MANAGEMENT SCIENCE
Vol. 30, No. 5, May 1984
Printed in U.S.A.

A REVISION OF THE POWER APPROXIMATION FOR COMPUTING (s, S) POLICIES*

RICHARD EHRHARDT AND CHARLES MOSIER
School of Business and Economics, University of North Carolina, Greensboro, North Carolina 27412
School of Management, Clarkson College, Potsdam, New York 13676

A revision of the Power Approximation for computing (s, S) inventory policies is presented. The revision incorporates modifications which (1) ensure homgeneity in the units chosen to measure demand and (2) ensure the proper limiting behavior of $S - s$ when the variance of demand is extremely small. Computational experience shows that the revision has operating characteristics that are typically within a few percent of optimal, which is nearly as accurate as the original Power Approximation.
(INVENTORY/PRODUCTION—APPROXIMATIONS; INVENTORY/PRODUCTION—PERIODIC REVIEW MODELS)

1. Introduction

The Power Approximation of Ehrhardt (1979) is an approximately optimal (s, S) inventory policy that has been shown to be very accurate over a wide range of parameter settings. The policy has two flaws, however, concerning its behavior when the units chosen to measure demand are varied and/or when the demand variance is extremely small. We report a revision of the policy that corrects the flaws using methods similar to those used to derive the original Power Approximation.

As in Ehrhardt (1979), we consider a single-item inventory system where unfilled demand is backlogged, there is a fixed lead time L between placement and delivery of an order, and demands during review periods are independently and identically distributed, having a mean μ and variance σ^2. Replenishment costs are composed of a setup cost K and a unit cost c. At the end of each review period, a cost h or p is incurred for each unit on hand or backlogged, respectively. The criterion of optimality is minimization of the undiscounted expected cost per period over an infinite horizon.

2. The Power Approximation

Under the assumptions, an (s, S) policy is optimal (Iglehart 1963); whenever the inventory position y (on hand plus on order minus backlog) is less than or equal to s, an order of size $S - y$ is placed. The Power Approximation is a computationally simple algorithm for computing approximately optimal values of s and S. It requires for demand information only the mean and variance and typically performs within a few percent of optimal total cost for a wide variety of demand distributions and parameter settings. At the core of the algorithm are equations (14) through (16) of Ehrhardt (1979):

$$D_p = 1.463\mu^{0.364}(K/h)^{0.498}\sigma_L^{0.138}, \quad \text{and} \tag{1}$$

$$s_p = \mu_L + \sigma_L^{0.832}(\sigma^2/\mu)^{0.187}(0.220/z + 1.142 - 2.866z), \qquad \text{where} \tag{2}$$

$$z = \{D_p/[(1 + p/h)\sigma_L]\}^{0.5}, \quad \mu_L = (L+1)\mu \quad \text{and} \quad \sigma_L^2 = (L+1)\sigma^2. \tag{3}$$

*Accepted by Warren H. Hausman, former Departmental Editor; received August 12, 1981. This paper has been with the authors 4 months for 1 revision.

0025-1909/84/3005/0618$01.25

In most situations the Power Approximation sets $s = s_p$ and $S = s_p + D_p$.

Expressions (1), (2), and (3) were derived by adjusting the approximations of Roberts (1962). He showed that the optimal policy parameters s^* and $D^* = S^* - s^*$ are given by

$$D^* = \sqrt{2K\mu/h} + o(D^*) \quad \text{and} \tag{4}$$

$$\int_{s^*}^{\infty} (x - s^*)\, d\Phi(x; L + 1) = D^*/(1 + p/h) + o(D^*), \tag{5}$$

where $\Phi(\cdot\,; n)$ is the n-fold convolution of the demand cumulative distribution function and $o(D^*)/D^*$ converges to zero as D^* becomes infinite. Expressions (4) and (5) were used to design regression models which were fitted to a grid of 288 inventory items with known optimal policies. Therefore, (1), (2), and (3) are the result of a numerical fine-tuning of (4) and (5).

The grid of 288 parameter settings is given in Table 1 of Ehrhardt (1979). The same grid is used in deriving the revised Power Approximation. Three types of demand distributions are used: Poisson, and negative binomial with variance-to-mean ratios of 3 and 9. Each demand distribution is given four mean values, 2, 4, 8, and 16. Three values, 0, 2, and 4, are assigned to lead time. Since the cost function is linear in the parameters K, p, and h, the value of the unit holding cost is a redundant parameter which is set at unity. The unit penalty costs are 4, 9, 24, and 99, and the setup cost values are 32 and 64. The unit replenishment cost c is unspecified because it does not affect the computation of an optimal policy for an undiscounted, infinite horizon. All combinations of these parameter settings are included in the grid, yielding 288 items.

An optimal policy is computed for each of the items utilizing the algorithm of Veinott and Wagner (1965). The optimal policies are used as data for the dependent variables of our regressions.

3. Reasons and Methods for Revision

As noted in Ehrhardt (1979, p. 786), it is likely that the accuracy of the Power Approximation will suffer when the variance of demand is very small, especially if K is large. The problem is that expression (1) for D_p vanishes as σ^2 approaches zero. The accuracy of (1) should not be seriously affected unless the demand variance is extremely small, since its exponent in (1) is only 0.069. In fact, (1) should be quite accurate for most, if not all, realistic values of demand variance. It is possible, however, that extremely small variances will occasionally arise when (1) is applied in a statistical environment where the demand moments are periodically estimated from a limited history of realized demands. In such a setting, an unusually low variance estimate in (1) could lead to a poor value of D_p. The proper low-variance limiting behavior is assured by designing a regression model of the form

$$\hat{D}_p = a\mu^{\alpha}(K/h)^{\beta}\left(1 + \sigma_L^2/\mu^2\right)^{\gamma}, \tag{6}$$

where a, α, β, and γ are constants to be fitted to the optimal policy data.

A second deficiency of the Power Approximation arises when the units of demand are changed. If demand units are rescaled by a factor, say f, then s_p and D_p should be transformed similarly. That is, if $\mu' = f\mu$ and $\sigma' = f\sigma$, then we should have $D' = fD$ and $s' = fs$. Notice that the Power Approximation (1), (2), and (3) does not have this property.

We remedy the situation by constraining the regressions for D and s. In (6) we simply set $\alpha = 1 - \beta$. Then, if $\mu = f\mu$ and $\sigma = f\sigma$, we have $K = K$ and $h = h/f$, ensuring that $D = fD$. The regression for s is easily modified as well. We use the model

$$\hat{s}_p = a_0 \mu_L + \sigma_L(a_1/z + a_2 + a_3 z), \tag{7}$$

where z is given by

$$z = \left[D_p/(\sigma_L p/h) \right]^{1/2}, \tag{8}$$

and D_p is determined by first fitting (6) to the optimal policy data. Since z is dimensionless, $\hat{s}_p$ is homogeneous in demand units. Notice that (8) differs from (3). This is because (8) was found to yield the best fit among several candidate forms for z. See Mosier (1981) for a detailed account of the fitting procedure.

4. The Revised Power Approximation

We fit the regression models (6) and (7) to the grid of 288 parameter settings. The following expressions, which replace (1), (2), and (3), yield an excellent fit to the data:

$$D_p = 1.30\mu^{0.494}(K/h)^{0.506}\left(1 + \sigma_L^2/\mu^2\right)^{0.116}, \tag{9}$$

$$z = \left[D_p/(\sigma_L p/h) \right]^{1/2}, \quad \text{and} \tag{10}$$

$$s_p = 0.973\mu_L + \sigma_L(0.183/z + 1.063 - 2.192z). \tag{11}$$

When D_p/μ is sufficiently small, say less than 1.5, the empirical modification of Wagner, O'Hagan, and Lundh (1965) is applied in the same manner as for the original Power Approximation (Ehrhardt 1979).

5. Policy Performance

We proceed with an analysis of the performance of the revised Power Approximation. We show that the policy given by (9)–(11) performs nearly as well as the original policy [(1)–(3)].

Consider the grid of 288 parameter settings used to derive the policies (see §2 of this note or Table 1 of Ehrhardt 1979). Let C and C^* denote the expected total cost per period for an item when using a Power Approximation and an optimal policy, respectively. Our measure of performance for a single item is $\Delta = 100\%\ (C - C^*)/C^*$, namely, the percentage by which the Power Approximation exceeds the optimal total cost. Our results for the 288 parameter settings of §2 are summarized in Table 1, which

TABLE 1

Error Frequencies and Cumulative Frequency (%) of Errors for Original 288 Parameter Settings

	Original Policy		Revised Policy	
Δ	Number of Items	Cumulative Percentage of Items	Number of Items	Cumulative Percentage of Items
[0.0, 0.1)	151	52%	118	41%
[0.1, 0.5)	102	87.8%	108	78.5%
[0.5, 1.0)	21	95%	25	87.2%
[1.0, 2.0)	11	99%	18	93.4%
[2.0, 3.0)	3	100%	11	97.2%
[3.0, 4.0)	0	100%	4	98.6%
[4.0, 5.0)	0	100%	3	99.7%
[5.0, 6.0)	0	100%	1	100%

TABLE 2

Single Parameter Extrapolations
Base Case: Negative Binomial Demand ($\sigma^2/\mu = 5$),
$\mu = 9$, $L = 2$, $p/h = 49$, $K/h = 48$

Extrapolated Value	Cost Accuracy, Δ	
	Original Policy	Revised Policy
$\sigma^2/\mu = 20$	0.0%	1.06%
$\mu = 20$	0.10%	0.01%
30	0.21%	0.14%
40	0.18%	0.16%
$K = 20$	0.11%	0.00%
15	0.28%	0.16%
9	0.63%	0.43%
$p = 132$	0.15%	0.02%
199	0.50%	0.18%
$L = 10$	0.02%	0.03%

lists the number of items having Δ in various ranges. Notice that the revised Power Approximation performs nearly as well as the original policy. The average value of Δ is 0.35% for the original Power Approximation and 0.47% for the revised policy.

We have also compared the policies using parameter settings that are interpolated and/or extrapolated from the 288-item grid. Table 2 lists the parameter settings and cost performance of each policy when parameter settings are chosen in the same manner as in Table 3 of Ehrhardt (1979). A base case is chosen for comparison. The parameter settings of the base case are near the midpoints of the ranges used in the 288-item grid (negative binomial demand, $\sigma^2/\mu = 5$, $\mu = 9$, $L = 2$, $p/h = 49$, and $K/h = 48$). Notice that all the parameter settings except for the lead time L are different from those used in the 288-item grid. We see that in all cases the revised policy yields low costs that are typical of the original policy. In fact, the revised policy outperforms the original policy for all extrapolations with the exception of the variance-to-mean value of 20. The results are particularly impressive when one considers that the extreme values of Table 2 differ from those in the 288-item grid by more than a factor of two.

Recall that the Power Approximation was originally designed as a refinement of a Normal Approximation (Ehrhardt 1979), which basically amounted to using a Normal demand distribution in (4) and (5). It is only natural to ask how the revised Power Approximation compares with the Normal Approximation. Since the revised policy performs nearly as well as the original Power Approximation, it compares with the Normal Approximation in much the same manner as the original Power Approximation. That is, it is slightly better than the Normal Approximation in most cases, and much better from a worst-case point of view. Notice in Table 1 that the revised Power Approximation has 4 items with expected cost between 3% and 4% above optimal, 3 items in the 4% to 5% range, and 1 item in the 5% to 6% range. In contrast, the Normal Approximation has 8 items in the 3% to 4% range, 3 in the 4% to 5% range, and 14 items with expected cost greater than 5% above optimal. The largest cost error for the Normal Approximation is 15% above optimal. If the policies were compared in a setting where demand parameters are statistically estimated, we expect that the results would be very similar to those in Table 4 of Ehrhardt (1979). That is, the revised

Power Approximation would be significantly superior to the Normal Approximation when penalty costs are large, lead times are large, and/or mean demands are small.

Ehrhardt and Wagner (1982) describe how the original Power Approximation can be generalized to systems with nonstationary demand, correlated demand, or stochastic lead times. We expect that the revised policy can be applied to these systems as well with only a modest degradation in total cost performance.

References

EHRHARDT, R., "The Power Approximation for Computing (s, S) Inventory Policies," *Management Sci.*, 25 (1979), 777–786.

——— AND H. M. WAGNER, "Inventory Models and Practice," *Advanced Techniques in the Practice of Operations Research*, American Elsevier, New York, 1982.

IGLEHART, D., "Optimality of (s, S) Policies in the Infinite Horizon Dynamic Inventory Problem," *Management Sci.*, 9 (1963), 259–267.

MOSIER, C., "Revised (s, S) Power Approximation," ONR Technical Report 18, School of Business Administration, University of North Carolina at Chapel Hill, February 1981, 18 pp.

ROBERTS, D., "Approximations to Optimal Policies in a Dynamic Inventory Model," *Studies in Applied Probability and Management Science*, by K. Arrow, S. Karlin, and H. Scarf (Eds.), Stanford University Press, Stanford, California, 1962, 207–229.

VEINOTT, A. F., JR. AND H. M. WAGNER, "Computing Optimal (s, S) Inventory Policies," *Management Sci.*, 11 (1965), 525–552.

WAGNER, H. M., M. O'HAGAN AND B. LUNDH, "An Empirical Study of Exactly and Approximately Optimal Inventory Policies," *Management Sci.*, 11 (1965), 690–723.