor is paid $10,800. Procurement policy influences cost and cost, therefore, becomes relevant for decisionmaking.

COMPUTING THE COSTS

We can illustrate how costs are computed by considering the simulation in table 1. The policy represented here is that of a 12-month reorder point-fixed order quantity for a $25 item with the following cost assumptions:

Carrying cost = 36 percent per year per item cost
Shortage cost = unknown, with a 90 percent fill rate
Replenishment cost = $100 per order

Table 1

Cost Computation Simulation

Month	Beginning inventory	Demand	Ending inventory	Carried	Short	Ordered
1	90	30	60	75.0	0	0
2	60	10	50	55.0	0	1
3	50	50	0	25.0	0	0
4	0	10	-10	0	5.0	0
5	80	30	50	65.0	0	1
6	50	30	20	35.0	0	0
7	20	10	10	15.0	0	0
8	100	30	70	85.0	0	0
9	70	20	50	60.0	0	1
10	50	20	30	40.0	0	0
11	30	50	-20	9.0	4.0	0
12	70	40	30	50.0	0	1
Average monthly amount				42.83	0.75	0.33
Average monthly cost				$32.13	$5.06	$33.33
Total cost: $70.52						

Figure 4

Average Amount Carried for Period 11

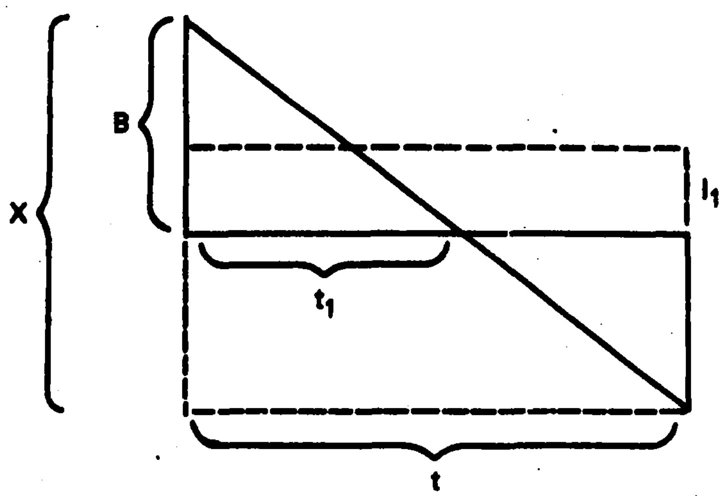

In addition, it is assumed that inventories are reviewed only once a month so that all calculations are expressed in monthly terms.

The carrying cost for the item is computed by determining the average amount carried per month or per period (42.8 units) and multiplying this amount by the carrying cost per item ($0.75, or $25 times 0.03). Except for period 11, computing the average amount carried is fairly straightforward. Since inventories are reviewed only once a month and there is no knowledge of the rate of withdrawal, we assume a linear withdrawal rate. Thus, for period 1, beginning with 90 units on hand and ending with 60 units, the linear average amount carried is 75 units--that is, the midpoint between 90 and 60.

Computing the average amount carried in period 11 is not as easy, but for this example it is fairly straightforward if we use some simple geometric relationships. Figure 4 shows B as beginning inventory, X as demand during time period t, t_1 as the time in which some inventory is carried, and I_1 as the average amount carried. Using the notion of similar triangles and the formulas for the area of a triangle and a rectangle, we can derive the formulas

$$\frac{t_1}{t} = \frac{B}{X} \qquad \text{and} \qquad I_1 t = \frac{B^2 t}{2X}$$

which in turn yield

$$I_1 = \frac{B^2}{2X}$$

We find, then, that the beginning inventory was 30 units and the demand during the period was 50. With the formula, we learn that on the average 9 units were carried during period 11.

We compute the shortage cost in much the same way as the carrying cost. First we determine the average amount short per period and then we multiply by the cost of being short one time. The cost of being short in this particular example will have to be valued implicitly with the 90 percent fill rate and the formula on page 12. We find that

$$C_2 = \frac{0.9}{1 - 0.9}(0.75) = \$6.75$$

Since the average amount short during the entire simulation is 0.75 units, multiplying 0.75 by $6.75 will yield the average cost of being short during this simulation, or $5.06.

Computing the cost of ordering or replenishing inventory is quite easy. Four orders were made during the simulation for an average of 0.333 per month. The average replenishment is thus $33.33 per month.

The most important amount in this example is the total cost of $70.52. Could different operating rules have produced a lower figure? Remember that it is not enough to say "These are the rules." We must be able to add ". . . and this is how much it is going to cost."

COST CHECKLIST

If managers are to be able to control the costs attributable to decisions about when to replenish stock and by how much, it is evident that these costs must be identified and measured. Categorizing inventory system costs as the cost of carrying stock, the cost of running out of stock, and the cost of replenishing stock, we can identify their components if we use a checklist something like this one:

Carrying costs

1. Capital, or interest on investments
   --In inventory
   --In land and buildings to hold inventory
   --In inventory handling and control equipment

2. Storage space
   --Rent on buildings
   --Taxes and insurance on buildings
   --Depreciation on buildings
   --Depreciation on warehouse installations
   --Maintenance and repairs
   --Utility charges for heat, light, and water
   --Salaries of security and maintenance personnel

3. Shrinkage
   --Obsolescence of inventory
   --Physical deterioration of inventory
   --Losses from pilferage
   --Insurance on inventory

4. Inventory service
   --Labor in handling and maintaining stocks
   --Clerical expense in keeping records
   --Employee benefits for warehouse and administrative personnel
   --Taxes on inventory

5. Handling equipment
   --Taxes and insurance
   --Depreciation
   --Fuel expense
   --Maintenance and repairs

Shortage costs

1. Back order
   --Overtime
   --Special clerical and administrative work

2. Lost sales or consequences to the customer

Replenishment costs

1. Dependent on quantity
   --Transportation
   --Receiving
   --Inspection
   --Contract administration

2. Independent of quantity
   --Order handling (paper, postage, telephone, and so on)
   --Setting up production
   --Computer time to update records and the like

CHAPTER 4

THE NATURE AND MEASUREMENT OF INVENTORY DEMAND

Inventories are kept to meet demands for orders--that is, to satisfy demands--but managers rarely have sufficient knowledge about what generates demand to predict any patterns with certainty. Some demand patterns show enough regularity that, for purposes of adequate approximations, they can be treated as being known with certainty. Often, however, demand is more random and we have to describe it probabilistically. In this chapter, we explain the probability of demand and the demand forecast as two requisites for making optimal decisions about when and how much to order.

THE PROBABILITY OF DEMAND

Demand for units in stock can seldom be predicted with certainty. Both the time between demands and the number of units demanded typically vary within a range of values and can therefore be thought of as "random variables." That is, their specific places in the possible range of values are determined by chance. Thus, the chance or likelihood or "probability" of a given time or quantity can be found by studying a probability pattern or the "distribution" of the random occurrences. Here we show how to identify the type of probability distribution that demand data most nearly fit.

The probability distribution

Probabilities are usually expressed as percentages or proportions and are computed by dividing the total number of items, values, events, or whatever in a given group or universe by the total of all possible types of item, value, event, or whatever in the same universe. For example, in a universe of a total of 1,000 vouchers made up of 250 receiving vouchers, 700 shipping vouchers, and 50 inventory adjustment vouchers, the probability that any one voucher selected at random will be an inventory adjustment voucher is 0.05, or 50 divided by 1,000.

A listing of the possible values of a variable and their associated probabilities is called a "probability distribution." When all the possible probabilistic values are summed, the total will equal 1.00. The probability distribution for the universe of vouchers in the example above is

Type of voucher	Probability
Shipping	0.70
Receiving	0.25
Inventory adjustment	0.05
	1.00

Table 2

A Frequency Distribution for Weekly Demand

Weekly demand	Frequency	Probability of demand
Up to 46	1	0.008
46-55	1	0.008
56-65	3	0.025
66-75	7	0.058
76-85	11	0.092
86-95	21	0.175
96-105	28	0.234
106-115	16	0.133
116-125	22	0.183
126-135	7	0.058
136-145	1	0.008
146+	2	0.017
	120	1.000

To make a reasonable guess or hypothesis about the distribution of a random variable, it is necessary to collect and analyze data, whether this is historical or experimental. Data that have been collected are usually summarized in a chart called a "frequency distribution" such as that shown in table 2. Notice that the range of values in the table has been broken into equal intervals, or classes, and that the frequency within each interval or class has been recorded. This is a common practice with inventory demand data when the range of possible values is large. Using the frequency table helps us interpret the probability of an event as the proportion of the time in which similar events will occur in the long run. Thus, for the data in table 2, we can expect weekly demand to be between 96 and 105 over approximately 23 percent of the time.

We can also show the distribution graphically as in figure 5 on the next page. When we do, we find a characteristic shape to the distribution. Knowing the mean and the standard deviation of demand can help us predict what shape a distribution of demand values will take.

The mean and standard deviation

The mean is the arithmetic average of a set of numbers and it is widely used as a measure of "central tendency." We can compute the mean by adding the values of all cases in a distribution and dividing that sum by the total number of cases. We can state this algebraically as

$$\text{Mean} = \frac{\text{Sum of values}}{\text{Number of cases}} \quad \text{or} \quad \bar{X} = \frac{\Sigma X}{N}$$

Figure 5

A Frequency Distribution
for Weekly Demand

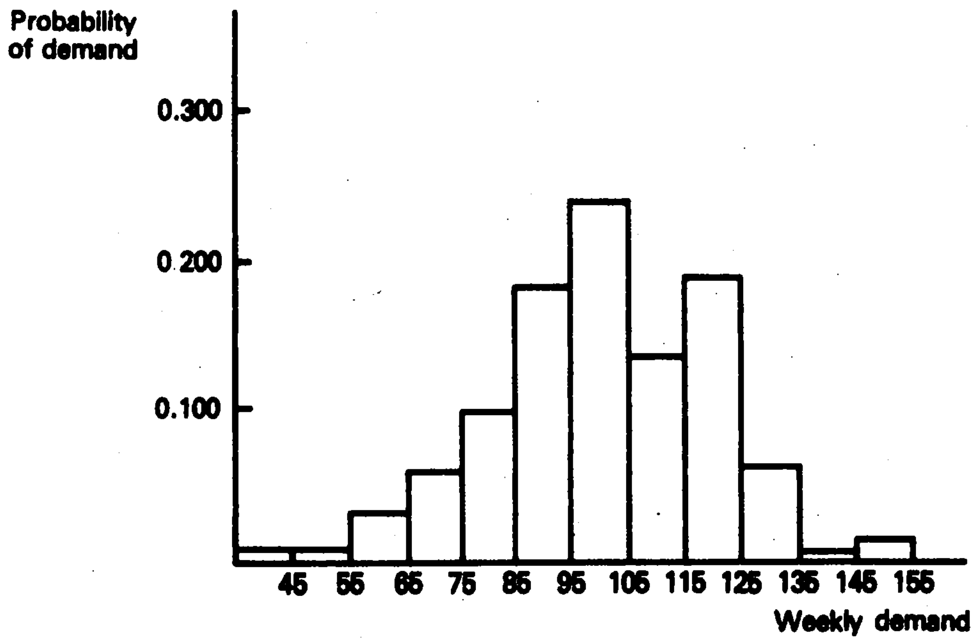

Figure 6

Mean and Standard Deviation
for a Single Distribution

X	$(X - \bar{X})$	$(X - \bar{X})^2$
3	−2	4
4	−1	1
5	0	0
6	1	1
7	2	4
25	0	10

Computation:

$$\bar{X} = \frac{\Sigma X}{N} = \frac{25}{5} = 5$$

$$S = \sqrt{\frac{\Sigma (X - \bar{X})^2}{N}} = \sqrt{\frac{10}{5}} = 1.4$$

For example, the set of values 1, 2, 3, 4, 5, sum to 15, and when we divide this sum by 5, or the number of cases, we find a mean, $\bar{X}$, of 3.

The standard deviation is the degree of spread in a series of numbers or a measure of dispersion. It is computed by extracting the square root of the average of the squared deviations of the individual values from their mean. Algebraically, it is denoted as

$$S = \sqrt{\frac{\Sigma(X - \bar{X})^2}{N}}$$

Figure 6 illustrates how the standard deviation is computed using the set of numbers 3, 4, 5, 6, and 7. 1/

Two sets of values can have the same mean but different standard deviations. For example, the set of numbers 1, 4, 3, 15, and 2 has the same mean as the set 3, 4, 5, 6, and 7 (the set in figure 6) but a different standard deviation--5.1 rather than 1.4. This difference reflects the fact that the values in 1, 4, 3, 15, and 2 are more widely dispersed around their mean than are the values in the other set. Knowing the amount of dispersion about the mean is especially important in inventory decisionmaking.

Some theoretical distributions

The frequency of observed data often compares well with some theoretical frequency distribution. Accordingly, mathematicians have formulated a number of theoretical probability distributions that approximate demand data found in inventory control situations. We should also take note of the recent observation that knowing the precise form of a demand distribution is not essential for determining optimal decisions. (Naddor, 1978, pp. 1769-72) In many cases, an optimal decision depends on the mean and standard deviation of demand but not on the specific form of the distributions. Nevertheless, we discuss briefly three theoretical distributions that are sometimes encountered.

1/The sum of the squares is sometimes divided by N - 1 rather than N when a sample of values is being studied rather than all values. In such cases, the standard deviation of the population, since it is unknown, must be estimated. McAllister (1975, p. 62) offers this reflection: "Previous samples . . . showed that . . . extreme population values . . . were seldom chosen. Hence, the variability of items about the sample mean will be smaller on the average than the variability of the population items about the population mean. If the numerator [for the standard deviation] is too small, then a correspondingly small denominator will compensate for this effect."

The normal distribution

The normal distribution is the theoretical distribution used most often in statistics because it represents a wide variety of actual distributions in nature and because it simplifies a number of statistical calculations. It is characterized as a "bell-shaped curve," being symmetrical about the mean, as can be seen in figure 7. One of its unique features is that it enables us to make an explicit determination of the probability of a certain value of customer demand per unit time being exceeded. For example, if demand is distributed normally (or approximately so), the mean value will be within 2 standard deviations 95.44 percent of the time. We can put this another way by saying that the mean value plus 2 standard deviations will, on average, be exceeded only 2.28 percent of the time. This concept is illustrated in figure 7 for various multiples of the standard deviation. Further values can be obtained from normal distribution function tables readily available in almost every statistics textbook.

According to C. D. Lewis, the normal distribution is often used to provide an approximate fit to the demand distribution at the factory level because the averaging that naturally takes place as customer orders are aggregated in the upward process from retailer to wholesaler to factory tends to produce this type of distribution. (Lewis, 1975, p. 20) However, use of the normal distribution may imply the existence of a "negative" demand, as in stock being returned, which may not be appropriate in many government situations. For example, if the mean demand is 10 units and the standard deviation is 5 units, then the normal distribution would indicate that demand is less than 0 (10 units minus 2 standard deviations) 2.28 percent of the time. This may not be a reasonable assumption.

The gamma distribution

Several practitioners have found the gamma distribution particularly useful in describing demand values, especially when demand cannot assume negative values. (Lewis, 1975, p. 24) The distribution is expressed by the following equation:

$$P(X) = \beta^{\alpha} X^{(\alpha-1)} e^{-\beta X} / (\alpha - 1)!$$

Unlike the normal distribution, the gamma distribution does not have a single characteristic shape, being defined in terms of two parameters, α and β, where α is the shape parameter and β is the scale parameter. These parameters are defined in terms of the mean and the standard deviation. The gamma distribution is illustrated in figure 8.

The Poisson distribution

In many applications involving failure statistics, such as for spare parts items, the Poisson distribution has been found

Figure 7

The Normal Distribution

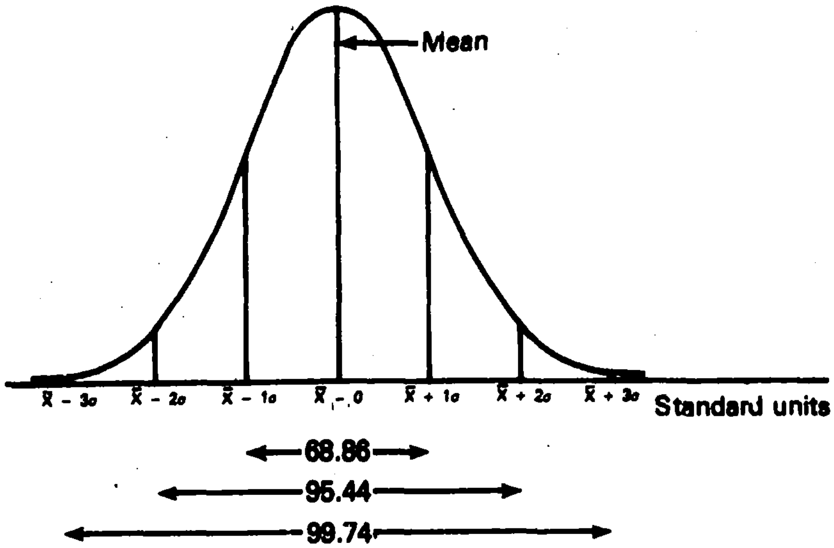

Percentage of area under the normal curve

Figure 8

The Gamma Distribution

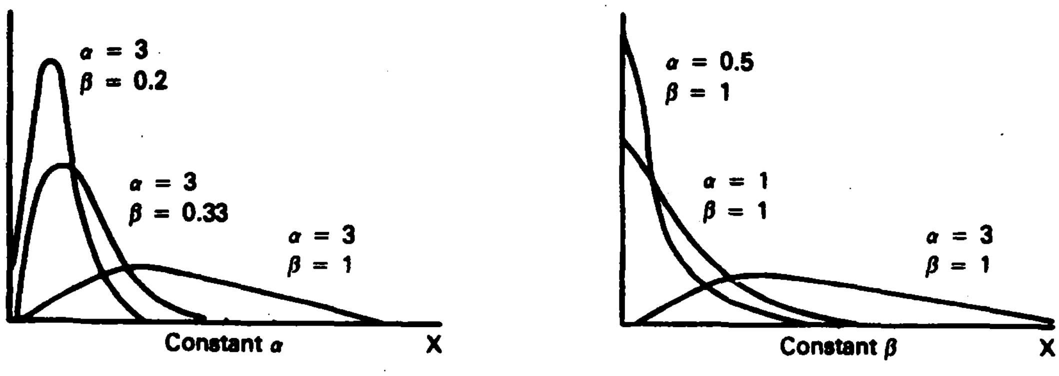

Mean $\bar{X} = \dfrac{\alpha}{\beta}$

Standard deviation $S = \sqrt{\dfrac{\alpha}{\beta^2}}$

Source: Adapted from R. E. Shannon, Systems Simulation: The Art and Science (Englewood Cliffs, N.J.: Prentice-Hall, 1975), p. 364.

Figure 9

The Poisson Distribution

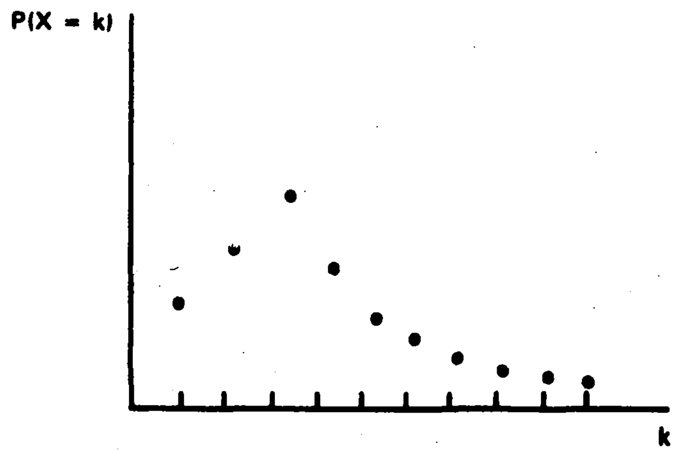

to be the simplest and most convenient to use. (Rudwick, 1969, p. 268) It contains only one parameter--the average failure rate. The standard deviation is equal to the square root of the mean. Mathematically, the Poisson distribution is expressed by the following equation:

$$P(X) = e^{-\lambda t} \cdot (\lambda t)^X / X!$$

This function gives an expression for the probability, P(X), of an event X occurring within a time period t with an average rate of λ per unit time. Unlike the normal and the gamma distributions, the Poisson is what is called a "discrete" distribution. That is, values for event X can be only integers.

The characteristic shape, shown in figure 9, is skewed to the left--that is, the bulk of the distribution is to the left of the mean. Thus, one way of detecting whether a distribution is likely to be approximately Poisson is to see whether it is skewed to the left (and whether the standard deviation is approximately equal to the square root of the mean).

Tables of values of the Poisson distribution do not generally include values for a mean above 20. For higher values of the mean, the distribution becomes very nearly normal. Thus, the Poisson distribution is usually characteristic only for a fairly narrow range of low, average values.

Other distributions

Demand data can be approximated from other theoretical distributions. The three we have discussed above are among the more frequently used. We avoided others that are mathematically more complex. The handbook of statistical distributions by N. A. J. Hastings and J. B. Peacock (1964) is a good reference for other distributions. We repeat, however, that the precise form of demand distribution is generally not essential for determining optimal decisions. Optimal decisions frequently depend only on the mean demand and its standard deviation.

THE DEMAND FORECAST

A demand forecast is an attempt to predict the mean of an assumed probability distribution. (Recall that the standard deviation also helps determine the shape of a distribution.) It is the link between the external and uncontrollable movements of the environment and the internal and uncontrollable affairs of an organization. Reliable, inexpensive demand forecasts are essential for planning soundly designed and operated inventory systems. In their absence, managing inventories is futile. Forecasting is generally based on collective opinion, on historical data, trends, and patterns, or on related information.

Forecasting from collective opinion

This method of forecasting is basically qualitative analysis built from the opinions of experts such as managers, wholesalers and retailers, customers, and economists. Market research, panel consensus, and the Delphi method are three approaches. 1/ The method might be used when data are scarce, as when a new product is being introduced to a market. It might be useful for building forecasts by product, customer group, or geographical territory. Regardless of the availability of other forecasting methods, it is useful for providing approximate estimates, and it can be used to verify or modify a forecast developed by other means.

Forecasting from historical data

This method of forecasting is based on the assumption that although change is going to take place, what has gone on before will continue. Its techniques are the statistical techniques of time series analysis and projection--or moving averages, exponential smoothing, and trend projections, as we discuss later in the chapter. Acceptable degrees of reliability in its use depend on there being sufficient data about the inventoried item.

To understand the forces that make for change and to understand and predict their future implications, the demand or use rate of an inventoried item may be described as subject to the several forces that operate on it simultaneously. Thus, one basic approach to forecasting demand from historical data is to decompose demand into its basic components. These are average demand, trends in the average, seasonal patterns, cyclical patterns, and random variations. The data plotted in figure 10 have been decomposed into seasonal pattern and random variations together with their linear trend as shown in figure 11.

1/The Delphi method is an attempt to obtain group consensus through revised estimations after additional group information is circulated in written form anonymously within a group.

Figure 10

A Plotting of Hypothetical Data

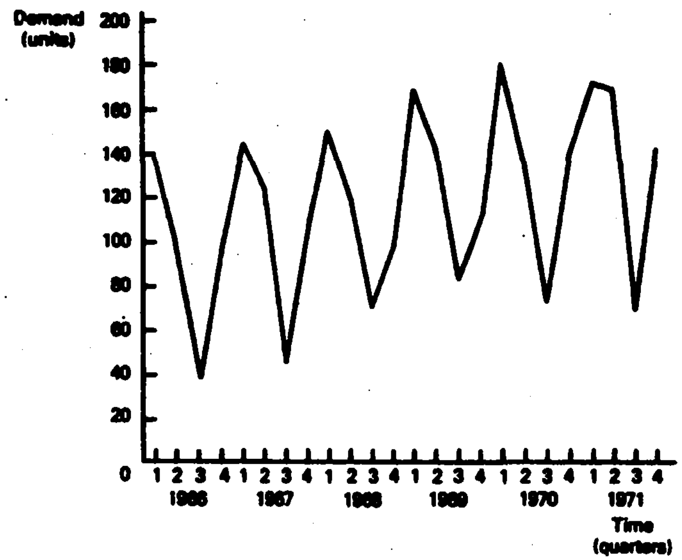

Quarter		Units
1966	1	138
	2	100
	3	40
	4	98
1967	1	145
	2	125
	3	44
	4	106
1968	1	152
	2	120
	3	68
	4	100
1969	1	169
	2	140
	3	79
	4	108
1970	1	180
	2	133
	3	72
	4	138
1971	1	168
	2	166
	3	70
	4	140

Figure 11

The Hypothetical Data Decomposed

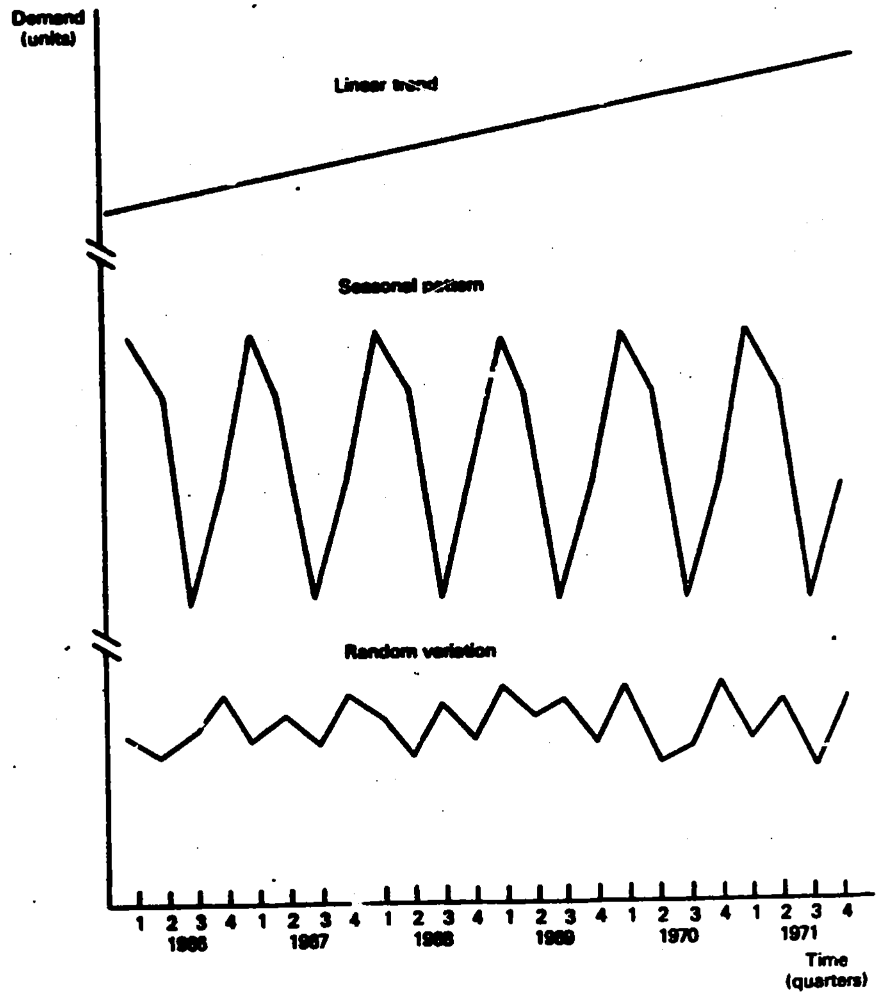

In forecasting from such components, we would analyze and predict each one and then combine them to construct a forecast.

Time series analysis and projection techniques are especially useful for forecasting the near future, but their reliability declines rapidly the further into the future we attempt to see with them. The two fundamental questions that must be answered in using historical data for forecasting are whether they are available and whether they are relevant. If the item in question has been inventoried for several years, demand rates, service levels, delivery patterns, and the like may all be available but often such data have not been accumulated in the forms that are the most appropriate for forecast decisionmaking. That is, the historical data may not have been clearly identified by specific items, at specific times and locations, or by specific classes of customer.

Even when historical data are available in usable forms, they may not be relevant for estimating the future. Some forecasting techniques are based on the assumption that existing patterns will continue into the future. If significant changes are taking place in the inventory system--as when new products are being designed or old ones are being improved, when customers are new or changing, when various buying behaviors differ, and so on--the similarity between the past and the future will diminish.

Forecasting from related information

Forecasts based on related information attempt to make explicit the relationships between the factor to be forecast and other factors. The forecaster tries to explain and understand the forces affecting future sales levels, either because it is easier to forecast these forces directly than it is to forecast sales or because they lead sales and are therefore known in advance of sales. Some of the more popular methods for doing this use regression models, econometric models, and input-output models. As a whole, these methods are more costly and time-consuming than other methods.

The most relevant techniques

An inventory system may need thousands of individual item forecasts, and the cost of developing them is an important consideration. A significant problem in designing forecasting systems for inventory control is to achieve an appropriate balance between the forecast's accuracy and its cost. Many inventory systems require forecasting procedures that are simple, computer-based routines that provide reasonable, short-range accuracy, but computational requirements and costs generally increase rapidly as refinements are added to the forecasting system. The point at which marginal savings from increased forecasting accuracy begin to exceed marginal costs is usually reached with relatively simple, automatic forecasting procedures.

Automatic forecasting procedures are

"objective computational routines that measure past patterns in the data series and generate new data with similar patterns for future periods. The causal factors underlying most product demand series, and thus, the pattern of these series, undergo changes through them. Forecasting systems based upon extending past patterns into the near future should therefore be responsive to these changes. New data should be utilized as soon as they are available and, since past observations become increasingly less reliable for establishing model parameters, greater weight should be given to more recent observations in analyzing the patterns of past data. While the forecasting system should respond quickly to permanent changes in the demand pattern, it should not be overly responsive to purely random variations. The appropriate level of responsiveness changes not only for different product demand series, but for the same product through time." (Groff, 1970, p. 257)

The techniques we discuss in the remainder of this chapter meet these requirements to some degree.

Moving average

Average demand can be simply the average of all past data. The hypothetical data we looked at in figure 10, for example, show a simple average of 116.7 units per quarter. This figure probably would not serve as a very good basis for a projection, however, since the data obviously express some trend. Because of the apparent trend, the average of quarterly demand is said to be "moving" over time.

The concept of a moving average enables us to smooth out the data so that the underlying trend becomes easier to see and analyze. To calculate a moving average, data from several periods numbering at least three are added together and divided by the number of periods. After that, data from the next period can be added while data from the earliest period are dropped. Figure 12 shows a three-quarter and a four-quarter moving average for the data presented in figure 10.

Exponentially smoothed average

As demand takes place, actual use may deviate from what had been projected, partly or entirely because of random effects or because the average level of demand has changed. Exponential weighted averaging, or "smoothing," is a means of adjusting a moving average of demand by incorporating the most recent information without dropping any earlier information, as is done in simple moving averaging. The equation for the exponential weighted average is

$$S_2 = S_1 + \alpha(y_2 - S_1)$$

Figure 12

A Plotting of a Hypothetical Moving Average

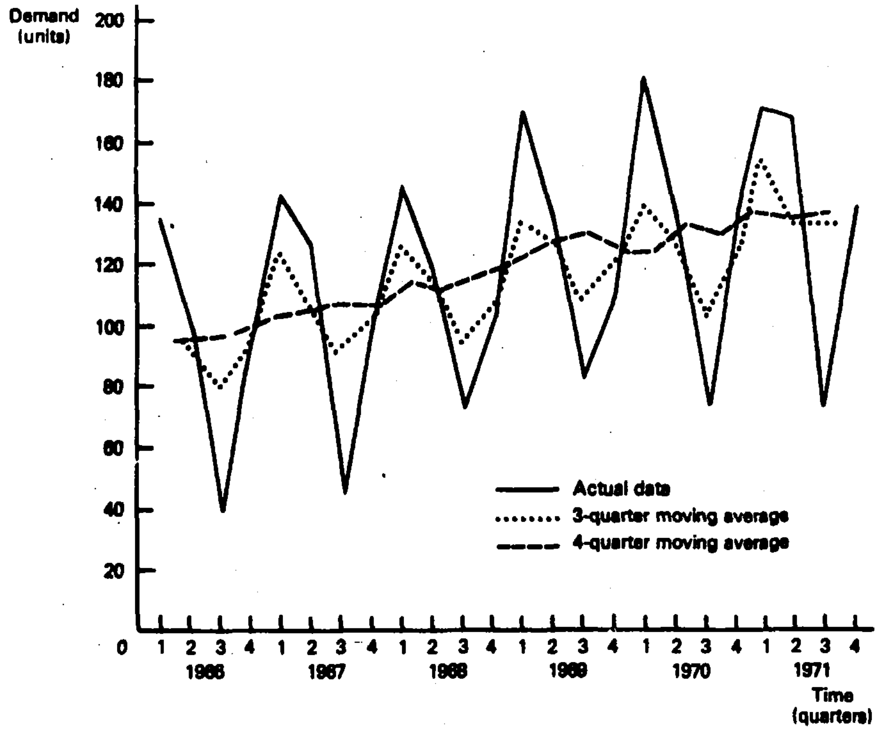

Quarter		Moving Average 3 Q	4 Q
1966	1	–	–
	2	93	94
	3	79	96
	4	94	102
1967	1	123	103
	2	106	105
	3	92	107
	4	101	108
1968	1	126	112
	2	113	110
	3	96	114
	4	112	119
1969	1	136	122
	2	129	124
	3	109	129
	4	122	125
1970	1	140	124
	2	128	131
	3	114	128
	4	126	136
1971	1	157	135
	2	135	136
	3	125	–
	4	–	–

where S_1 = average forecast in period 1, S_2 = average forecast in period 2, y_2 = actual demand in period 2, and α = exponential smoothing constant (0 to 1).

The forecast average--that is, the smoothed demand--for period 2 is equal to the forecast average for period 1 plus or minus some fraction of the deviation between the actual demand in period 2 and the forecast average for period 1. This fraction, the alpha value, can range from 0 to 1. It is usually set at 0.10 to 0.20. To make the computation easier, the terms can be rearranged to

$$S_2 = \alpha y_2 + (1 - \alpha)S_1$$

For any succeeding time period t, the smoothed value S_t is found by computing

$$S_t = \alpha y_t + (1 - \alpha)S_{t-1} \text{ where } 0 \leq \alpha \leq 1$$

For example, returning to the data in figure 10, we can compute the following smoothed values with an alpha of 10 percent and S_0 equal to 100:

$$\begin{aligned} S_1 &= \alpha y_1 + (1 - \alpha)S_0 \\ &= (0.1)(138) + (0.9)(100) \\ &= 13.8 + 90.0 \\ &= 103.8 \end{aligned}$$

Figure 13

A Plotting of Exponentially Smoothed Averages

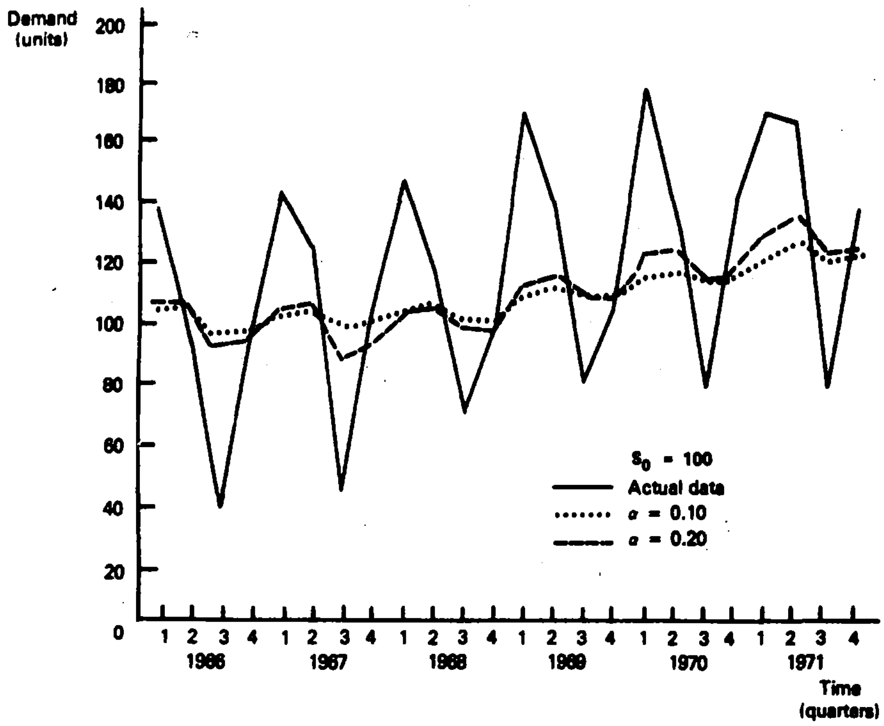

Quarter		Units	α = 0.10	α = 0.20
1966	1	138	103.8	107.6
	2	100	103.4	106.1
	3	40	97.1	92.9
	4	98	97.2	93.9
1967	1	146	102.8	104.1
	2	125	105.0	108.3
	3	44	98.9	95.4
	4	106	100.7	98.5
1968	1	152	106.8	102.0
	2	120	107.2	105.6
	3	68	103.3	97.8
	4	100	103.0	98.2
1969	1	169	109.6	112.3
	2	140	112.6	117.8
	3	79	109.2	110.0
	4	108	109.0	109.6
1970	1	180	116.1	123.6
	2	133	117.7	125.4
	3	72	113.1	114.7
	4	138	115.5	119.3
1971	1	168	120.7	129.0
	2	166	125.2	136.4
	3	70	119.6	123.1
	4	140	121.6	126.5

$$
\begin{aligned}
S_2 &= \alpha y_2 + (1-\alpha)S_1 \\
&= (0.1)(100) + (0.9)(103.8) \\
&= 10.0 + 93.4 \\
&= 103.4
\end{aligned}
$$

In figure 13, the rest of the S_t values are calculated and also plotted for alphas of 10 and 20 percent.

The fundamental idea of this and other exponentially smoothed models is that of filtering out "noise" or random variations. When actual demand changes gradually, the forecasting system can track the changes rather well. If, however, demand changes suddenly and permanently, a forecasting system using small values for the smoothing constants will lag substantially behind the actual change.

Recently, several schemes have been proposed for the adaptive control of exponential smoothing constants. In many of these schemes, it is assumed that forecasted values that lag behind actual demand values will cause biased error when error is defined as the difference between forecast and actual demand in each period. When the forecasting system is tracking changes in demand correctly, the positive errors in some periods are balanced by negative errors in others, so error is not biased one way or another. Adaptive systems detect any tendency toward bias and correct the smoothing constants to remove it.

Other techniques

Many forecasting techniques that are relevant for inventory management require sophisticated statistical analysis, but it is possible to describe some of them in brief and general terms. We mention six of the more promising models.

The exponentially weighted moving average (EWMA) is a technique appropriate for generating forecasts of an item whose demand has a pronounced seasonal effect. The EWMA model operates by separately estimating at each point in time the smoothed process average, the process trend, and the seasonal factor and then combining these components to build a forecast. The sinusoidal forecasting model is a least-squares technique using sine and cosine functions of time as the independent or predictor variables. It detects seasonal effects over time that may exist within the time series. The autoregressive forecasting model uses a linear technique that accounts for correlations between adjacent observations in a time series. Lagged process values assume the role of the independent or predictor variables. (Mendenhall and Reinmuth, 1971, pp. 413-20, 399-408)

Exponential smoothing is a special case of the Box-Jenkins model. The time series is fitted with an optimal mathematical model whose parameters must be estimated. The X-11 technique, developed by the Bureau of the Census, simultaneously removes seasonal data from raw data and fits a trend-cycle line to the data. It provides details on seasons, trends, the accuracy of the seasonal data and the trend cycle fit, and a number of other measures. (Chambers, Mullick, and Smith, 1971, pp. 71-72)

Curve-fitting techniques fit a trend line or curve to a mathematical equation and project into the future with it. The simplest type of trend is a linear function, although more complex curvilinear functions, such as polynomials or growth equations, are used.

EVALUATION CONSIDERATIONS

Ideally, inventory managers should have detailed information on the demand for each item inventoried over extended periods, but it is rarely available in a form that is useful for analysis. When it is available, sometimes there is more of it than can be handled efficiently. The usefulness of data depends on their ability to display the true characteristics of customer-demand variations. Some prevalent data sources are (1) orders showing the dates of receipt and shipment, (2) consolidated demand figures by day or week, (3) consolidated shipment records by day or week, and (4) records of inventory balances by day, week, or some other normal review period. Data in the first form are the most useful; usefulness declines in the order listed.

Keeping this in mind, an inventory manager might select a sample of items and develop the details of their demand characteristics. The manager might then study these characteristics however crudely arranged, to form some idea of average levels of demand and random variations, trends, seasonal variations, and so on. From this information, frequency charts and graphs may be prepared so that judgments may be made about the type of probability distribution that best approximates the data. Computer programs are available for simplifying this task.

In evaluating the forecasting aspects of inventory management, managers should attempt to keep the analysis conceptual and avoid getting lost in details. To this end, several important questions can be asked.

--Does the system attempt to create updated forecasts of expected demand for inventoried items?

--What is the average time period between successive forecasts?

--On what bases are forecasts made?

--How is the accuracy or reliability of forecasts measured?

--What is the record of the historical accuracy of the forecasting system being used?

--How much time and effort go into forecasting?

--What savings or advantages result from the forecasting system?

--Is there a pronounced seasonality to demand? If so, do forecasting techniques systematically take it into account?

--Have the forecasting techniques been automated or do they require substantial human computation and discretion?

--Do forecasting techniques require much computer storage capacity and time?

--Can the forecasting techniques being used be adapted to permanent changes in user demand levels and patterns?

Although concrete information is not likely to result from any one of these questions alone, overall patterns should emerge. Data evaluators can generally study the appropriateness of a forecasting system in terms of the reliability, timeliness, and economy of operation that are needed.

CHAPTER 5

THE NATURE AND MEASUREMENT

OF REPLENISHMENT LEAD TIME

In this chapter, we define some important properties of replenishment lead time and the peculiarities that make it a problem for inventory decisionmakers. We also describe the variability of lead time and some of the problems that make it difficult to forecast lead time duration.

LEAD TIME CHARACTERISTICS

Total lead time, also known as "order and ship" time, "procurement lead" time, and "lag" time, is the time between an order's being placed and the receipt of the new stock followed by its storage in a warehouse. Usually stated in terms of days, weeks, or months, it can also be seen as internal and external lead time.

Internal lead time consists of all or any of the following elements. It may be the time it takes to become aware that the stock position is low. This is usually the time between review cycles. It may be the time it takes to process request documents or the time of a delay in obtaining an approval of a request. It may be the time required to advertise a procurement or process an order. Internal lead time may be negotiation or advertising time or the time required to award a contract. It may be the time required to process the receipt of material, place it in storage, and make it ready for issue. (Kuhlman, 1969, p. 74)

External lead time is, as the term implies, the time delay in filling a request that is external to the supply manager's organization and, to some extent, beyond the manager's control. External lead time may consist of all or any of the following. It may be mail or transmittal time. It may be the time required to process orders or set up an operation. It may be manufacturing time. It may be inspection delay. It may be preservation and packing time or shipping time. (Kuhlman, 1969, p. 74)

When a supply manager can replenish stock from a nearby depot, delay may be as little as a day or two. If instantaneous delivery could be assured, lead time would be zero. If a stock item must be manufactured commercially, delay may extend for months or a year or more, depending on the complexity of the item. Thus, lead time can be identified only individually for each item, and what it will be depends on the facts of the supply source. Once it has been established, lead time cannot be ignored. It must be reviewed constantly under the dictates of actual experience and changes in supply sources.

Aggregate internal and external lead time should be considered in establishing inventory levels. For example, if the

Table 3

The Components of Pipeline Inventory

Movement	Average transit or delay time (days)	Average movement inventory (500 units per day)
Factory-to-factory warehouse	1	500
Delay at factory warehouse	2	1,000
Warehouse to distribution center	6	3,000
Delay at distribution center	2	1,000
Distribution center to user	3	1,500
Total	14	7,000

total lead time is 3 months, it may be desirable to establish a reorder point such that when it is reached, stock on hand is enough for at least 3 more months or, alternatively, to set the order level sufficiently high that a 3-month supply will be available whenever an order is placed. Stock held for these purposes is frequently referred to as "pipeline" inventory or "transit" or "movement" inventory.

At any one time, an organization may need to have thousands of items in pipeline inventories, representing an investment of millions of dollars. The average pipeline inventory can be calculated from transit times, delays, handling times, and use rates for the inventory system. The simple example in table 3 illustrates the calculation of pipeline inventories for an organization using an average of 3,500 units of an item each week--that is, 500 units each day.

In many situations, it is not necessary to hold great quantities of stock if suppliers' delivery times are strictly adhered to. Unfortunately, delivery times vary considerably, making the holding of some stock essential. Some managers seeking to replenish stock think that suppliers control lead time, particularly when deliveries are late. On closer examination, however, delay can often be attributed to the purchasing organization or other agents as well as the supplier.

C. D. Lewis gives an interesting illustration of the constituents of delay. His scenario begins at the earliest point in time that the purchasing organization realizes a replenishment order should be placed. The first delay is at the purchasing organization in transmitting the replenishment information to its own buying department. The second delay occurs when the purchasing organization's buying department delays in compiling the order. Next, the post office delays in sending the order to the supplier. The supplier creates the fourth, fifth, and sixth delays in processing the order, then in manufacturing the items ordered or in taking them from stock, and finally in packaging and shipping the order. Seventh is a transit delay, whether postal, rail, or road. The eighth delay occurs in the purchasing organization's "goods inwards" department in accepting and unpacking the order, the ninth in its inspecting and controlling the quality of goods received, and the tenth in its getting the replenished stock into stores ready for issue. Finally, after clearing the paperwork, the replenished item is ready to be issued. (Lewis, 1975, p. 101) Of these ten delays, five--the first two and the last three--lie within the control of the purchasing organization. It is rare that five delays make up 50 percent of the total. Nevertheless, by no means all delays in overall lead time can be attributed to suppliers.

LEAD TIME VARIABILITY

Lead time is generally not constant. The time to fill an order at its source, the shipping time, the time required for paperwork, and the time for various other activities can vary from one order to another. Even though lead time duration varies, organizations rarely analyze it to estimate its mean and standard deviation.

In chapter 4, we pointed out that demand data are often not available. This is even more true with lead time data. When lead times are fairly long and orders are placed infrequently, data are rarely sufficient to yield a probability distribution for lead time. Sometimes about the best that can be done is to obtain crude estimates of what the maximum and minimum lead times are, average them to find the mean lead time, and estimate the standard deviation as the range divided by six (assuming that a normal distribution of six standard deviations, three on either side of the mean, will include essentially all the possible values of the distribution). (Hadley and Whitin, 1964, p. 419)

In addition to the difficulty of trying to estimate lead time distribution, it may be that it is not stationary, since lead times tend to change continually. Furthermore, some orders may be split and not shipped at one time. Some orders may be expedited if it seems they will go out of stock. Because lead time and demand forecasts predict the future, and because conditions may change, the best of information may change too and the forecasts become erroneous.

Another important but unpublished observation is that lead time variability will not affect optimal decision rules significantly if the system can generate a "good" estimate of the lead time mean or average. 1/ That is, optimum decisions on when and how much to order will not vary significantly in the long run if an accurate forecast of the mean replenishment lead time is available, even though lead time varies.

EVALUATING LEAD TIME ESTIMATES

Most questions pertaining to forecasting demand also pertain to analyzing how an organization forecasts lead time. Some additional considerations that can be used as a guide for evaluation are:

--In a comparison of established lead times with actual lead times, what are the reasons for differences?

--Have both internal and external elements been considered in estimating lead time?

--How much stock is pipeline inventory? Can this be reduced?

--In an evaluation of the actual physical movement of a replenishment action, what aspects of it does it seem the decisionmaker can control?

--Could lead time be shortened if several commodities were ordered from the same supplier simultaneously? (Recall that this practice saves on replenishment cost.)

--In a consideration of the effects of revised order quantities on lead time, do large quantities, for example, have a longer or a shorter replenishment cycle?

--Has a priority system been established for expediting more essential, costly items in replenishing stock or is lead time longer for these items because of more intensive management?

--Is lead time treated as constant or as variable in deciding when and how much to order?

1/The observation was made by Levy and Naddor and derives from their work in Levy, 1979a, and Levy and Naddor, 1979b.

CHAPTER 6

ASSIGNING OPTIMAL VALUES

TO POLICY VARIABLES

Now that we have examined the aspects of inventory systems that have to be considered in making proper decisions about when and how much to order, we can discuss how to use this information to arrive at optimal decision values. The usual means is mathematical analysis under one of three approaches--the analytic, simulation, and heuristic approaches.

Before discussing these techniques in detail, it might be beneficial to reiterate the properties of inventory systems that influence how decisions are made. First are the three distinct decisionmaking policies--the reorder point-fixed order quantity policy, the fixed interval-order level policy, and the reorder point-order level policy--one of which must be chosen. Further, if costs are to be minimized, a decision about when and how much to order must be made in view of conflicting cost functions, variations in demand, and variations in replenishment lead time. Thus, inventory systems are characterized by policy (reorder point or inventory review interval, order quantity or order level), cost (carrying, shortage, replenishing), demand (constant, variable), and lead time (insignificant or significant, constant or variable) among other factors. These are summarized in table 4. Managers should consider all factors when making inventory control decisions if they want to minimize costs.

Table 4

Inventory System Properties

POLICY--WHEN AND HOW MUCH

1. Reorder point-fixed order quantity
2. Fixed interval-order level
3. Reorder point-order level

COST

1. Carrying
2. Shortage
   a. Explicit
   b. Implicit
3. Replenishment
   a. Constant
   b. Variable

DEMAND

1. Constant
2. Variable or probabilistic

LEAD TIME

1. Constant
   a. Virtually zero
   b. Positive
2. Variable

The literature on optimizing inventory control decisions is quite extensive. Many analysts find the problems of optimizing inventory control policy variables mathematically interesting because there is a seemingly infinite number of possible models. Furthermore, the same basic principle may be applied in the social and political sciences as in determining what to do in a situation in which doing either too much or too little will result in excessive costs or unduly low benefits. (Nagel and Neef, 1976 and 1979)

Our intention is not to synthesize the copious literature on the subject. Instead, we give an overview of three techniques for arriving at decisions that will tend to minimize costs. Evaluating precisely how well a government organization is making inventory decisions will, in all likelihood, require some expertise in quantitative methods. 1/ Nevertheless, the pertinent theories, assumptions, approximations, and judgments can be understood in general terms.

THE ANALYTIC APPROACH

In the analytic approach, operating rules that minimize costs are determined by constructing a mathematical model of the system. This method has as its primary purpose to develop optimal decision rules and determine the minimum total cost of systems. It can also be used to analyze the sensitivity of results and to compare different inventory policies.

In operations research literature, the concept of a mathematical model is presented as a functional relation between some appropriate measure of effectiveness or utility and a set of controllable and uncontrollable variables. (Ackoff, 1962, p. 111) Such a characterization has the form $V = f(X_i, Y_j)$, where V is the measure of value or utility for the system under study (cost, for example), X_i is the set of controllable variables (when and how much to order), and Y_j is the set of uncontrollable variables that affect performance (the cost of the item, demand factors, certain lead time factors); f signifies the functional relationship between the independent variables and constraints, X_i and Y_j, and the dependent variable V.

While this may be considered the prototype of all models of problem situations, an actual model may contain several equations and inequalities. Because there is considerable flexibility in model construction, few principles can assist analysts in this phase of work. In fact, the process by which one derives a model of a system has been described as an intuitive art. Any set of rules for developing models has limited usefulness at best and can serve only as a suggested framework.

1/GAO evaluators should refer to chapter 11 of the GAO Project Manual.

The few models we discuss are simple and we present them in order to provide some idea of the optimization formulas generally in use. The "EOQ" formula is famous and the simplest of all models because of its underlying restrictive assumptions. After we show these, we discuss how they can be altered in two other formulas for determining minimum system costs.

The EOQ model

Ford Harris of the Westinghouse Corporation used mathematical methods in inventory analysis as early as 1915. (Hadley and Whitin, 1963, pp. 2-4; Naddor, 1966, pp. 16-17) His formula for deriving the most "economical order quantity" (EOQ) has had more applications than any other single result in inventory systems analysis. The formula can be applied to systems with the following characteristics--carrying cost and replenishment cost are constant; no shortages occur and demand is always met and, thus, shortage cost is not considered; the precise rate of demand is known and constant; the replenishment lead time is known and constant.

Inventory managers who use the formula can compute the order quantity that will balance the cost of carrying inventory with the cost of replenishing stock. The formula, whose derivation is explained in most standard textbooks on inventory theory, is as follows:

$$Q = \sqrt{\frac{2^{*}D^{*}C_3}{C_1}}$$

where

Q = the economical order quantity
D = demand, usually in units per year
C_1 = unit carrying charge, usually in dollars per year
C_3 = the cost of replenishment in dollars per replenishment

It is important to note that the formula answers only "How much to buy?" "When to buy?" is a separate issue. Because of the restrictive assumptions upon which the formula is based, however, "When to buy?" is relatively easy to answer. Orders are placed whenever the amount on hand equals the expected demand during the lead time period. This will insure that some stock will always be on hand when a demand occurs (no shortages being allowed or expected). As this implies, the EOQ model is usually used in connection with reorder point policies, inasmuch as the reorder point, R, equals the lead time demand, Lt*D--that is, R = Lt*D.

The total expected cost (TEC) of managing items in this particular system can be computed by using the following formula:

$$TEC = \frac{Q}{2} * C_1 + \frac{D}{Q} * C_3$$

where

$\frac{Q}{2} * C_1$ = yearly cost to carry the item in inventory, Q/2 being the average amount of inventory on hand

$\frac{D}{Q} * C_3$ = yearly cost of replenishment, D/Q being the number of orders placed per year

Let us consider an example. Suppose

Demand	= 500 units per year
Lead time	= 6 months, or 0.5 years
Replenishment cost	= $10.00 per order
Carrying cost	= $0.10 per unit per year

Then

$$EOQ = \sqrt{\frac{2*500*10}{0.10}} = 316$$

$$\begin{aligned} \text{Lead time demand} &= 500(0.5) \\ &= 250 \\ R &= 250 \\ TEC &= \frac{316}{2}(0.10) + \frac{500}{316}(10) \\ &= \$15.81 + \$15.81 \\ &= \$31.62 \end{aligned}$$

Thus, when the inventory level reaches 250 units, an order for 316 units is placed. The order will be received in 6 months, at which time the inventory level will be zero. The cost of managing the item--that is, carrying it in inventory and placing orders--is about $31.62. The cost to carry and the cost to order are equal--that is, they are balanced.

The stockouts permitted model

Permitting stockouts implies a shortage cost. All other assumptions are the same as in the EOQ model. Shortage cost may be difficult to measure precisely, as we discussed in chapter 3, but management-imposed service goals or fill rates imply that there will be one. Thus, a 90 percent fill rate, for example, suggests that it is acceptable to be out of stock 10 percent of the time. This reduces the cost of carrying an item by 10 percent over what it would be if stockouts were not permitted, valuing the cost of maintaining 1 unit on back order at 9 times the cost of carrying 1 unit in inventory for the same length of time.

Using this concept, we can find the optimal order quantity (Q), reorder point (R), and total expected cost (TEC) from the following expressions:

$$Q = \sqrt{\frac{2*D*C_3}{C_1*(F/100)}}$$

$$R = D*Lt - Q*(100 - F)/100$$

$$TEC = \frac{S^2}{2Q}*C_1 + \frac{(Q-S)^2}{2Q}*C_2 + \frac{D}{Q}*C_3$$

where

D = annual demand
F = fill rate percentage
C_1 = unit carrying cost per year
C_2 = unit shortage cost; recall that $C_2 = [C_1*F(1 - F)]$
C_3 = replenishment cost
Lt = replenishment lead time
S = expected inventory level after an order has been received-- in this case, $S = Q + R - D*Lt$

Consider the following example, in which it is given that

D = 500 units per year
Lt = 0.5 year
C_1 = \$0.10 per unit per year
F = 95 percent
C_2 = \$1.90 per unit per year, computed from C_1 and F
C_3 = \$10 per order

Then we can compute as follows:

$$Q = \sqrt{\frac{2*500*10}{(0.1)*(0.95)}} = 324$$

$$R = 500*0.5 - 324*0.05 = 234$$

$$TEC = \frac{308^2}{2*324}*0.1 + \frac{16^2}{2*324}*1.9 + \frac{500}{324}*10 = 14.64 + 0.75 + 15.43 = \$30.82$$

Thus, when the inventory level reaches 234 units, an order for 324 units is placed. The cost of managing the item under these conditions is \$30.82 per year. No other combination of reorder point and order quantity will give a lower cost.

The probabilistic demand model

In the probabilistic demand model, we change one more of the four original assumptions, so that demand will vary, taking on a probabilistic nature. Thus, the systems assumptions are now that carrying and replenishment costs are constant, stockouts are allowed and thus there is a constant shortage cost, demand varies, and replenishment lead time is known and constant.

With these assumptions, the mathematics becomes more complicated. Continuing this trend, the next type of model would show not only demand varying but also lead time. We conclude our discussion on the analytic approach with an illustration of the probabilistic demand model, because the level of mathematics needed to go any further becomes greater than intended for this document. Readers interested in the formulas and their manipulation may refer to the bibliography in appendix II.

The model we discuss here is sometimes referred to as the "ship provisioning model" because of its applicability in determining optimum stock levels for spare parts when ships are being constructed. These parts are generally expensive but they are bought even though it is unlikely they will be used, because it is considerably more expensive to buy them when they are needed than to keep them on hand.

Since there is only one period of concern--the life of a ship--deciding when to order is moot. The only decision is how much to order. Hence, the objective is to store enough parts to meet future demands efficiently, which fluctuate because of equipment failure. Lead time--that is, production lead time--is not important either, since presumably the spare parts are being produced along with the construction of the ship. Lead time is zero.

Two possible situations can be encountered, as depicted in figure 14. In situation A, too many parts have been purchased for the period of concern, resulting in an unnecessary investment cost, even though the leftover parts may have some salvage value that will partially offset this. In situation B, the stock has been depleted before the end of the period, and a penalty must be paid either for having to construct and transport the part that is needed or for having the ship (or aircraft, tank, combat unit, or whatever) go completely out of commission.

Situations A and B both indicate that three costs must be considered--(1) the buying price (B) for one item, (2) a salvage value (S) for each item left at the end of the period of concern, and (3) a penalty cost (P) for each item needed but not in stock. Since the model incorporates demand that is known not exactly but only probabilistically, decisions can be made only in terms of the "expected" cost of stocking a certain number of items. Expected values are computed by multiplying the probability of

Figure 14

Inventory Positions in the Probabilistic Demand Model

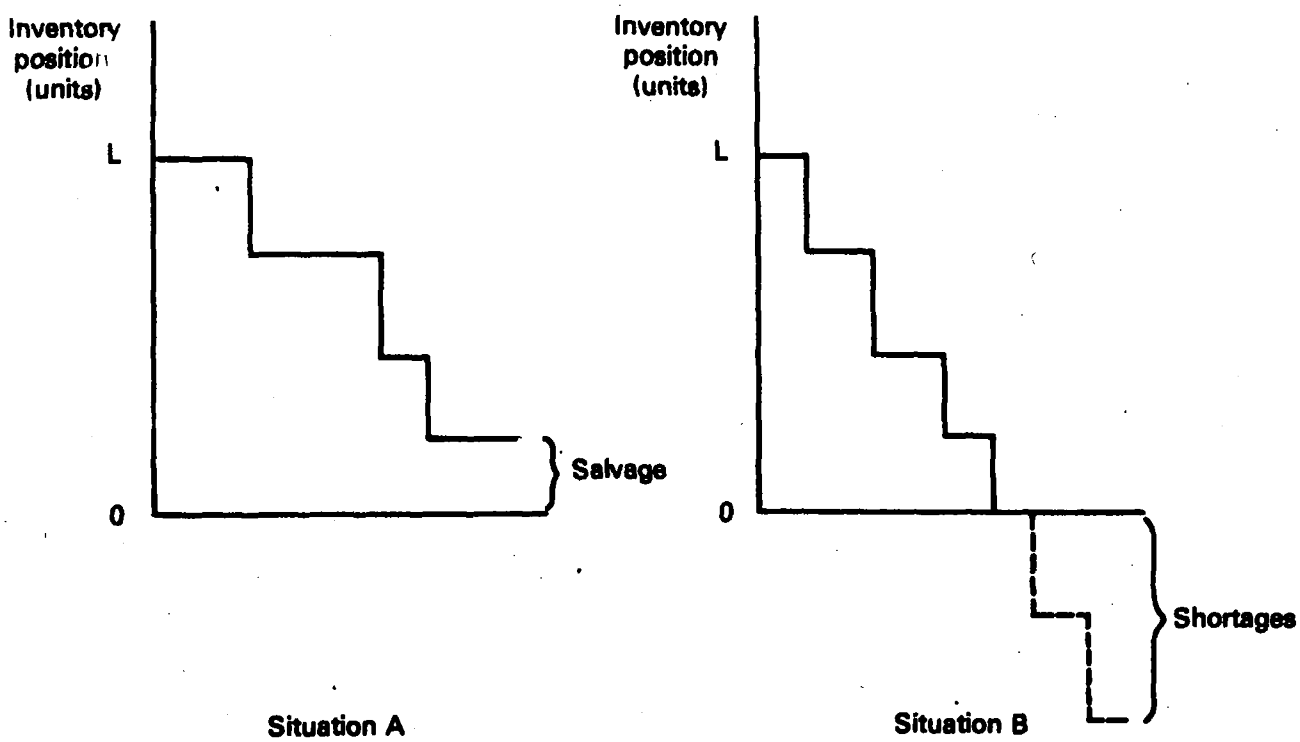

an outcome by the value of the outcome if it does occur and then summing the products obtained.

To obtain the optimum stock level (L) we could tabulate the expected cost of various stock levels. The level that gives the lowest expected cost is the optimum amount to stock. A formula derived from this concept gives the optimal L without computing the expected cost for all values. That is, the optimal value of L is one for which

$$Pr(X \leq L - 1) \leq (P - B)/(P - S) \leq Pr(X \leq L)$$

in which the righthand term is "the probability that demand (X) is less than or equal to L."

Suppose, for example, that a submarine is being provisioned for a polar voyage. One of the items to be carried is a very specialized bilge pump, and we have to determine the number of reserve pumps to carry. They cost $3,700 each and have a salvage value of $1,000 if they are not used. If the initial provisioning is inadequate, however, the cost of manufacturing and delivering a replacement pump will be $15,000, including production, delivery, and downtime costs. The distribution of demand, determined from past history of similar activities, is as follows:

Demand X	Probability Pr(X)
0	0.10
1	0.15
2	0.20
3	0.30
4	0.20
5	0.05

The probability that demand will be less than or equal to level L is then

Level L	$Pr(X \leq L)$
0	0.10
1	0.25
2	0.45
3	0.75
4	0.95
5	1.00

Using the concept of expected value, we can construct the following cost table: 1/

Level	Cost	Expected penalty	Expected salvage	Total cost
0	$ 0	$37,500	$ 0	$37,500
1	3,700	24,000	100	27,600
2	7,400	12,750	350	19,800
3	11,100	4,500	800	14,800
4	14,800	750	1,550	14,000
5	18,500	0	2,500	16,000

Thus, the optimal decision would be to procure 4 pumps, since this choice has the lowest expected cost. Using the optimization formula above, we find

$$(P - B)/(P - S) = (15{,}000 - 3{,}700)/(15{,}000 - 1{,}000) = 0.807$$

and, since $Pr(L = 3) = 0.75 < 0.807 < Pr(L = 4) = 0.95$

the optimal level L is 4.

1/For example, if one pump is bought but never needed, there will be a return of $1,000 in salvage. However, since there is only a 10 percent chance of this occurring, the expected salvage value is only $100(0.10*1000) when the pump is purchased. All figures for expected values in the table are computed in this manner.

THE SIMULATION APPROACH

Simulation is the development and use of models to aid in evaluating ideas and studying systems or situations. The essential characteristic of simulation is understood in the observation that "a model represents a phenomenon, but that simulation imitates it." (Ackoff, 1962, p. 346) It allows experimentation with systems that would otherwise be impossible or impractical. Simulation is not unique to the analysis of inventory or supply systems, having numerous applications in almost every social, economic, technological, and humanistic endeavor. This concept is both simple and intuitively appealing.

Simulation is, however, based heavily on computer science, the mathematics of probability, and statistics. Thus, simulation models of real systems are usually computerized, though this is not always true. Experiments with them are for the purpose of either evaluating various strategies for the operation of a system or understanding the behavior of the system. In the rest of this section, we illustrate both purposes with examples that keep the mathematics as simple as possible.

Evaluating alternative decisions

Simulation is used frequently to compare new operating rules with others whose experience is known. This is generally done for a system for which no valid analytical model has yet been developed; it can be simulated once with its new rules and again with old ones, employing the same pattern of demand in both. Thus, a comparison can be made by computing the costs that would be incurred under the two operating doctrines over the length of time for which the simulations were carried out.

The following example illustrates this with a relatively simple computer simulation model. Suppose that demand is probabilistic with the following distribution:

Monthly demand	Probability of demand
10	0.2
20	0.2
30	0.3
40	0.2
50	0.1

where

C_1 = \$0.75 per item per month
F = 90 percent
C_2 = \$6.75 per stockout per month
C_3 = \$100 per order
Lt = 2 months

Table 5

12-Month Simulation of Reorder Point-Order Level Policy--Reorder Point = 50 and Order Level = 120

Month	Beginning inventory	Demand	Ending inventory	On hand, on order	Ordered	Received
1	70	30	40	40	80	0
2	40	10	30	110	0	0
3	30	50	-20	60	0	80
4	60	10	50	50	70	0
5	50	30	20	90	0	0
6	20	30	-10	60	0	70
7	60	10	50	50	70	0
8	50	30	20	90	0	0
9	20	20	0	70	0	70
10	70	20	50	50	70	0
11	50	50	0	70	0	0
12	0	40	-40	30	90	70

	Carried	Short	Orders
Average monthly amount	31.7	2.1	0.4
Average monthly cost	$23.79	$14.44	$41.67

Total cost: $79.90

Table 6

12-Month Simulation of Reorder Point-Order Level Policy--Reorder Point = 40 and Order Level = 100

Month	Beginning inventory	Demand	Ending inventory	On hand, on order	Ordered	Received
1	50	30	20	20	80	0
2	20	10	10	90	0	0
3	10	50	-40	40	60	80
4	40	10	30	90	0	0
5	30	30	0	60	0	60
6	60	30	30	30	70	0
7	30	10	20	90	0	0
8	20	30	-10	60	0	70
9	60	20	40	40	60	0
10	40	20	20	80	0	0
11	20	50	-30	30	70	60
12	30	40	-10	60	0	0

	Carried	Short	Orders
Average monthly amount	22.7	2.3	0.4
Average monthly cost	$17.06	$15.70	$41.67

Total cost: $74.43

Now suppose that an organization uses the reorder point-order level policy, setting the reorder point at 50 units and the order level at 120. Table 5 shows a 12-month simulation of such a situation. The average monthly cost of managing this item is $79.90. Rules of 40 and 100, respectively, would cost only $74.43, as shown in table 6, and thus would save approximately 7 percent in operating costs. The optimum rules using an analytic model are 60 and 150, respectively, and would total only $73.10 in system operating costs.

Simulation can also be used in evaluating strategies to study parameter variations or to make sensitivity analyses, which are difficult to do analytically. For example, it may be interesting to study what happens to operating cost estimates when lead time is reduced or increased, when the demand pattern is changed slightly, when different cost figures are used, and so on.

Understanding system behavior

Simulation often gives useful insights into a system's operation, sometimes by revealing inadequacies or inconsistencies in the operating rules that might not come to light short of implementing them in the real world. It is a valuable and relatively efficient tool for this purpose. GAO has used it, as reported in Alternatives Available for Reducing Requirements for Spare Aircraft Engines (LCD-77-418, October 12, 1977). The following example is drawn from that report.

Spare aircraft engines are routinely needed to replace engines undergoing periodic maintenance or major overhaul, the purpose being to keep a $20 million aircraft operable at the expense of an extra $2 million engine. In other words, a high cost is placed on being out of stock. Requirements for spare engines are governed by predicted flying hours, engine removal rates, pipeline times in transporting and repairing or overhauling the engines, and the number of aircraft bases to be supported.

Separate requirements are computed for the central depot and each base. Spare engines are stocked at the depot and sent to a base whenever an aircraft's engine is removed and shipped to the depot for overhaul or repair. Enough spares are stocked at each base to cover the time required for receiving replacement engines from the depot or for repairing engines at the base. The continual movement of engines from aircraft to overhaul and repair facilities and back is shown in figure 15 on the next page.

Total requirements are determined analytically with a model very similar to the ship provisioning, or probabilistic demand model, illustrated in the section above. In the case at hand, however, the Navy included quantities of engine stock to support aircraft operations at land bases and also aboard aircraft carriers. In our analysis, we deemed this to be duplicative, the problem arising because the Navy computed its requirements for

Figure 15

Spare Engine Supply System

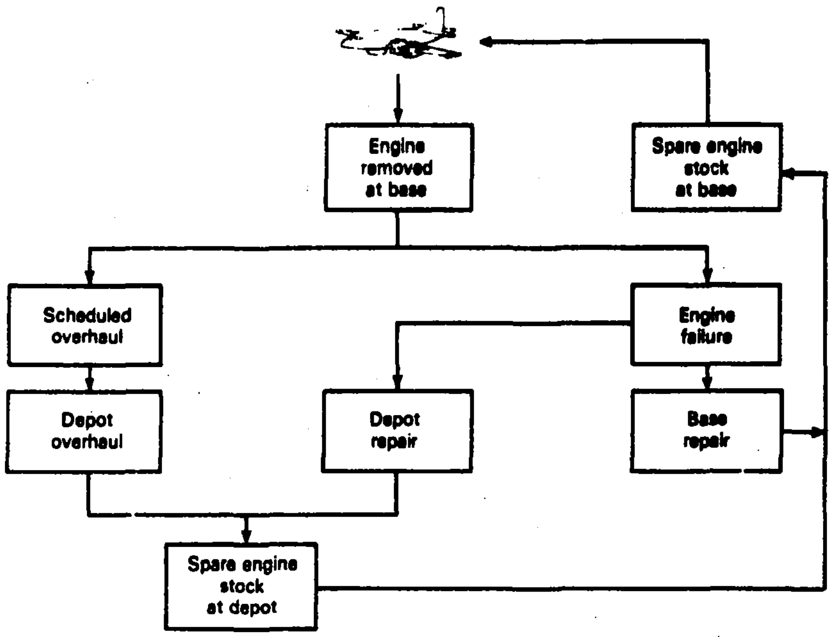

bases as if aircraft operated at them year-round, when in fact many aircraft are stationed at bases but also aboard carriers for one-third to one-half the time. The Navy supplied a quantity of spares aboard each carrier to support the same aircraft being supported by spare engine stock at their land bases.

When we suggested that this practice was "not consistent with the pipeline concept which is designed to determine aircraft operational needs and distribute assets considering depot overhaul capability and shipping times," agency personnel expressed concern. (U.S. GAO, 1977, p. 19) They feared that large reductions would hamper their ability to meet contingency surges in demand for spare engines. Consequently, a computer simulation of the spare engine support system was developed to evaluate how the system would behave under various conditions.

The results of one simulation run for one particular engine, the Navy's F-14 aircraft engine, is shown in table 7. It is constructed in 10-day intervals so that the flow of engines in and out of the various pipelines may be seen as well as the amount of spares on hand. GAO's estimate of the agency's requirements was 191, compared to the Navy's claim that 262 engines were needed. The simulation demonstrates the effect on the amount of stock on hand over a period of time as the engines going into the pipeline begin to offset the engines leaving it. Thus, the feasibility of eliminating the duplication in spares can be seen, given that shortages occur only rarely (the depot was

Table 7

Spare Engine Supply System Simulation

	Days elapsed	Number of engines: Introduced into pipeline: Depot overhaul	Number of engines: Introduced into pipeline: Base repair and resupply	Number of engines: Returned from pipeline: Depot overhaul	Number of engines: Returned from pipeline: Base repair and resupply	Number of serviceable spares: Depot	Number of serviceable spares: Bases and carriers	Total
Peacetime	0	0	0	0	0	69	122	191
	10	8	15	0	0	61	107	168
	20	8	23	0	1	53	85	138
	30	9	24	0	17	44	78	122
	40	11	29	0	21	33	70	103
	50	9	26	8	23	32	67	99
	60	12	23	8	30	28	74	102
	70	9	25	9	28	28	77	105
	80	12	20	11	21	27	78	105
	90	8	21	9	26	28	83	111
Wartime	100	12	35	12	19	28	67	95
	110	21	43	9	21	16	45	61
	120	17	47	12	39	11	37	48
	130	15	40	8	43	4	40	44
	140	14	33	12	42	2	49	51
	150	14	44	21	41	9	46	55
	160	16	43	17	32	10	35	45
	170	7	42	15	44	18	37	55
	180	2	47	14	43	30	33	63

out for a short period between day 140 and day 150) and that, in a contingency, a sufficient quantity is available to maintain the system's operability. At day 90, the ending of peacetime, the table shows 111 spares on hand. Other areas in the system where cost savings could be achieved were also shown in the simulation and are discussed in detail in the 1977 report.

Because it is sometimes difficult to represent reality in complete detail, simplifications and approximations become necessary. In making them, we should be careful that the resulting system represents enough of reality to adequately test whatever is being examined. It is, of course, difficult in a simulation to subject a system to all possible types of stress and interaction. Thus, it is often necessary to give parameters numerical values that are difficult to estimate. The outcome may or may not be sensitive to them. This is especially likely when there are a number of such parameters, because it is very difficult to make sensitivity analyses of them all. Problems like these make it difficult in many cases to derive clear-cut conclusions from the results of a simulation.

THE HEURISTIC APPROACH

As we have seen, the computational effort required in the numerous analytical models available for deriving optimal inventory decision rules sometimes prohibits their practical application. Frequently, they are difficult to understand and their data requirements are cumbersome. When the computation of an optimal policy requires complete specification of the probability distribution of demand, for example, the information required for the model may be unrealistic for many practical settings it might be applied to. Consequently, another approach to inventory decisionmaking is growing rapidly in importance. This is the heuristic approach.

The usual meaning of "heuristic" in problem-solving defines an act of learning that may lead to further discoveries or conclusions but provides no proof of whether outcomes are correct or not. (Wiest, 1975) This meaning has been expanded somewhat to include any device or procedure that can be used to reduce a problem-solving effort. In short, a heuristic may be a rule of thumb.

Without using the term, however, many government inventory systems are actually managed from heuristic models or rules of thumb. For example, the EOQ formula might be used to answer the question "How much?" but when its assumptions that demand and replenishment lead time do not vary are not, in practice, met (as they seldom are), then some sort of rule of thumb might be used for deciding when to order. That is, orders are usually placed when the inventory level falls below a quantity equal to the replenishment lead time demand plus some "safety stock." There is no proof that this procedure is the least costly. Indeed, costs are never even computed in this method. It is, however, simple to use and easy to program in a computer. As long as the system service level meets management's approval (usually stockouts are to be kept to a minimum), the system is believed to be operating satisfactorily.

Heuristic rules for when and how much to order have recently been developed that compare very favorably with optimal rules in terms of cost yet are very simple to use, requiring no significant computer time. (Ehrhardt, 1979; Freeland and Porteus, 1980; Naddor, 1975a; Nahmias, 1979) Use of them comes very close to balancing and, thus, minimizing the costs of carrying inventory, placing orders, and running out of stock.

To illustrate, we can consider the following data on a $2.40 item in a system with variable demand properties but a relatively stable replenishment lead time of four months. The policy is the reorder point-order level with values of 700 and 1,500, respectively. Costs are known to be as follows--carrying cost is $0.04 per month, shortage cost is $3.96 per stockout per month (implied from the carrying charge and the agency goal of 99 percent service level), and replenishment cost is $100 per order. The mean demand is 83.5 per month, while the standard deviation of demand is 50.3 per month.

Using the heuristic rules presented in appendix I for a system with these characteristics and an optimization model

Source	Reorder point	Order level	Monthly operating costs Carrying	Shortage	Order	Total	Projected service level
Agency	700	1,500	$30.60	$ 0	$10.00	$40.60	100.0
Heuristic	550	1,200	21.60	0.90	12.30	34.70	99.7
Analytic	500	1,150	19.60	2.00	12.30	33.80	99.4

developed by Naddor (1975a and 1975b), results were obtained as shown in the accompanying display. 1/ The agency could reduce operating costs by about 15 percent by using the heuristic decision rules because they show less stock carried. Orders are placed more frequently, but the extra replenishment cost is not enough to offset the savings in inventory carrying costs. Costs could be reduced slightly more by using the analytic solution, but this particular model is very complex, being highly mathematical, and takes considerably more computer time; indeed, solving for the rules manually is unthinkable.

While a heuristic approach may not always lead to the "best" solution, experience has proved it generally useful in finding good solutions with a minimum of effort. In some cases, it can result in computer savings. (Naddor, 1975b) More importantly, it is easy to understand. Actually, the basic notion of heuristic problem-solving is not new. Recent sophisticated extensions of this basically simple concept, when combined with the power of a computer, can enable decisionmakers to consider many complex situations successfully--among them some that have resisted solution by other techniques.

CONCLUSION

It has been said that inventory system management is based largely on intuitive judgment and experience. (Hadley and Whitin, 1964, p. vi) Further, we have pointed out that because there are few principles to assist in the development of mathematical models, inventory management frequently depends on intuition. Thus, all that can be recommended for constructing models of inventory systems that will derive optimal decision rules is that modelers and evaluators be aware of the four properties we outlined in table 4--inventory policy, cost, demand, and lead time. They must be given due consideration.

From an evaluative point of view, however, it may not be necessary to model a given system. It may be necessary only

---

1/For the analytic solution, a probabilistic model was used in which the gamma distribution was assumed (see chapter 4). This model was also used to compute the monthly operating costs for all three sets of decision rules.

to check on the assumptions underlying policy, cost, demand, and lead time. If some assumptions are found to be out of line with reality, it may be sufficient to report this. Since such a report could not give interested parties any idea about the monetary effect of assumptions or communicate a procedure that would improve the system, it might be desirable to make specific recommendations about what to do to improve decisionmaking for purposes of minimizing costs. In such cases, expertise in building models becomes a necessity. It is not enough to say "These are the rules." It is equally important to be able to state ". . . and this is how much it is going to cost."

HEURISTIC DECISION RULES

FOR THE REORDER POINT-ORDER LEVEL POLICY

Eliezer Naddor (1975b) has proposed the following heuristic decision rules as providing excellent results for the reorder point-order level policy in probabilistic systems.

$$(1)\ R = (Lt + w/2)\bar{X} - (Q + u)/2 + N\sqrt{(Lt + w/3)\sigma^2 + [(w\bar{X})^2 + Q^2 + u^2]/12 - u^2P^w/6}$$

$$(2)\ L = (Lt + T/2)\bar{X} + N\sqrt{(Lt + T/3)\sigma^2 + (T\bar{X})^2/12 + u^2(1 - P^T)/6}$$

where

R = reorder point, to the nearest multiple of u

L = order level, to the nearest multiple of u

Lt = lead time in units of time

N = number of standard deviations of a standardized normal distribution

P = probability of no demand in a unit of time. Expressions containing P are of theoretical interest only and may be ignored. If this probability is unknown, set P = 0.

Q = lot size or order quantity

σ = standard deviation of demand in a unit of time

T = scheduling period in units of time. This is the length of time between consecutive decisions about replenishments; it is the fixed interval in the fixed interval-order level policy.

u = unit quantity in which demand data are considered. (See text discussion at table 2.)

w = reviewing period in units of time. This is the unit of time in which inventory levels are determined. Generally, w = 1. Where inventory levels are reviewed bimonthly, for example, but all other variables are in monthly units, w = 2.

$\bar{X}$ = mean or average demand in a unit of time

It is assumed that inventory is reviewed every reviewing period w. If the amount on hand and on order is equal to or less than R, an order is placed to raise the inventory on hand and on order to the Level L. If there are shortages in the system, it is assumed that they are on back order.

The variables N, Q, and T are derived from

(3) $N = A^{-1}(F)$

(4) $C_2 = C_1F/(1 - F)$

(5) $T = \sqrt{2C_3/(C_1F\bar{X})}$ to the nearest positive multiple of w

(6) $Q = \sqrt{2C_3\bar{X}/(C_1F)}$ to the nearest multiple of u

where

F = fraction of time during which inventory is available; this is the fill rate

C_1 = carrying cost per unit per unit of time

C_2 = shortage cost per unit per unit of time

C_3 = cost of replenishing inventory

A = cumulative normal distribution

For F = 1, a suitable value for N in most applications is 3. In deterministic systems, when $\sigma = 0$, $N = (2F - 1)(3)^{1/2}$. If we work through an example, we can assume the following:

<u>Costs</u>

C_1 = \$0.10 per unit per month

F = 90 percent

C_2 = \$0.90 per unit per month (using 90 percent fill rate goal)

C_3 = \$100 per month

<u>Demand</u>

$\bar{X}$ = 100 units per month

σ = 20 units per month

u = 10

<u>Lead time</u>

Lt = 4 months

<u>Other variables</u>

w = 1

P = 0

To begin computing this example, we would approximate the variables T and Q--the "fixed interval" and "fixed order quantity" as described in chapter 2--with formulas 5 and 6 from above, as follows:

$$T = \sqrt{2*100/0.1*0.9*100} = 4.7, \text{ or 5 months, to the nearest multiple of } w$$

$$Q = \sqrt{2*100*100/0.1*0.9} = 471.4, \text{ or 470 units, to the nearest multiple of } u$$

The number, N, of standard deviations corresponding to F = 90 percent is approximately 1.28. N may be found by using formula 3 with the aid of normal distribution tables.

Substituting into the formulas, we obtain the following results. Note that for computational ease, every term may first be divided by u = 10 if the final answer is then multiplied by u = 10.

$$R = (4 + 1/2)100 - (470 + 10)/2 + 1.28\sqrt{(4 + 1/3)20^2 + (100^2 + 470^2 + 10^2)/12}$$

$$= 210 + 1.28(144.856)$$

$$= 395.416, \text{ or 400 units, to the nearest multiple of } u$$

$$L = (4 + 5/2)100 + 1.28\sqrt{(4 + 5/3)20^2 + (5*100)^2/12}$$

$$= 650 + 1.28(151.987)$$

$$= 844.54, \text{ or 840 units, to the nearest multiple of } u$$

While the formulas appear to be considerably more complex than the EOQ formula, they are not difficult to use manually and are easily programmed into a computer, where they require an almost insignificant amount of time. They have been shown to be very nearly optimal. Thus, organizations can manage inventories at lower cost while simultaneously maintaining some desired service level with other heuristic decisionmaking policies already in place. We have shown this recently, in a 1981 report, Better Investment Decisions Can Save Money at GSA and FAA, PLRD-81-30, June 5, 1981.

BIBLIOGRAPHY

Ackoff, R. L. Scientific Method: Optimizing Applied Research Decisions. New York: John Wiley & Sons, 1962.

Barden, G. The Accounting Basis of Inventories, Accounting Research Study No. 13. New York: American Institute of Certified Public Accountants, 1973.

Chambers, J. D., K. Mullick, and D. Smith. "How to Choose the Right Forecasting Technique." Harvard Business Review, July-August 1971, pp. 45-74.

Ehrhardt, R. "The Power Approximation for Computing (s,S) Inventory Policies." Management Science, 25 (1979), 777-86.

England, W. B., and M. R. Leenders. Purchasing and Materials Management. Homewood, Ill.: Richard D. Irwin, Inc., 1975.

Freeland, R. R., and E. L. Porteus. "Evaluating the Effectiveness of a New Method for Computing Approximately Optimal (s,S) Inventory Policies." Operations Research, 28 (1980), 353-64.

Freund, J. E. Introduction to Probability. Encino, Calif.: Dickenson Pub. Co., 1973.

Groff, V. "An Evaluation of Recursive, Discounted Least-Squares Methods for Short-Range Forecasting." Paper delivered at the 13th Annual International Conference of the American Production and Inventory Control Society, October 7-9, 1970.

Hadley, J. G., and T. M. Whitin. "A Review of Alternative Approaches to Inventory Theory." Memorandum RM-4185-PR, Rand Corporation, Santa Monica, Calif., September 1964.

-----. Analysis of Inventory Systems. Englewood Cliffs, N.J.: Prentice-Hall, 1963.

Hastings, N. A. J., and J. B. Peacock. Statistical Distributions: A Handbook for Students and Practitioners. New York: John Wiley & Sons, 1975.

Institute of Internal Auditors. Internal Audit of Inventory Control and Management, Research Committee Report No. 16. New York: 1970.

Kuhlman, H. C. Supply Management. Washington, D.C.: Industrial College of the Armed Forces, 1969.

Levy, Y. "Inventory Systems with Independent Stochastic Leadtimes." Ph.D. dissertation, Johns Hopkins University, 1979a.

-----, and E. Naddor. "Optimal Heuristic Decisions for the s,S Policy in Inventory Systems with Probabilistic Leadtimes."

Technical Report No. 315, Johns Hopkins University, Department of Mathematical Sciences, Baltimore, Md., July 1979b.

Lewis, C.D. Demand Analysis and Inventory Control. Lexington, Mass.: Lexington Books, 1975.

McAllister, H. E. Elements of Business and Economic Statistics: Learning by Objectives. New York: John Wiley & Sons, 1975.

McCarthy, J., and J. Morison. Controlling and Accounting for Supplies, Publication No. 7576. New York: National Association of Accountants, 1975.

Magee, J. F., and D. M. Boodman. Production Planning and Inventory Control. New York: McGraw-Hill, 1967.

Mendenhall, W., and E. Reinmuth. Statistics for Management and Economics. Belmont, Calif.: Buxbury Press, 1971.

Naddor, E. "Sensitivity to Distributions in Inventory Systems." Management Science, 24 (1978), 1769-72.

-----. "Optimal and Heuristic Decisions in Single- and Multi-item Systems." Management Science, 21 (1975a), 1234-49.

-----. "Optimal Heuristic Decisions for the s,S Inventory Policy." Management Science, 21 (1975b), 1071-72.

-----. Inventory Systems. New York: John Wiley & Sons, 1966.

Nagel, S. S., and M. Neef. Policy Analysis in Social Science Research. Beverly Hills, Calif.: Sage Publications, 1979.

-----. Operations Research Methods, Sage University Paper Series on Quantitative Applications in the Social Sciences No. 07-001. Beverly Hills, Calif.: Sage Publications, 1976.

Nahmias, S. "Simple Approximations for a Variety of Dynamic Lead Time Lost-Sales Inventory Models." Operations Research, 27 (1979), 904-25.

National Association of Accountants. "Techniques in Inventory Management." Research Report No. 40. New York: February 1, 1964.

Rudwick, B. H. Systems Analysis for Effective Planning: Principles and Cases. New York: John Wiley & Sons, 1969.

Shannon, R. E. Systems Simulation: The Art and Science. Englewood Cliffs, N.J.: Prentice-Hall, 1975.

U.S. General Accounting Office. Better Investment Decisions Can Save Money at GSA and FAA, PLRD-81-30. Washington, D.C.: June 5, 1981.

U.S. General Accounting Office. Alternatives Available for Reducing Requirements for Spare Aircraft Engines, LCD-77-418. Washington, D.C.: October 12, 1977.

Wiest, D. "Heuristic Programs for Decision-Making." In Managerial Economics and Operations Research, Edwin Mansfield, ed. New York: W. W. Norton & Co., 1975.

(973351)

Re: <u>Evaluating Inventory Control Decisions</u>

Request: Please respond to the following questions and return this sheet to

Eleanor Chelimsky, Director
Institute for Program Evaluation

1. How do you rate this document's job-related usefulness?

	Little or none	Some	Moderate	Substantial	Very great
Inventory policymaking	☐	☐	☐	☐	☐
Cost considerations	☐	☐	☐	☐	☐
Inventory demand	☐	☐	☐	☐	☐
Inventory lead time	☐	☐	☐	☐	☐
Assigning optimal values to policy variables	☐	☐	☐	☐	☐
Document as a whole	☐	☐	☐	☐	☐

2. Has the document been useful for any of your immediate job tasks? Can you give an example? ______________________________

____

3. How helpful will it be for designing and understanding future evaluation studies? Can you give an example? ______________

____

4. What suggestions do you have for improving this document? ___

____

____

5. What topics would you like to see discussed in future evaluation methodology transfer papers? ___________________________

____

____