INVENTORY THEORY

BY

HERBERT E. SCARF

COWLES FOUNDATION PAPER NO. 1036

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281
2002

INVENTORY THEORY

HERBERT E. SCARF

Sterling Professor of Economics, Yale University, 88 Blake Road, Hamden, Connecticut 06517, herbert.scarf@yale.edu

Where to begin? I became involved with inventory theory in 1955, but the proper starting point for this memoir is some years earlier, in the fall of 1951, when I began my graduate training in the Department of Mathematics at Princeton University. I had lived at home during my undergraduate years at Temple University and was totally unprepared for the remarkable features of life at Princeton in the early 1950s. I had a room in the splendidly gothic Graduate College, and in a short time I met my classmates Ralph Gomory, Lloyd Shapley, John McCarthy, Marvin Minsky, Serge Lang, and John Milnor. John Nash and Harold Kuhn had left Princeton the year before, but I saw a good bit of them during their regular returns. Martin Shubik was then a graduate student in the Department of Economics, passionately engaged with Oscar Morgenstern in the early development of game theory. Some 50 years later, Martin and I have offices in the same building, the Cowles Foundation, at Yale University, and Ralph and I meet regularly.

I wrote my Ph.D. thesis under the direction of Salomon Bochner, a student of Erhard Schmidt, who was himself a student of David Hilbert. Albert Tucker was the chairman of the department; other faculty members were Soloman Lefschetz, William Feller, Emil Artin, and Ralph Fox. I got to know John Tukey during a daily commute to Bell Labs in the summer of 1953, where I would occasionally have a glimpse of Claude Shannon. I heard von Neumann lecture on computers and the brain and would often see Einstein and Gödel during their regular afternoon walks near the Institute for Advanced Study.

I left Princeton for the RAND Corporation in June of 1954. 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 that 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 the 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.

I moved into a simple office, far from my previous colleagues in mathematics, and sat 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

Subject classifications: Inventory/production: personal reflections. Professional: comments on.
Area of review: Anniversary Issue (Special).

0030-364X/02/5001-0186 $05.00
1526-5463 electronic ISSN

level y that maximizes the minimum expected profit for *all* probability distributions 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. My office was in a charming building known as Serra House, sitting in a grove of eucalyptus trees at the edge of the Stanford campus. I was on the second floor of the building along with Kenneth, Hiro Uzawa, and Patrick Suppes. Richard Atkinson was on the first floor, and Leo Hurwicz and Bill Estes were frequent visitors.

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 entitled "Studies in the Mathematical Theory of Inventory and Production," published in 1958. I will describe two papers that appeared in the monograph, in addition to the previously cited min-max solution.

The first of these, entitled "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 that 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. I remember an early conversation on this topic with Harry Markowitz, which took place in a rowboat, in the ocean off the coast of Santa Barbara, as we watched a seal sunning itself on a marker buoy.

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

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

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

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

The classic work of the three authors Arrow, Harris, and Marschak, whose names appear in the title of the paper, is more relevant to the second selection from this volume. Their work, "Optimal Inventory Policy," published in *Econometrica* in 1951, 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

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

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.

When a particular ordering policy is implemented, the level of inventories will typically follow a Markov process, which can be analyzed to determine the long-run expected cost associated with that policy. If the discount factor $\alpha = 1$, the stationary probability distribution of the inventory level at the beginning of the period will permit us to calculate expected holding and shortage costs, the quantity ordered and the expected number of periods between orders. If $\alpha < 1$, the stationary distribution is replaced by a discounted sum of probability distributions of inventory levels at the beginning of each period: a function that can be found by solving a suitable integral equation.

Working with fixed inventory policies, and their associated Markov processes, gives one a good deal of flexibility in the design of the model to be studied. In the paper, "Stationary Operating Characteristics of an Inventory Model with Time Lags," which also appears in the volume with Arrow and Karlin, I assume that demands are discrete, with an arbitrary probability distribution for the time between successive demands. An (S, s) policy is used, and when an order is placed the delivery time is a random variable, governed by its own distribution. The underlying stochastic processes are similar to those arising in queuing theory and in the design of telephone networks. A subsequent paper in the volume, written with Karlin, generalizes this analysis and relates it to an important class of stochastic processes known as renewal processes.

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 began my teaching career at Stanford with an undergraduate course built around linear programming, two-person zero-sum games, and the elements of the theory of convex sets. At various times, I taught a graduate course in inventory theory and related stochastic processes, and became aware for the first time of the importance of being able to lecture on one's current research. I began to work with a fine set of graduate students who were to complete their theses in the next several years. And I started to learn something about economics from discussions with Kenneth, Leo Hurwicz, and Hiro Uzawa and by attending the seminar on mathematics in the social sciences which Kenneth, Sam, and Pat Suppes had organized.

I am not sure that I got much research done during that busy year. I remember writing one paper entitled "Bayes Solutions to the Statistical Inventory Problem." The paper studies a conventional dynamic inventory problem, in which the purchase cost is strictly proportional to the quantity purchased, so that the optimal policy is defined in period n by a single critical number $\bar{x}_n$. The innovation in the paper is to allow the density of demand $\varphi(\xi, \omega)$ to depend on an unknown parameter, ω, about which we have a prior distribution with density, $f(\omega)$. As time evolves, the sequence of realized demands generates holding, shortage, and purchase costs, but in addition, we learn more about the true value of the underlying parameter.

For the analysis to be manageable, the demand distribution is assumed to take the form

$$\varphi(\xi, \omega) = \beta(\omega) e^{-\xi\omega} r(\xi).$$

With this specification, if we enter the nth period with a knowledge of the current stock level, x, and a history of past demands, $\xi_1, \ldots, \xi_{n-1}$, the entire history may be summarized in the sufficient statistic,

$$s = \frac{\sum_1^{n-1} \xi_i}{n-1},$$

so that the dynamic programming formulation depends only on the two-state variables x, s. Much effort is spent in the paper on demonstrating monotonicity of the critical numbers $\bar{x}_n(s)$, which now depend on s, and on determining their asymptotic behavior as n tends to infinity. The research underlying this paper taught me much about the elements of Bayesian analysis; it was very likely a genuflection in the direction of the department in which I was now employed.

I met Tjalling Koopmans in 1957 during a return visit to RAND. Tjalling was concerned with water storage policies in a hydroelectric system, and we talked about the relationship between this problem and the optimal management of inventories. Tjalling invited me to visit the Cowles Foundation during the academic year 1959–1960, and was surely responsible for the offer I received from Yale in 1963. We became close friends during the long period of our joint tenure at Yale. We played chess, swam, and canoed together, travelled to various conferences, and spent much time discussing our mutual research interests.

I spent the summer of 1958 working with Andrew J. Clark in a research group of the General Electric Corporation located in Santa Barbara. Andy was a gifted scholar, extremely knowledgeable about inventory problems and their applications. But for me, one of his most important virtues was his willingness to calculate optimal inventory policies numerically, on whatever primitive computers were available at that time. Andy showed me the results of his computations in models of varying complexity; we pored over the optimal policies for the simple dynamic model with instantaneous delivery, with an ordering cost composed of a set-up cost K and a constant marginal cost c. Andy mentioned to me that he had never seen a case of this problem in which the optimal policy was not of the (S, s) sort, ..., well *almost* never, he said. I was lucky that I didn't quite hear this final phrase.

As I look back over my research style, I realize that every now and then I become obsessed by some particular problem, that I worry about endlessly and find difficult to dismiss. In the fall of 1958, I turned to finding a proof of the optimality of (S, s) policies with this sense of obsessive involvement. The standard procedure for establishing the structure of optimal policies was a backwards recursion, based on a dynamic programming formulation. The optimal value function for an n period inventory problem, $f_n(x)$, satisfies the recursive relationship

$$f_n(x) = \operatorname*{Min}_{y \geqslant x}\left\{c(y-x) + L(y) + \alpha \int_0^\infty f_{n-1}(y-\xi)\varphi(\xi)\,d\xi\right\}.$$

In the absence of a set-up cost, the characterization of the optimal policy depends on a recursive demonstration that the function

$$cy + L(y) + \alpha \int_0^\infty f_{n-1}(y-\xi)\varphi(\xi)\,d\xi,$$

is convex, assuming that the single-period costs, $L(y)$, are themselves convex. But $f_n(x)$ cannot be convex when the ordering cost includes a set-up cost, and it was not clear precisely how to proceed. Some special cases had been studied by Bratton and Karlin in which optimality was obtained by making quite restrictive assumptions on the costs and on the demand density itself, but these arguments were not generally applicable.

I remember quite vividly, after many a fruitless month, turning in desperation to the special case in which demand was known with certainty to find some property of the value functions that could be carried along inductively and that was sufficient to demonstrate the optimality of (S, s) policies. After stumbling about with some tentative constructions, I realized that the value functions in this special case did indeed satisfy a condition, which I called *K-convexity*, that could be extended to the general problem as well. A function, $f(x)$, is called *K-convex* if the secant line connecting any two points on the graph of the function, when extended to the right, is never more that K units above the function, or in analytical terms, if

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

for $a, b > 0$ and all x.

If a definition can be said to have a smile associated with it, this is such a case. *K-convexity* is such an odd departure from the marginal conditions that characterize optimality for convex optimization, but it fits the inventory problem perfectly. If $f_{n-1}(x)$ is *K-convex*, then it is easy to show that

$$cy + L(y) + \alpha \int_0^\infty f_{n-1}(y-\xi)\varphi(\xi)\,d\xi$$

is *K-convex* as well and that the optimal policy in period n is an (S_n, s_n) policy with a pair of numbers appropriate to that period. It is then elementary to show that $f_n(x)$ is also *K-convex* and the induction continues. The demand densities are arbitrary and the single-period expected costs, $L(y)$, can be a general convex function; both the densities and these costs can vary over time. The only constraints on the parameters of the problem relate to the sequence of set-up costs, say, K_n. A function that is *K-convex* is also *K'convex* if $K' \geqslant K$, but not necessarily for smaller K'; it follows that if the set-up costs vary over time, they must decrease with increasing time for the inductive argument to be valid. It is easy to construct more complex optimal policies if the set-up costs increase over time; these are the occasional departures from the optimality of (S, s) policies that Andrew Clark had noticed in his numerical examples.

The paper was presented at the Stanford Symposium on Mathematical Methods, held at Stanford in June 1959. At the symposium were Gerard Debreu, Lionel McKenzie, Michio Morishima, Paul Samuelson, Bob Solow, Hans Theil, David Kendall, and Jascha Marschak—extraordinary scholars whom I met on this occasion for the first time. Some of these acquaintances turned into close professional and personal friendships, lasting over many years. Paul Samuelson and I actually became relatives in 1982, when my daughter Martha married his son, Paul R. Samuelson.

Clark and I published two research papers together, one of which, "Optimal Policies for a Multi-Echelon Inventory Problem," deserves special mention. The term multi-echelon was invented by Clark and describes 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 that 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 our second joint paper we examine a more general network of installations for which such a decomposition is not possible. We provide a bound on the increase in total cost associated with policies obtained by approximating the correct value function by a well-chosen sum of functions of single variables.

My research interests moved to other topics after these papers with Clark. But I did return at one point to write a survey paper entitled "A Survey of Analytical Techniques in Inventory Theory," which assembles a considerable amount of information that was known about inventory problems in 1963.

I was fortunate to have had a number of fine graduate students who wrote their Ph.D. theses on topics related to inventory theory. Donald M. Roberts derived elegant approximations to the optimal values of S, s when the set-up cost and the shortage cost tend to infinity in a suitable fashion. His work is reported in the paper "Approximations to Optimal Policies in a Dynamic Inventory Model," which appears in the volume *Studies in Applied Probability and Management Science* (1962 Arrow, Karlin, Scarf, eds.). Donald Iglehart's thesis involved the limiting behavior of the value function, $f_n(x)$, and optimal policies, (S_n, s_n), as n tends to infinity in the difficult case in which the discount factor $\alpha = 1$. His thesis "Dynamic Programming and Stationary Analysis of Inventory Problems," was published in the volume *Multistage Inventory Models and Techniques* (1963 Scarf, Gilford, Shelly, eds.).

In his thesis, entitled "Polya Type Distributions in Renewal Theory with an Application to Inventory Theory," Frank Proschan discussed a complex system, composed of many subunits, each of which had an independent, random failure time. If a subunit did fail and no replacement was available in stock, the entire system would no longer function. The issue was to determine the inventory of parts to be purchased with a fixed budget, so as to minimize the probability of failure, in a given time, of the complex system. Proschan and his collaborator, Richard Barlow, were major contributors to the Mathematical Theory of Reliability. Their work was recognized by the award of the Von Neumann Prize by the Institute of operations research and the Management Sciences in 1991.

At Yale, my students were concerned with quite different topics: cooperative game theory, fixed point theory, the computation of economic equilibria, applied general equilibrium theory, and indivisibilities in production. But in 1983, one of my graduate students, Andrew Caplin, wrote an important thesis applying inventory theory to a central problem in macroeconomics. Inventory management, as it appears in the literature of operations research, typically deals with a single firm or a small conglomerate of firms. Business cycle theorists, on the other hand, are primarily concerned with the behavior of inventories at the national level. One possible line of analysis linking these two concerns is to describe the economy-wide behavior of inventories by the aggregation of a vast number of individual optimizing decisions.

The aggregation of inventory decisions does not seem like an easy task, even if we restrict our attention to a single item that is stocked by a large number of small retailers. Under optimimizing behavior, the inventory levels of each retailer will be a stochastic process, governed, say, by the particular (S, s) policy adopted by that entity. Each such process is a Markov process of the sort discussed in the original paper by Arrow, Harris, and Marschak, but it is far from clear how to make a significant statement about their sums. This difficulty persists even if we focus not on the full stochastic process, but on the long-run stationary distribution of inventories at each firm. For a specific retailer, the stationary distribution can be calculated using renewal theory, but its precise form depends crucially on the distribution of demand that applies to that particular firm.

Caplin's thesis contains an ingenious resolution of this difficulty. Consider a single firm following a particular (S, s) policy with $Q = S - s$, and whose inventory levels are monitored periodically. Let us assume the orders can be placed *continuously* throughout the monitoring cycle, so that if the inventory level at the beginning of the period is $S - y$, with $0 \leqslant y < Q$, and the cumulative demand in the period is ξ, then

$$\lfloor (y+\xi)/Q \rfloor \text{ orders of size } Q$$

will be placed during the period and the resulting inventory at the beginning of the succeeding period will be $S - y'$ with

$$y' = y + \xi - Q\lfloor (y+\xi)/Q \rfloor.$$

It is then elementary to argue that a stationary distribution for a Markov process with this transition rule is the *uniform* distribution on the interval, $[0, Q]$, *regardless* of the distribution assumed for the random variable ξ. More generally, if there are n firms with a joint density of demand given by $\phi(\xi_1, \xi_2, \ldots, \xi_n)$, then the joint distribution in which the inventory levels are *independent* across firms and *uniform* for each firm is a stationary distribution for the resulting Markov process. Aggregation is easy to carry out.

I returned to inventory theory myself several years ago. Two colleagues at Yale, George Hall and John Rust, established a close relationship with a Connecticut company whose primary business activity is the purchase, storage, and eventual sale of a variety of steel products to local manufacturers in northeastern United States. Rust and Hall were kind enough to invite me to visit the company and discuss their procedures for inventory management, thereby reintroducing me to a research topic that I had left almost 40 years ago.

These discussions suggested a variation of the classical inventory model, in which the inventory manager may elect to meet a fraction of the demand if the sequence of costs and revenues make such a choice profitable. Such a situation might arise if the cost, sales, and demand parameters vary substantially over time, possibly in a stochastic fashion. In the Cowles Foundation Discussion Paper, "Optimal Inventory Policies When Sales Are Discretionary," written in 2000, I showed that under classical conditions the optimal policy is again of the (S, s) form. The argument makes use of a property of *K-concave* functions that I had never seen before. I am grateful to Guillermo Gallego, who took the time to read a draft of this paper and made a number of useful suggestions.

ACKNOWLEDGMENTS

The material of this account is a reworking of an essay to appear in a selection of my research papers to be published by the M.I.T. Press.

REFERENCES

Arrow, Kenneth J., Theodore Harris, Jacob Marshcak. 1951. Optimal inventory policy. *Econometrica* 250–272.

Caplin, Andrew S. 1985. the variability of aggregate demand with (S, s) inventory policies. *Econometrica* **53**(6).

Iglehart, Donald. 1963. Dynamic programming and stationary analysis of inventory problems. H. E. Scarf, D. M. Gilford, M. W. Shelley, eds. *Multistage Inventory Models and Techniques.* Stanford University Press, Stanford, CA.

Roberts, Donald M. 1962. Approximations to optimal policies in a dynamic inventory model. K. J. Arrow, S. Karlin, H. E. Scarf, eds. *Studies in Applied Probability and Management Science.* Stanford University Press, Stanford, CA.

Scarf, Herbert E. 1959. Bayes solutions of the statistical inventory problem. *Ann. Math. Statist.* **30**(2).

——. 1995. The optimality of (S, s) policies in the dynamic inventory problem. *Mathematical Methods in the Social Sciences.* Stanford University Press, Stanford, CA. Also in E. Sheshinski, Y. Weiss, eds. *Optimal Pricing, Inflation, and the Cost of Price Adjustment.* MIT Press, Cambridge, MA, 1993.

——. 1963. A survey of analytic techniques in inventory theory. H. E. Scarf, D. D. Gilford, M. W. Shelley, eds. *Multistage Inventory Models and Techniques.* Stanford University Press, Stanford, CA.

——. 2000. Optimal inventory policies when sales are discretionary. Cowles Foundation Discussion Paper No. 1270, August.

——, A. J. Clark. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* **6**(4) 475–490.

——, ——, 1962. Approximate solutions to a simple multi-echelon inventory problem. K. J. Arrow, S. Karlin, H. E. Scarf, eds. *Studies in Applied Probability and Management Science*, Stanford University Press, Stanford, CA.

——, S. Karlin. 1958. A min-max solution of an inventory problem. K. J. Arrow, S. Karlin, H. E. Scarf, eds. *Studies in Applied Probability and Management Science.* Stanford University Press, Stanford, CA.

——, ——. 1958. Inventory models and related stochastic processes. K. J. Arrow, S. Karlin, H. E. Scarf, eds. *Studies in Applied Probability and Management Science.* Stanford Unversity Press, Stanford, CA.

——, ——. 1958. Inventory models of the Arrow-Harris-Marschak type with time lag. K. J. Arrow, S. Karlin, H. E. Scarf, eds. *Studies in the Mathematical Theory of Inventory and Production.* Stanford University Press, Stanford, CA.

——, ——. 1958. Stationary operating characteristics of an inventory model with time lag. K. J. Arrow, S. Karlin, H. E. Scarf, eds. *Studies in Applied Probability and Management Science.* Stanford University Press, Stanford, CA.