A Dynamic Inventory Model with Random Replenishment Quantities

By: H. Baker and R. Ehrhardt

Baker, H. & Ehrhardt, R. (1995). A dynamic inventory model with random replenishment quantities. *Omega*. pp. 23(1), 109-116.

Abstract:
A periodic-review, random-demand inventory model is analyzed under the assumption that replenishment quantities are random fractions of the amounts ordered. Results of a previous study of a single-period model are generalized to form an easily computed heuristic adaptation of the *(s, S)* policy for use in this environment. The heuristic is based on the simple practice of scaling down pipeline inventories to estimate the inventory position, and scaling up the order quantity in anticipation of an average replenishment yield. Simulation experiments are used to estimate the most cost-efficient *(s, S)* policies and to estimate the performance of heuristic policies in environments where replenishment randomness ranges from mild (0-20% defectives) to moderate (0-50% defectives). The heuristic is shown to perform quite well, with expected total costs typically within a few percent of the best (s, *S)* costs. The results tend to support common practice in industry which is similar to the approach studied here. Although the heuristic is naive in the sense that it ignores the degree of randomness in the replenishment quantity, the simulation results support the speculation that unless the target service level is extremely high, the replenishment process must be extremely random for its variability to be a significant explicit factor in the selection of a practical, cost-effective policy.
Key words—inventory control, replenishment, random replenishment quantities

Article:
Consider a problem that might, for example, confront the production manager of an electronics manufacturing firm. The manager is responsible for producing a highly sophisticated component that must meet very stringent quality specifications. Each production run yields a batch having only a random portion of the components which pass inspection. The component is produced to stock in response to continuing demands of random size. The problem at hand is to specify the inventory replenishment policy for the component which will drive the production process. It is likely that an optimal policy will compensate by prescribing an order for more units than are actually desired.

The problem is an example of a situation that is not addressed by traditional inventory management models. That is, one cannot prespecify the exact size of an inventory replenishment at the time of ordering. Rather, one is limited to choosing from a set of ordering levels that influence the probability distribution of a random replenishment quantity. While it is certainly true that all real inventory systems have this property, many systems have replenishment processes that are sufficiently predictable to assume that the quantities are deterministic. This is often the case when stocking a simple, standard item that is in abundant supply. On the other hand, it may not be appropriate to assume deterministic replenishment quantities in situations when quality requirements are unusually high or when the replenishment process itself is inherently random, such as in agriculture, blood donations and semiconductor chip production.

Topics in the areas of random replenishment inventory systems and variable yield production systems have received considerable attention in recent years, yet little is known about optimal policies in a multi-period, random demand environment other than the fact that they do not possess many of the more elegant properties of their deterministic-replenishment counterparts. Silver [17], Sepheri *et al.* [15], Lee and Yano [10], and Shih [16] analyze models with deterministic demand. Shih [16], Karlin [9], Noori and Keller [13] and Ehrhardt and Taube [4] discuss single-order problems with random demand and various assumptions about the ordering process, cost structure and demand distributions.

Ehrhardt and Taube [4] show that when the replenishment quantity is a random fraction of the amount ordered, an optimal single-period ordering policy can be found with a simple generalization of the traditional newsvendor result. They also show that a simple scaling-up heuristic is an effective approximation to optimal performance. The heuristic computes an order size by starting with the order size that would be optimal with deterministic replenishment, and dividing it by the expected value of the replenishment yield fraction.

Moinzadeh and Lee [11] study a continuous review system under Poisson demand, and analyze ordering policies confined to the approximately-optimal *(r, Q)* form, where the reorder point *r* is based on the sum of inventory on hand plus that on order. They perform an exact Markov-chain analysis of the system and find that the best values of *r* and *Q*, while difficult to compute, are nicely approximated by an intuitively appealing adaptation of the heuristic of Hadley and Whitin [6]. The adaptation, similar in spirit to Ehrhardt and Taube's [4], consists of merely scaling up the mean of the Poisson process by dividing it by the expected value of the replenishment yield fraction. Gerchak *et al.* [5] analyze a multi-period model with linear ordering cost and random demand. They establish the existence of order points, and find that optimal policies are in general difficult to compute and non-myopic in character. A significant theoretical contribution was made by Henig and Gerchak [7], who discuss single- and multi-period models with more general assumptions about the random replenishment distribution and the cost structure. They show that for a single-period model there exists, under very general conditions, an optimal order point whose value is independent of replenishment randomness. They also analyze a zero-leadtime multi-period model having replenishment quantities that are random fractions of the order sizes, and linear order costs. They know that optimal policies are not of the order-up-to type, and that the optimal order quantity and reorder point both tend to increase with replenishment randomness, although the strength of the relationships are not characterized. They also establish convergence of value functions to their infinite horizon counterparts, and that infinite-horizon optimal order points are no smaller than when replenishment is deterministic. Recently Bollapragada and Morton [1] analyzed a multiperiod model with normally distributed demand, linear cost functions, and replenishment quantities that are random fractions of the quantities ordered. They show that the problem is nearly myopic under certain assumptions, and present several accurate heuristics. Their model is very effectively analyzed using myopic and near-myopic methods, but it does not, however, include an order setup cost or a delivery lead time.

In the remainder of this paper, we analyze a multi-period model with independently distributed demands and an expected-cost optimality criterion. Specifically, in the next section we specify a dynamic model having linear holding and shortage costs, an ordering setup cost, and a rather simple mechanism for generating replenishment quantities. An Appendix is also provided to place the model in a more general context, and to discuss the theoretical reasons for difficulties in computing exactly optimal ordering policies. In the following section we introduce an easily-computed heuristic *(s, S)* policy that is inspired by results from a previously published single-period version of the model. Then we present the results of a simulation experiment comparing the performance of the heuristic policy with that of the best policy of the *(s, S)* form. Finally, we briefly draw conclusions.

A DYNAMIC INVENTORY MODEL

In this section we describe the model that we analyzed in detail with a simulation experiment. We develop our reasons for confining attention to a model this specific, and for resorting to simulation rather than mathematical analysis, in Appendix A, where we briefly present a more general theoretical discussion of random replenishment dynamics.

We consider a periodic review inventory system in which demands in successive periods are independent and identically distributed. The sequence of events within a period is order, receipt of any replenishment, and demand. There is a fixed lead time $L \geq 0$ between the period in which an order is placed and the period in which it is delivered. We assume that any unsatisfied demand in a period is completely backlogged.

Let the amount of stock ordered be denoted by *z*, and the amount that is actually delivered by *a (z)*, a random variable with cumulative distribution function B_z. One would typically expect that *a (z)* possesses certain general properties. For example, it is reasonable to assume that if $u \leq v$, then $a(u) \leq a(v)$ in a stochastic sense, i.e. $B_u(w) \geq B_v(w)$ for all *w*. The remainder of this paper is based on the assumption that $a(z) = Az$, where *A*, the fraction of the order that is actually delivered, is a random variable with support [0, 1] and mean μ_a. We shall refer to μ_a as the *average replenishment yield.*

The cost of ordering stock is composed of a setup cost *K* and a unit purchase cost *c* for each unit of stock delivered. Although this implies that only the delivered units are purchased, the modeling approach also encompasses applications where all units ordered are purchased. There is a holding cost *h* for each unit of inventory on hand at the end of a period, and a backlog penalty cost *p* for each unit backlogged at the end of a period. The management objective is to minimize the long-run average undiscounted total cost per period.

A MODIFIED *(s, S)* HEURISTIC POLICY

We note that the heuristic policy defined below is sufficiently robust to be applied in more general environments (nonstationary demand, convex cost structures, etc.), but we evaluate its performance only in an infinite horizon setting with stationary demand distributions and linear holding and shortage costs.

The heuristic policy we define is based upon the deterministic-replenishment version of the model, and relies on the decision maker's knowledge of μ_a to (1) scale up the desired order size when ordering, and (2) to scale down the on-order quantity when assessing the inventory position. The optimal policy form for the deterministic-replenishment version of the model is *(s, S)*. That is, if the inventory position *w* (stock on hand plus on order) is less than or equal to *s*, then order an amount $z = S - w$; otherwise order nothing.

The heuristic policy is defined as follows. Let x_t and z_t be the starting on-hand inventory and the order size in period *t*, and let (s^*, S^*) be a policy that is optimal for the deterministic-replenishment version of the model. We first define the inventory position in period *t* (before ordering and delivery) as

$$w_t = x_t + \mu_a[z_{t-L} + z_{t-L+1} + z_{t-L+2} + \dots + z_{t-1}].$$

Then the prescribed order size in period *t* is zero unless $w_t \leq s^*$, in which case it is given by

$$z_t = (S^* - w_t)/\mu_a.$$

In practice, the optimal deterministic-replenishment policy can be replaced by a good, easily computed approximation (see [2, 3, 12, 14]). If so, the total cost performance should not suffer by more than a percent or so, given the accuracy of the approximations in the deterministic replenishment environment. If a theoretically exact calculation of (s^*, S^*) is desired, a good approximation can be improved by using the efficient algorithm of Zheng and Federgruen [21].

The logic of the heuristic is simple: take credit only for the expected value of the amount of outstanding orders, and order so that the expected value of the order size matches the deterministic-replenishment order size. Aside from the common sense appeal of the heuristic, it is motivated by the computational results of Ehrhardt and Taube [4] who found that the same scaling-up approach was quite effective in the single period version of the model with no setup costs. They compared the scaling-up policy with exactly optimal policies in a very random replenishment environment (uniform on [0 1]), and found that it was within 0.8% of optimal total cost when the holding-to-shortage cost ratio *p/h* was 4 or 9. When *p/h* was 24, the heuristic averaged within 1.4% of optimal total cost except for one item which was 9.1 % above optimal. It clearly outperformed another, more complicated heuristic which was designed to adjust the ordering policy for both the mean and variance of the replenishment yield variable *A.*

We note that the heuristic could be considered naive in that it ignores the degree of randomness in the replenishment distribution. Indeed, it would prescribe the same policy for a replenishment yield distributed uniformly on [0.4, 0.6] as for one distributed uniformly on [0, 1]. In view of the fact that such a naive approach was very effective in a single-order environment [4], we investigate its potential in a dynamic environment despite the established result [7] that it is not exactly optimal. We estimate the effectiveness of the heuristic using a simulation experiment described in the following section.

A SIMULATION EXPERIMENT

We focus on the long-run behavior of the model. To characterize the performance of the heuristic *(s, S*) policy, we compare it with the best known *(s, S)* policy, estimated using a simulation experiment. The simulation, written in Pascal, searches for the *(s, S*) policy having the smallest average total cost per period. For each item simulated, we record the sample means and 95% confidence intervals of the following performance characteristics: holding cost per period, order cost per period, shortage cost per period, total cost per period, and backlog frequency. We describe details of the experimental design and summarize results below.

We simulated a full factorial permutation of the parameter settings listed in Table 1, resulting in 256 distinct inventory items. Without loss of generality, we have set the unit holding cost rate *h* equal to 1, and the unit revenue *r* and unit purchase cost *c* equal to zero. The replenishment yield distributions are uniform on [0.5, 1.0] and [0.8, 1.0] for the mean values 0.75 and 0.90, respectively. For each of the 256 items (parameter combinations), we computed the optimal deterministic replenishment policy *(s*,S*)* utilizing the algorithm of Veinott and Wagner [19], which we then modified to a heuristic policy using the method described in the previous section.

Table 1. Parameter values of simulated items

Parameter	Symbol	Values
Unit holding cost	h	1
Unit shortage cost	p	4, 9, 24, 99
Order setup cost	K	32, 64
Replenishment lead time	L	0, 2
Mean replenishment yield	μ_a	0.75, 0.90
Mean demand	μ	2 4, 8, 16
Demand distribution		Poisson ($\sigma^2/\mu = 1$)
		Negative binomial ($\sigma^2/\mu = 3$)

Table 2. Frequencies of percent differences in mean total cost

Range (%)	No. of items	Cumulative % of items
0.0–0.5	130	50.8
0.5–1.0	64	75.8
1.0–2.0	36	90.0
2.0–3.0	12	94.5
3.0–4.0	5	96.5
4.0–5.0	2	97.3
5.0–7.0	3	98.4
7.0–10.0	3	99.6
10.0–20.0	1	100

The heuristic policy was the center of a range of policies we simulated to assess cost performance. We determined the range by first varying the quantity $D = S - s$ by 10% above and below the heuristic value, and then for each value of *D* by varying *s* above and below its heuristic value by the larger of 10% or 10. In those instances where the lowest-cost policy was found to be at the end of the range of either *D* or *s,* we extended the range until a local minimum was established. We refer to the policy found by this simulation response surface method as the *best (s, S) policy.* Approximately 54,600 policies were simulated in studying the 256-item collection of parameter combinations. The number of policies simulated for an individual item ranged from a minimum of 33 [for an item having a heuristic policy of (2, 13)] to a maximum of 1073 [for an item having a heuristic policy of (88, 134)]. For every policy simulated, we ran 101 independent replications of 1000 periods. The first run of 1000 periods was discarded before collecting data. The program computed mean values and 95% confidence intervals for the total cost per period as well as its components and the backlog frequency. The confidence intervals are within ± 1.8% of the mean value, on average.

The total cost performance of the heuristic policy is quite close to that of the *best (s, S) policy* in our simulation experiment. In Table 2, we list frequency statistics for the estimated percentage difference between the mean total costs of the *best (s, S) policy* and the heuristic. Notice that the heuristic is within 0.5% of the *best (s, S) policy's* total cost for 50.8% of the 256 items, and within 2.0% for 90.0% of the items. The average percent difference in mean total cost is 0.93% for all 256 items. The largest percent difference is 19.7%, for the item having Poisson demand, $\mu = 16$, $L = 2$, $\mu_a = 0.75$, $K = 64$ and $p = 99$. The heuristic policy for this item is (55, 95), while the *best (s, S) policy* is (64, 99). Although five of the seven items having a cost difference greater than 5% have $p = 99$, it is difficult to explain the larger difference for this particular item. The same item with μ_a increased to 0.9 has a cost difference of only 1.38%. It appears as if managers should be wary if they wish to design extremely high-service inventory systems in the presence of significantly random replenishment quantities.

Table 3. Effects of replenishment variability on average costs

Policy	μ_a	Holding	Shortage	Setup	Total
Heuristic	0.75	18.8	6.3	10.1	35.3
	0.90	18.2	5.9	10.1	34.0
Best	0.75	19.5	5.2	10.2	34.8
	0.90	18.5	5.1	10.2	33.8

Table 4. Effects of μ_a and p/h on backlog protection

			Mean backlog protection	
p/h	μ_a	$p/(p+h)$	Best	Heuristic
4	0.75	0.800	0.819	0.788
	0.90	0.800	0.816	0.791
9	0.75	0.900	0.911	0.885
	0.90	0.900	0.911	0.892
24	0.75	0.960	0.965	0.952
	0.90	0.960	0.965	0.956
99	0.75	0.990	0.992	0.987
	0.90	0.990	0.992	0.989

We found that the percentage difference in total cost between the *best (s, S) policy* and the heuristic is in general not very sensitive to the parameter settings. The average percentage difference in total cost between the two policy rules is no more than 1.57% when aggregated over all items having the same value of any single parameter setting. The variance in total cost percentage differences is significant at the 0.05 level only when comparing the two demand distributions, mean demand values of 2 and 8, the two levels of mean replenishment yield, unit shortage cost values of 4 and 99, or unit shortage cost values of 9 and 99.

It is interesting to examine how the two policy rules differ in the values of their cost components. In Table 3 we list the average values of total cost per period and the components of total cost per period when aggregated over all items having the same value of the average replenishment yield μ_a. Notice that the *best (s, S) policies* tend to hold more inventory, have smaller shortages, and have slightly larger setup cost per period than the heuristic policies. The differences are more pronounced for the smaller replenishment yield value. It appears that the *best (s, S) policies* are holding more safety stock and ordering slightly smaller batches than the heuristic policies. This observation is supported when the actual policy numbers are examined. Although the differences in shortage cost appear significant, we judge this simple heuristic to be surprisingly effective in its overall (total cost) performance.

Finally, we examine the backlog performance of the two policies by computing their values of *backlog protection*, which is the fraction of periods in which a backlog does not exist. Managers are often interested in such a measure because the backlog penalty cost parameter is often not based on real costs, but rather used as a tool to strike a reasonable balance between carrying costs and customer service. We tabulate backlog protection values of the two policy rules, and compare them to their theoretical bounds in a perfect replenishment environment. Recall that in a periodic review system with deterministic replenishment, and optimal *(s, S)* policy has backlog protection equal to $p/(h + p)$ if demand is continuously distributed. If demand is discrete, then the optimal backlog protection is equal to or slightly greater that this figure [19, p. 539]. In Table 4 we list average simulated backlog protection aggregated over all items with the same value of p/h and average replenishment yield μ_a. Notice that the heuristic policies have somewhat lower backlog protection, which is consistent with our observations about the cost components listed in Table 3. The *best (s, S) policies* tend to behave more as optimal *(s, S)* policies would in a perfect replenishment environment, in that they give slightly higher protection than the theoretical target value. The heuristic policies, appear to be just as close to the target value on average. Notice that these results parallel the analytic result [7] for $K = 0$ that random replenishment tends to increase the optimal reorder point above its deterministic replenishment value. These results would seem to indicate, however, that such an effect is rather weak for the parameter values used in this study.

CONCLUSIONS

We have formulated a dynamic model of an inventory system having random replenishment quantities. Although the mathematics of computing optimal policies are quite cumbersome, we have shown that an easily-computed modified *(s, S)* heuristic performs well compared with other policies within the modified *(s, S)* class when there is mild to moderate randomness in replenishment quantities.

The results should offer some encouragement to inventory managers, since the heuristic approach (scaling-up replenishment lots) parallels a common industry practice for items known to have variable replenishment quantities, and is easily adapted to other policy forms, *(Q, r)* for example. Although the scaling-up heuristic tends to order lots that are slightly too large and hold safety stocks that are slightly too small, its good total cost performance and its closeness to traditional backlog protection targets provide quantitative evidence that the practice is a sound one. We note that although we did not simulate environments with extremely random replenishment yields, our parameter set focused on relatively low values of demand variance, which were the more challenging cases for this type of heuristic in an earlier study [4]. We expect that the heuristic would perform at least as well as noted here if larger values of demand variance were simulated. It may turn out, however, that the heuristic is not as efficient for extremely random replenishment yields. Limited evidence in this study also points to the possibility that very high service level targets (values of p/h in the neighborhood of 99) may also be associated with sensitivity of optimal policies to randomness in replenishment quantities.

From a theoretical perspective, there is room for improvement in the heuristic due to the fact that its safety stocks are consistently smaller than those of the best policies found, even though the difference is not large. It is also true that this study does not compare the heuristic with truly optimal policies. Although significantly better total cost performance may be possible with true optimal policies, the enormous complexity of computing and implementing them would seem to make them of interest primarily to mathematical analysts. That is not to say that additional research is not warranted. It appears, though, that the most promising avenues of investigation are those that confine attention to convenient policy forms [11].

APPENDIX A

A Mathematical Model of the Dynamic Inventory System

In this Appendix we discuss a general version of the random replenishment inventory model. The discussion illuminates the difficulties surrounding the computation and implementation of theoretically optimal policies, and motivates the need for effective heuristic policies that are easily implemented.

We consider a multi-period model of length T periods in which time-related quantities are labelled with the subscript t, but such subscripts will be suppressed in the interest of simplicity of exposition wherever clarity permits. The sequence of events within a period is order, receipt of any replenishment and demand. Assume initially that replenishment orders are delivered immediately (the leadtime $L = 0$). Demand D_t in each period is independent and identically distributed. Let the amount of stock on hand before ordering be denoted by x, the amount of stock ordered by z, and the amount that is actually delivered by $a(z)$, a random variable with cumulative distribution function B_z. Also, let $y_t = x_t + a(z_t)$ denote the amount of stock available to meet demand in period t, and assume that all unsatisfied demand is backlogged. One would typically expect that $a(z)$ possesses certain general properties. For example, it is reasonable to assume that if $u \leq a(v)$, then $a(u) \leq a(v)$ in a stochastic sense, i.e. $B_u(w) \geq B_u(w)$ for all w. The model that follows is based on the assumption that $a(z) = Az$, where A is a random variable with support [0, 1] and mean μ_a.

We follow the general pattern of stochastic inventory model analysis presented by Heyman and Sobel [8, pp. 306-327]. All cost functions and parameters are assumed to include the necessary discount factors so that they represent values discounted to the first period of decision making. Let $c_t(u)$ be the ordering cost in period t when u units of stock are delivered. Although this implies that only the delivered units are purchased, the model is easily adjusted to applications where all units ordered are purchased. We assume, unless otherwise noted, that the ordering cost has a linear component $c_t u$ as well as a setup cost K,. For example, if these derive from a constant unit purchase cost c and a constant setup cost K, then, when $u > 0$,

$$c_t(u) = K_t + c_t u = \alpha^{t-1}(K + cu),$$

α being the discount factor per period. We express the remaining costs in period t as the function $g_t(y, D_t)$, where y is the amount of stock available to meet demand, that is, the initial stock plus the replenishment quantity received in that period. (A negative value of y signifies a persisting amount of unsatisfied demand, which we assume to be completely backlogged.) For example, if h and p are the respective unit costs of inventories and backorders, and r is the price, then

$$g_t(y, d) = \alpha^{t-1}[h(y-d)^+ + p(d-y)^+ - rd]$$

if all revenue is received at the time goods are demanded. Finally, let $g_T(y_T, D_T)$ include all costs and revenues associated with salvage of excess stock or liquidation of an backlog at the end of the decision-making horizon.

Let B denote the present value, at the beginning of period 1, of all costs that are incurred during periods 1, . . . , T + 1. Then

$$B = \sum_{t=1}^{T} [K_t\delta(z_t) + c_t a(z_t) + g_t(y_t, D_t)]$$

$$B = \sum_{t=1}^{T} \{K_t\delta(z_t) + c_t a(z_t) + g_t[x_t + a(z_t), D_t]\},$$

where $\delta(x) = 1$ *if* x is positive and zero otherwise. It is clear that the following dynamic programming recursion is appropriate for minimizing the expected value of B.

$$\begin{aligned} &f_{T+1}(x) = 0, \\ &f_t(x) = \inf\{K_t\delta(z) \\ &\qquad + E[c_t Az + g_t(x + Az, D_t) \\ &\qquad + f_{t+1}(x + Az - D_t)]: \quad z > 0\}, \\ &t = 1, \ldots, T. \end{aligned}$$

We identify $f_t(x)$ as the minimum expected discounted (to period 1) total cost for periods t through T, starting with an inventory level of x.

At this point it is clear that the model cannot be transformed in the traditional way to one in which the decision variables are the amounts of stock available to meet demand $\{y_t, t = 1, \ldots, T\}$, since choosing a value for z_t, *does not fix the value of* y_t; it merely determines its probability distribution. It follows that the structure of optimal policies that has been developed for deterministic replenishment models [myopic optima, *(s, S)* policies, etc.] will not transfer directly to our model. In fact, it has been shown [7] that an order-up-to policy will not generally be optimal. It appears that true optimal policies are generally both complicated to compute and, once computed, would be tedious to implement.

Additional complications arise when the replenishment lead time L is positive. In this case we define a state vector having L + 1 components. The first component is the stock on hand at the beginning of a period (before receipt of any outstanding order and before demand), and the remaining components are the sizes of outstanding replenishment orders in order of delivery. Therefore, the state in period t is given by

$$(x_t, z_{t-L}, z_{t-L+1}, \ldots, z_{t-1})$$

which is transformed to

$$[x_t + a(z_{t-L}) - D_t, z_{t-L+1}, z_{t-L+2}, \ldots, z_t]$$

after one period. The state space cannot be collapsed into a scalar inventory position, as is the case in more standard backlogging models. Therefore, the model has the same state space as the lost-sales model, and also shares the same curse of dimensionality. Optimal policies are significantly more difficult to compute than for the zero-leadtime model described above, and they would be even more tedious to implement. This is why we turn our attention to heuristic policies that are both easily-computed and easily-implemented, and why we resort to simulation rather than analytic computation.

REFERENCES

1. Bollapragada S and Morton TE (1993) The periodic review inventory problem with random yield: near-myopic properties, heuristics and testing. WP # 199311, Graduate School of Industrial Administration and The Robotics Institute, Carnegie Mellon University.
2. Ehrhardt R (1979) The power approximation for computing *(s, S)* inventory policies. *Mgmt Sci.* 25, 777-786.
3. Ehrhardt R and Mosier C (1984) A revision of the power approximation for computing *(s, S)* policies. *Mgmt Sci.* 30, 618-622 (1984).
4. Ehrhardt R and Taube L (1987) An inventory model with random replenishment quantities. *Int. J. Prod. Res.* 25, 1795-1803.
5. Gerchak Y, Vickson RG and Parlar M (1988) Periodic review production models with variable yield and uncertain demand. *IIE Trans.* 20, 144-150.
6. Hadley G and Whitin TM (1963) *Analysis of Inventory Systems.* Prentice-Hall, Englewood Cliffs, N.J.
7. Henig M and Gerchak Y (1990) The structure of periodic review policies in the presence of random yield. *Opns Res.* 38, 634-643.
8. Heyman DP and Sobel MJ (1984) *Stochastic Models in Operations Research, Vol. II.* McGraw-Hill, New York.
9. Karlin S (1958) One-stage models with uncertainty. In *Studies in the Mathematical Theory of Inventory and Production* (Edited by Arrow KJ, Karlin S and Scarf H). Stanford University Press, Stanford, Calif.
10. Lee HL and Yano CA (1988) Production control in multistage systems with variable yield losses. *Opns Res.* 36, 269-278.
11. Moinzadeh K and Lee HL (1987) A continuous-review inventory model with constant resupply time and defective items. *Naval Res. Log.* 34, 457-468.
12. Naddor E (1975) Optimal and heuristic decisions in single- and multi-item inventory systems. *Mgmt Sci.* 21, 1234-1249.
13. Noori AH and Keller G (1986) One-period order quantity strategy with uncertain match between the amount received and quantity requisitioned. *INFOR* 24, 1-11.
14. Porteus E (1985) Numerical comparisons of inventory policies for periodic review systems. *Opns Res.* 33, 134-152.
15. Sepehri M, Silver EA and New C (1986) A heuristic for multiple lot sizing for an order under variable yield. *HE Trans.* 18, 63-69.
16. Shih W (1980) Optimal inventory policies when stock-outs result from defective products. *Int. J. Prod. Res.* 18, 677-685.
17. Silver EA (1976) Establishing the order quantity when the amount received is uncertain. *INFOR* 14, 32-39.
18. Veinott AF (1966) The status of mathematical inventory theory. *Mgmt Sci.* 12, 745-777.
19. Veinott AF and Wagner HM (1965) Computing optimal *(s, S)* inventory policies. *Mgmt Sci. 11,* 525-552.
20. Wagner HM (1975) *Principles of Operations Research,* 2nd edition. Prentice-Hall, Englewood Cliffs, N.J.
21. Zheng Y and Federgruen A (1991) Finding optimal *(s, S)* policies is about simple as evaluating a single policy. *Opns Res.* 39, 654-665.