7

A Survey of Analytic Techniques in Inventory Theory

HERBERT E. SCARF, *Stanford University*

During the last ten years many articles and books have been written about inventory theory. The purpose of this paper is to summarize some of the results that have been reported, and to describe some of the techniques that have turned out to be fruitful.

There is, of course, a certain ambiguity in the choice of topics to be included under the heading of inventory theory. Broadly speaking, virtually any topic in operations research can be thought of as dealing with the economical management of stocks of commodities. In addition, large areas of economic theory are concerned with problems similar to problems of operations research. Any dynamic problem in economic theory is necessarily concerned with stocks, whether these be interpreted as capital, manufacturers inventories, bank reserves, or the accumulated savings and assets of an individual consuming unit.

In this paper I shall restrict my attention to those topics in inventory theory that can most naturally be thought of as operations research rather than economic theory. This is by no means a sharp distinction. I shall, for example, include some discussion of the problem of minimizing production and storage costs in a dynamic model, but shall not include any remarks on the problems of allocating assets into current consumption and savings, many of which can be treated by virtually identical techniques. My choice of topics to be included in this survey is necessarily somewhat arbitrary, and should not be misinterpreted as a judgment on the relative merits of these topics.

In organizing this paper, I have found it convenient to maintain the categories that were used by Arrow, Karlin, and Scarf [3]. I shall begin with a discussion of the general characteristics of inventory problems. The second section will be on deterministic models, in which all costs, prices, and demands are known with certainty. This will be followed by a section on stochastic inventory models, emphasizing the dynamic programming approach to problems in which demand is a random variable. In the fourth section I shall discuss the analysis of infinite-period inventory models from a variety of points of view. The final section deals with multi-echelon models.

Reprinted from MULTISTAGE INVENTORY MODELS AND TECHNIQUES edited by Herbert E. Scarf, Dorothy M. Gilford, and Maynard W. Shelly. Published by Stanford University Press.

The emphasis in this survey will be primarily theoretical and analytical; however, this is not meant to suggest that inventory theory has been devoid of practical application. In the Navy alone, a modest but nevertheless impressive number of items are in fact being controlled on the basis of results obtained from inventory theory. The fact that these items are chiefly routine housekeeping items of relatively low cost should come as no surprise. The results of inventory theory will never supplant skilled judgment in dealing with items of high value and relatively low turnover rate, for which demand is very difficult to predict. The purpose of research in inventory theory must be considerably more modest, that is, to rationalize the large number of inventory decisions for routine, low-value items.

1. General Characteristics of Inventory Problems

Inventory theory is concerned with the analysis of several types of decisions relating primarily to the problem of when to buy and how much to buy of a given item. This analysis may involve consideration of when the item should be manufactured, problems of transportation and distribution of stock, and questions of repair, maintenance, and obsolescence. At times it may be appropriate to consider several items simultaneously, e.g., in the event that manufacturing or procurement costs cannot be factored into separate costs.

We are, of course, interested in obtaining inventory policies that are reasonable or optimal according to some appropriate measures of effectiveness. A measure of effectiveness is essentially a method of summarizing in terms of costs, probabilities, or related variables, the conflicting motives for holding stock. Therefore, I shall begin with a list of reasons, either for or against holding stocks, that have been discussed in the literature on this problem.

(a) Inventories are frequently held because of economies of scale in production or procurement. If the average cost of purchasing stock decreases as larger quantities are purchased, clearly it is economical to purchase in relatively large quantities. The result is the accumulation of stock prior to actual need. Consider a single piece of equipment or item of stock such that z units can be purchased for a total cost of $c(z)$. The average cost is $c(z)/z$, and if this is a decreasing function of z, there will be a motive for maintaining stock in advance of requirements.

A cost function of this type arises, for example, when there is a set-up cost or an administrative cost associated with procurement, in addition to a cost proportional to the quantity purchased. In this example we would have

$$c(z) = \begin{cases} K + c \cdot z & z > 0\,, \\ 0 & z = 0\,, \end{cases}$$

where K is the set-up cost and c is the proportional or unit cost. This type of cost function has appeared very frequently in the literature of inventory theory not necessarily because of its realism, but because it provides one of the few examples of cost functions with a decreasing average cost for which the analysis of inventory policies is relatively easy. In section 2, we shall examine some problems that involve a more general decreasing average cost function (such as that obtained by superimposing various "price breaks").

(b) The requirements for the items may vary substantially over time, and this, itself, may serve as an incentive for holding stock. To make this point somewhat clearer, consider the case of increasing average cost of procurement. In this case, given a constant flow of requirements, there will be no advantage to purchasing in large quantities, and hence no inventories will be maintained. (This is, of course, in the absence of other motives for holding stock.) In other words, given this type of cost function, purchases will be made at a constant rate to satisfy requirements for the item. On the other hand, if requirements vary sufficiently, a policy of this sort may result in a very high marginal cost. It may be advantageous to procure the item before it is needed at a lower marginal cost, thus contributing to the formation of inventories. This motive for holding inventories will be reinforced if the cost function displays decreasing average cost.

(c) Another motive for holding stock, in addition to the fluctuation of requirements over time, is that the costs may themselves be functions of time. For example, commodities may be held for the sole reason that the anticipated rise in price will more than cover purchase costs plus the cost of maintaining stock [more will be said about this latter aspect of cost under (e)].

If the fluctuations in cost are predictable in advance, a good deal can be said about the resulting inventory policies. The more important problem in which future prices and costs are unknown has received relatively little attention in the operations research literature, possibly because much of this literature is concerned with the control of military items, where this problem is not particularly significant.

(d) Uncertainty of future requirements is also a strong motive for holding inventories. For an item that is essential or has considerable utility, an uncertain future demand may result in higher levels of stock than would be desirable if future demand were relatively predictable.

It has been customary when demand uncertainty is an essential part of the problem to describe future demands by means of a stochastic process, that is, in terms of the joint probability distribution of demand for sets of future time instants. This, of course, requires us to assume a substantial amount of current information about the parameters describing future demands—means, variances, correlations, etc.—and may involve excessively strong assumptions for items with little or no demand history. Some attempts have been made to introduce *a priori* distributions that are successively refined as more and more demand data become available (see Scarf [50], [51], and Dvoretzky, Kiefer, and Wolfowitz [16]), though this approach has not yet proved to be particularly fruitful.

If stockage decisions are made at regularly spaced intervals, such as at the beginning of each week or month, then it is most natural to consider the cumulative demand during such periods. In other words, we would be concerned with a discrete sequence of random variables $\xi_1, \xi_2, \ldots,$ where ξ_n represents the cumulative demand during the nth decision period. The simplest and most frequent assumption is that these demands are independently distributed from period to period, and are specified by, say, a cumulative distribution Φ_n. The cumulative distribution, or its density $\varphi_n(\xi)$, may, of course, change over time, reflecting either periodic fluctuations or more general fluctuations in the level of demand.

Independence of demand from period to period is undoubtedly a restrictive assumption, especially if the periods in question are of short duration. It is probably correct to say, however, that the loss in realism is relatively slight compared to the increase in ability to calculate inventory policies based on this assumption.

Some work has been done recently in which the assumption of independence has been relaxed. This work may be divided into two categories. In one of these, demands are assumed to arise from a Markov process, so that the level of demand, or "state" of the demand is dependent upon the state in the previous period (Karlin and Fabens [35], Iglehart and Karlin [27]). In the other approach, demands are assumed to arise from a stationary stochastic process and few or no structural assumptions are placed on the demand process. The increase in generality is balanced by the fact that the results obtained are valid under a rather restrictive class of cost functions (Holt, Modigliani, Muth, and Simon [24]). The techniques involved in this latter approach are related to "control theory" and will be discussed by Ronald A. Howard in chapter 6 of this volume.

If demands or future requirements are uncertain, then regardless of the stockage policy adopted there is generally some probability that available stock levels will be insufficient to meet demand. Any of a number of courses of action might be followed in this event. One possibility is to satisfy the request by some alternative method of procurement, by a substitute part, or by a higher level of assembly. Another possibility, if the need is not very urgent, is to wait until sufficient stock becomes available through the normal functioning of the supply system to satisfy the request. Still another possible action is to delay the request for a fixed period of time and then use an alternative method of procurement if adequate stock is not yet available. In the formulation of a specific inventory model the treatment of this problem necessitates a number of arbitrary assumptions.

We are naturally led to make assumptions that will facilitate the analysis, and produce simple policies, with the hope (which of course can be verified by other methods) that the resulting distortion will not be too great. As we shall see later, the most convenient assumption is that excess demands or shortages are "backlogged," i.e., permitted to wait until adequate stock is available. Technically this means introducing the mathematical fiction of negative stock levels to refer to the cumulative excess of unsatisfied demands over the current stock position.

An inventory policy that permits large shortages to occur persistently will generally be of little value. The advantages of keeping shortages at a tolerable level must be weighed against the cost of maintaining high levels of stock. Hence it would be convenient to describe shortages numerically in terms of costs or of probabilities. For the latter type of description we might restrict our attention to those policies for which the probability of at least one shortage is kept below a certain preassigned level. If, as is customary, demands for the item persist over a long period of time, we would then have to select probability levels for each of a number of periods. If the demand distributions were changing substantially over time, the procedure might involve a considerable degree of arbitrariness. Although this approach has occasionally been utilized (see, for example, Proschan [47]), it has been more customary to assume that the effect of shortages can be described in terms of a shortage cost.

In some inventory models the appropriate shortage cost may be suggested quite naturally. If, for example, shortages are satisfied by emergency shipment, it would seem natural to use as a shortage cost the difference in cost between emergency shipment and routine purchase. In a model in which the items are being sold for certain prices, and in which sales are lost if they cannot be met out of current inventory, the appropriate shortage cost would be related to the loss in profit. In both of these instances, shortages produce a real increase in cost or a decrease in revenue that can be translated into an appropriate cost. There are models, however, in which the cost implications of a shortage are by no means as immediate, and the selection of the correct shortage cost becomes more difficult. If, for example, excess demands are backlogged until stock becomes available, the shortage cost must be determined by an analysis of the implications of waiting for the particular item. In some instances the shortage may not be a real cost, and may reflect only one's judgment as to the urgency with which an item is required when a demand is presented.

Consider the case in which the stock position is reviewed at the beginning of a number of regularly spaced time intervals. The shortage cost is generally assumed to be some function of the excess of demand over supply at the end of a typical period. If the excess is u ($u \geqq 0$), the shortage cost to be charged will be, say $p(u)$. The functions that have most frequently been suggested are either a cost proportional to the size of the shortage, or a fixed cost to be charged at the end of any period in which there is a shortage, regardless of the size. If the first of these functions is used with the model in which excess demands are backlogged, shortage costs will reflect not only the size of the shortage but also the actual duration.

Before concluding this discussion of random demands as a motive for holding stock, I should like to say something about models in which stock decisions are made continuously over time. In models of this type the course of future demand must be described by a continuous-time stochastic process. If we carry over the convenient assumption made in the discrete-time model, i.e., that demand is independent from period to period, we are led to consider stochastic processes with independent increments. The Poisson process is an example of this type of stochastic process. A more general example would be obtained by assuming that the points of time at which demands are made are given by a Poisson process, but that the size of demand at any of these time points is random. Most of the elementary treatments of inventory models involving continuous time are quite ambiguous on the point of just what stochastic process is being assumed.

Occasionally other descriptions of the continuous-time demand process are used. We may characterize the Poisson process by saying that demands arrive one at a time, and that the times between successive demands are independently distributed according to an exponential distribution. We may generalize by assuming that demands arrive one at a time, and that the times between demands have an arbitrary distribution. This type of assumption is most conveniently made for the stationary analysis described in section 4.

(e) Now let me turn to the motives for *not* holding inventories. If items are to be held in inventory, they must be stored somewhere, and holding or storage charges may be assessed against the stock. These charges may be based on the cost of maintaining the stock in usable condition, the cost of obsolescence, the cost of renting storage space, or the cost of constructing new storage space. Holding costs may also be derived costs, for example, costs implied by weight or volume limitations aboard a ship.

In a discrete-time inventory model, holding costs are generally assumed to be calculated as some function $h(u)$ of stock on hand either at the beginning of the period, at the end of the period, or averaged throughout the period. The most frequent assumption is that holding costs are proportional to the quantity of stock held, though more general functions have occasionally been used.

An interest rate provides a motive for not holding inventories if it accurately reflects the return available from alternative uses of funds. Here the advantages of investing in inventories today must be compared with those of investing these funds in some alternative way and using these funds augmented by interest payments for purchases of stock at some future date. In this respect the interest rate acts as a holding cost. In fact, a number of studies have neglected holding costs completely, with the interest rate providing the dominant motive for keeping inventories down.

In an inventory problem that lasts for some length of time, costs will generally be incurred at various moments of time, and we have the problem of summarizing these costs in a single number so that alternative policies can be compared. If funds can actually be borrowed and lent at a rate of interest of i per period, then policies should be compared on the basis of the discounted value of their costs, with a discount factor $\alpha = 1/(1+i)$.

(f) In most inventory problems the actual costs that influence the selection of an appropriate policy are those described above, i.e., production and procurement costs, holding costs, and shortage costs. There are, however, several additional motives for holding stock that are not directly translatable into cost figures. For example, higher inventory levels would generally be held if there is a substantial period of time between placing an order for stock and its eventual delivery. This motive would not be operative if future requirements were known precisely; orders could then anticipate requirements. However, if future demands are random, the possibility that shortages will occur during a lead time must be incorporated in the analysis.

As another, somewhat similar motive for holding stock, consider a situation in which inventories are stored in a warehouse and at several using activities. Assume that stocks are sent from the warehouse to the using activities, but that direct shipment among using activities is prohibited. In this situation, high levels of stock might be kept at the warehouse in order to postpone their commitment to the using activities.

2. Deterministic Inventory Models

In this section I shall discuss some of the work that has been done on inventory problems in which all of the parameters—costs, demand rates, and so on—are assumed to be known in advance.

The most fundamental deterministic inventory model is the familiar one that yields the Wilson lot-size formula as the solution (see Whitin [59] for a discussion of the early development and applications of this result). This is a continuous-time model in which purchases may be made at any moment of time according to the cost function

$$c(z) = \begin{cases} K + c \cdot z & z > 0\,, \\ 0 & z = 0\,. \end{cases}$$

The set-up cost K, which provides the motive for large, infrequent purchases, must be weighed against the cost of holding inventories in choosing the minimum-cost policy. Holding costs are assumed to be proportional to the size and duration of inventories, and h is the proportionality constant (in dollars per unit of stock, per unit of time).

Demand for the item (which is required to be satisfied) is assumed to be known in advance and to occur at a uniform rate of m per unit of time. In comparing policies, interest charges are usually neglected (or incorporated directly into the holding cost), and the comparison is made on the basis of average cost per time unit.

It is easy to verify that the only policies that need be considered are those which place orders when there is no available stock (neglecting lead times in delivery as is appropriate for this model), and which order the same quantity whenever an order is placed. A single parameter Q, the size of the order, is sufficient to specify the policy, and as a function of Q, the average cost per time unit will be

$$cm + \frac{mK}{Q} + \frac{hQ}{2}.$$

The policy that minimizes this cost is the Wilson lot-size formula

$$Q = \sqrt{\frac{2Km}{h}}.$$

The model on which this calculation is based is highly simplified and neglects a good number of the important reasons for maintaining inventories. On the other hand, the formula given above provides a remarkably good approximation to "optimal policies" in considerably more elaborate and realistic models.

The Wilson formula represents a balance between economies of scale in purchasing, and costs associated with maintaining inventories. It does, however, make use of a particular cost function, and it is of interest to examine optimal policies when this purchase cost is replaced by a general cost function exhibiting economies of scale. Some work has been done (see, for example, Churchman, Ackoff, and Arnoff [10]) using a cost function composed of several linear segments, while maintaining the other assumptions given above.

In [58], Wagner and Whitin have described an algorithm for the solution of a considerably more general problem, in which the purchase and holding cost functions are general concave functions, possibly changing over time, and where the requirements, instead of being constant as in the derivation of the Wilson formula, are permitted to vary in an arbitrary fashion. In their model, decisions are assumed to be made at the beginning of a sequence of regularly spaced intervals, which are numbered chronologically as periods $1, 2, \ldots, N$. The purchase cost function relevant to period j is $c_j(z)$, and the holding cost to be charged on stock at the beginning of the period j is $h_j(y)$. Requirements in the jth period are denoted by r_j. The objective is to find that policy which satisfies the requirements and minimizes the sum of purchase and holding costs (if there is an interest rate, the discounted sum can be introduced with no additional complexity).

This problem can be treated as a special case of the stochastic inventory models considered in section 3 (with the demand distributions degenerate to a point, and with an extremely high shortage cost), and hence a solution may be found by means of the dynamic programming procedures usually used in this more general problem. The algorithm given by Wagner and Whitin, while itself a dynamic programming algorithm, represents a considerable simplification based on the following result: If c_j and h_j are concave, there exists an optimal policy that makes purchases only at those times when there is no stock. This theorem permits us to focus our attention on the moments of time at which purchases are made, rather than on the quantities of stock purchased. The actual algorithm considers numbers $C_1, \dots, C_N$, with C_n defined as the minimum cost for a problem which lasts a total of n periods and in which the requirements and costs are identical with those in the first n periods of our original problem. These numbers satisfy a very simple recursive relation of the dynamic programming type, from which the optimal purchase or production plan may be obtained.

One of the important generalizations of this problem, which has recently been receiving some attention, is to consider several items, each with its own pattern of requirements. The problem can be factored into the separate consideration of each of the items if there are no joint constraints in production. If, on the other hand, a single machine with a capacity limitation is used to produce all of the items, then the individual production plans obtained by the use of the Wagner-Whitin algorithm may not fit within the capacity limitations, and new computational procedures must be devised. Discussions of this problem may be found in Manne [40], and in Gilmore and Gomory [22].

In the inventory models discussed above, the dominant motive for holding inventories has been economies of scale in production. As I pointed out in the introduction, inventories will also be held with increasing marginal costs of production if the requirements are sufficiently fluctuating over time. The problem becomes one of smoothing production rather than one of taking advantage of economies of scale.

The basic production-smoothing problem is that in which a sequence of requirements are to be met, and the production cost functions $c_j(z)$ and the holding cost functions $h_j(y)$ are convex (increasing marginal cost). This problem has been studied by a number of authors, using techniques ranging from the calculus of variations to linear programming. The fundamental paper on this problem is due to Johnson [28], who develops an exceptionally simple algorithm.

In describing Johnson's algorithm it is useful to assume that requirements and possible production levels can all be expressed as multiples of some common unit. The algorithm proceeds as follows: Let the requirements of the first period be satisfied by production in that period, with the understanding that if r_1 units are required in the first period, the smallest r_1 marginal costs will be assigned to their production. Take the remaining marginal costs arranged in increasing order and add to them, one by one, the marginal costs of storage, arranged in increasing order. This provides a list of marginal costs available for satisfying requirements in period 2. Of course, requirements in period 2 may also be satisfied by production in the second period, providing us with an alternative list of increasing marginal costs. If the two lists are combined and arranged in increasing order, we obtain the actual production function available for satisfying requirements in period 2 and subsequent periods. The process is then repeated; the smallest r_2 marginal costs are eliminated from this list, marginal holding costs are added, and the list is combined with the marginal costs of producing in period 3, and so on.

A simple modification of this argument will cover the case in which arbitrary quantities of the firm's output can be sold in a competitive market at fixed prices, so that the problem becomes one of scheduling production and sales. The procedure has also been extended by Wagner [55] to cover the case of a monopolist setting prices in addition to making production decisions, under an appropriate assumption as to the elasticity of demand.

Johnson's procedure depends crucially on the assumption of increasing marginal costs of production and storage. It will not apply, for example, to the problem in which there is a cost associated with changing the level of production from one period to the next. This latter problem, with an arbitrary schedule of requirements, linear production cost and storage cost functions, and a cost proportional to the increase in production (no cost is charged if the production level decreases), has been examined by Johnson and Dantzig [29]. The problem is posed in linear programming terms, and is solved by means of a simplified version of the simplex method. A modification of this problem in which the item is perishable and consequently cannot be stored has been discussed in a discrete-time version by Karush and Vazonyi [38] and in a continuous-time version by Arrow and Karlin [2].

3. Stochastic Inventory Models

In this section I shall begin to discuss inventory models that explicitly take into consideration the possibility of uncertain future demand. For the moment I shall restrict my attention to models in which decisions are made at the beginning of each of a number of periods, and in which the demand distributions are independent from period to period. The cumulative distribution for demand in period n will be denoted by $\Phi_n(\xi)$, with the subscript omitted if demands are identically distributed. There is a certain ease in exposition if it is assumed that the demand distributions have density functions $\varphi_n(\xi)$, and I shall consistently make this assumption even though the calculations of optimal policies, which are invariably done on a digital computer, involve using a discrete distribution.

The simplest model of this type involves a single period. Assume that y units of stock are purchased at the beginning of the period for a total cost of $c \cdot y$, and that a random demand ξ is given at the end of the period. The quantity sold will be the smaller of y and ξ, and if the selling price is p, expected profit will be

$$-c \cdot y + p \int_0^y \xi\varphi(\xi)\, d\xi + py \int_y^\infty \varphi(\xi)\, d\xi \, .$$

It is easy to verify that this function is concave as y ranges from zero to infinity; the maximum is therefore obtained by setting the derivative equal to zero (if $p > c$), and we see that the optimal value of y should satisfy the equation

$$\frac{c}{p} = \int_y^\infty \varphi(\xi)\, d\xi \, .$$

This solution may be found in Whitin [59] and in Arrow, Harris, and Marschak [1].

This model is, of course, superficial. It does, however, provide a point of departure for the discussion of more elaborate inventory models. The negative of the above expression may be considered an expected cost (in the sense that we would like to minimize it), and may be written as

$$c \cdot y + p \int_y^\infty (\xi - y)\varphi(\xi)\, d\xi - pm \, ,$$

where m is the mean of the demand distribution. Only the first two terms are relevant in minimizing this expression. The first term is the cost of purchasing y units; the second term is equal to the selling price multiplied by the expected excess of demand over supply. In other words, the second term may be looked upon as an expected shortage cost, with the shortage cost function proportional to the number of shortages. With this interpretation, our original problem becomes one of minimizing the expected cost of purchasing and shortages.

Several generalizations come to mind immediately. There may be situations other than the one described above in which a different shortage cost function is appropriate. This would mean replacing the second term by $\int_y^\infty p(\xi - y)\varphi(\xi)\,d\xi$, with a possibly non-linear shortage cost function.

On the other hand, it might be reasonable to charge other costs, such as holding costs, during the period. If the holding costs are charged at the beginning of the period, an additional cost $h(y)$ would be added. If the holding cost were charged at the end of the period, on the basis of the excess of supply over demand, the additional term would be $\int_0^y h(y - \xi)\varphi(\xi)\,d\xi$.

It is convenient to introduce the notation $L(y)$ to represent the expected holding cost plus shortage cost during a period if the initial stock is y. The problem is then to select y so as to minimize $c \cdot y + L(y)$, and if sufficient regularity conditions are assumed, the solution is given by

$$c + L'(y) = 0\,.$$

In the case in which both the holding cost and the shortage cost are linear (holding cost charged at the end of the period), this equation becomes

$$1 - \Phi(y) = \frac{c + h}{p + h}\,.$$

Single-period inventory problems are of relatively little importance by themselves. There are, however, a number of instances in which the single-period problem is looked upon as the last period in a dynamic model. The important distinction is that we no longer consider the initial inventory before placing an order to be equal to zero, but permit it to be an arbitrary size x. If y represents the level to which stock is raised, the expected cost will now be given by

$$c(y - x) + L(y)\,.$$

I shall at this point and for the remainder of the paper consider the case in which y must be greater than or equal to x. The other possibility, which involves the disposal of stock, has been considered by Fukuda [18] and others, in both the dynamic and single-period models.

The problem is then to select $y \geqq x$ so as to minimize $c(y - x) + L(y)$. The optimal choice of y will, of course, depend on x, and those assumptions on both L and c which give a simple form for the optimal policy have been extensively investigated.

If the purchase cost function is linear, and if $L(y)$ is convex (this will occur if the holding and shortage cost functions are linear, and in many other cases as well), the optimal policy will be of the particularly simple form characterized by a single "critical number" $\bar{x}$. The value of $\bar{x}$ is determined from the equation $c + L'(\bar{x}) = 0$, and the policy is to raise the stock to $\bar{x}$ if it is initially below this level, and not to order if the stock is above $\bar{x}$. Examples have been given in which considerably more complex policies occur if the convexity assumption of $L(y)$ is dropped. For example, if a constant shortage cost is charged independently of the size of the shortage, $L(y)$ may not be convex and the resulting policy may involve several regions in which orders are alternately placed and not placed.

The next problem, as far as complexity is concerned, is that in which the purchase cost function involves a set-up cost K in addition to the linear cost function $c \cdot x$. Again, if the expected holding and shortage cost function is convex, the resulting optimal policy has a simple form based on two numbers S and s. The policy is to raise the stock to S if it is initially below s, and not to order if the stock is initially above s. Consistent with the economies of scale in purchasing, small orders will not be placed. The minimum order size is $S - s$, which will generally be greater than zero unless $K = 0$, in which case the policy is identical with that given above. For the single-period model, S is determined by the equation $c + L'(S) = 0$, and s by $cs + L(s) = K + cS + L(S)$.

For example, if the holding and shortage cost functions are linear and if the demand distribution is exponential with mean m, the equations for S and s become

$$e^{-S/m} = \frac{c + h}{p + h}, \quad \text{and} \quad e^{Q/m} = \frac{K}{m(c + h)} + \frac{Q}{m} + 1,$$

where $Q = S - s$. If K is small, then Q will be small and the left-hand side of the latter equation may be approximated by the first several terms of its Taylor series expansion, with the result that

$$Q \sim \sqrt{\frac{2Km}{c + h}}.$$

This suggests a relationship between the minimum order size and the Wilson lot-size formula, which occurs repeatedly in inventory theory.

If the convexity assumption on $L(y)$ is dropped, the optimal policy may be considerably more complex than an (S, s) policy. Karlin [33] thoroughly analyzes the types of optimal policies implied by various forms of the holding and shortage cost functions and demand distributions.

It is also possible, in the single-period problem, to determine the form of the optimal policy if more general purchase cost functions are considered. The difficulty is that the results do not apply to multiperiod models, whereas the results for the two cost functions discussed above can, in fact, be generalized. If one takes the point of view, as I do, that the single-period problem is insignificant compared to the dynamic problem, there is little point in discussing these extensions.

Let me now turn to those models which continue for a number of periods, and in which purchases can be made at the beginning of each of the periods. The following assumptions, a number of which will be relaxed later, will be made in the present discussion.

1. A single item is being considered. The item may be ordered at the beginning of various time periods; the purchase cost function is given by $c(z)$.

2. Demands for the item are independent from period to period, and their density function is $\varphi(\xi)$.

3. Delivery of stock is immediate.

4. Excess stock is backordered—this is, of course, a minor assumption when there is no lead time in delivery.

5. The expected single-period costs, if the stock level at the beginning of the period is y, will be $L(y)$.

6. No disposal is possible.

7. The discount factor is α, between zero and one.

8. The problem will persist for a total of N periods. It is convenient to label the periods in the reverse order from the chronological order.

An inventory policy is a set of rules that define the quantity to be purchased at the beginning of each period as a function of whatever information has accumulated up to that time. Because of the nature of the problem and the assumptions that have been made, it is sufficient to consider purchasing rules that depend only upon the stock size at the moment of time in question. More elaborate policies are required if some of the assumptions are relaxed.

Any specific purchasing policy will generate a random sequence of costs, the randomness being caused by the uncertainty of future demand. We are interested in calculating those policies (optimal policies) which minimize the expectation of the present value of these costs.

By appropriately identifying and relabeling the quantities involved, we can usually transform programming problems that involve time into problems in which only one decision is made. The possibility of converting dynamic problems into static problems is implicit in the literature of economic theory, and has been recognized explicitly in game theory and statistical decision theory. Especially when uncertainty is involved, the resulting static problems are quite complex and almost useless for calculating optimal policies or even for discussing their general form. On the other hand, the static problem may be quite useful for the purpose of proving existence theorems and examining related questions. For the inventory problem, this approach was taken by Karlin in [30].

The alternative approach is to utilize the fact that sequential decision problems may be split into two parts—the current decision and all subsequent decisions—and to obtain an iterative sequence of functional equations whose solution yields the optimal policies. This technique, which Bellman [7] has used to form the basis of dynamic programming, also has a history in economic theory, game theory, and statistical decision theory. In the context of the inventory problem, the technique was first introduced by Massé [42] and by Arrow, Harris, and Marschak [1] (with a restriction to (S, s) policies), and was then examined with complete generality by Dvoretzky, Kiefer, and Wolfowitz [15].

The dynamic programming technique for the solution of inventory problems proceeds in the following way. We consider an inventory problem that begins with an inventory of x units of stock and lasts for a total of n periods. If an optimal policy is used, the minimum of the expected discounted costs will be a function of x and n, which I shall denote by $C_n(x)$. This cost may be divided into two parts, the costs incurred in period n, and the expected future costs. If a decision is made to raise stock from x to y in period n, there will be a purchase cost of size $c(y - x)$ and a single-period cost $L(y)$. At the beginning of period $n - 1$ the inventory will be $y - \xi$; if we proceed optimally from this period onward, the expected future cost will be

$$\alpha \int_0^\infty C_{n-1}(y - \xi)\varphi(\xi)\, d\xi .$$

The purchasing decision y should be made so as to minimize cost, and we obtain the well-known functional equation of inventory theory

$$C_n(x) = \min_{y \geq x} \left\{ c(y - x) + L(y) + \alpha \int_0^\infty C_{n-1}(y - \xi)\varphi(\xi)\, d\xi \right\} .$$

This equation (or variants of it), which may be programmed on a digital computer, provides a very efficient procedure for calculating optimal policies. The algorithm is not only rapid, it is also quite flexible in the sense that all costs and demand distributions may vary over time (this would involve introducing subscripts on these parameters in the functional equation).

There are several cases in which the functional equation can be analyzed and the form of the optimal policies determined. There is, of course, considerable appeal in knowing that a simple policy, such as a single critical number or an (S, s) policy, is optimal; simple policies are, naturally enough, easy to implement. In addition, the time necessary to compute optimal policies may be cut substantially if it is known beforehand that the policy has a particular form.

The first analysis of the inventory equation whose purpose was to determine the form of the optimal policy was given by Bellman, Glicksberg, and Gross in [8]. In that paper the purchase cost is assumed to be linear, there is no set-up cost, and the holding and shortage cost functions are also linear (actually the last requirement may be replaced by the considerably weaker assumption that $L(y)$ is convex). It is then demonstrated that the optimal policy is defined by a sequence of critical numbers $\bar{x}_1$, $\bar{x}_2$, ...; if the stock level at the beginning of period n is below $\bar{x}_n$, an order for the difference is placed. No ordering is done during this period if the stock level exceeds $\bar{x}_n$. The proof is quite simple and depends only on verifying, by induction, that the functions $C_n(x)$ are convex. The main tool is the simple fact that if $f(x)$ is convex and bounded from below, $g(x) = \min_{y \geqq x} \{f(y)\}$ is also convex.

The argument is quite general in the sense that the demand distributions and costs may vary over time. If, however, these are assumed to be unchanging over time, it is quite easy to show that $\bar{x}_1 \leqq \bar{x}_2 \leqq \bar{x}_3 \ldots$, so that the critical numbers are increasing monotonically as we move away from the end of the program. It is also possible to show that $\bar{x}_n \leqq \bar{x}$, with $\bar{x}$ defined as the solution of the equation

$$c(1 - \alpha) + L'(\bar{x}) = 0 .$$

The number $\bar{x}$ has a specific interpretation, which I shall defer until later in this section.

If a set-up cost in purchasing is introduced, the situation becomes somewhat more complex. Recall that for the single-period problem the optimal policy would be of the (S, s) type if $L(y)$ were assumed to be convex. This result could easily be demonstrated for the dynamic problem if it were possible to demonstrate inductively that the functions $C_n(x)$ were convex, even in the presence of a set-up cost in ordering. Unfortunately, this is never the case, and a different argument is required. In [52] Scarf introduced the class of K-convex functions. A function is K-convex if the secant line, when extended to the right, is never more than K units above the function, or in analytical terms if

$$f(x + a) - f(x) - a\left[\frac{f(x) - f(x - b)}{b}\right] + K \geqq 0$$

for $a \geqq 0$, $b \geqq 0$, and all x. For $K = 0$, this reduces to the ordinary definition of convexity. It is then shown inductively that each of the $C_n(x)$ is K-convex, and it follows immediately from this result that the optimal policy in period n is defined by a pair of numbers (S_n, s_n). The main tool of the proof is the fact that if $f(x)$ is K-convex, then

$$g(x) = \min_{y \geqq x} \left\{ K\delta(y - x) + f(y) \right\}$$

is again K-convex, where $\delta(u) = 1$ if $u > 0$, and $\delta(u) = 0$ if $u = 0$.

Again this argument is quite general and permits the costs and probability distributions to change over time. The only restriction is that the set-up costs do not decrease as we move away from the end of the program. It is possible, given decreasing set-up costs, to obtain optimal policies of a more complex sort.

Very little is known about the relationship of the policies (S_n, s_n) in different periods, even if it is assumed that all costs and demand distributions are unchanged over time. The fragmentary results that are known are due to Iglehart [26]; they suggest that the relationship will by no means be simple.

In introducing the basic functional equation, explicit use was made of the assumption that orders are delivered instantaneously. The assumption is a questionable one, since in reality there will generally be a lag in delivery. Perhaps the most common case is one in which the delivery lag is random. It is easier, however, to begin with the assumption that the delivery lag is of fixed length and is equal to an integral number of periods. Specifically, let us assume that an order placed at the beginning of the present period will be delivered at the beginning of the period, λ periods from now. It is possible to write a functional equation as before—the difficulty is that the functions involve λ variables, current stock x, and orders y_j $(j = 1, 2, \ldots, \lambda - 1)$ due in the subsequent $\lambda - 1$ periods. Even with a lead time of four or five periods, the recursive calculations involved in the calculation of optimal policies would be prohibitively long.

Karlin and Scarf [36] showed that if excess demands are backlogged, the functional equation can be reduced to one involving only a single variable. More specifically, if $C_n(x, y_1 \ldots y_{\lambda-1})$ represents the expected discounted cost for an n-period problem if an optimal policy is followed, then

$$C_n(x, y_1 \ldots y_{\lambda-1}) = L(x) + \alpha \int_0^\infty L(x + y_1 - \xi)\varphi(\xi)\, d\xi$$

$$+ \ldots + \alpha^{\lambda-1} \int_0^\infty L\left(x + \sum_1^{\lambda-1} y_j - \xi\right)\varphi_{\lambda-1}(\xi)\, d\xi + g_n\left(x + \sum_1^{\lambda-1} y_j\right),$$

where $g_n(u)$ satisfies the functional equation

$$g_n(u) = \min_{y \geq u} \left\{c(y - u) + \alpha^\lambda \int_0^\infty L(y - \xi)\varphi_\lambda(\xi)\, d\xi + \alpha \int_0^\infty g_{n-1}(y - \xi)\varphi(\xi)\, d\xi\right\}$$

(we are using the notation φ_j to represent the density of demand over j periods), and moreover the optimal policy is determined from this latter equation, with u representing stock on hand *plus* total stock on order. Other than replacing $L(y)$ by $\alpha^\lambda \int_0^\infty L(y - \xi)\varphi_\lambda(\xi)\, d\xi$, this equation is identical with the zero lead time equation, and therefore all of the theory developed for the special case of zero lead time is applicable in general, if we backlog.

The backlog assumption enters into this simplification by providing us, at any moment of time, with a horizon with two properties: (i) No action taken now will influence any costs before the horizon, and (ii) all costs after the horizon are functions of stock on hand *plus* stock on order.

There are other lead-time models for which horizons with these properties could be found, and which would therefore result in a corresponding simplification of the functional equation. One example is the case in which the lead times are random, but at most one outstanding order is permitted at any moment of time. Another involving random lead times is the case in which outstanding orders "queue up" before being delivered. The case of no backlog does not fall into this category, nor do cases in which the lead times for the various orders are obtained by independent randomizations, with some possibility that orders will be delivered in a sequence other than that in which they were placed.

This treatment of delivery lead time has a number of drawbacks. One of them is that backlogging is required, and it may be quite unrealistic to assume that customers are willing to wait indefinitely until their requests are satisfied. There may be some hope in analyzing a modified problem in which excess demands stay on the books for a fixed number of periods, after which they are canceled. Another deficiency is that this treatment does not permit the possibility of purchasing shorter lead times at higher cost. In chapter 2 of this volume, Daniel analyzes a model that involves a routine lead time of a single period, and an instantaneous emergency alternative that supplies items at a higher unit cost.

Now I shall turn to a problem considered by a number of the authors cited above; that is, the inventory problem in which all costs and probability distributions are unchanged over time and the length of the program tends to infinity. It will be convenient in this section to restrict attention to the case in which the discount factor is strictly less than one. Corresponding questions when $\alpha = 1$ will be considered in the next section.

In an infinite-period inventory problem of this type, the situation faced at the beginning of a period will differ from that at the beginning of the next only in the size of current stock levels; we shall, of course, be no closer to the end of the program. This suggests that we consider only one minimum cost function $C(x)$ with no subscript indicating dependence upon the length of the program, and that $C(x)$ should satisfy the functional equation

$$C(x) = \min_{y \geq x} \left\{ c(y - x) + L(y) + \alpha \int_0^\infty C(y - \xi)\varphi(\xi)\, d\xi \right\}.$$

This equation is more complex than the ones previously discussed, inasmuch as the same function C enters on both sides. The questions that come to mind are: Is there a unique solution to this functional equation, and how is it related to the functions $C_n(x)$; what are the optimal policies associated with this equation, and how do they relate to the policies previously discussed? The answers to these questions have a certain technical difficulty, since the variable x ranges over the infinite interval. If it were possible to restrict both x and y, or merely the difference between x and y, to a finite interval, the general techniques described by Bellman [7] and Karlin [30] would be adequate to answer the question of the existence and uniqueness of $C(x)$. If we do not choose to make this *a priori* restriction, the following type of analysis is available. First of all, it is easy to verify that the functions $C_n(x)$ are monotonically increasing and bounded from above, so that they converge to some function $f(x)$. It is also simple to verify that $f(x)$ satisfies the inequality

$$f(x) \leqq \inf_{y \geqq x} \left\{ c(y-x) + L(y) + \alpha \int_0^\infty f(y-\xi)\varphi(\xi)\, d\xi \right\}.$$

In order to establish the reverse inequality, we generally need some additional information concerning the optimal policies for the finite-period problem—namely, the existence of some number X such that no ordering takes place at the beginning of any finite-period problem if the initial stock is above X, and that $C_1(x) - f(x) \geqq -A$ for all $x \leqq X$, with some finite constant A. For the case in which $L(y)$ is convex and the ordering cost is composed of a set-up cost plus a unit cost, these results have been established by Iglehart [26]. It is then a relatively simple matter to verify inductively that

$$C_n(x) - f(x) \geqq -\alpha^{n-1}A \qquad (x \leqq X).$$

This result is sufficient to show that $f(x)$ satisfies the equation

$$f(x) = \min_{y \geqq x} \left\{ c(y-x) + L(y) + \alpha \int_0^\infty f(y-\xi)\varphi(\xi)\, d\xi \right\}.$$

Then, under suitable assumptions, we may demonstrate uniqueness of this solution and then identify $f(x)$ with $C(x)$.

If the functions $C_n(x)$ are convex, $C(x)$ will be convex; if they are K-convex, $C(x)$ will also be K-convex. Therefore, if the ordering cost is composed of a set-up cost plus a unit cost, the optimal policy for the infinite-period problem will be an (S, s) policy, which degenerates to a single critical number $\bar{x}$ in the event that $K = 0$. It is a simple matter to show that $\bar{x}$ is given by the solution of the equation

$$c(1-\alpha) + L'(\bar{x}) = 0,$$

which will be unique if L is strictly convex (say, with a strictly positive second derivative).

If the finite-period calculations are thought of as being an approximation to the infinite-period calculations, then it is of interest to relate the optimal policies of the finite-period model to those described above. In practice, optimal policies seem to converge quite rapidly. Our theoretical information on these points is, however, quite scanty. If $K = 0$, it is possible to demonstrate that

$$\overline{\lim_{n\to\infty}} \left| \frac{\bar{x} - \bar{x}_n}{\alpha^n} \right| \leqq \frac{2c}{L''(\bar{x})},$$

so that convergence of the critical numbers is geometric. (This result should be compared with the corresponding one in section 4 with an interest rate of zero.) If the set-up cost is positive, the analysis is hampered by the fact that we have no information on the behavior of S_n and s_n as functions of time. It is not even known, for example, whether these numbers converge as n tends to infinity, though numerical calculations suggest that it is reasonable to conjecture that they do.

All of these results extend directly to the case of a time lag in delivery, if excess demand is backlogged. In particular, if the set-up cost is zero, the optimal policy for the infinite-period problem may be obtained directly by solving the equation

$$c(1 - \alpha) + \alpha^{\lambda} L'_{\lambda+1}(\bar{x}) = 0,$$

where $L_{\lambda+1}(y)$ is the expected holding and shortage cost over $\lambda + 1$ periods of demand.

There are examples other than those discussed above in which an infinite-period inventory model has a relatively simple solution. Iglehart and Karlin [27] consider the following demand process: At the beginning of each period the system is assumed to be in one of n states, with the transition from period to period governed by a Markov process. If at the beginning of a period the system is in state i, the demand density for the period will be given by $\varphi_i(\xi)$. By varying the transition law for the Markov process, one may obtain a variety of inventory models, including, for example, cyclic fluctuations in the demand distributions. An algorithm for the optimal policy that involves the recursive solution of at most $n!$ equations is presented. (Generally far fewer than $n!$ equations will suffice.)

I should like to mention one additional example before concluding this section. Karlin [34] has noticed that if the demand distributions are stochastically increasing over time ($\int_0^x \varphi_n(\xi)\, d\xi \geqq \int_0^x \varphi_{n+1}(\xi)\, d\xi$, with the subscript referring to the demand distribution relevant in period n counting from the beginning of the problem), the optimal policy in the first period will be identical with the first-period optimal policy for a problem in which the demand density is φ_1 throughout. This idea is used to establish a number of interesting results involving demand distributions that are stochastically fluctuating over time. In chapter 4 of this volume, Veinott gives a number of rather surprising results of this same sort. For example, if the demand distributions are all translates of a given distribution, the optimal policy in the first period is independent of precisely which translates are used in the subsequent periods.

4. Stationary Inventory Models

Section 3 described the use of recursive calculations or dynamic programming techniques in obtaining optimal policies for the finite-horizon inventory problem. These techniques have much to be said in their favor; they are quite flexible, and if knowledge of the specific forms of the optimal policies is incorporated in the computing codes, they are surprisingly rapid. For example, the computing time involved in the calculation of optimal policies for a twenty-period inventory problem, where the demand distribution in each period can take on, say, forty possible values, will be somewhat over one minute even on a slow machine such as the 650.

In a certain sense the speed of these calculations means that the inventory problem involving a single item stocked at a single installation has been solved. On the other hand, the dynamic programming approach provides us with no information about the dependency of optimal policies on the many parameters involved in the model or about the sensitivity of costs as a function of the policies. This type of information may often be obtained from the probabilistic approach to be discussed.

Another reason for the techniques discussed in this section is that frequently optimal policies are not really required. For low-cost, high-turnover, routine items (and these are the items to which inventory theory can most successfully be applied), it is generally sufficient to obtain some relatively simple analytic approximations to optimal policies; and these may be obtained more naturally from a stationary approach than by the analysis of functional equations.

As I shall also indicate in this section, many of the apparently different techniques involved in the stationary analysis of inventory problems are closely connected.

The fundamental observation, which permits us to analyze inventory problems from a probabilistic point of view, is that, given the assumptions that we have generally been making and the use of a fixed policy, the stock levels will frequently be described by a Markov process with stationary transition probabilities (see Massé [42], [43]). Any number of examples will illustrate this point. I shall begin with an example close to the models described in the previous section.

Consider a single item which can be ordered at the beginning of each period and for which delivery is immediate. Demands for the item will be independently and identically distributed in various periods with the common density function $\varphi(\xi)$, and excess demands are to be backlogged. Now let us assume that a specific (S, s) policy has been adopted for use in all periods. If the stock size at the beginning of the first period is some specific value, then, because of the randomness of demand, stock sizes at the beginning of subsequent periods will form a sequence of random variables x_1 , x_2 , ... If x_n is known, it is easy to obtain the possible values for x_{n+1} and their associated probabilities. Let ξ be the random demand during the nth period. If $x_n > s$, no order will be placed during this period, and $x_{n+1} = x_n - \xi$. On the other hand, if $x_n < s$, the stock level will first be raised to S, so that $x_{n+1} = S - \xi$.

The sequence of stock levels is therefore a Markov process, and since we are assuming that demands arise from the same distribution in different periods, the transition probabilities will be stationary in time. The stock size at the beginning of the nth period will have a density function $f_n(x)$ related to f_{n-1} by the typical recursive relationships of Markov processes. Generally, this sequence of densities will converge, as n tends to infinity, to a limiting density $f(x)$. The interpretation is that if the stock size is examined at the beginning of a particular period, it will be a random variable with the density function $f(x)$, if the process has been going on sufficiently long.

Purchase costs and holding and shortage costs will be accumulating from period to period. If the interest rate is zero, so that the discount factor is one, the present value of these costs during the first n periods will be equal to the total cost during these periods and will tend to infinity as n tends to infinity. On the other hand, the average cost will tend to a limit that depends not only on the parameters describing the costs, but also on the specific (S, s) policy used, and it may therefore be used as a basis for comparing various (S, s) policies.

It is quite easy to calculate the long-run average cost per period once the limiting density function is known. For any given stock level x, we write the expected purchase and holding and shortage costs during the period, and then take the expectation of these costs with respect to the distribution of initial stock x. For example, suppose that the purchase cost is given by $K + c \cdot z$, and the expected holding and shortage cost during the period by $L(y)$. If $x > s$, no ordering will be done and the expected total cost during the period will be $L(x)$. On the other hand, if $x < s$, an order of size $S - x$ will be placed; hence the total expected cost during the period is $K + c(S - x) + L(S)$, and we obtain the following formula for the long-run cost per period:

$$\int_s^S L(x)f(x)\,dx + [K + cS + L(S)] \int_{-\infty}^s f(x)\,dx - c \int_{-\infty}^s xf(x)\,dx\,.$$

This approach can be used in many different formulations of the inventory problem, formulations involving, for example, lead times in delivery which may or may not be random, continuous-time rather than discrete-time models, and policies that are not (S, s) policies. In all cases the problem is to determine $f(x)$, the limiting density for stock size.

For the particular model that I have been discussing, the determination of $f(x)$ makes use of that branch of probability theory known as renewal theory, and it will be appropriate to describe some of the elements of renewal theory at this point. For a more elaborate exposition the reader is referred to Karlin [31], [32], and to some of the references listed in these papers.

Renewal theory is concerned with the following type of problem: Consider a piece of equipment that has a random failure time with a cumulative distribution $\Phi(t)$. The equipment is installed at time zero and as soon as it fails it is replaced by a substitute whose failure time is independent of that of the first item, and is identically distributed. The process is continued with replacements being installed whenever a failure occurs.

In the inventory problem the failure-time distribution is to be identified with the demand distribution in a fixed period. A typical problem in inventory theory is to determine the distribution of the number of periods between successive orders, which is the same as determining the distribution of the number of observations from Φ until the sum first exceeds $Q = S - s$. The related problem in renewal theory is to determine the distribution of the number of replacements in a fixed length of time t. Another problem in inventory theory, related to the calculation of shortage costs, is to find the distribution of the amount that accumulated demands cause stock to fall below s the first time they do so. In renewal theory the related problem is to determine the "excess" distribution, or the distribution of the time by which the first failure after t exceeds t.

The probability of n failures between 0 and t, $p_n(t)$, may be written as

$$p_n(t) = \Phi^{(n)}(t) - \Phi^{(n+1)}(t),$$

where $\Phi^{(0)}(t) = 1$ and $\Phi^{(n)}$ is the n-fold convolution of $\Phi(t)$. In the particular case in which failure times are exponentially distributed with a mean m, the number of failures between 0 and t has a Poisson distribution with a mean of t/m. In the general case, the mean of the number of failures in $(0, t)$, which we denote by $M(t)$, is *not* given by the simple form t divided by the mean time between failures.

The quantity $M(t)$ is of particular importance in renewal theory, and may be calculated in several ways. One procedure, which is generally not too effective, is to use the relationship

$$M(t) = \sum_{n=1}^{\infty} \Phi^{(n)}(t).$$

A somewhat better procedure is to observe that $M(t)$ satisfies the fundamental equation of renewal theory, i.e.,

$$M(t) = \Phi(t) + \int_0^t M(t - \xi)\, d\Phi(\xi).$$

If Φ is itself a convolution of exponential distributions (or in the case corresponding to discrete stock sizes, a convolution of geometric distributions), the renewal equation may be converted to a finite system of linear differential (or difference) equations with constant coefficients, and may therefore be solved quite readily. If the failure-time distribution has a density $\varphi(t)$, it is frequently useful to introduce $m(t) = M'(t)$, which satisfies the equation

$$m(t) = \varphi(t) + \int_0^t m(t - \xi)\varphi(\xi)\, d\xi .$$

A third procedure for solving the renewal equation is to use Laplace transforms, which is quite natural, since the renewal equation involves convolutions.

Although $M(t)$ is generally not equal to t/m (m is the mean time between failures, or the mean demand per period in the related inventory problem), this relationship is asymptotically correct. In other words,

$$\lim_{t\to\infty} \frac{M(t)}{t} = \frac{1}{m} .$$

An even stronger theorem, due to Blackwell [9], states that

$$\lim_{t\to\infty} M(t + h) - M(t) = \frac{h}{m} ,$$

as long as Φ is not a "lattice distribution." Very useful generalizations of Blackwell's theorem are known, but I shall not consider them in this paper.

The general implication of these theorems is that $M(t)$ is approximately linear for large t. For inventory theory this means that if the set-up cost, and consequently Q, are large, optimal policies may be expected to approximate those in which the demand density is exponential. I shall come back to this point after a word about the excess distribution.

Let $h(t, x)$ represent the density of the excess distribution, i.e., the time by which the first failure past t exceeds t. The variable of the distribution is x, and t is thought of as being a parameter. If the renewal function $M(t)$ is known, $h(t, x)$ may be expressed conveniently as

$$h(t, x) = \varphi(t + x) + \int_0^t \varphi(t + x - \xi)m(\xi)\, d\xi .$$

As t tends to infinity, the renewal density tends to a limit given by

$$\lim_{t\to\infty} h(t, x) = \frac{1 - \Phi(x)}{m} .$$

Now let me turn to the application of renewal theory to the inventory problem. As I remarked before, the important problem is that of determining the limiting distribution of stock size $f(x)$. I have indicated above the transition probabilities for stock size viewed as a Markov process. As is customary in the study of Markov processes, this permits us to write a recursive relationship for f, which in this case becomes

$$f(u) = \varphi(S-u)\int_{-\infty}^{s} f(x)\,dx + \int_{s}^{S} f(x)\varphi(x-u)\,dx\,.$$

The solution of this equation can be expressed in terms of the renewal function as follows:

$$f(x) = \begin{cases} \dfrac{m(S-x)}{1+M(Q)} & s \leqq x \leqq S\,, \\[2ex] \dfrac{\varphi(S-x) + \int_0^Q m(\xi)\varphi(S-x-\xi)\,d\xi}{1+M(Q)} & x \leqq s\,. \end{cases}$$

Note the similarity between the second part of this expression and the corresponding formula for the excess distribution.

Using this expression for the limiting distribution of stock size, we can write the long-run average cost per period as

$$cm - \frac{K + L(S) + \int_0^Q L(S-x)m(x)\,dx}{1+M(Q)}\,,$$

which may then be minimized as a function of Q and S to provide us with the optimal policy. This formula may be found in Arrow, Harris, and Marschak [1]. A large number of numerical examples based on a discrete stock analogue may be found in Wagner [57].

I shall now indicate what these formulas become in the special case of an exponential demand distribution $[\varphi(\xi) = 1/m(e^{-\xi/m})]$. The limiting distribution for stock size is given by

$$f(x) = \begin{cases} \dfrac{1}{m+Q} & s \leqq x \leqq S\,, \\[2ex] \dfrac{e^{(x-s)/m}}{m+Q} & x < s\,, \end{cases}$$

and the long-run average cost per period, assuming linear holding and shortage costs, the former charged at the beginning of the period and the latter at the end, is

$$cm + \frac{1}{m+Q}\left[Km + hsm + hQm + hsQ + \frac{hQ^2}{2} + mpe^{-s/m}\right].$$

In order to find the optimal policy we set the two derivatives, with respect to Q and s, equal to zero, and obtain the following remarkably simple expressions:

$$Q = \sqrt{\frac{2Km}{h}}, \qquad e^{-s/m} = \frac{h}{p}\left(1 + \frac{Q}{m}\right).$$

The first of these expressions is, of course, the Wilson lot-size formula, and I, for one, have always found it surprising that the Wilson formula should appear in so many different areas of inventory theory. Another interesting point is that the expression for Q is independent of the shortage cost.

If we depart from the exponential distribution, or from the zero lead time problem, the expressions given above are no longer correct. I have mentioned, however, that for large values of the set-up cost, the limit theorems of renewal theory suggest that the distributions may be approximated by exponential distributions, and we may therefore expect to find approximations similar to the exact expressions given above. This point has been examined by Roberts [48], using a general demand distribution, linear holding and shortage costs, and an arbitrary lead time λ with the backlog model.

In Roberts' study it is necessary to take large values of both the set-up cost and the shortage cost. It is shown that Q and s may be approximated by

$$Q \sim \sqrt{\frac{2Km}{h}} \qquad \text{and} \qquad p\int_s^{\infty} (\xi - s)\varphi_{\lambda+1}(\xi)\, d\xi \sim \sqrt{2Kmh}\,.$$

Again we see that Q is given approximately by the Wilson lot-size formula, independently of the shortage cost. If a more refined approximation is desired, it is necessary to add to the expression for Q a correction term involving s, and therefore involving the shortage cost to a slight degree.

These approximations are remarkably good, even for moderately small values of K and p; moreover they seem to be good approximations to the limiting optimal policies in the dynamic programming calculations. This is true even if there is an interest rate involved in the dynamic programming formulation; as long as the interest rate is not too large, the above formulas are applied, with the holding cost increased by the interest rate times the unit cost.

Table 1 compares the value of Q obtained by means of dynamic programming with that obtained from the Wilson lot-size formula. In these examples a geometric distribution is used with an interest rate of zero, and holding and unit costs of one. The agreement is remarkable, even without the corrective second order term for Q.

The comparison for s, although not as close as that for Q, is sufficiently good to suggest using these approximations if the dynamic programming calculation involves a long horizon.

TABLE I

Mean	Lead time	Shortage cost	Set-up cost	Dynam. prog. Q	Wilson Q
4	0	30	4	6	5
4	6	1000	100	29	28
1	6	1000	100	15	14
4	6	1000	4	7	5
4	6	100	100	31	28
1	4	100	100	16	14
4	2	100	100	30	28
1	4	30	100	15	14
4	6	100	4	8	5
1	4	30	4	4	2
.25	6	100	4	2	1

These calculations are based on the steady-state model of this section, which has been selected as the stationary analogue of the dynamic programming model of section 3. Other stationary models that involve continuous review of stock, discrete stock levels, random lead times, and policies other than (S, s) policies have been suggested by a number of authors. I shall say a bit about some of these other models, but first I should like to indicate another aspect of the close connection between the model of this section and the dynamic programming model. As we shall see, these remarks are related to some other techniques for solving inventory problems.

In the stationary approach we select a particular (S, s) policy, calculate the long-run costs based on this policy, and then select the policy variables so as to minimize long-run cost. Let the minimum cost be denoted by k. In the dynamic programming approach, the technique depends on the minimum cost functions $C_n(x)$ and the functional equations of section 3. If the interest rate is zero, then as n becomes infinite, $C_n(x)$ will tend to infinity. It seems plausible that there will be some connection between $\lim_{n\to\infty} [C_n(x)]/n$ and k. Note that this is a rather delicate mathematical question. If the optimal policies obtained from dynamic programming were, in fact, independent of n, then standard ergodic theorems could be used to demonstrate that these two costs are the same. However, if the set-up cost is positive, we know nothing about the long-run behavior of (S_n, s_n), not even that the numbers converge. Contrary to the opinion that seems to be held by a number of people, there are no general theorems about the convergence of optimal inventory policies. Even if there were, it would not be easy to demonstrate the equivalence of the two methods of calculating long-run cost.

This is precisely what Iglehart has done in chapter 1 of this volume. He demonstrates the general theorem that

$$\lim_{n\to\infty} \frac{C_n(x)}{n} = k ,$$

and several related and stronger theorems. His method, which is based upon some previous work of Bellman [5], [6], depends upon constructing a solution to the modified functional equation

$$\psi(x) + k = \min_{y \geqq x} \left\{ c(y - x) + L(y) + \int_0^\infty \psi(y - \xi)\varphi(\xi)\, d\xi \right\} ,$$

and then demonstrating inductively that

$$| C_n(x) - nk - \psi(x) | \leqq A ,$$

for all $x \leq X$; the constant A depends upon X, which may be selected as being arbitrarily large. Iglehart also demonstrates the interesting result that if $K = 0$, then

$$\bar{x} - \bar{x}_n = 0\left(\frac{1}{\sqrt{n}}\right) ,$$

which should be compared with the geometric rate of convergence in the case of a positive interest rate.

The method of proof used by Iglehart has a remarkable connection with the technique of "policy iteration" discussed by Bellman [7] and Howard [25]. To fix the ideas, let me return for a moment to the case of an infinite-period problem with a positive interest rate. The associated functional equation is

$$C(x) = \min_{y \geqq x} \left\{ c(y - x) + L(y) + \alpha \int_0^\infty C(y - \xi)\varphi(\xi)\, d\xi \right\} .$$

One procedure for solving this equation, which Howard calls "value iteration" and Bellman the "method of successive approximations," is to follow the method of the previous section, i.e., to calculate the functions $C_n(x)$ and take their limit as n tends to infinity. This method may be modified somewhat, since we are interested only in the infinite-period problem, by taking as the initial choice of $C_0(x)$ some function other than the one identically zero. As we know, the cost functions will converge, and therefore the optimal policy for the infinite-period problem will be obtained.

There is a modification of this approach, halfway between the value iteration and policy iteration techniques, which utilizes the fact that for certain cost functions we do know the form of the optimal policy. For example, if we know that the optimal policy is (S, s), we may select an arbitrary choice of these two numbers and solve the following equation for $C(x)$:

$$C(x) = \begin{cases} K + c \cdot (S - x) + L(S) + \alpha \int_0^\infty C(S - \xi)\varphi(\xi)\, d\xi & x \leqq s\,, \\ L(x) + \alpha \int_0^\infty C(x - \xi)\varphi(\xi)\, d\xi & x > s\,. \end{cases}$$

This equation is a modification of the standard renewal equation and may be solved by any number of methods. The resulting function $C(x)$ will, of course, depend upon the particular policy used. The optimal policy will then be obtained by minimizing, say, $C(0)$, with respect to S and s. This procedure, which is obviously similar to the stationary analysis described above, was first suggested by Arrow, Harris, and Marschak [1].

A third method for the infinite-period problem with a discount factor is that of policy iteration. The difference between this method and the previous one is somewhat subtle; we begin in the same way by selecting the policy and calculating the function $C(x)$. Since the policy is not optimal, this is not the correct infinite-period cost. The next step is to obtain an improvement over the original policy by selecting that policy $y^1(x) \geqq x$ which minimizes

$$c(y - x) + L(y) + \alpha \int_0^\infty C(y - \xi)\varphi(\xi)\, d\xi\,.$$

On the basis of this policy, a new cost function $C^1(x)$ is calculated and then a second policy $y^2(x)$ is selected, using an expression like that above with C replaced by C^1. The process continues until the same C is reached on two successive stages.

To my mind, policy iteration has a serious drawback for the inventory problem, even though it is undoubtedly an excellent procedure for other problems, especially those in which the number of state variables is finite. The drawback is that although (S, s) policies may be optimal, the C functions appearing in the middle of the calculation will not be K-convex, and the intermediary policies will therefore be of a more complex sort. Calculating the subsequent C's will be more difficult than solving a renewal equation. For finite-state problems the C's may be calculated by inverting a matrix; in the inventory problem, however, there is considerable information about the structure of the problem that should not be omitted in the calculation of optimal policies.

Of the three techniques described above, the first two go over immediately to the case of a zero interest rate. For policy iteration, however, there is no analogue of the function $C(x)$. Howard suggests (in another context and in another notation) that the calculation be based on the functions $\psi(x)$ and the related equation.

The technique of policy iteration has much in common with a recent suggestion for solving infinite-period inventory problems by means of linear programming (see Manne [41], Wagner [56], and d'Epenoux [14]). I shall begin with the ideas of d'Epenoux. Since the problem is to be discussed by means of linear programming, we must assume a finite number of stock sizes at the beginning of the period, and this is most conveniently done by assuming that excess demand is *not* backlogged, that stock sizes are integral multiples of some common unit, and that no order is placed which raises stock to a level above, say, N.

Instead of the functions $C(x)$, we shall be dealing with a discrete sequence c_i, the minimum discounted cost if the problem begins in state i. The functional equation satisfied by these costs will be

$$c_i = \min_{i \leqq j \leqq N} \left\{ l_{ij} + \alpha \sum_{k=0}^{j} c_{j-k} p_k + \alpha c_0 \sum_{k=j+1}^{\infty} p_k \right\},$$

where l_{ij} represents the purchase, holding, *and* shortage cost if we begin in state i and raise stock to state j. The fact that there are two terms involving the demand distribution $(p_0, p_1, \ldots)$ depends upon the assumption that we are not backlogging.

To convert this problem to a linear programming problem we consider $c_0, c_1, \ldots, c_N$ to be the variables of the problem, unrestricted in sign, and satisfying the $N(N+1)/2$ linear inequalities

$$c_i \leqq l_{ij} + \alpha \sum_{k=0}^{j} c_{j-k} p_k + \alpha c_0 \sum_{k=j+1}^{\infty} p_k .$$

There is an inequality for every pair (i, j) with $j \geqq i$. We then select an arbitrary set of non-negative weights $u_0, \ldots, u_N$ (not all zero) and consider the problem of maximizing $\Sigma\, u_i c_i$ with respect to all of those $\{c_i\}$ satisfying the above linear equalities. Assuming that the optimal inventory policy is unique (there are, otherwise, some problems of degeneracy), we may show that the same solution is obtained regardless of the values of the u_i. Moreover, the solution will have the following important property: for every i there will be a single value of $j \geqq i$ for which there is equality for the corresponding linear inequality. This means, of course, that the solution of the linear programming problem satisfies the optimal inventory equation

$$C_i = \min_{i \leqq j \leqq N} \{l_{ij} + \ldots\} .$$

This seems like an entirely different approach to the solution of inventory problems. It is, however, relatively easy to see that if we solve the dual of this linear programming problem, we are essentially using the method of policy iteration. The dual problem involves non-negative variables x_{ij} with $j \geqq i$. The dual constraints (one for each i) are

$$\sum_{j \geqq i} x_{ij} - \alpha \sum_{j \geqq i} p_{j-i} \left(\sum_m x_{mj} \right) = u_i \qquad (i \neq 0),$$

and

$$\sum_j x_{0j} - \alpha \sum_j \left(\sum_{k=j}^{\infty} p_k \right) \left(\sum_m x_{mj} \right) = u_0 ,$$

and we are to minimize $\Sigma\Sigma_{j \geqq i}\, l_{ij} x_{ij}$.

It is simple to verify that the dual constraints imply that $\Sigma\Sigma_{j \geqq i}\, x_{ij} = \Sigma\, u_i/(1 - \alpha)$. Since the only condition on the u_i is that they be non-negative with a sum different from zero, we may replace this problem by that of minimizing $\Sigma\Sigma_{j \geqq i}\, l_{ij} x_{ij}$ subject to

$$\sum_{j \geqq i} x_{ij} - \alpha \sum_{j \geqq i} p_{j-i} \left(\sum_m x_{mj} \right) \geqq 0 \qquad (i \neq 0),$$

$$\sum_j x_{0j} - \alpha \sum_j \left(\sum_{k \geqq j}^{\infty} p_k \right) \left(\sum_m x_{mj} \right) \geqq 0 ,$$

and

$$\sum_{j \geqq i} \sum x_{ij} = 1 .$$

Now in the policy iteration technique we begin with a specific inventory policy giving j as a function of i. Let us see how this policy may be used to obtain a feasible solution to the dual problem. If this policy is used, the stock on hand at the beginning of a period forms a Markov process, and we may calculate $\pi_i(n)$, the probability that the system is in state i at the beginning of period n, given any particular initial distribution $\pi_i(0)$. We then define x_{ij}:

$$x_{ij} = (1 - \alpha) \sum_{n=0}^{\infty} \alpha^n \, \pi_i(n)$$

if the policy under question involves raising the stock from i to j, and $x_{ij} = 0$ otherwise. It may be verified by direct substitution that these quantities satisfy the dual constraints.

The next step in policy iteration is to calculate the C functions based on this policy, or, in our context, to calculate the numbers $c_0, c_1, \dots, c_N$ that satisfy the primal constraints with equality for those pairs (i, j) involved in the policy. But these numbers will obviously be the prices associated with the particular dual feasible solution we have obtained. The next step in the simplex method is to obtain another dual feasible solution based on the prices that have been calculated, and it is at this point that the simplex method and policy iteration diverge slightly. Both methods lead to an alternative dual feasible solution; in the simplex method the new basis is adjacent to the old one, whereas in Howard's method the new basis may involve a substantial change.

In the case of a zero interest rate (the case discussed by Manne), the dual variables x_{ij} may be identified with the joint stationary probabilities of being in state i and raising stock to state j, based on the particular policy being used, and the problem is then to minimize stationary expected cost. The primal problem involves quantities analogous to the ψ functions previously discussed. Again there is a great similarity between Howard's procedure and the simplex method, with the important difference noted above.

It is now time to leave these matters and turn to a discussion of the different inventory models that have been suggested for use with the stationary approach. Many operations research texts discuss a simple-looking continuous-time inventory model that yields the Wilson lot-size formula and a corresponding equation for s. The model is never discussed from the point of view of finding the stationary distribution for a class of policies and then minimizing cost with respect to the policy variables. Often the approach taken obscures the fact that the model as usually presented contains a large number of contradictions that can be avoided only by rather restrictive assumptions, such as that demand arises from a Poisson time series, and the like. Perhaps the clearest treatment of this particular model is to be found in Schlaifer [53].

The techniques for analyzing continuous-time inventory models correctly are very similar to those of queueing theory. Consider, for example, the following model: items are demanded one at a time according to a Poisson time series. When an order is placed the lead time distribution is assumed to be exponential, with independent lead times for different orders; for the moment let us assume that if a demand cannot be satisfied out of current stock, it will be lost.

In order to apply the stationary approach to this problem, it is necessary to specify an inventory policy. Since excess demands are not backlogged and there is a random lead time that may overlap orders, it is no longer correct that simple policies such as (S, s) policies will be optimal. There is no reason, however, why we should not restrict our attention to policies such as (S, s) policies and attempt to select the one that minimizes long-run cost.

Let us begin with the special case of an (S, s) policy based on stock on hand plus stock on order in which $s = S - 1$. This means that a total of S items are to be kept on hand and on order; whenever a demand is satisfied, an additional order is placed. With these assumptions, the number of outstanding orders is a continuous-time Markov process. The stationary distribution for the number of outstanding orders proves to be

$$\pi_m = \frac{\alpha^m/m!}{\sum_{j=0}^{S} \alpha^j/j!},$$

where α is the expected lead time multiplied by the average demand per unit of time.

This model is identical with the machine repair problem, which is well known in queueing theory (see Feller [17, p. 416]). In this latter problem a single repairman services S machines, whose failure times are independently and exponentially distributed. The repair time for a machine is assumed to be exponential, and machines queue for service if several have failed. The number of machines currently operating is identical with the number of outstanding orders in the inventory problem.

For the particular policy described above, the limiting distribution proves to be independent of the specific form of the service-time distribution, and depends only on its mean. This result is also correct if excess demands are backlogged. A corresponding result holds if we consider an (S, s) policy with $s = 0$. In this case the relevant variable is the number of items in current inventory, and the limiting distribution is

$$p_0 = \frac{\alpha}{S + \alpha}, \qquad p_m = \frac{1}{S + \alpha} \qquad (m > 1).$$

This model, assuming Poisson demand, exponential lead time, and a general (S, s) policy, has been analyzed by Scarf [49] in both the backlog and the non-backlog form, and also by Morse [45] and by Galliher, Morse, and Simond [19]. As might be expected, results are considerably easier to obtain in the backlog case; and it is also possible to generalize the model to one in which there is an arbitrary distribution of time between successive demands and an arbitrary lead time distribution. (See Karlin and Scarf [37] for an analysis of this problem that leans heavily upon some previous work of Takács [54].)

Occasionally the analysis is made somewhat simpler by the use of other types of policies. For example, Morse [45] considers the following version of an (S, s) policy: When inventory falls to s, an order of size $S - s$ is requested; if inventory on hand is depleted to zero while this order is outstanding, an additional order of size s is placed.

Another variation is considered by Gaver [20]. When stock on *hand* falls to s, an order of size $(S - s)$ is requested; only one outstanding order is permitted at any given moment of time. A general lead-time distribution is considered, excess demand is not backlogged, and a compound Poisson process is assumed for demands.

A third variation is discussed by Karlin and Fabens [35] and by Hadley and Whitin [23] in a discrete-time inventory model of the type considered at the beginning of this section. Orders are placed whenever stock falls below s, but the order size is a multiple of $S - s$. The largest multiple that keeps the stock size less than or equal to S after ordering is selected. The virtue of this type of policy is that the stationary distribution of stock size (after placing the order) will be uniform between s and S, regardless of the demand distribution. For (S, s) policies of the conventional sort, this result would be correct only for an exponential demand distribution.

All of the work reported in the last two sections has treated supply and demand in a somewhat symmetrical fashion. Demands have been assumed to be random and exogenously given, in contrast to supply, which has been assumed to be controllable. There are any number of situations, however, in which it is required to store an item whose supply is not completely under our control. The most important example of this type of situation is in the problem of water storage, where the supply of water, possibly random, is given as a datum.

The problems of reservoir storage have been analyzed by means of techniques very similar to those already described for the inventory problem, although a number of modifications must be made in incorporating exogenous supply. A thorough treatment of this area would require at least as many pages as I have already given to inventory theory. It seems more appropriate to indicate several references that the reader may wish to pursue himself.

For the deterministic version, the paper by Koopmans [39] may be consulted, and for the dynamic programming approach, the article by Gessford and Karlin [21], and also the book by Massé [42]. A vast amount of work has been done on the stationary analysis of water storage problems under the heading of the Theory of Dams. Much of the work done in this field prior to 1959 is discussed by Moran in [44].

5. Multi-Echelon Inventory Problems

In the last several sections I have surveyed the major techniques that have been used to analyze decisions about the control of a single item stocked at a single installation. The relatively simple inventory policies obtained from this type of analysis are probably sufficient for the control of low-cost routine items. On the other hand, if we are concerned with high-cost items of considerable importance, with relatively small demands, it may be useful to consider a more elaborate type of analysis.

One of the assumptions implicit in our previous treatment of the inventory problem was the infinite availability of stock. Delivery of an order was assumed to occur with some fixed or perhaps random lead time independently of the size of the order. The lead time may be thought of as the time required to manufacture the item, the time necessary to transport the item, or some combination of the two. If the former interpretation is taken, our previous assumption of infinite availability of stock is equivalent to assuming away any possible limitations on the rate at which the item may be produced. Some work has been done on the calculation of optimal policies for the inventory problem with random demands, with the costs of modifying the rate of production explicitly considered (Beckmann [4] and Orr [46]). As might be expected, both the functional equations and the optimal policies are of a more complex sort than those previously obtained.

If the lead time in delivery is considered to be primarily the time required to transport the item, our previous analysis is equivalent to assuming the existence of a large supply of the item stocked at some alternative location, say a warehouse. For low-demand items, however, warehouse supply should not be taken for granted. The warehouse may itself run out of stock, resulting in longer lead times for the delivery of stock to the using activity. It would seem appropriate, in this type of situation, to consider the problems of stocking at the warehouse and at the using activity together, as a single inventory problem.

The warehouse may, of course, be supplying stock to several installations. One of the problems would then be that of allocating available warehouse stock among the various installations sending in requests. This problem would be most pronounced in those periods in which warehouse stock was insufficient to satisfy all incoming requests. On the other hand, if stocks are maintained at several installations, at the beginning of a period one installation may be considerably overstocked in comparison with the others. In this event, it may be economical to transship stock among the various installations, in addition, perhaps, to delivering stock from the warehouse. If the stockage and transportation decisions are considered simultaneously, it should be clear that optimal policies will be of a rather complex form.

In chapter 3 of this volume Gross considers the single-period problem with one warehouse, and an arbitrary number of activities supplied by the warehouse. Optimal policies are obtained for two using activities, and an iterative procedure is given for the general case. In chapter 5 of this volume, Hadley and Whitin suggest various simple policies for the multi-period version of this problem.

The dynamic, multi-installation inventory problem can be examined from the point of view of dynamic programming. As usual, we define a sequence of cost functions whose independent variables indicate the state in which the system may be. The cost functions will satisfy a recursive equation, which may, at least in theory, be solved so as to obtain optimal purchasing and transshipment rules. The difficulty is, however, that the number of possible states will be exceptionally large. The disposition of stock at all of the installations will have to be indicated, along with a description of quantities on order and being shipped. With so large a number of possible states, dynamic programming calculations take a long time.

The same type of difficulty appeared in our previous discussion of optimal policies in the presence of a lead time in delivery. We overcame the delivery lead time problem by assuming that excess demand is backlogged, and on the basis of this assumption transforming the functional equation to one involving a single variable. For certain rather extreme types of multi-echelon problems, a similar type of approach is possible (see Clark [11] and Clark and Scarf [12]).

Consider the case of N installations arranged in series, so that installation 1 receives stock from installation 2, 2 from 3, and so on. Stock enters the system through the highest installation, with a fixed delivery lead time, and is shipped through the subsequent installations, stock being removed at each step to satisfy random demands. If the following assumptions are made, then the optimal policies may be calculated easily.

1. The cost of shipping from a single installation to the next one is proportional to the quantity shipped.

2. Excess demand is backlogged.

3. The expected holding and shortage costs to be charged to each installation during a single period are functions of the stock at that installation *plus* all stock in transit and on hand at lower installations.

Of the three assumptions, the first is the crucial one; the third may easily be satisfied by an appropriate accounting of holding and shortage costs.

Let me consider a special case in order to indicate how the method works. There will be two installations, with installation 1 receiving stock from installation 2. Both the shipment time from 2 to 1 and the lead time to 2 will be assumed to be a single period. Demand for the item occurs only at the lower installation with a density $\varphi(\xi)$. The expected holding and shortage cost functions will be $L_1(x_1)$ and $L_2(x_2)$, respectively, with x_2 representing stock at *both* installations.

At any moment of time the state of the system is described by the pair (x_1, x_2), and we may therefore define a sequence of minimum cost functions $C_n(x_1, x_2)$ that satisfy

$$C_n(x_1, x_2) = \min_{\substack{x_1 \leq y \leq x_2 \\ 0 \leq z}} \Big\{ c(z) + c_T \cdot (y - x) + L_1(x_1) + L_2(x_2) + \alpha \int_0^\infty C_{n-1}(y - \xi, x_2 + z - \xi)\varphi(\xi)\, d\xi \Big\},$$

where $c(z)$ is the cost of purchasing z units, and c_T the unit transportation cost. We begin our analysis of this equation by considering the lower installation by itself, and the functional equation that would apply if there were no limitation on available stock, i.e.,

$$C_n(x_1) = \min_{y \geq x_1} \left\{ c_T \cdot (y - x_1) + L_1(x_1) + \alpha \int_0^\infty C_{n-1}(x_1 - \xi)\varphi(\xi)\, d\xi \right\}.$$

If L_1 is convex, the optimal policy for this equation, which as yet has no relevance to the original problem, is determined by a sequence of critical numbers $\bar{x}_1, \bar{x}_2, \bar{x}_3, \ldots$, with the interpretation that at the beginning of each period stock is raised to the corresponding level. The problem, of course, is that there may not be adequate stock at the second installation to satisfy these requests. It may be shown, however, that the optimal system policy is to satisfy as many of these requests as possible.

This takes care of the optimal policy at the lower installation. The remaining problem is to determine the ordering policy for system stock. The rather surprising conclusion is that the system stockage policy is a function of x_2 alone; moreover, the optimal policy may be obtained from the recursive calculation of a sequence of functions of one variable. The important idea is that the system shortage cost involved in L_2 must be augmented by an additional shortage cost resulting from the fact that system stock may be insufficient to supply the requests from installation 1.

The actual determination of the additional shortage cost is relatively simple. Suppose that at the beginning of period n, the stock on hand at the lower installation is $u < \bar{x}_n$. An order of size $\bar{x}_n - u$ would then be placed, resulting in a cost of

$$c_T \cdot (\bar{x}_n - u) + L_1(u) + \alpha \int_0^\infty C_{n-1}(\bar{x}_n - \xi)\varphi(\xi)\, d\xi .$$

If the system stock x_2 is less than $\bar{x}_n$, it will be impossible to satisfy this request, and the resulting cost will be

$$c_T \cdot (x_2 - u) + L_1(u) + \alpha \int_0^\infty C_{n-1}(x_2 - \xi)\varphi(\xi)\, d\xi .$$

The second cost is larger than the first, and the difference in cost is precisely the incremental shortage cost to be charged to the system. This extra cost may be written as

$$\Lambda_n(x_2) = \begin{cases} c_T \cdot (x_2 - \bar{x}_n) + \alpha \int_0^\infty [C_{n-1}(x_2 - \xi) - C_{n-1}(\bar{x}_n - \xi)]\varphi(\xi)\, d\xi & (x_2 \leqq \bar{x}_n), \\ 0 & (x_2 > \bar{x}_n). \end{cases}$$

The optimal system policy is then calculated on the basis of the functional equation

$$q_n(x_2) = \min_{z \geq 0} \left\{ c(z) + L_2(x_2) + \Lambda_n(x_2) + \alpha \int_0^\infty q_{n-1}(x_2 + z - \xi)\varphi(\xi)\, d\xi \right\}.$$

In the case of several installations linked in series, the procedure may be repeated with an additional shortage cost added at each step because of the inability to satisfy requests from below. It is important to realize that we are, in fact, obtaining optimal policies. Although the argument given above has a fanciful quality, it may be made quite rigorous by means of the result

$$C_n(x_1, x_2) \equiv C_n(x_1) + q_n(x_2) .$$

The arrangement of installations involved in the preceding discussion is unusual. The more customary situation would be one in which a warehouse serves several installations rather than just one. In this case the procedure suggested by Clark is to calculate the critical numbers at each of the lower levels independently. If at the beginning of a period system stock is inadequate to supply all requests, the available stock is rationed among the various installations so as to minimize the sum of the additional shortage costs.

This policy is generally not optimal. In order to obtain an optimal policy it would be necessary to add an additional shortage cost that depends not only on system stock, but on the actual distribution of stock among the lower levels. If this were done, however, we would no longer be calculating system purchases on the basis of functions of a single variable.

There is one possible suggestion that has not yet been thoroughly explored. That is to attempt to bound the additional shortage cost from above and below by functions of x_2 alone. If this were done, we would obtain simple bounds both above and below for the true minimum cost, and also a policy whose cost falls between these bounds. If the bounds were close, the policy could be used.

This type of approximation has not yet been done for the case of a warehouse supplying several installations. It has, however, been discussed for a different multi-echelon problem that is not capable of being factored. The problem is that of two installations arranged in series, with a set-up cost in transportation between them (Clark and Scarf [13]).

References

[1] Arrow, K. J., T. Harris, and J. Marschak. Optimal Inventory Policy. *Econometrica*, 1951, **19** (3), 250–72.

[2] Arrow, K. J., and S. Karlin. Production Planning without Storage. Chap. 6 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in the Mathematical Theory of Inventory and Production*, Stanford, Calif.: Stanford Univ. Press, 1958.

[3] Arrow, K. J., S. Karlin, and H. Scarf. *Studies in the Mathematical Theory of Inventory and Production*. Stanford, Calif.: Stanford Univ. Press, 1958.

[4] Beckmann, M. J. Production Smoothing and Inventory Control. *Operations Res.*, 1961, **9**, 456–67.

[5] BELLMAN, R. On a Quasi-Linear Equation. *Canadian J. Math.*, 1956, **8**, 198–202.
[6] BELLMAN, R. A Markov Decision Process. *J. Math. and Mech.*, 1957, **6**, 679–84.
[7] BELLMAN, R. *Dynamic Programming*. Princeton, N.J.:Princeton Univ. Press, 1957.
[8] BELLMAN, R., I. GLICKSBERG, and O. GROSS. On the Optimal Inventory Equation. *Management Sci.*, 1955, **2**, 83–104.
[9] BLACKWELL, D. A Renewal Theorem. *Duke Math. J.*, 1948, **15**, 145–50.
[10] CHURCHMAN, C. W., R. L. ACKOFF, and E. L. ARNOFF. *Introduction to Operations Research*. New York: Wiley, 1957.
[11] CLARK, A. J. "A Dynamic, Single-Item, Multi-Echelon Inventory Model." RM-2297, Santa Monica, Calif.: The RAND Corporation, 1958.
[12] CLARK, A. J., and H. SCARF. Optimal Policies for a Multi-Echelon Inventory Problem. *Management Sci.*, 1960, **6**, 475–90.
[13] CLARK, A. J., and H. SCARF. Approximate Solutions to a Simple Multi-Echelon Inventory Problem. Chap. 5 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in Applied Probability and Management Science*, Stanford, Calif.: Stanford Univ. Press, 1962.
[14] D'EPENOUX, F. Sur un Problème de Production et de Stockage dans l'Aléatoire. *Revue Française de Recherche Opérationelle*, 1960, **4**, 3.
[15] DVORETZKY, A., J. KIEFER, and J. WOLFOWITZ. The Inventory Problem: I. Case of Known Distributions of Demand. *Econometrica*, 1952, **20** (2), 187–222.
[16] DVORETZKY, A., J. KIEFER, and J. WOLFOWITZ. The Inventory Problem: II. Case of Unknown Distributions of Demand. *Econometrica*, 1952, **20** (3), 450–66.
[17] FELLER, W. *An Introduction to Probability Theory and its Applications*. 2d ed., New York: Wiley, 1957. Vol. 1.
[18] FUKUDA, Y. "Optimal Disposal Policies." R-152, Los Angeles, Calif.: Planning Research Corporation, March 1960.
[19] GALLIHER, H. P., P. M. MORSE, and M. SIMOND. Dynamics of Two Classes of Continuous Review Inventory Systems. *Operations Res.*, 1959, **7**, 362–84.
[20] GAVER, D. P., Jr. Renewal Theoretic Analysis of a Two-Bin Inventory Control Policy. *Naval Res. Logist. Quart.*, 1959, **6**, 141–63.
[21] GESSFORD, J., and S. KARLIN. Optimal Policy for Hydroelectric Operations. Chap. 11 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in the Mathematical Theory of Inventory and Production*, Stanford, Calif.: Stanford Univ. Press, 1958.
[22] GILMORE, P. C., and R. E. GOMORY. A Linear Programming Approach to the Cutting Stock Problem. *Operations Res.*, 1961, **9** (6), 849–59.
[23] HADLEY, G., and T. M. WHITIN. A Family of Inventory Models. *Management Sci.*, 1961, **7**, 351–71.
[24] HOLT, C. C., F. MODIGLIANI, J. F. MUTH, and H. A. SIMON. *Planning Production, Inventories and Work Force*. Englewood Cliffs, N.J.: Prentice-Hall, 1960.
[25] HOWARD, R. A. *Dynamic Programming and Markov Processes*. New York: Wiley, 1960.
[26] IGLEHART, D. L. Optimality of (S, s) Policies in the Infinite Horizon Dynamic Inventory Problem (to be published in *Management Science*).
[27] IGLEHART, D. L., and S. KARLIN. Optimal Policy for Dynamic Inventory Process with Nonstationary Stochastic Demands. Chap. 8 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in Applied Probability and Management Science*, Stanford, Calif.: Stanford Univ. Press, 1962.

[28] JOHNSON, S. M. Sequential Production Planning over Time at Minimum Cost. *Management Sci.*, 1957, 3, 435–37.

[29] JOHNSON, S., and G. DANTZIG. A Production Smoothing Problem. *Proceedings of the Second Symposium in Linear Programming*, Vol. 1. Washington, D.C.: 1955, 151–76.

[30] KARLIN, S. The Structure of Dynamic Programming Models. *Naval Res. Logist. Quart.*, 1955, 2 (4), 285–94.

[31] KARLIN, S. Steady State Solutions. Chap. 14 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in the Mathematical Theory of Inventory and Production*, Stanford, Calif.: Stanford Univ. Press, 1958.

[32] KARLIN, S. The Application of Renewal Theory to the Study of Inventory Policies. Chap. 15 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in the Mathematical Theory of Inventory and Production*, Stanford, Calif.: Stanford Univ. Press, 1958.

[33] KARLIN, S. One Stage Inventory Models with Uncertainty. Chap. 8 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in the Mathematical Theory of Inventory and Production*, Stanford, Calif.: Stanford Univ. Press, 1958.

[34] KARLIN, S. Dynamic Inventory Policy with Varying Stochastic Demands. *Management Sci.*, 1960, 6 (3), 231–58.

[35] KARLIN, S., and A. FABENS. A Stationary Inventory Model with Markovian Demand. Chap. 11 in Arrow, K. J., S. Karlin, and P. Suppes (eds.), *Mathematical Methods in the Social Sciences, 1959*, Stanford, Calif.: Stanford Univ. Press, 1960.

[36] KARLIN, S., and H. SCARF. Inventory Models of the Arrow-Harris-Marschak Type with Time Lag. Chap. 10 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in the Mathematical Theory of Inventory and Production*, Stanford, Calif.: Stanford Univ. Press, 1958.

[37] KARLIN, S., and H. SCARF. Inventory Models and Related Stochastic Processes. Chap. 17 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in the Mathematical Theory of Inventory and Production*, Stanford, Calif.: Stanford Univ. Press, 1958.

[38] KARUSH, W., and A. VAZSONYI. Mathematical Programming and Service Scheduling. *Management Sci.*, 1957, 3, 140–48.

[39] KOOPMANS, T. C. Water Storage Policy in a Simplified Hydroelectric System. *Proceedings of the First International Conference on Operational Research*, Baltimore, Md.: Oper. Res. Soc. of Amer., 1957.

[40] MANNE, A. S. Programming of Economic Lot Sizes. *Management Sci.*, 1958, 4 (2), 115–35.

[41] MANNE, A. S. Linear Programming and Sequential Decisions, *Management Sci.*, 1960, 6, 259–67.

[42] MASSÉ, P. *Les Réserves et la Régulation de l'avenir dans la vie Économique*. 2 vols. Paris: Hermann, 1946.

[43] MASSÉ, P. *Le Choix des Investissements*. Paris: Dunod, 1959.

[44] MORAN, P. A. P. *The Theory of Storage*. London: Methuen, 1959.

[45] MORSE, P. M. *Queues, Inventories and Maintenance*. New York: Wiley, 1958.

[46] ORR, D. "Production Stability and Inventory Variation." Res. Memo. No. 15, Econometric Research Program, Princeton, N.J.: Princeton Univ. May, 1960.

[47] PROSCHAN, F. Optimal System Supply. *Naval Res. Logist. Quart.*, 1960, 7 (4), 609–45.

[48] ROBERTS, D. M. Approximations to Optimal Policies in a Dynamic Inventory Model. Chap. 13 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in Applied Probability and Management Science*, Stanford, Calif.: Stanford Univ. Press, 1962.

[49] SCARF, H. Stationary Operating Characteristics of an Inventory Model with Time Lag. Chap. 16 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), *Studies in the Mathematical Theory of Inventory and Production*, Stanford, Calif.: Stanford Univ. Press, 1958.

[50] SCARF, H. Bayes Solutions of the Statistical Inventory Problem. *Ann. Math. Stat.*, 1959, **30** (2), 490–508.

[51] SCARF, H. Some Remarks on Bayes Solutions to the Inventory Problem. *Naval Res. Logist. Quart.*, 1960, **7** (4), 591–96.

[52] SCARF, H. The Optimality of (S, s) Policies in the Dynamic Inventory Problem. Chap. 13 in Arrow, K. J., S. Karlin, and P. Suppes (eds.), *Mathematical Methods in the Social Sciences, 1959*, Stanford, Calif.: Stanford Univ. Press, 1960.

[53] SCHLAIFER, R. *Probability and Statistics for Business Decisions*. New York: McGraw-Hill, 1959.

[54] TAKÁCS, L. On Secondary Stochastic Processes Generated by Recurrent Processes. *Acta Math. Acad. Sci. Hung.*, 1955, **6**, 363–80.

[55] WAGNER, H. M. A Postscript to "Dynamic Problems in the Theory of the Firm." *Naval Res. Logist. Quart.*, 1960, **7** (1), 7–12.

[56] WAGNER, H. M. On the Optimality of Pure Strategies. *Management Sci.*, 1960, **6**, 268–69.

[57] WAGNER, H. M. *Statistical Management of Inventory Systems*. New York: Wiley, 1962.

[58] WAGNER, H. M., and T. M. WHITIN. Dynamic Version of the Economic Lot Size Model. *Management Sci.*, 1958, **5**, 89–96.

[59] WHITIN, T. M. *The Theory of Inventory Management*. 2d ed., Princeton, N.J.: Princeton Univ. Press, 1957.