An inventory model with random replenishment quantities

By: Richard Ehrhardt and Larry Taube

Ehrhardt, R. & Taube, L. (1987). An inventory model with random replenishment quantities. *International Journal of Production Research*, 25(12), pp. 1795-1803.

Abstract:
A single-period random demand inventory model is analysed under the assumption that the replenishment quantity is a random function of the amount ordered. The structure of optimal policies is characterized for linear ordering cost functions, both with and without a setup cost. Generalizations of base-stock policies and (s, S) policies are shown to be optimal. Closed form analytic expressions for optimal policies are obtained for the special ease of linear ordering cost and uniformly distributed demand. Optimal policies are compared with two heuristics in a more general demand environment. It is shown that a very simple heuristic policy performs quite well.

Article:
1. Introduction
Consider a problem that might, for example, confront the production manager of an electronics manufacturing firm. The manager must order semiconductor chips for assembly into a highly sophisticated component. Unfortunately, no supplier can reliably provide chips that meet the stringent performance specifications required for use in this component. The manager must order more chips than are actually desired, hoping that a reasonable number of them will be usable. But how many chips should be ordered?

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 among 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 such as:

(a) choosing the size of a production run that is to replenish the stock of a high-quality product to be sold to clients;

(b) attempting to replenish the supply of a rare blood type in a, blood bank; and

(c) deciding on the amount of land to devote to growing a particular crop.

The terminology used in this paper is that, of inventory models. The reader should remain aware that, as is true of most papers about inventory models, applications include a variety of scenarios. In particular, when the term replenishment quantity is used it merely refers to the amount of usable material that received from the ordering or production process. That is, the model encompasses situations where the amount received or produced is deterministic but the amount that passes inspection is random. The term perfect replenishment will be used when referring to the standard model in which the amount delivered is always equal to the amount ordered.

Literature in the topic area is sparse. Silver (1976), Sepehri *et al. (1986),* and Shih (1980) analyse models with deterministic demand and simple cost structures. Shill (1980) also briefly analyses a model with continuously distributed random demand, but stops after deriving an optimality condition for the single-period model with a simple cost structure. Karlin (1958) considers a single-period problem with random demand and general cost structure, but limits the ordering decision to only two alternatives: (a) don't order and (b) order with a given random result. A range of order sizes is not permitted. Recently, Noori and Keller (1986) have obtained results characterizing optimal order quantities for the specific eases of uniformly and exponentially distributed demand. There are apparently no results in the literature concerning a multi-period model with random demand.

In the remainder of this paper we analyse a single-period model with random demand and an expected-cost optimality criterion. We characterize optimal policies and compare them with easily-computed heuristic policies. Specifically, in §2 we discuss a single-period model with general forms for the cost, functions, demand distribution and replenishment distribution. We present properties of optimal policies for the general model and the important special case of linear cost functions. In §3 we present detailed analytic results that apply when an assumption of uniformly distributed demand is added to the model. We discuss numerical results for negative binomial demand distributions in §4, and draw conclusions in §5.

2. A single-period model

We consider a model in which a single ordering decision is to he made so as to minimize expected total cost. The model applies to situations where a final order is placed for an item that will soon be obsolete, and to dynamic systems where the time between ordering decisions is so dong that either lost sales or salvaging takes place after demand is realized. The sequence of events is: order, delivery and demand. The actual times at which events occur needn't be specified, since the expected values of discount factors can be included in the cost functions.

Let the amount of stock on hand before ordering be denoted by x_o, 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. One would typically expect that *a(z)* possesses certain general properties. For example, it is reasonable to assume that if $u \geq v$, then $a(u) \geq a(v)$ in a stochastic sense; that is $B_u(x) \leq B_v(x)$ for all *x*.

Let *c(u)* be the ordering cost 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. The remaining costs are expressed as the function *g(x),* where *x* is the amount of stock remaining after demand is satisfied. (A negative value of *x* signifies an amount of unsatisfied demand.) The function *g(x)* includes the costs and revenues from holding and/or salvaging any excess stock when $x \geq 0$, and the penalty costs of unsatisfied demand when $x < 0$. Let $C(z; x_0)$ be the expected total cost when *z* units are ordered (and *a(z)* units are delivered) with an initial inventory of x_0. We have

$$C(z;x_0)=\int_0^\infty\left[c(u)+\int_0^\infty g(x_0+u-t)dF(t)\right]dB_z(u)$$

where *F* is the cumulative distribution function of demand.

Now suppose that *g(•)* is convex, and that both c(•) and a(•) are linear, that is $c(u) = cu$ and $a(z) – Az$, where *c* is a given constant and *A* is a random variable with support [0, 1], distribution function *H* and mean μ_a. Incorporating these assumptions, we have

$$C(z;x_0)=\int_0^1\left[cza+\int_0^\infty g(x_0+az-t)dF(t)\right]dH(a) \qquad (1)$$

Now $C(z; x_0)$ is convex on $z \geq 0$, and an optimal policy may be found by examining first order conditions. Since eqn. (1) differs from the perfect replenishment version only in the expectation with respect to the distribution function H, optimal policies are found with essentially the same techniques. It is normally standard practice in inventory theory when assuming a linear ordering cost to transform variables so that the inventory position (stock on hand plus on order), rather than the order size, is of interest (Veinott 1966). This procedure is not applicable here because the random variable A premultiplies z, but not x_0 in the function g. Hence, optimal policies will generally not be of the traditional base-stock or (s, S) forms.

For the commonly considered special case of linear holding and shortage costs, namely

$$g(x) = h \cdot (x)^+ + p \cdot (-x)^+$$

expression (1) becomes

$$C(z; x_0) = (\mu - x_0)p + (c - p)\mu_a z + (h + p)\int_0^1 \int_0^{x_0 + az} (x_0 + az - t)\,dF(t)\,dH(a) \qquad (2)$$

where p is the mean of the demand distribution. As in the standard newsboy problem, linear revenues from satisfying demand and/or salvaging excess stock are easily incorporated by adjusting the values of h and p (Heyman and Sobel 1984, p. 21). Using the same general approach, we make an additional simplification in notation by setting $c = 0$ without loss of generality. Under this convention, the model is easily adjusted to situations with $c > 0$ by transforming the holding and shortage cost parameters. All policy-dependent terms of $C(z; x_0)$ and all expressions for optimal policies are recaptured if c is added to h and subtracted from p.

Defining the dimensionless parameter $\pi = p/h$, we observe that when $C(z; x_0)$ is given by eqn. (2), the optimal order size $z^*(x_0)$ has the simple form:

$$z^*(x_0) = \begin{cases} 0, & \text{if } F(x_0) \geqslant \pi/(1+\pi) \\ z_0, & \text{otherwise} \end{cases}$$

The value of z_0 is given by the smallest z satisfying

$$\int_0^1 aF(az + x_0)\,dH(a) \geqslant \mu_a \pi/(1+\pi) \qquad (3)$$

which always has a solution when $F(x_0) < \pi/(1 + \pi)$. If demand is continuously distributed, eqn. (3) will be solved as an equality.

This optimal policy specification differs from the more complicated form of eqn. (19) in Shih (1980), which can be obtained from eqn. (3) by (*a*) assuming that demand is continuously distributed, *(b)* assuming that $x_0 = 0$, and (*c*) integrating by parts. Expression (3) is especially interesting in its given form because it is a generalization of the standard newsboy result (Wagner 1975, p. 804):

$$F(z + x_0) \geqslant \pi/(1+\pi) \qquad (4)$$

Notice that eqn. (3) reduces to eqn. (4) if the replenishment random variable A equals one with probability one.

It is of interest to consider the form of the optimal policy when a setup cost K is added to the ordering cost function. In this case we modify eqn. (2) by adding a constant term K when $z > 0$. Then it will not be optimal to order unless the setup cost can be offset by savings in other cost components. That is, we must have:

$$F(x_0) \leqslant \pi/(1+\pi) \qquad (5)$$

and

$$C(0; x_0) \geqslant C(z_0; x_0) + K \qquad (6)$$

where C is given by eqn. (2) and z_0 by eqn. (3). Let s be the largest value of x_0 that satisfies eqns. (5) and (6). If s exists, then the optimal order size is given by

$$z^*(x_0) = \begin{cases} 0, & \text{if } x_0 \geq s \\ z_0, & \text{if } x_0 < s \end{cases}$$

This policy is the random-replenishment analogue of the optimal (s, S) policies in perfect replenishment settings (Kardin 1958).

For t he remainder of this paper, we simplify the exposition by assuming that the initial inventory $x_0 = 0$. Hence, the optimal order size z* will be the same as the optimal stock level after ordering and before demand. Although one can generalize the ensuing results for $x_0 > 0$, they will become more complicated in form. We do not believe, however, that the thrust of our general conclusions is affected by the assumption.

3. Uniform demand—analytic results

We find additional structure by considering the ease of uniformly distributed demand with support [0, b]. As discussed in §2, we analyse the model with both *c* and x_0 set to zero. Recall that $\pi = p/h$ is the ratio of unit shortage cost to unit holding cost. Then, if the replenishment quantity were perfect, the optimal order size would be given by

$$\tilde{z} = b\pi/(1+\pi) \qquad (7)$$

When the replenishment quantity is random, the optimal order size is found by solving eqn. (3), which, for uniformly distributed demand is given by.

$$\int_0^1 a(az^*/b)\,dH(a) = \mu_a \pi/(1+\pi)$$

Therefore

$$z^* = [b\pi/(1+\pi)][\mu_a/(\sigma_a^2+\mu_a^2)] = \tilde{z}[\mu_a/(\sigma_a^2+\mu_a^2)] \qquad (8)$$

where σ_a^2 is the variance of the random variable A. Notice that z^* does not depend upon the entire distribution *H*, only its first two moments. This is not necessarily the case when $x_0 > 0$.

Finally it is interesting to examine the properties of a heuristic policy that ignores the variability of the replenishment quantity and merely attempts to correct for the average level of replenishment:

$$\bar{z} = z/\mu_a \qquad (9)$$

Notice that $\bar{z} \leq z^* \leq \bar{z}$, and that z^* converges to $\bar{z}$ as σ^2_a approaches zero.

Analytic expressions for the expected total cost of each policy can be developed by substituting eqns. (7)-(9) for *z* in eqn. (2). After considerable algebraic manipulation, we find the following expressions in terms of the dimensionless parameters π and $\alpha = \sigma^2_a / \mu^2_a$.

$$C(\tilde{z}) = (pb/2)\left[1 - \left(\frac{\pi}{1+\pi}\right)(2\mu_a - \mu_a^2 + \sigma_a^2)\right]$$

$$C(z^*) = (pb/2)\left[1 - \left(\frac{\pi}{1+\pi}\right)\left(\frac{1}{1+\alpha}\right)\right]$$

$$C(\bar{z}) = (pb/2)\left[1 - \left(\frac{\pi}{1+\pi}\right)(1-\alpha)\right]$$

A great deal of information can be gleaned from these expressions, but unfortunately they are not valid for all interesting values of the model parameters. In particular, the optimal policy expressions [z* and *C*(z*)] are valid only when

$$\mu_a + \sigma^2_a / \mu_a \geq / (1 + \pi) \qquad (10)$$

and the expression for $C(\bar{z})$ is valid only when $\bar{z} \leq b$. Another expression for $C(\bar{z})$ was derived for $\bar{z} > b$. It is considerably more complicated and is not reproduced here.

Table 1 display's the cost performance of the policies $\bar{z}$, z* and $\bar{z}$ for a typical set of parameter values, which were chosen to satisfy eqn. (10). The replenishment distributions were assumed to be uniform over the upper-range of replenishment quantities. That is, the random variable *A* was assumed to be distributed uniformly on the interval [(2μ_a — 1). 1]. Therefore the variance of the replenishment distribution is given by $\sigma^2_a = (1 - \mu_a)^2/3$. Note that the perfect replenishment policy can be costly in the random environment. The apparently naive policy does quite well in all cases except when μ_a, the mean fraction of the order that is actually delivered, is very small.

		Optimal cost	% above optimal cost	
π	μ_a	$C(z^*)$	$C(\bar{z})$	$C(\bar{z})$
2	1/2	4·0	11	8
2	5/8	3·2	13	2
2	3/4	2·9	9	0
2	7/8	2·7	3	0
3	3/4	3·3	13	0
3	7/8	3·1	4	0
5	7/8	3·4	7	0
7	7/8	3·7	9	0

Table 1. Cost performance of uniform-demand policies [$b=8$, $h=1$, $\sigma_a^2=(1-\mu_a)^2/3$].

4. Other demand distributions—computational results

The uniform demand ease is interesting because it provides closed form solutions for the optimal policy and expected costs. A useful model, however, must produce results for more realistic demand distributions and a wider range of parameter values. We momentarily depart from the assumption that *a(z) = Az*. and examine a model with integer-valued demand following the negative binomial distribution.

Consider a random replenishment process that orders from an infinite source that has a known fraction *x* of defective items. Then the number of items delivered *a(z)* out of an order of size *z* follows the binomial distribution, that is

$$P\{a(z) = \mu\} = \binom{z}{u} r^{\,-a}(1 - r)^a, u = 0,1,\ldots,z. \qquad (11)$$

Next suppose that the fraction *r is* unknown, and is replaced by a random variable *R.* For the sake of mathematical convenience, suppose that *R* follows a beta distribution having density

$$F_R(r) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} r^{\alpha-1}(1-r)^{\beta-1}, 0 \leq r \leq 1 \qquad (12)$$

Although this assumption is made merely to facilitate the computation of results, it is not very limiting to the model. The beta distribution encompasses a wide variety of density functions on [0, 1]. Using eqn. (11) with (12), we find that the number of items delivered, *a(z)*, has the probability distribution

$$P\{a(z) = u\} = \binom{z}{u} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \frac{\Gamma(z+u+\alpha)\Gamma(u+\beta)}{\Gamma(z+\alpha+\beta)}$$

We report sample calculations for the special case of $\alpha = \beta = 1$, which results in the discrete uniform replenishment distribution

$$P\{a(z)= u\}=1/(1+z),\ u = 0,1,2,\ldots, z$$

Therefore, if we choose in this special case to formulate the replenishment process as $a(z) = Az$, the random variable A has the uniform distribution

$$P\{A = a\} = 1/(1 + z),\ a = 0,\ 1/z,\ 2/z,\ \ldots,\ (z - 1)/z,\ 1 \qquad (13)$$

This distribution is utilized in eqn. (2) (the c.d. f. H) to compute the expected cost of an ordering policy.

We retain our assumption (discussed in §2) that $x_0 = c = 0$, and evaluate the expected cost of three heuristic ordering policies. The first is the newsboy policy given by the general version of eqn. (7)

$$\bar{z} = \min \{u : F(u) \geq \pi/(1 + \pi)\}$$

This policy would be optimal if the amount delivered always equalled the amount ordered. The second policy is the heuristic given by

$$\bar{z} = \bar{z}/\mu_a$$

which corrects for the expected value of the amount delivered but not for its variability. The third policy is inspired by eqn. (8) and given by

$$z^{+} = \bar{z}\mu_a /(\sigma^2_a + \mu^2_a)$$

As shown in §3, this policy is optimal when demand is uniformly distributed and eqn. (10) holds. We choose to examine this policy because it would appear to be a good approximation for the optimal policy z^*. This conclusion is based on the frequently observed phenomenon that optimal inventory policies tend to be insensitive to the form of the demand distribution (beyond its mean and variance) when an expected-cost optimality criterion is specified. It would be an attractive formula because of its closed form and the fact that only the first two moments of the replenishment distribution (rather than the entire distribution) are needed.

We compute the expected costs of the heuristic policies by substituting eqn. (13) for the distribution H in eqn. (2). The optimal policy z^* is also found by evaluating eqn. (2) in the neighbourhood of the heuristics until the minimum expected cost is found.

We illustrate the numerical results in Table 2 for 24 items having negative binomial demand. The results are for a full factorial design with mean demands of 2, 4, 8 and 16, demand variance-to-mean ratios of 3 and 9, and unit shortage cost ratios of 4, 9 and 24. Replenishment quantity distributions are uniform with a mean of one-half the amount ordered, which should be a difficult test of the less sophisticated approximation $\bar{z}$.

Notice that the perfect replenishment policy $\bar{z}$, gives very poor performance, up to 75%, above optimal cost . The best heuristic is clearly $\bar{z}$, with costs typically within a few percent of optimal. Although the supposedly more sophisticated heuristic z^{+} performs nearly as well for low shortage costs, it is clearly inferior for high shortage costs. It appears that the fact that $\bar{z} > z^{+}$ gives $\bar{z}$ an added cushion against high penalties when shortages are costly.

σ^2/μ	μ	π	z^*	$C(z^*)$	% Above optimal cost		
					$\bar{z}$	z'	$\tilde{z}$
3	2	4	6	5·0	9·6	0·5	0·0
		9	10	8·4	14·7	1·4	0·0
		24	16	14·8	31·8	7·4	1·1
	4	4	11	8·0	14·1	1·9	0·4
		9	17	13·4	18·7	1·7	0·3
		24	27	23·8	43·8	12·1	2·5
	8	4	21	13·4	16·6	1·2	1·6
		9	30	22·5	28·6	4·4	0·0
		24	48	40·5	57·2	19·4	4·8
	16	4	39	23·5	26·9	3·0	0·5
		9	56	40·1	43·0	9·8	0·8
		24	88	72·9	74·9	28·1	9·1
9	2	4	5	6·8	1·6	0·0	0·8
		9	11	12·0	5·9	0·6	0·3
		24	22	21·6	14·2	2·1	0·0
	4	4	11	11·0	4·1	0·0	1·4
		9	20	18·6	10·6	1·0	0·3
		24	36	32·8	19·5	3·5	0·0
	8	4	22	17·6	9·4	0·5	1·0
		9	36	29·2	16·7	2·0	0·2
		24	59	51·5	31·0	7·0	0·4
	16	4	43	28·8	15·7	1·6	0·9
		9	63	47·9	25·4	4·4	0·0
		24	101	85·5	45·0	13·0	2·0

Table 2. Cost performance of four policies with negative binomial demand.

5. Conclusions

We have analysed a single-period, random demand inventory model under the assumption that the replenishment quantity is a random function of the amount ordered. Simple generalizations of base-stock policies and (s, S) policies are optimal when a linear ordering cost function is assumed either without or with a setup cost. We have found closed form analytic expressions for optimal policies and their costs for the special case of linear ordering cost and uniformly distributed demand.

The most significant result is that a very simple heuristic policy, one that accounts for the expected value of the replenishment quantity but not its variability, performs quite well for both uniformly distributed demand and negative binomial demand. Although a common approach in practice, that is the first evidence of its effectiveness in a random demand environment. The result is especially interesting in light of results obtained by Sepehri *et al.* (1986) showing poor performance of this heuristic in a multi-period model deterministic demand.

References

HEYMAN, D.P., and SOBEL, M.J., 1984. Stochastic Models in Operations Research. Vol. 11 (New York: McGraw-Hill)
KARLIN, S., 1958. One stage models with uncertainty, Studies in the Mathematical Theory of Inventory and Production. Arrow et al. (editors) (Stanford, California: Stanford University Press).
NOORI, A. H., and KELLER, G., 1986. One-period order quantity strategy with uncertain match between the amount received and quantity requisitioned. INFOR. 24.
SEPEHRI, M., SILVER, E.A., and NEW, C., 1986. A heuristic for multiple lot sizing for an order under variable yield. IIE Transactions. 18, 63-69.
SHIH, W.,1980. Optimal inventory policies when stockouts result from defective products. International Journal of Production Research. 18, 677-685.
SILVER, E.A., 1976. Establishing the order quantity when the amount received is uncertain. INFOR. 14, 32-39
VEINOTT, A.F., JR., 1966. The status of mathematical inventory theory. Management Science. 12, 745-777.
WAGNER, H.M., 1975. Principles of Operations Research, 2nd edition (Englewood Cliffs, N.J.: Prentice Hall).