2

THE NATURE AND STRUCTURE OF INVENTORY PROBLEMS

KENNETH J. ARROW
SAMUEL KARLIN
HERBERT SCARF

1. Decision Problems

The last decade has seen a rapid spread of the use of scientific disciplines in the rationalization of management decisions, a movement symbolized by the terms "operations research" and "management sciences." There has been a steady interplay between the more careful formulation of business problems and the development of mathematical tools for their solution, such as linear programming, information theory, queueing theory, and various new concepts in statistics.

The problems of business management come under the general heading of *decision problems*. A decision problem has typically four parts: (1) a *model*, expressing a set of assumed empirical relations among a set of variables; (2) a subset of *decision variables*, whose values are to be chosen by the firm or other decision-making entity; (3) an *objective function* of the variables in the model, a function such that a higher value represents a more desirable state of affairs from the viewpoint of the firm; and (4) *computing methods* for analyzing the effects of alternative values of the decision variables on the objective function. Ideally, we would like to have computational methods which lead to an *optimal* solution—that is, to the determination of those values of the decision variables which maximize the objective function subject to the constraints implied by the model. Where no optimal solution can be found (as often happens in the present state of decision theory), we seek at least to determine the value of the objective function for any given set of values of the decision variables. We may refer to this sort of solution as a *descriptive* solution. We shall have occasion in this book to develop both optimal and descriptive solutions for different decision problems.

In the present work we are interested in descriptive solutions primarily as an aid to determining good policies. Such solutions may at least enable us to accept or reject certain values of the decision variables as leading to satisfactorily high or unsatisfactorily low values of the objective function, or to compare a limited number of alternative policies. In this way, descriptive solutions are a partial substitute for optimal solutions. Moreover, in other contexts the economist, interested in describing the economy as a whole rather than in giving advice to individual decision-makers, will be able to make use of the techniques developed here for descriptive solutions whenever he is in a position to make assumptions about the decisions of the entrepreneurs. Such studies have been made, usually in connection with business cycle analysis (see, for example, Abramovitz [1], chap. i and pp. 127–31; Whitin [16], chaps. ii and v), but usually only with respect to grossly oversimplified models.

The term "computing methods" is, of course, to be interpreted broadly as the mathematical specification of algorithms for arriving at a solution (optimal or descriptive), rather than in terms of precise programming for specific machines. Nevertheless, we want to stress that solutions which are not effectively computable are not properly solutions at all. Existence theorems and equations which must be satisfied by optimal solutions are useful tools toward arriving at effective solutions, but the two must not be confused. Even iterative methods which lead in principle to a solution cannot be regarded as acceptable if they involve computations beyond the possibilities of present-day computing machines.

The acceptability of a solution is thus relative to computing technology. With modern developments, straightforward simulation of intricate situations with little or no loss in essential features is sometimes possible, so that descriptive solutions are possible with little abstraction or preliminary mathematical analysis. Some use has been made of these procedures in the study of inventory problems. A procedure of this sort involves a priori selection of what is considered to be a reasonable class of inventory policies. Each possible policy is then tested in the computer's simulation of the model, and the appropriate policy is selected according to the objective function.

In most practical problems of any complexity, however, if the number of reasonable strategies is at all large, the machine time of simulation is apt to be prohibitively costly, if indeed it is at all possible. It is at this point that theoretical discussions of decision problems become essential. In a descriptive solution it may be possible, given initial assumptions about the elements of the model, to give a quantitative description of its workings by theoretical analysis. The role of the computing machine is then reduced to the evaluation of formulas, which will usually be far less costly than direct simulation. There will still remain the problem of choosing among alternative policies by carrying out the computations for each. But in some cases we can actually arrive at optimal, not merely descriptive, solutions, and this is the best situation of all.

Immediate application to practical problems is not the sole justification of theoretical work. We do not expect that the results of this book will always or even frequently be applicable directly to practical problems. We do feel that they represent progress in understanding the phenomena involved. The situation is analogous to that in many branches of applied mathematics, in which theoretical constructions furnish the conceptual framework for practical decisions rather than a complete procedure for decision-making in all cases. In many cases, the model is not a complete representation of the true situation in all its complexity, but a simplified version amenable to analysis. The solution to the model will not be strictly applicable in all its quantitative features; but it will call attention to the chief qualitative features of the appropriate policy, the form it will take, and the directions in which it can be expected to vary with changes in the underlying parameters. This qualitative information can be very useful in improving the intuition of the decision-maker, helping him to organize his data and his direct knowledge, and reducing the field of alternative policies to manageable proportions.

2. Inventory and Production Problems

We shall not attempt to give a precise definition of the area to be included under the heading of inventory and production problems. Only through a study of typical models can one acquire a sense of the scope of the subject. An inventory is a stock of goods which is held or stored for the purpose of future sale or production. Some of our problems deal only with inventories, others with situations in which the possibility of holding inventories has repercussions on production policy by loosening the relation between production and sales. In some cases studied, it would be rather difficult to identify any component of the problem as being an inventory in any strict sense. But there is a similarity of structure in all these problems in that there is an element of interaction between the present and future policies; decisions made today affect the limits of and profitability of future decisions.

An inventory problem might, for example, involve deciding how much typing paper to stock each month for an office, or how many spare parts to keep on hand for a given machine. When production is involved, the inventory problem might require determining how much wheat to plant per year or how much gasoline of a certain variety to have blended. How much water to release from a dam for electricity and irrigation purposes is an inventory problem; how many workers to hire for a given labor force is another. Inventory problems may involve scheduling, production, determining efficient distribution of commodities in certain markets, finding proper replacement policies for old equipment, determining proper prices for goods produced, or combinations of these elements.

3. Cost and Revenue Considerations

Basically, inventories constitute an alternative to production in the future. To have available one unit of product tomorrow, we may either produce (or purchase) it then or produce it today and store it until tomorrow. The choice between the two procedures depends upon their relative profitability. Holding inventories usually involves storage costs; it also means that capital is tied up which might be invested elsewhere. On the other hand, it may be that production tomorrow is more expensive than production today for any one of several reasons. If unit production costs increase with the scale of operations and if more is being planned for tomorrow than for today, it will be cheaper to produce an additional unit today. If, on the other hand, unit production costs decrease with the scale of operation, it may be profitable to produce a great deal today to derive the benefits of the economies of scale and store the product for tomorrow. Changing the rate of production may be so costly that it pays to produce more today than is needed today and the opposite tomorrow, holding inventories to bridge the gap. It may even be impossible to produce tomorrow. In general, a failure to meet demands will result in some loss by the firm. We will now examine these various cost elements in somewhat more detail.

The Cost of Ordering or Producing. In stocking any commodity, there will be a cost $c(z)$ to ordering or producing a given amount z of the commodity. Various assumptions about the cost function $c(z)$ appear reasonable in different circumstances. The simplest is that the cost of ordering is directly proportional to the amount ordered. Another common situation is that $c(z)$ is a concave function of z, which means that each additional unit costs less. A special case of this would obtain when $c(z)$ is composed of a cost proportional to the amount ordered plus a set-up cost which is constant for z positive and zero for $z = 0$. In production, the set-up costs can be due to the preliminary labor and other expenses of starting a production run (e.g., setting a lathe); in ordering, the set-up term is due to the administrative expenses of processing the order. Concave cost functions also arise whenever there are economies to large-scale production, usually by making profitable the use of more expensive machinery. If the firm is ordering rather than producing, any economies of scale in the firm from which the ordering takes place will be reflected in quantity discounts.

On the other hand, $c(z)$ may be a convex function of z when additional output requires hiring additional workers and purchasing additional equipment without increasing the size of the plant as a whole, so that production becomes less efficient. In a specific problem, a realistic cost function may be made of pieces, each of which is either concave, convex, or linear.

Storage Costs. A second set of costs is associated with the stock of inventories on hand (the cumulated excess of supply over demand). Storage or handling costs may be incurred by the actual maintenance of stocks or the rent of storage space ; or, in a more generalized form, they may be a measure of obsolescence or spoilage.

For example, a firm's fleet of trucks gives rise to storage costs in the forms of maintenance and of replacement because of aging or obsolescence. Another firm produces electronic equipment that becomes defective if stored without use. In this case the value of a certain fraction of the current inventory could be considered as a storage cost ; an alternative procedure would be to consider the cost of repairing a defective item as a storage cost.

It will usually be assumed here that storage costs are proportional to the size of the stock of inventory, and for many purposes this approximation is adequate. However, other cases should be considered. For example, if storage takes place in a warehouse, the unit storage cost may jump from zero (when the amount stored does not exceed the capacity of the warehouse) to a large positive number (representing the cost of hiring space elsewhere).

Discount Rate. If a firm produces or orders and then stores, it is laying out money which will not return to it for a while. It could, instead, have invested this money elsewhere, say in government bonds which are absolutely secure and yield a return of 4 per cent. The firm thus has the alternative of receiving \$1.04 a year hence for each dollar invested today. Therefore, it should regard any return of \$1.04 a year hence in its production and sales operations as equivalent to \$1.00 today ; or, put alternatively, a dollar profit a year hence is to be evaluated as equal to a dollars today, where $a = 1/1.04$. The quantity a is referred to as the *discount rate.* Thus in arriving at the net benefits of an inventory policy, the net receipts of the current period should be added to a times the net receipts a year hence. The discount rate for dollars two years hence would be a^2, since if the bond investment alternative were adopted, the entire \$1.04 available one year hence could be reinvested. Similar considerations apply for further periods.

The period of one year is, of course, chosen arbitrarily for illustration. However, it is clear that for short periods, a may become very close to 1 ; hence, in an inventory problem where the relevant time horizon is very short, discounting is unimportant, and for practical purposes we may let $a = 1$. The discount rate is of most consequence when the planning period of the firm extends into the far future. The discounted returns of the very far future become negligible.

We have introduced the concept of discounting with regard to the alternative of investment in government bonds, but other alternative investments may be available to the firm. For example, a firm with several types of processes may have within it an alternative investment which has a higher rate of return than that on government bonds; in that case, the discount rate should be computed from the best alternative investment available.

Penalty Costs. The storage cost arises because supply (including both current output and accumulated stocks from the past) exceeds demand; the penalty cost arises when demand exceeds supply. It may be impossible (or possible, but too costly) to guarantee that demand will be met under all circumstances, especially when future demands are uncertain. The failure to meet demands generates costs, though in different ways under different circumstances. A simple example of penalty costs occurs in stocking spare parts for a given machine. When the parts are not available, the machine becomes inoperative and its output is lost to the firm.

Although the form that penalty costs take is related to the structure of the model, which will be discussed below, two extreme cases may be visualized here. One possibility is that if a demand occurs beyond the available inventory, it will be met by a priority shipment. In this case, the penalty cost would be measured by the difference between the cost of priority shipment and the cost of routine delivery or production. Another extreme case is that any demand which cannot be satisfied out of stock is backlogged and satisfied when the commodity becomes available. In this case, the penalty cost would be the loss of the customer's goodwill and his possible future unwillingness to do business with the firm. Such a penalty cost is real but may be very hard to measure in any precise sense.

As in the case of the cost function for ordering or production, a variety of assumptions about the shape of the penalty cost function are possible in varying circumstances. The simplest, again, is that penalty costs are proportional to the amount of shortage. In industrial applications, it may be that the penalty cost should be a convex function—i.e., small shortages are of little consequence but larger ones create more than proportionately great difficulties for the customers. In military applications, it is sometimes held that a given amount of material must be available for success in a mission and that any shortage at all will cause failure; in that case, the penalty cost function is discontinuous, being a positive constant for positive shortages and zero for a zero shortage.

Revenues. In the models studied in this book and indeed in most of those in the literature, it will be assumed that both the price and the demand for the product are independent of the firm's control. Under this assumption, the stream of revenues from sales is the same regardless of the inventory policy of the firm, and hence can be disregarded in further analysis except for the situation in which the firm is unable to meet a demand. This situation can be handled by including in the penalty cost the price which would have been obtained by sales.

Other assumptions about price and demand are possible, as we have seen in Chapter 1, Section 3. The cases studied here are of the constant-price variety, but one may also consider the perfectly competitive and imperfectly competitive cases : in the first, prices are given to the firm (though possibly changing over time), but sales are completely within its control ; in the second, both price and sales are decision variables, but sales are a given function of price.

Cost of Changes in the Rate of Production. In some cases, a change in the rate of production leads to costs which are different from those associated with continued production at the new rate. They are associated with the derivative of the rate of production, not with the rate itself. Partly they are what is referred to in economics as "the excess of short run over long run costs." Thus an increase in production may require a rapid increase in some variable factors for a period in which other factors, such as equipment, cannot be increased, and therefore a temporary increase in total costs will occur which will be reduced as relatively immobile factors become adapted to the new production level. There are other elements in the cost of changing the rate of production : the hiring of inexperienced personnel, the need for learning new organizational methods appropriate to a higher production rate, the breaking-in of new equipment. There may even be costs to reducing the rate of production, such as those involved in separation of personnel (intensified by guaranteed wage plans) or in making special provisions for the care of inactive equipment. Ordinarily, the costs involved in reducing the rate of production would be expected to be much smaller than those involved in increasing it.

Salvage Costs. Suppose we analyze an inventory problem only with respect to a single time period. Then at the end of the period we may have stocks left over. These stocks have a value ; at the very least they can be sold for some price. A salvage value must therefore be included in analyzing the problem ; the negative of the salvage value is referred to as a salvage cost. The salvage cost may be positive if there is a cost to disposal of the surplus. For most purposes, it is more appropriate to assign the inventories a salvage value which is equal to their value to the firm itself in its future operations. In this sense, the salvage cost is a fictitious cost which would not appear in the problem if an analysis were made of all future time periods.

4. Demand

Inventories are, of course, held for the ultimate purpose of satisfying demands. A number of alternative assumptions are possible and appropriate in different circumstances.

As we have already noted, virtually all work in inventory theory assumes that the demand is independent of the firm's control, although other assumptions are certainly reasonable in some conditions. Even with this restriction, we may still have the choice of assuming that demand can be perfectly forecast or that it cannot, and that demand conditions do or do not remain stable in future periods.

The case where the demand for the commodity in subsequent time periods is regarded as known may be referred to as the *deterministic* case. In most of the interesting analyses in this field, the demand is not assumed constant over time. We thus have a sequence of future time periods for each of which the demand is assumed known, though it will in general vary from period to period. The deterministic case is studied in Part II of this book.

If the demand is not assumed to be known ahead of time, it is very convenient and sometimes justifiable to assume that the demand in each future period is a random variable with a known probability distribution. Although in principle this distribution may be known to change from time to time (e.g., when there are seasonal fluctuations or a long-term trend), in practice we are confined to assuming the distribution of demand to be the same in each future time period. A variation of this assumption is that the size of each demand is fixed but that the times at which the successive demands come are random variables. Under this assumption demand is considered as a continuous time stochastic process, and this is frequently more appropriate when it is possible to place orders at any moment of time. In most examples in which this concept is applied, the assumption is made that the times between successive demands are independent identically distributed random variables.

Even under these restrictive assumptions, the determination of optimal policies remains difficult. Because analysis of the stochastic case has been confined to the situation of identically distributed demands, the deterministic case is sometimes a better approximation when demand conditions are in fact changing over time.

Even a fuller study of the stochastic case would not exhaust the realistic possibilities. It is not always reasonable to suppose that the probability distribution of possible demands is known to the firm. One possibility is to apply criteria for decision-making under uncertainty which do not presuppose knowledge of a probability distribution—e.g., the min-max rule (see Chapter 12). In an inventory problem which is extended over time, there is the possibility of estimating the probability distribution of demand from the successive observations and using the estimates to arrive at successively improved inventory problems. The combination of estimation and decision-making in an inventory problem is a complicated sequential decision problem which has not yet been explored.

5. Deliveries

Another important element in the mechanism of the inventory process is the lag in delivery of the commodity after an order is placed or a decision is made to produce. If demand is assumed to be known with certainty, this lag is of no consequence; all that is required is that the orders be placed correspondingly earlier. In the case of uncertain demand, however, the assumptions about delivery are important, since the information about the amount needed for ordering changes with time.

The existence of lags in delivery is an essential element of inventory holding in so far as it serves as protection against uncertainty, for if deliveries could be made instantaneously without extra cost, the firm could place its orders after knowing what the demand is rather than before and thus avoid all possibility of penalty costs.

Delivery lags may enter the models in several different ways. In some cases it is appropriate to assume a fixed lag between order and delivery. In others the lag is a random variable with a known distribution. Finally, one may permit more generality by admitting two kinds of shipments: one a routine shipment with a lag, the other a priority shipment without lag but at a higher price.

6. The General Structure of Inventory Models

Discrete vs. Continuous Time. In some models, we assume that all orders and deliveries take place on a succession of equally spaced time points. In some cases this can be regarded as reasonable from the point of view of normal business practices. In other cases it must be regarded as a simplification for analytic convenience, notably in the case of random demands, where the assumption is made that the demands in successive periods are independent random variables; this assumption is very unlikely to be valid when the periods are sufficiently short. The transition to continuous time in such models will require formulation of demands as a continuous time stochastic process. We believe that such a reformulation will turn out to be not merely more realistic but also ultimately simpler to handle analytically, but it has not yet been accomplished.

There are two situations in which time has been treated as a continuous variable. In one, it is assumed that demands occur as a continuous or at least piecewise continuous function of time, and that the corresponding amount ordered or produced is a similar function. In view of the preceding remarks, it is not surprising to find that such models so far assume deterministic demands (see the studies in Part II). In the past literature, similar models have been treated by assuming time as discrete, but as might be expected in the light of experience in other branches of applied mathematics, the continuous approximation leads to notable simplification. There is indeed the initial difficulty that somewhat deeper mathematics is needed (calculus of variations instead of ordinary calculus, in effect); but the simplification of the resulting algorithms and proofs, and in particular the ease of obtaining qualitative conclusions, seems to us an overwhelming argument.

The second situation is that in which the demands come discontinuously at time points which are not equally spaced but occur at random. In such cases of discontinuous demand, it is frequently natural to assume that orders can be placed at such points. Then time is a continuously varying parameter, but ordering and demand are discontinuous functions of it.

Stock-Flow Identities and Nonnegativity Conditions. If the commodity dealt with is not in any way perishable, there is the obvious identity that the stock on hand at the end of any period equals the stock on hand at the beginning of the period plus the amount delivered to the firm less the amount sold by it. Let x_t be the initial stock level in a period, y_t the stock level after ordering, z_t the amount ordered, and r_t the amount sold. The stock-flow identity can then be written as

$$z_t = y_t - x_t \, , \tag{1}$$

$$x_{t+1} = y_t - r_t \, . \tag{2}$$

Equations (1) and (2) hold in each time period. These identities are valid if time is taken as a discrete variable, or if time is taken as a continuous variable but with discontinuous demands and orders. If time is taken as continuous, and demand and ordering or production as continuous or piecewise continuous functions of time, then the distinction between x_t and y_t disappears, and (1) and (2) are replaced by

$$\frac{dy}{dt} = z(t) - r(t) \, , \tag{3}$$

or, in integral form,

$$y(t) = y(0) + \int_0^t [z(\tau) - r(\tau)] \, d\tau \, . \tag{4}$$

The above identities raise the important question of the relation between sales r_t and demand ξ_t. The two cannot be identical in all circumstances because the demand may be such as to make the inventory x_{t+1} negative. Since this is physically impossible, some assumptions must be made and several alternatives are reasonable under different circumstances.

(a) It is almost invariably assumed that demand will be met if physically possible—that is, that in the case of discrete time or discontinuous demands,

$$(5) \qquad r_t = \xi_t \quad \text{if} \quad \xi_t \leqq y_t \; .$$

Although one can conceive of cases where it would pay to leave a demand unsatisfied and incur the corresponding penalty, usually it is reasonable to suppose that this is not the case.

(b) One possible assumption is to insist on policies for which x_{t+1} is never negative—that is, that demands are always met so that

$$(6) \qquad y_t \geqq \xi_t \quad \text{for all } t \; .$$

Such an assumption is usually made when demands are deterministic (see Chapters 4–6), though it is not logically necessary ; a firm may find it more profitable not to meet the demand (an example is given in Chapter 7). When demand is uncertain, it may be very costly or even impossible to require that (6) hold for all possible values of the demand.

(c) When we admit the possibility that demand can exceed the inventory on hand, there will be a shortage $\xi_t - y_t$. Whenever this quantity is positive, there will be a penalty p, which we take to be a function $p(\xi_t - y_t)$ of the shortage. There are several possible assumptions about the firm's behavior in dealing with the shortage. One is to order the necessary goods for immediate delivery, the firm paying a premium over the usual price for this service. In this case, the customers are satisfied, and the penalty is interpreted as the premium cost of immediate delivery. In that case, (2) is modified to

$$(7) \qquad x_{t+1} = \max\,(0,\, y_t - \xi_t) \; ,$$

and is necessarily nonnegative. It then follows that y_t is nonnegative since $z_t \geqq 0$.

(d) A second possible situation when a shortage occurs is that the unsatisfied demand is never met. Thus the customer may go elsewhere for his goods or he may have an immediate need which cannot be postponed. The firm again has a penalty, which in this case reflects the loss of customer goodwill rather than the cost of priority shipment. Also (7) holds in this case. Thus the two cases are mathematically identical, though the interpretation of the penalty is different. We shall refer to both of these as the *non-backlog* case, as opposed to the case to be discussed.

(e) A frequent policy in case of shortage is to leave an unfilled order on the books and satisfy it as soon as possible. Again we must assume that there is some penalty reflected in price, or at least in customer goodwill, for failure to satisfy the order immediately. In this case, however, a negative inventory has meaning, since it represents the cumulated total of unfilled orders. Such a negative inventory is usually referred to as a *backlog*. The penalty function has the same role as in the non-backlog case, but the stock-flow relation (7) is replaced by

(8) $$x_{t+1} = y_t - \xi_t \, .$$

Static Models. Let us now turn to the time structure of inventory models. The simplest case has but one time period. A production or ordering decision is made at the beginning of the period, so that the inventory y is determined; then the demand occurs and the revenues and costs are determined by the demand and the inventory. Much of the earlier work in inventory theory has dealt with such models, and some are studied here. Static models form a reasonable approximation when the item is perishable, or when it obsolesces rapidly because of technological change or fluctuations in styles (e.g., women's dresses). In such cases, whatever is left over at the end will not be used further as an inventory. Nothing at all may be left; if something is left, there is a salvage cost of eliminating it, which may be negative if it can be disposed of for some other use.

Of the costs listed in Section 3, the discount rate will not be relevant in the present case. We will use the subscript t to designate the time period considered even though we are considering a static model, in order to make use of the formulation in dynamic extensions. The profit π_t will be the revenue less the costs. The revenue equals price, r, times sales, which in accordance with the previous section equals $\xi_t - \max(0, \xi_t - y_t)$. The ordering cost is $c(z_t)$, the penalty cost is $p(\xi_t - y_t)$ and is zero unless the argument is positive, the storage cost is assumed to depend upon the inventory y_t and is designated as $h(y_t)$, the salvage cost is $v(y_t - \xi_t)$, and the cost associated with changing the rate of production is $G(z_t - z_{t-1})$. Then we have

(9) $$\pi_t = r[\xi_t - \max(0, \xi_t - y_t)] - c(z_t) - p(\xi_t - y_t) - h(y_t) - v(y_t - \xi_t) - G(z_t - z_{t-1}).$$

If the demand ξ_t is known, then the problem is to choose z_t so as to maximize (9); of course z_{t-1} is given, and y_t is determined from (1). If the demand is considered to be a random variable, then z_t is chosen to maximize the expected value of (9).

Since in any case ξ_t is independent of the firm's control, the term $r\xi_t$ can be ignored when comparing policies. The term $r \max(0, \xi_t - y_t)$ can then be absorbed in the penalty function. The problem in the static model can then be restated as that of minimizing the loss,

(10) $$\mathscr{L}_t(z_t|x_t) = c(z_t) + p(\xi_t - y_t) + h(y_t) + v(y_t - \xi_t) + G(z_t - z_{t-1}) \, ,$$

or, if ξ_t is a random variable, to minimize the expected value of $\mathscr{L}_t$. For a particular static situation z_{t-1} is given, and the last term can be absorbed into the first.

In many, perhaps most, cases, the value of the inventory left over is precisely that it can be used in the future as a new initial inventory. This is not properly accounted for in the static model. However, the static model is an approximation to the dynamic in two senses: First, there exists a suitable salvage cost for the inventory which would correctly represent its value for future use; and whereas the exact determination of this value would be equivalent to solving the full dynamic problem, in many cases an approximation based on judgment and past experience would provide an acceptable basis for determining the optimal value of z_t in the current period. Second, because of this property there is some similarity between the mathematical form of the solutions to the static and the dynamic problems, and therefore experience in solving the former is of great help in suggesting the approach to the latter.

Dynamic Models. If we consider more than one period, we have a loss such as (10) for each future period. As we have seen earlier, the future losses have to be discounted to make them comparable to present losses. Thus in a two-period model, the discounted stream of losses becomes

$$\lambda(z_1, z_2|x_1) = \mathscr{L}_1(z_1|x_1) + \alpha\mathscr{L}_2(z_2|x_2) , \tag{11}$$

where we put the dependence on the initial stock of inventories, x_1, into evidence. Notice that $\mathscr{L}_2$ is dependent on the decision z_1 and the sales ξ_1 in the first period since it depends on x_2, which in turn is determined by z_1 and ξ_1 through relation (7) or (8). If demand in the two periods is deterministic, we wish to choose z_1 and z_2 to minimize (11) subject to the constraints (1) and (7) or (8).

When demands are random, the minimization takes a more complicated form, one which is of the greatest importance. It is not necessary or advisable to choose a definite value for z_2 at time 1. At time 2, more information will be available, in that a specific value of the random variable ξ_1 will have been observed and the stock level at time 2 will be known. What is required is that whatever decision is made at any time be a function of all the information available at that time. Hence we need to choose a policy which prescribes how we shall choose z_2 under all possible alternative values of the observation ξ_1. The minimization problem is to choose a *number* z_1 and a *function* $z_2(\xi_1)$ so as to minimize (11) subject to the restraints. The idea that we choose not a definite set of actions for the future but a strategy specifying how to act under all possible contingencies is of fundamental importance in all problems involving behavior over time where uncertainty is present.

Before generalizing the dynamic models to more time periods, let us observe that the salvage cost in (10) should appear only in the last time period; for other time periods it is assumed that any leftover inventories appear as the initial inventories x_{t+1} of the next period. Only at the end of the entire horizon considered are they ready for scrap. From now on, then, we assume

$$\mathscr{L}_t(z_t|x_t) = c(z_t) + h(y_t) + p(\xi_t - y_t) + G(z_t - z_{t-1}) . \tag{12}$$

To take account of salvage cost in the last period, we rewrite (11):

(13) $$\lambda(z_1, z_2|x_1) = \mathscr{L}_1(z_1|x_1) + \alpha\mathscr{L}_2(z_2|x_2) + V(y_2 - \xi_2) ,$$

where V is salvage cost discounted to the initial period. If we go beyond time period 2, the loss in time period 3 has a discount factor α^2, and so forth. Then for T time periods, (13) is generalized to

(14) $$\lambda(z|x_1) = \sum_{t=1}^{T} \alpha^{t-1} \mathscr{L}_t(z_t|x_t) + V(y_T - \xi_T) .$$

There is really no logical horizon, though for practical purposes we may act as if there were. In principle, there is no difficulty in contemplating an infinite sequence of time periods, in which case there would be no salvage term, and the total discounted loss is given by

(15) $$\lambda(z|x_1) = \sum_{t=1}^{\infty} \alpha^{t-1} \mathscr{L}_t(z_t|x_t) .$$

The symbol z as an argument of the function λ stands for the whole finite or infinite sequence of z_t's in (14) or (15). If demand is deterministic, then the problem is to choose the z_t's so as to minimize (14) or (15). If demand is taken as a random variable, then for $t > 1$, z_t is chosen, not as a number but as a function of all the random variables which have been observed before time t. Such a specification of functions constitutes an *inventory policy*, and it is desired to choose an inventory policy which minimizes the expected value of (15).

Similar remarks apply if the demands are considered as arriving continuously. However, a precise formulation of this case does not exist in the literature except in the circumstance where the demand process is deterministic. In this case, as indicated above, it is customary to add the constraint that inventories must meet demands, so that $\xi_t \leqq y_t$, and the penalty term in (12) disappears. As noted in Chapter 1, Section 4, such problems have been studied in the earlier literature, with G identically equal to zero; and a definitive solution for this case, with the additional assumption of increasing marginal costs, has been given in the important paper of Modigliani and Hohn [14]. The cost of changing the rate of production was introduced as a consideration by Hoffman and Jacobs [10], and solutions under special circumstances given by them and others.

When we consider time as a continuous variable, the infinite sum in (14) is replaced by an integral. The finite difference as the argument of G then becomes a derivative, and we have

(16) $$\lambda(z|x) = \int_0^T \left\{c[z(t)] + h[y(t)] + G\left(\frac{dz}{dt}\right)\right\} dt .$$

In Chapters 4–6 we present algorithms for minimizing (16) for various special cases of the functions entering into it. In the stochastic case, where the demands in successive periods are considered as independent random variables, the function G has been ignored in every study we are aware of, including those in this volume. This model was first formulated and studied by Arrow, Harris, and Marschak [3], though, as we have seen in Chapter 1, Massé had earlier formulated a special case along similar lines. Arrow, Harris, and Marschak did not, however, consider all possible inventory policies but restricted attention to the class of so-called two-bin or (s, S) type policies.

7. Types of Analysis

Optimization. The highest hope for analysis of a model is to find a constructive method for determining the optimal strategy. As we have explained, there are two variants of this aim: one is to find a procedure which will provide a practicable computing procedure for determining the optimal strategy, given all the parameters of the problem; the other is to characterize the solution sufficiently sharply so that its qualitative characteristics can be studied closely. These two aims are distinct and may not be realized in the same way. The second aim is important if we accept a model, as we frequently do, not as literally true but as true enough to suggest the nature of the solution.

We are usually interested in finding the solution as a function. Even in a static model, a theoretical analysis usually seeks to determine the ordering or production, z_t, as a function of the initial stock x_t, even though it is a given magnitude in any concrete situation. In dynamic models, as we have seen, z_t, for each t greater than 1, is a function of the random variables observed up to time t. Let us now examine the optimization in a dynamic model a little more closely by considering the case $T = 2$, as specified in equation (11).

For any given z_1 and ξ_1, x_2 is determined. Then, to minimize λ, we have to choose the remaining variable, z_2, which only occurs in the last term. Hence, we choose z_2 to minimize the last term, taking x_2 as given (if ξ_2 is a random variable, we minimize the expected value of the last term, but again x_2 is taken as given, since it is known at the time the choice of z_2 is to be made). We thus have a function $z_2(x_2)$. If we substitute the function for z_2 in (11), we have left a single decision variable z_1, which now enters not only the first term but the second, through x_2. We then have to minimize with respect to z_1.

Let us apply the same approach to the case of an infinite number of time periods. In this case, because we are assuming that the various losses have the same functional form, a very interesting simplification emerges. At time 2, the decision-maker is again facing an infinite future of the same structure as that faced at time 1; the discount factor applied to the losses in period t is α^{t-2}, which differs only by a constant factor from the discount factor applied in period 1. The only difference is the magnitude of the initial stock, which is now x_2 instead of x_1.

Equation (15) may be written

$$\lambda(z|x_1) = \mathscr{L}_1'(z_1|x_1) + a\lambda(\bar{z}|x_2) , \tag{17}$$

where $\bar{z}$ stands for the infinite sequence of z_t's beginning with z_2. Consider the choice of z_2 at time 2; the range of alternatives and the losses are the same as at time 1 except in so far as the initial stock of inventory is different, and hence the value of z_2 chosen would be the same as the value of z_1 if x_1 were to equal x_2—that is, the optimal policy must have the form of requiring that z_2 be the same function of x_2 that z_1 is of x_1. The same argument extends to all future times. The optimal policy then takes the form of specifying a function φ and then requiring that $z_t = \varphi(x_t)$ for all t. Notice that x_t depends, among other things, on all the demands $\xi_1, \ldots, \xi_{t-1}$ in periods preceding time t; hence, implicitly, z_t is a function of these demands, as we observed was necessary in the stochastic case (see p. 29 above).

By an obvious change in notation, let $\lambda(\varphi|x)$ be the loss if the optimal policy is defined as just shown by a function φ and the initial inventory is x. For such a policy, (17) can be written

$$\lambda(\varphi|x_1) = \mathscr{L}_1'(z_1|x_1) + a\lambda(\varphi|x_2) .$$

Since for an optimal policy z_1 must be the optimal choice, we can write

$$\lambda(\varphi|x_1) = \min_{z_1} [\mathscr{L}_1'(z_1|x_1) + a\lambda(\varphi|x_2)] . \tag{18}$$

Equation (18) is a functional equation which must be satisfied by the function $\varphi(x)$ which defines the optimal inventory policy.

It is important to realize that an assumption has been tacitly made which permits us to summarize all of the relevant information at the start of a period in the knowledge of the stock level alone. This assumption is that there is no time lag in delivery; if this were not the case, the size and dates of all orders which have been placed but not yet delivered would be included in the summary of relevant information. The introduction of time lags makes it impossible to reduce the optimal policies, or for that matter any effective policy, to the simple form $z_t = \varphi(x_t)$ as described above. For example, if there is a time lag of λ periods of time, then the relevant information at the beginning of the nth time period includes not only the stock level at the beginning of this period, but also a specification of all orders placed during the preceding $(\lambda_1 - 1)$ time periods. Stockage policies are therefore sequences of functions of λ variables rather than merely one variable. Of course, if the inventory problem is symmetric in the sense that for each period the decision-maker faces the same problem, the sequence of functions may be collapsed into one function of λ variables. It is occasionally possible to make additional simplifications if more restrictive assumptions on the model are made (see Chapter 10).

It is sometimes of interest to consider the time lag as not fixed but random. If the values that the random time lag may assume with positive probability are sufficiently large, it may be essential to include all past orders along with the present stock level in the specification of information.

A functional equation frequently occurs as an analytical expression of a recursive situation. It is the symmetries and the recursive character of the dynamic inventory problem that permit the derivation of relationship (18). Recursion relations are traditionally used in examining a process over time ; they are at the heart of many analytical studies of stochastic processes, as well as the basis of the theory of differential equations. It is therefore not surprising that relationship (18) plays a fundamental role in the analysis of dynamic inventory problems.

As we have seen in Chapter 1, the approach defined in (18) was applied in special cases by Shaw and by Hart in the 1930's, and more systematically by Massé, and the reasoning behind it was used by Wald in the development of sequential analysis. (Wald's own expositions made such specific use of the statistical problems concerned that the use of the functional equation is somewhat obscured ; it appears more clearly in the exposition and development of Wald's work in Arrow, Blackwell, and Girshick [2].) The functional-equation approach was applied to the inventory problem in the form used here by Arrow, Harris, and Marschak [3], who, however, restricted attention to functions φ of a specified form. A general formulation was given by Dvoretzky, Kiefer, and Wolfowitz [8], who proved the existence of a solution $\lambda(x)$ to (18). The proof consists of choosing an initial $\lambda_0(x)$ and defining $\lambda_n(x)$ recursively from $\lambda_{n-1}(x)$ by (18). Because $a < 1$ the functions $\lambda_n(x)$ converge exponentially fast to a solution $\lambda(x)$ in accordance with the classical methods of successive approximation

Thus the functional-equation approach can be used to establish the existence of solutions to (18) and to provide a method for computing the solution in principle. The existence theorem can, however, be stated and demonstrated more simply and directly ; see Karlin [11]. The possibility of recursive computing methods has been examined by Bellman [4] and applied to many different fields. However, the practical range of these computing methods is limited to relatively simple problems. What is perhaps even a more serious limitation is that they usually cannot be used to determine the qualitative features of the solution ; they only provide a numerical method for determining the function φ for a given set of cost assumptions.

The underlying principle of the functional-equation approach, that a policy which is optimal over the entire period considered is optimal starting from any point in the interval, is an important component of the solution to any concrete problem, but it is not a sufficient basis for analysis. What we would like are theorems which indicate the qualitative nature of the function φ, as determined by qualitative features of the cost functions. Such qualitative information not only is interesting in itself as a guide to decision-making, but also, by restricting the range of policies which needs to be examined, greatly improves the possibility of successful computing.

It is worth emphasizing that an explicit functional equation in itself is not the solution to the problem. By analogy, the claim that a physical problem is resolved if the differential equations which describe the phenomena are known is never acceptable unless the detailed qualitative description of the solution is ascertained, or efficient numerical approximations to the solution are available. The same applies with regard to the study of (18). In order to deduce the deeper qualitative properties of the optimal policy, as distinguished from establishing its existence, it seems necessary to bring to bear more powerful mathematical tools. More exact knowledge of the solution in most cases requires further assumptions about the nature of the cost functions and the form of the demand distribution. Determining the relationship of the form of the solution to the parameters of the model entails a formidable problem of classification, one of the most challenging problems of inventory theory.

The theory of maximization, whether in the form of ordinary calculus or in that of the calculus of variations, plays an important role. Many new tools are needed and some are presented in the following chapters. In some cases, where suitable convexity properties of the functions involved can be assumed, the tools of nonlinear programming are of great assistance in qualitative analysis (see Koopmans [12] and Chapter 7 below). If all cost functions are linear, then the problem may formally be regarded as one in linear programming. A straightforward reduction to a standard simplex procedure is almost always computationally impractical, but in some cases a simplex method which exploits the special properties of the problem can be used (see Dantzig [7], Bowman [6], Manne [13]). But especially in problems involving stochastic demands, none of the standard methods are applicable, and new ones must be developed; see the studies in Part III. Here attention is concentrated on finding conditions that will guarantee that the optimal policy has some simple form, such as the two-bin policy (see below).

Nonnegativity and other boundary conditions are used extensively in the following analyses. Roughly speaking, the optimal policies are frequently characterized by alternating phases, in some of which one or another boundary constraint is operative while in others the solution is an interior one. The interior segments can frequently be found to have a simple form, and the problem is reduced to determining when one phase ends and another begins.

One interesting approach which ignores these boundary conditions has been developed by Simon [15] for the stochastic case. He assumes that all cost functions are quadratic. The restrictions implied by the boundary conditions are supposed to be expressed in his model by rapidly increasing costs. Then it can be shown that the first step in the optimal policy, z_1, is the same as it would be if all random variables were in fact assumed to be known equal to their expected values. This permits a considerable simplification of the analysis; indeed, the solution becomes independent of the probability distribution of demands. We mention this approach as an interesting alternative to the direction emphasized here, but it is probably a valid approximation only in a limited range of cases.

Optimization with Restricted Ranges of Strategies. Because of the difficulty of choosing an optimal policy out of all possible functions $\varphi(x)$, we may use intuition part way by assuming that the functional form of φ is known except for a finite set of parameters. The functional forms are chosen as those types of policies used in practice and demonstrated to be optimal in cases where such analysis has proved possible. Perhaps the most frequently studied policy of this type is the two-bin or (s, S) policy, which is implemented as follows: Order or produce only if the present stock level falls below some given value s. When ordering is done, the stock is increased to a second value S. The study of Arrow, Harris, and Marschak referred to above confined its attention to choice among such policies—that is, to choosing the optimal values of s and S. The authors show that the total expected cost incurred from use of an (s, S) policy satisfies a renewal equation, which is solved. The actual underlying renewal process describing the flow of stocks is directly investigated in Chapter 15. The methods of Chapter 15 extend to the analysis of other kinds of inventory policies.

An ordering rule closely related to the (s, S) policy, based on a single critical level, proceeds as follows: Whenever any demand comes in, ordering is immediately effected to replenish the stock consumed, so that a constant stock level is always aimed at. The decision problem here is to choose the optimal level of stock. In terms of our previous notation, the (s, S) policy can be described by a function $\varphi(x)$ which is 0 if $x \geqq s$ and $S - x$ if $x < s$. The second policy is the special case where $s = S$; it may be useful to remark here that this policy seems best when there is no cost to ordering which is independent of the magnitude of the order (no set-up cost).

Bellman, Glicksberg, and Gross [5] determined the optimal policy for the case in which the ordering and penalty cost are both linear. Extensions of these results are contained in Chapter 9 of this volume. Dvoretzky, Kiefer, and Wolfowitz give some sufficient conditions for establishing that the optimal policy is an (s, S) policy for the single-stage inventory problem (they assume that the penalty cost is a fixed constant whenever suffered, and that the ordering cost is composed of a linear term plus a set-up cost).

Operating Characteristics of Inventory Policies. Instead of optimizing, we may turn the problem around and ask what the effects of a given inventory policy will be. A given policy, together with a given specification of random demands, determines a stochastic process involving inventories and production. We may then ask what the expected costs would be. If we do not want to specify the cost functions, we may find instead the probability distribution of the magnitudes which enter the cost functions, such as the stock of inventory, the amount of shortage, and the amount ordered over time. If we do specify the cost functions, we may be able to discuss choosing the policy parameters to minimize expected costs. Other criteria may be used for selecting the parameters of the policies once the performance characteristics of these policies are determined. For example, a reasonable ordering rule could be that policy minimizing the expected storage costs among all policies of a given class which guarantee not to exceed a prescribed maximum probability of run-outs (shortage) per period.

Processes generated by simple policies such as the two-bin policy or the process determined by random demand and delivery pose very interesting mathematical problems. The processes are related to many of those studied in recent years, such as the processes arising in queueing and counter problems.

Stationary Distributions. An important role in the theory of stochastic processes is played by the equilibrium or stationary distribution for a specified inventory policy. The distributions of the stock level, shortage, or amount ordered will usually tend to converge in time to some limit. These distributions are easier to study than the distributions corresponding to a finite section of the process, and it may be argued that in the long run they are more important. Again we can use these distributions to choose the parameters of the policy so as to minimize expected loss —in this case, long run expected loss.

These stationary costs are closely related to the solution of the dynamic model where the discount rate a is 1. The precise statement of this result is as follows. If $\lambda^{(a)}(x)$ (where a is the discount rate) represents the total discounted expected costs accrued when effecting the prescribed policy, then

$$\lim_{a\to 1} (1 - a)\lambda^{(a)}(x)$$

is exactly the long run expected loss. An equivalent way of computing this stationary cost is by averaging the losses for n periods with n growing arbitrarily large. It should also be observed in this case that long run costs are independent of the initial stock level.

Knowledge of the stationary distribution is fundamental, in that its determination involves no specific assumptions concerning the cost elements of the inventory model. The stationary distributions, whenever they exist, are a function only of the given policy and the nature of supply and demand. In this respect equilibrium distributions may be regarded as a description of the fluctuations in stock size resulting from the given policy in operation for a long period of time.

REFERENCES

[1] Moses Abramovitz, *Inventories and Business Cycles*. New York: National Bureau of Economic Research, 1950.

[2] Kenneth J. Arrow, David Blackwell, and M. A. Girshick, "Bayes and Minimax Solutions of Sequential Decision Problems," *Econometrica*, XVII (1949), 213-44.

[3] Kenneth J. Arrow, Theodore E. Harris, and Jacob Marschak, "Optimal Inventory Policy," *Econometrica*, XIX (1951), 250-72.

[4] Richard Bellman, *Dynamic Programming*. Princeton: Princeton University Press, 1957.

[5] Richard Bellman, Irving Glicksberg, and Oliver Gross, "On the Optimal Inventory Equation," *Management Science*, II (1955), 83-104.

[6] E. H. Bowman, "Production Scheduling by the Transportation Method of Linear Programming," *Operations Research*, IV (1956), 100-103.

[7] George B. Dantzig, "Upper Bounds, Secondary Constraints, and Block Triangularity in Linear Programming," *Econometrica*, XXV (1957), 124-83.

[8] Aryeh Dvoretzky, J. Kiefer, and Jacob Wolfowitz, "The Inventory Problem," *Econometrica*, XX (1952), 187-222, 450-66.

[9] Aryeh Dvoretzky, J. Kiefer, and Jacob Wolfowitz, "On the Optimal Character of the (s, S) Policy in Inventory Theory," *Econometrica*, XXI (1953), 586-96.

[10] A. J. Hoffman and Walter Jacobs, "Smooth Patterns of Production," *Management Science*, I (1954), 86-91.

[11] Samuel Karlin, "The Structure of Dynamic Programming Models," *Naval Research Logistics Quarterly*, II (1955), 285-94.

[12] Tjalling C. Koopmans, "Water Storage Policy in a Simplified Hydroelectric System," to appear in *Proceedings of the International Conference on Operational Research* (Bristol, England: John Wright & Sons, 1958).

[13] Alan S. Manne, "A Note on the Modigliani-Hohn Production Smoothing Model," *Management Science*, III (1957), 371-79.

[14] Franco Modigliani and Franz Hohn, "Production Planning over Time and the Nature of the Expectation and Planning Horizon," *Econometrica*, XXIII (1955), 46-66.

[15] Herbert A. Simon, "Dynamic Programming under Uncertainty with a Quadratic Criterion Function," *Econometrica*, XXIV (1956), 74-81.

[16] Thomson Whitin, *The Theory of Inventory Management*. Princeton: Princeton University Press, 1953.