MANAGEMENT SCIENCE
Vol. 50, No. 12 Supplement, December 2004, pp. 1791–1793
ISSN 0025-1909 | EISSN 1526-5501 | 04 | 5012S | 1791

informs®

DOI 10.1287/mnsc.1040.0293

Comments on "Optimal Policies for a Multi-Echelon Inventory Problem"

Herbert Scarf

I joined the RAND Corporation in June of 1954 after receiving my Ph.D. in Mathematics from Princeton University. One of my reasons for choosing RAND rather than a more conventional academic appointment was my desire to be involved in applied rather than abstract mathematics. I could not have selected a better location to achieve this particular goal. George Dantzig had arrived recently and was in the process of applying linear programming techniques to a growing body of basic problems. Richard Bellman was convinced that all optimization problems with a dynamic structure (and many others) could be formulated fruitfully, and solved, as dynamic programs. Ray Fulkerson and Lester Ford were working on network flow problems, a topic which became the springboard for the fertile field of combinatorial optimization. Dantzig and Fulkerson studied the Traveling Salesman problem and other early examples of what ultimately became known, under the guidance of Ralph Gomory, as integer programming.

In 1955 the organization was visited by a budgetary crisis and I was asked if I would mind taking up temporary residence in the Department of Logistics. The Logistics Department was a junior subgroup of the Department of Economics at RAND, with a much more prosaic mission than that of its senior colleagues. The members of the Logistics Department were concerned with scheduling, maintenance, repair and inventory management, and not with the deeper economic and strategic questions of the Cold War. Andy Clark was in the department, and we may have met each other casually but we certainly didn't do any work together, and we probably never had a serious conversation about research while we were both at RAND. Later in these notes I will describe our collaboration on the multiechelon inventory problem in the spring of 1959. Since that time, some 45 years ago, Andy and I have not seen each other at all. It has been a great pleasure to renew our acquaintance during the last several weeks. We have been in frequent contact discussing our joint work and correcting at least one seriously faulty memory of mine. Another member of the Logistics Department was Herb Karr who was actively involved in computing optimal inventory policies using the primitive computers that were available at that time. Herb figures in this failure of memory.

I moved into a simple office, far from my previous colleagues in mathematics, and sat there for a few weeks wondering what I was meant to do. I don't remember receiving any specific instruction or being presented with any particular research topic, but at some point I learned about the most elementary inventory problem: the decision about the quantity of a single nondurable item to purchase in the face of an uncertain demand. In the terminology of my first paper on inventory theory, "A Min-Max Solution of an Inventory Problem," the marginal cost of purchasing the item is a constant c. If y units are purchased and the demand is ξ, then the actual sales will be $\min[y, \xi]$ and if the unit sales price is r, profits will be given by the random amount

$$r\min[y, \xi] - cy.$$

The standard treatment of the problem was to assume a known probability distribution for demand, with the cumulative distribution given by $\Phi(\xi)$, so that expected profits are

$$r\int_0^\infty \min[y, \xi]\, d\Phi(\xi) - cx.$$

It is trivial to set to 0 the derivative of expected profit as a function of y and obtain the optimal quantity to be purchased as the solution of the equation

$$1 - \Phi(y) = c/r.$$

In my paper the probability distribution of demand is assumed not to be fully known. I study the decision problem in which the inventory manager selects the inventory level y that maximizes the minimum expected profit for *all* probability distribu-

Note: These comments are a condensation of my essay on Inventory Theory that appears in the 50th anniversary issue of *Management Science*.

Key words: multiechelon inventory problems; optimal inventory policies

tions of demand with a fixed mean μ and standard deviation σ. There is no particular reason for thinking in advance that this version of the inventory problem would have anything but a clumsy solution, but in fact the answer turned out to be surprisingly simple. If we define the function

$$f(a)=1/2\frac{1-2a}{\sqrt{a(1-a)}},$$

then the optimal policy is to stock

$$y=\begin{cases}\mu+\sigma f(c/r) & \text{if } (1+\sigma^2/\mu^2)<r/c\\ 0 & \text{if } (1+\sigma^2/\mu^2)>r/c\end{cases}.$$

The paper continues with a comparison between this inventory policy and those associated with the Normal and Poisson distributions.

It was my good fortune to meet Samuel Karlin and Kenneth Arrow at RAND. They were both interested in inventory problems, and they kindly invited me to spend the academic year 1956–1957 with them at Stanford University. My natural home at Stanford would have been a department of Operations Research, but such a department had not yet been established and I was formally located in the Department of Statistics.

Kenneth and Sam became good friends and mentors; both had an enormous impact on my professional career. We worked intensively on inventory problems during this year, and our efforts resulted in a monograph titled *Studies in the Mathematical Theory of Inventory and Production*, published in 1958. I will describe a particular paper that appeared in the monograph that bears on the subsequent work with Andy on the multiechelon inventory problem.

The paper "Inventory Models of the Arrow-Harris-Marschak Type with Time Lag" was coauthored with Samuel Karlin. It was concerned with an important feature of inventory problems: that an order, when placed, may not be immediately delivered. Two distinct treatments of the problem caused by time lags in delivery are examined. Imagine, for concreteness, a retailer whose stock is replenished by purchases from a wholesaler. When a customer arrives at the retailer with a request for the item, the currently available stock may be insufficient to meet this demand, even though there may be an adequate quantity previously ordered and sitting in the pipeline. One treatment is to assume that the customer will take his trade elsewhere and that the sale will be lost. A second possibility is that the customer is willing to wait until the item is received by the retailer, possibly experiencing some disutility which is charged to the retailer as a shortage cost. In the former case the sale is *lost* and in the latter case the sale is *backlogged*.

The paper has two major results. The first is to show that when sales are backlogged, the optimal ordering policy is a function of the stock on hand plus the stock previously ordered and not yet delivered. The second result is to show that policies of this simple form are not optimal when sales are lost. A more detailed study of optimal policies is then presented for the case of lost sales, a delay in the receipt of orders of a single period, and a purchase cost strictly proportional to the quantity ordered.

The analysis of optimal policies in the case of lost sales is conducted by means of the standard dynamic programming formulation of the inventory problem. Let the inventory on hand at the beginning of the period be x, and suppose that the initial decision is to order up to y units at a cost of $c(y-x)$. If the time lag is a single period, the order will be delivered at the end of the period resulting in a new level of inventory $\max[y-\xi, 0]$, where ξ is the random demand for stock during the period, governed, say, by a probability distribution with density $\varphi(\xi)$. Let $L(y;x)$ be the expected costs experienced during the period, and α the discount factor. Then $f(x)$, the discounted expected costs associated with the series of optimal decisions, will satisfy the dynamic programing equation

$$f(x)=\operatorname*{Min}_{y\geq x}\Big\{c(y-x)+L(y;x)+\alpha\Big[f(0)\int_y^\infty\varphi(\xi)\,d\xi+\int_0^y f(y-\xi)\varphi(\xi)\,d\xi\Big]\Big\}.$$

If delivery were instantaneous, the expected costs during the period would be a function, $L(y)$, of the immediately available inventory, and if backlogging were permitted the stock level could become negative and the pair of integrals in braces would be replaced by the single integral

$$\int_0^\infty f(y-\xi)\varphi(\xi)\,d\xi$$

resulting in the somewhat simpler equation

$$f(x)=\operatorname*{Min}_{y\geq x}\Big\{c(y-x)+L(y)+\alpha\int_0^\infty f(y-\xi)\varphi(\xi)\,d\xi\Big\}.$$

The classic work of the three authors Arrow, Harris, and Marschak, whose names appear in the title of the paper was published in *Econometrica* in 1951. Their work, "Optimal Inventory Policy" studies many aspects of inventory theory, including problems in which demand is known with certainty, single-period models with random demand, and general dynamic inventory models. In their analysis of the dynamic problem, they made the specific assumption that the cost of purchasing stock is composed of two parts: a set-up cost K incurred whenever an order is placed, and a unit cost c proportional to the size of the order. This is the case in which the optimal inventory policy was suspected to be an (S,s) policy, defined by the pair of numbers S, s and taking the form

$$y = \begin{cases} S & \text{if } x < s, \\ x & \text{if } x \geq s. \end{cases}$$

The optimality of such a policy was not known when the Arrow, Harris, and Marschak paper was written. What they did instead was to restrict their attention to policies of this particular form, to calculate the discounted expected cost associated with each such policy and to discuss the selection of that pair S, s yielding the lowest cost.

My original invitation to Stanford was for a single year, but the invitation was extended and in the fall of 1957 I was appointed an assistant professor in the Department of Statistics. I remained at Stanford, aside from a year-long visit to the Cowles Foundation at Yale in 1959–1960, until my departure for Yale in 1963.

I spent several weeks in the summer of 1958 working with a research group of the General Electric Corporation located in Santa Barbara. In my earlier account of this visit I stated incorrectly that Andy was working at General Electric and that he and I had discussed the results of his numerical computations of optimal inventory policies. But in the review of our histories during the last several weeks, we realized that Andy has never been at Santa Barbara. Herb Karr, whom I have previously mentioned, was the source of my invitation to Santa Barbara, and it was he whose numerical computations lead me to the conjecture that (S, s) policies were indeed optimal for the dynamic inventory problem.

Clark and I published two research papers together, one of which, "Optimal Policies for a Multi-Echelon Inventory Problem," has been cited as one of the 10 most influential papers published in *Management Science* in the 50 years of the journal's existence. The term "multiechelon" was invented by Clark and describes, in its simplest form, a situation with N installations linked in series, with installation $i-1$ receiving stock only from installation i, for $i = 2, \ldots, N$. If installation $i-1$ places an order from installation i, the length of time for the order to be filled is determined not only by the natural delivery time but also on the availability of stock at installation i. The collective optimal policies for the N installations can be found by solving a dynamic programming recursion in which the value function depends on the stock levels at each installation and the orders from successive installations which have not yet been delivered. The large number of arguments in these functions compromises our ability to obtain explicit numerical solutions.

In the paper we demonstrate that the value functions can, under certain assumptions, be decomposed into functions of a single variable, each of which satisfies its own recursive equation which can be solved quite readily. The major assumptions are that demand at each installation is backlogged, and that the purchase cost at each intermediary installation is linear, aside from the first installation in which a set-up cost is permitted. In a latter section we study the case in which several installations receive stock from the same supplier and demonstrate that the optimal policy is quite complex. The analysis is conducted by means of a simple example in which two installations, A_1 and A_2, receive stock from a single supplier B. Let

$$C_n(x_1, x_2, x_3)$$

be the minimum discounted system cost if there are n periods remaining, and the stock at the beginning of the period consists of x_i units at installation A_i and x_3 units at the supplier B. We assume a time lag of a single unit in the shipment of goods from B to each of the destinations, and that demands are backlogged if the stock is not currently available. Each installation has its own independent demand distribution.

It is easy to write down the dynamic program that relates C_n to C_{n-1}, and permits us to solve recursively for optimal policies. But, of course, the immediate recursion will involve a number of variables equal to the number of installations, and may be very time consuming to carry out, unless the cost functions can be written as the sum of functions each depending on a single variable. But unfortunately this is not generally the case. While I have not followed the enormous literature of the last 50 years on multiechelon problems, I suspect that much of it is devoted to the determination of inventory policies that are optimal in some approximate sense, for this more complex network of suppliers. This literature is a remarkable tribute to Andy's extraordinarily fruitful formulation of the multiechelon problem.

Our second joint paper, "Approximate Solutions to a Simple Multi-Echelon Inventory Problem" was published in the monograph *Studies in Applied Probability and Management Science*, edited by K. J. Arrow, S. Karlin, and H. E. Scarf, Stanford University Press, 1962. In this paper we examine an example of two installations in which set-up costs appear at the intermediary installation as well. Again, the cost functions do not simplify, and the true optimal policy may be quite complex. We provide an easily computable approximate inventory policy and an explicit set of bounds on the losses associated with this policy.