MANAGEMENT SCIENCE
Vol. 25, No. 8, August 1979
Printed in U.S.A.

THE POWER APPROXIMATION FOR COMPUTING (s, S) INVENTORY POLICIES*

RICHARD EHRHARDT†

In this paper we present a new analytic approximation for computing (s, S) policies for single items under periodic review with a set-up cost, linear holding and shortage costs, fixed replenishment lead time, and backlogging of unfilled demand. The approximation formulae are derived by using existing results of asymptotic renewal theory to characterize the behavior of the optimal policy numbers as functions of the model parameters. These functions are then used to construct regressions with coefficients that are calibrated by using a grid of 288 known optimal policies as data. The resulting Power Approximation policies (formulae) are easy to compute and require for demand information only the mean and variance of demand over lead time. Extensive computational results show that the approximations yield expected total costs that typically are well within one percent of optimal. The approximation's robustness is exemplified by analyzing its performance when statistical estimates are used in place of the actual mean and variance of demand.
(INVENTORY/PRODUCTION—APPROXIMATIONS)

1. Introduction

We consider a periodic review, 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 independent and identically distributed, having mean μ and variance σ^2. Replenishment costs are comprised 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.

Under these assumptions an (s, S) policy is optimal (Iglehart [1]). That is, whenever inventory on hand plus on order y is less than or equal to s, an order of size $S - y$ is placed. Iterative methods for computing optimal policies are available (Veinott and Wagner [7]), but unfortunately, the computational effort required is prohibitive for practical implementation. Furthermore, the computation of an optimal policy requires the complete specification of the demand distribution, and this level of demand information is particularly unrealistic in practical settings. Most managers would be very fortunate if they had accurate knowledge of only the first two moments of the demand distribution.

A model of this system was analyzed by Roberts [5], who used a renewal-theoretic approach to derive approximations for the optimal values of s and S when the parameters K and p are large. The approximations are easy to compute, but they require specifying the form of the demand distribution. Wagner [9, p. 833] modified the Roberts' approximations by (i) substituting a normal distribution for the actual demand distribution, and (ii) incorporating an empirical modification for small values of setup cost. The resulting approximation, henceforth referred to as the Normal Approximation, requires specifying only the first two moments of the demand distribution. MacCormick [4] has shown that the Normal Approximation yields expected costs that are typically within a few percent of optimality. Costs in excess of 10% above optimal, however, are quite common when the Normal Approximation is

* Accepted by Bennett J. Fox; received March 29, 1978. This paper has been with the author 3 months for 1 revision.
† University of North Carolina at Chapel Hill.

0025-1909/79/2508/0777$01.25

used with certain parameter settings (large penalty cost p, large demand variance σ^2, very large or small demand mean μ).

In this paper we present an alternative to the Normal Approximation called the Power Approximation. The normality specification is dropped, and regression analysis is used to fit approximations of Roberts' form to a large grid of known optimal policies. The result is an easily-computed policy, requiring only the mean and variance of demamd, that is an excellent approximation to optimality for the range of parameter settings that we have examined.

In §2 we describe our methodology for improving upon Roberts' analytic results. In §3 we present the formulae used in executing the Power Approximation, and in §4 we analyze its performance. In §5 we consider aggregate costs in a hypothetical multi-item system. The Power Approximation is compared with the Normal Approximation when μ and σ^2 must be estimated from a limited historical sample of demands. Finally we draw conclusions in §6.

2. Methodology

In [5] Roberts used asymptotic renewal theory to characterize the limiting behavior of an optimal policy as K and p grow large. He obtained the following expressions for optimal policy parameters s^* and $D^* \equiv S^* - s^*$, as D^* grows large.

$$D^* = \sqrt{2K\mu/h} + o(D^*), \tag{1}$$

$$\int_{s^*}^{\infty} (x - s^*)\, d\Phi(x; L+1) = D^*/(1 + p/h) + o(D^*), \tag{2}$$

where $\Phi(\cdot\,; n)$ is the cumulative distribution function of the n-fold convolution of demand, and $o(D^*)/D^*$ converges to zero as D^* becomes infinite. Furthermore, when the demand distribution can be standardized, we can define $\Psi(\cdot)$ as the standardized distribution function, that is

$$\Phi(x; L+1) = \Psi\left\{\left[x - (L+1)\mu\right]/\left[\sigma\sqrt{L+1}\,\right]\right\},$$

and then equation (2) becomes

$$F(u) \equiv \int_{u}^{\infty} (x - u)\, d\Psi(x) = D^*/\left[(1 + p/h)\sigma\sqrt{L+1}\,\right] + o(D^*), \tag{3}$$

where $u \equiv [s^* - (L+1)\mu]/[\sigma\sqrt{L+1}\,]$.

Equations (1) and (3) motivate considering the following formulae for computing an approximately optimal policy

$$D = \sqrt{2K\mu/h}\,, \tag{4}$$

$$S = (L+1)\,\mu + \sigma\sqrt{L+1}\; G\left\{D/\left[(1 + p/h)\sigma\sqrt{L+1}\,\right]\right\}, \tag{5}$$

where $G(x) \equiv F^{-1}(x)$. The difficulty with this approach lies in the computation of the function $G(\cdot)$. One method [9, p. 833] is to assume a normal distribution for demand. The function $G(\cdot)$ can then be computed using an iterative procedure or a rational function approximation. Our approach is to use numerical analysis to fit a power series to $G(\cdot)$ using known optimal policies as data. In fact, we go one step further and use regression to adjust several parameters in functions having the general form of (4) and (5).

Optimal Policy Data

Before discussing the specific regression models, we present the data used in this study. A grid of 288 inventory items has been specified to generate data for the analysis; Table 1 lists the parameter settings. 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.

TABLE 1

System Parameters

Factor	Levels	Number of Levels
Demand distribution	Poisson ($\sigma^2/\mu = 1$)	3
	Negative Binomial ($\sigma^2/\mu = 3$)	
	Negative Binomial ($\sigma^2/\mu = 9$)	
Mean demand (μ)	2, 4, 8, 16	4
Replenishment lead time (L)	0, 2, 4	3
Replenishment setup cost (K)	32, 64	2
Unit penalty cost (p)	4, 9, 24, 99	4
Unit holding cost (h)	1	1

We have calculated an optimal policy for each of the 288 items using the algorithm of Veinott and Wagner [7]. The resulting sets of 288 values for s and S are the data utilized for regression fits. The optimal policies have expected values of period-end inventory ranging from 4 to 82, period-end backlogs from 0.006 to 0.9, backlog frequencies between 0.009 and 0.19, and ordering frequencies ranging from 0.08 to 0.39.

An Approximation for D^*

We first construct a regression model for D^*, the optimal value of $D = S - s$. We generalize expression (4) to the multiplicative form

$$\overline{D} = C\mu^{\alpha}(K/h)^{\beta}(L+1)^{\gamma}(\sigma)^{\delta}(p/h)^{\epsilon}, \tag{6}$$

where C, α, β, γ, δ, and ϵ are constants to be fitted. The variables μ and K/h appear in (6) with the same form as in (4). The remaining variables of our inventory model do not appear in (4), so we arbitrarily include them in (6) as simple multiplicative factors. We use $(L + 1)$ instead of just L in (6) because this is the way that leadtime appears in analytic expressions for expected cost [5].

We form a linear model by taking the logarithm of (6) and use least-squares regression to fit the model to our 288 values for D^*. Then we examine the results and refit a refined model as follows. The variable p/h is removed because the fitted value of ϵ is so close to zero that variations in p/h do not affect the value of $\overline{D}$ when rounded to the nearest integer. Also, we discard five outliers which are identified by a visual examination of plotted data. Finally, after refitting and consolidating similar terms, we obtain the approximation for D^*

$$D_p = 1.463\mu^{0.364}(K/h)^{0.498}\left[(L+1)\sigma^2\right]^{0.0691} \tag{7}$$

Note that K/h has nearly the same exponent as in the Wilson lot-size, namely 0.5. The constant 1.463 is also close to its Wilson lot-size counterpart, 1.414. The exponent of mean demand μ, however, is significantly lower than the Wilson value of 0.5, and we have new variables L and σ^2 in our expression. The regression fit (7) has a coefficient of determination R^2 equal to 0.98. We note, however, that in this study R^2 is merely a measure of how close our regression curve is fitted to the set of 288 deterministic data points. Its sole use is to screen fits to these points, and any conclusions about the robustness of the approximation must be based on the cost data in §§4 and 5.

An Approximation for s^*

We now address the problem of using (5) to construct a regression model for an approximation to the optimal reorder point s^*. Actually, at this point in the analysis we are no longer interested in an approximation for s^*, but rather, an expression for the best s to use in conjunction with D_p above. We generated a new set of data for this analysis by calculating for each of the 288 items a value of s that minimizes expected total cost per period when (7) is used to set $(S - s)$. We call this value s_p^*. This approach is consistent with expression (5), which was derived by computing the partial derivative of expected total cost with respect to s and setting it equal to zero. Hence, expression (5) is an approximation for the best value of s when using a particular value of D. (The best values of s to use in conjunction with given values of D are generated as intermediate results when using the Veinott-Wagner algorithm [7]. Values of s_p are therefore obtained with the same software package used to compute optimal policies.)

The first step in fitting (5) to our data is to choose a functional form for $G(\cdot)$. Let

$$y = D_p / \left[(1 + p/h)\sigma\sqrt{L+1} \right],$$

where D_p is given in (7). For each of the 288 items we compute a value for y and for

$$u_p^* = \left[s_p^* - (L+1)\mu \right] / \left[\sigma\sqrt{L+1} \right].$$

We seek to fit expression (5) to these data by using the model

$$u_p^* = G(y) + \epsilon,$$

where ϵ is an error term. We considered the following set of functions as a representation of $G(\cdot)$

$$G^{n,m}(y) = \sum_{i=n}^{m} A_i(n,m) y^{i/2},$$

where the $A_i(n,m)$ are coefficients to be determined by regression for each specification of n and m; we examined 24 models defined by the range $n = 0, -1, -2$ and $m = 1, 2, \ldots, 8$. The best fit resulted from setting $n = -1$ and $m = 1$; that is,

$$u_p^* = A_1/\sqrt{y} + A_2 + A_3\sqrt{y} + \epsilon$$

was the simplest model among those yielding the highest correlation with the data.

We now substitute $G^{-1,1}(\cdot)$ for $G(\cdot)$ in (5) to obtain a function for an approximately optimal s.

$$\bar{s} = (L+1)\mu + \sigma\sqrt{L+1}\left(A_1/\sqrt{y} + A_2 + A_3\sqrt{y} \right). \tag{8}$$

At this point we use (8) to specify a sequence of regression analyses. First we regress for the constants A_1, A_2, and A_3; then we adjust the multiplier $\sigma\sqrt{L+1}$, and finally we readjust the constants A_1, A_2, and A_3. A detailed discussion of the procedure follows.

Expression (8) motivates the regression model

$$s_p^* = C_0 + C_1[(L+1)\mu] + C_2[\sigma\sqrt{L+1}/z] + C_3[\sigma\sqrt{L+1}] + C_4[\sigma\sqrt{L+1}\,z] + \epsilon, \tag{9}$$

where $z = \sqrt{y}$ and ϵ is an error term. Let $a_0, a_1, \ldots, a_4$ be the values of C_0, $C_1, \ldots, C_4$ resulting from the least-squares regression of (9), yielding an approximately optimal s of the form

$$\bar{s}_1 = a_0 + a_1(L+1)\mu + \sigma\sqrt{L+1}\,(a_2/z + a_3 + a_4 z). \tag{10}$$

A detailed examination of the data reveals that the fit can be improved by adjusting the factor $\sigma\sqrt{L+1}$. The adjustment is performed by replacing the factor with a variable f and solving (10) for the value of f that would yield $\bar{s}_1 = s_p^*$, namely

$$f = [s_p^* - a_0 - a_1(L+1)\mu]/(a_2/z + a_3 + a_4 z).$$

We then seek an approximate expression for f of the form

$$f = \omega C (L+1)^{\alpha} \mu^{\beta} (\sigma^2/\mu)^{\gamma} (p/h)^{\delta} (K/h)^{\epsilon},$$

where ω is an error term, and $C, \alpha, \beta, \gamma, \delta$, and ϵ are constants. The expression is converted to a linear model by taking its logarithm, and least-squares regression is used to set the constants. After discarding insignificant variables and grouping similar terms we obtain our best fit

$$\bar{f} \propto [(L+1)\mu]^{0.416} (\sigma^2/\mu)^{0.603} \equiv q.$$

Expression (9) is then modified to the model

$$s_p^* = C_0 + C_1[(L+1)\mu] + C_2[q/z] + C_3[q] + C_4[qz] + \epsilon. \tag{11}$$

We use least-squares regression to fit (11) to our data, and, after discarding outliers and neglecting insignificant terms [as was done to obtain (7)], we refit to obtain the Power Approximation expression for s

$$s_p = (L+1)\mu + [(L+1)\mu]^{0.416} (\sigma^2/\mu)^{0.603} [0.220/z + 1.142 - 2.866z]. \tag{12}$$

The final regression fit has a coefficient of determination R^2 equal to 0.999. Recall, however, that in this study R^2 is merely a measure of how close our regression curve is fitted to the set of 288 deterministic data points. Its sole use is to screen fits to these points, and any conclusions about the robustness of the approximation must be based on the cost data in §§4 and 5.

3. The Power Approximation

Expressions (7) and (12) yield an approximation (s_p, S_p) for the optimal (s, S) policy. Recall that the approximation is based on theory which assumes large values for the parameters K and p. Wagner, O'Hagan and Lundh [8] observe that the Wilson lot size is a reasonably good approximation for D^* when K/h is large relative to μ, but that it does not approach zero as rapidly as D^* when K/h is allowed to become relatively small. Therefore, they devise an empirical modification to Roberts' policy which we adopt here for the Power Approximation as well.

When D_p/μ is sufficiently small, say less than 1.5, S_p is compared with a single critical number which would be approximately optimal if K were equal to zero. The smaller of these two numbers is then used as S in the policy, thereby reducing the separation between S and s. The single critical number we use is one which would be optimal if demand followed a normal distribution and K were equal to zero. Define S_0 as

$$S_0 = (L+1)\mu + v\sigma\sqrt{L+1}\ , \tag{13}$$

where v is the solution to

$$\int_{-\infty}^{v} \exp(-x^2/2)/\sqrt{2\pi}\ dx = p/(p+h).$$

The Power Approximation is defined as follows. Let $\mu_L = (L+1)\mu$ and $\sigma_L = \sigma\sqrt{L+1}$. Compute

$$D_p = 1.463\mu^{0.364}(K/h)^{0.498}\sigma_L^{0.138}, \tag{14}$$

$$z = \{D_p/[(1+p/h)\sigma_L]\}^{0.5}, \tag{15}$$

and

$$s_p = \mu_L + \sigma_L^{0.832}(\sigma^2/\mu)^{0.187}(0.220/z + 1.142 - 2.866z). \tag{16}$$

If D_p/μ is greater than 1.5, let $s = s_p$ and $S = s_p + D_p$. Otherwise, let $s = \text{minimum}\{s_p, S_0\}$ and $S = \text{minimum}\{s_p + D_p, S_0\}$. If demands are integer-valued, s_p, D_p, and S_0 are rounded to the nearest integer.

4. Policy Performance: Known Mean and Variance of Demand

We proceed with a thorough analysis of the performance of the Power Approximation. In this section we assume that we have accurate values for the mean and variance of demand. In the next section we assess the approximations when the mean and variance are estimated from a limited history of demand values. We show here that the policy as given by (13)–(16) performs very well when several parameter values are simultaneously interpolated or extrapolated from settings used to derive the policy. We also show that the Power Approximation [(13)–(16)] performs well when demands are derived from different probability distributions than are given in Table 1.

We first examine the set of 288 items that are used to derive the Power Approximation (see Table 1). Let C_p and C^* be the expected total cost per period for an item when controlled using the Power Approximation and the optimal policy, respectively. Our performance measure for a single item is

$$\Delta_p = 100\%(C_p - C^*)/C^*,$$

namely, the percentage by which the Power Approximation cost exceeds the optimal cost. Our results for the 288-item system of Table 1 are summarized in Table 2, which lists the number of items in the system having values of Δ_p in various ranges. Note that the Power Approximation yields expected total costs within 0.1% of optimality for 52% of the items and within 0.5% for 84% of the items. The average value of Δ_p in the system is 0.3%, with the largest values found for items having the largest unit penalty costs.

We next examine a 288-item system with parameter settings other than those used in deriving the Power Approximation. The system is based on the parameter settings of

TABLE 2

Frequencies of Δ_p for a 288-Item System

Range for Δ_p	Number of Items	Cumulative Percentage of Items
[0%, 0.1%)	151	52%
[0.1%, 0.5%)	102	84%
[0.5%, 1.0%)	21	95%
[1.0%, 2.0%)	11	99%
[2.0%, 3.0%)	3	100%

Table 1 with values of mean demand μ changed to 3, 7, 11, and 15, and values of unit penalty cost p changed to 3, 9, 27, and 81. In this system, the Power Approximation yields an average of 0.6% above optimal total cost per period, with individual items ranging up to 4.7% above optimality. The magnitude and distribution of values for Δ_p in this system are quite similar to those found in the original 288-item system. We conclude that the Power Approximation is very accurate over the range of parameter values spanned by Table 1.

We also consider extreme extrapolations of parameter values beyond the ranges used in deriving the Power Approximation. A single item with interpolated parameter settings is used as a base case (negative binomial demand, $\sigma^2/\mu = 5$, $\mu = 9$, $L = 2$, $h = 1$, $p = 49$, and $K = 48$). Table 3 lists the parameter settings and the resulting values of Δ_p for each extrapolation. We see that in all cases the Power Approximation yields total costs within 0.63% of optimal. The results are particularly impressive when one considers that the extreme parameter values of Table 3 differ from those of Table 1 by more than a factor of two.

TABLE 3

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

Extrapolated Value	Δ_p
$\sigma^2/\mu = 20$	0.0%
$\mu = 20$	0.10%
30	0.21%
40	0.18%
$K = 20$	0.11%
15	0.28%
9	0.63%
$p = 132$	0.15%
199	0.50%
$L = 10$	0.02%

We have seen that the Power Approximation is accurate for parameter values that are different from those used to derive the policy rule. We also consider the performance of the Power Approximation when it is used to control items having *demand distributions* different from those used to derive the policy, especially ones with considerable skew. (Of course, the Poisson and negative binomial distributions are themselves skewed to the right.) Consider two new systems, one having exponential demand distributions, and one with bimodal (compound negative binomial) demand distributions. In each case we compute optimal and Power Approximation policies [expressions (13) through (16)], and compare expected total costs per period.

We examine a 75-item system with exponentially-distributed demands. Replenishment lead time is zero for all 75 items since this is required by the analytic expression we use to evaluate costs for exponential demand distributions. The other parameters have a wide range of values, in most cases extending beyond those used in the 288-item system of Table 1. The values of mean demand are 1, 2, 4, 8 and 16; setup cost values are 16, 32 and 64; unit penalty costs are 9, 49, 99, 132 and 199. Unit holding cost is normalized at 1. Power Approximation costs are found to be very close to optimal. Values of Δ_p are less than 0.2% for 23 items, between 0.2% and 0.6% for 27 items, between 0.6% and 1.0% for 20 items, and between 1.0% and 1.5% for the remaining 5 items.

We examine a 72-item system with bimodal (compound negative binomial) demand distributions. Demand is taken to be zero with probability 0.25, and it is sampled from a negative binomial distribution with probability 0.75. Mean demand is set at 2, 4, 8 and 16, with a variance-to-mean ratio of 9. The other parameter settings are 0, 2 and 4 for lead time, 32 and 64 for setup cost, and 4, 9 and 99 for unit penalty cost. Unit holding cost is normalized at 1. In this system the Power Approximation has an average expected total cost per period of only 1.2% above optimal, a surprisingly robust result.

Thus, the Power Approximation [(13)–(16)] performs well for a broader family of demand distributions than those used in setting the coefficients.

5. Policy Performance: Estimated Mean and Variance of Demand

We have presented a policy which provides an excellent approximation to optimality when the mean and variance of demand are accurately specified. In a typical applied setting, however, demand parameters are estimated from a limited history of observed past demands. In this situation the concept of an optimal policy is not well defined. We suggest using an (s, S) policy that is computed by substituting estimates of the demand mean and variance in place of the actual mean and variance in expressions (13)–(16). We next analyze the performance of the Power Approximation, and also compare it with the Normal Approximation, when demand estimates are used in place of actual values. Each policy is also compared with the optimal policy which could be computed if the demand distribution were completely specified.

Specifically, we assume that a history of n demands is used in setting the policy to be employed over the subsequent n periods, and that during this interval of $2n$ periods the demand distribution parameters remain unchanged. (In other words, we assume it is warranted to use the past n observed demands to estimate the mean and variance of demand for the next n periods.) The demand history is used to calculate a sample mean and sample variance which are substituted in place of μ and σ^2 in expressions (13)–(16) to compute a Power Approximation policy (or in the Normal Approximation formulae to compute a Normal Approximation policy).

The mathematical complexity underlying this procedure of policy determination necessitates our using a simulation program to evaluate policy performance. Specifically, we make 200 replications of this policy computation for each group of parameter settings: we examine a 72-item system with negative binomial demand distributions having a large variance-to-mean ratio of 9. Values of mean demand are 2, 4, 8 and 16, setup costs are 32 and 64, unit penalty costs values are 4, 9 and 99, and lead time is set at 0, 2 and 4. Demand history length values n are 13, 26 and 52 (corresponding to a quarter, half, and full year of weekly data). The simulation program computes estimates of expected cost components for each item in the system, and aggregates the costs to produce estimates of system-wide performance.

Table 4 shows estimates of percentage increases in expected total cost per period when statistical control is compared with optimal control given full information. The Power Approximation yields lower total costs than the Normal Approximation for all the systems in Table 4, with the most noticeable differences occuring for small demand history length, high unit penalty cost, and low mean demand. Notice that the Power Approximation yields total costs within 1% of the Normal Approximation for unit penalty costs of 4 and 9. The Power Approximation yields significantly lower costs, however, for the subsystems with a unit penalty cost of 99. The difference reaches a maximum of 14% for the system with a 13-period demand history. Percentage differences are even larger for subsystems of items having two or more troublesome parameter settings. For example, the subsystem of six items having a revision history length of 13 periods, $p = 99$ and $\mu = 2$ yields for the Normal Approximation costs 54% above optimal with full information as compared to 35% for the Power Approximation, that is, a 1.14-fold increase.

TABLE 4

Percentages above Optimal Full Information Total Costs for a 72-Item System Under Statistical Control

Decision Rule and Demand History Length	Total Aggregate Cost	Costs Aggregated by Parameter Value											
		Penalty Cost			Setup Cost		Lead Time			Mean Demand			
		4	9	99	32	64	0	2	4	2	4	8	16
Power Approximation													
13 Periods	20	9	12	30	21	19	9	19	28	24	23	18	18
26 Periods	12	5	7	18	12	11	6	11	16	18	14	10	9
52 Periods	6	3	4	10	7	6	3	6	9	10	7	6	5
Normal Approximation													
13 Periods	27	10	12	44	28	25	11	26	38	34	31	26	22
26 Periods	17	6	7	28	19	15	8	15	24	27	20	15	12
52 Periods	10	4	4	17	11	9	6	9	13	18	13	9	6

Several conclusions emerge from the data in Table 4. The Power Approximation is clearly preferable to the Normal Approximation, especially when unit penalty costs are high. Good performance is evidenced when a moderately large demand history is used to estimate demand parameters, but significant degradation over optimal costs is to be expected for small demand histories. Thus, in the context of Table 4, with a year's worth of weekly demand history and a penalty cost not exceeding 9, the *combined* impact of using statistical information and an approximation for the policy computations may yield less than 5% degradation in total costs over having complete information. These conclusions are supported by similar data for systems with Poisson demand distributions and negative binomial distributions with $\sigma^2/\mu = 3$. The data are presented in [2], along with detailed information about the components of total cost and other operating characteristics such as backlog frequency.

6. Concluding Remarks

We have presented an approximately-optimal policy which (1) is easily computed, (2) limits the required demand information to only the mean and variance, and (3) provides an excellent approximation to optimality for a wide range of parameter settings. The approximation is robust when model assumptions are relaxed, including the assumption of accurately specified mean and variance of demand. Additional results, reported by Kaufman [3] and Schultz [6], indicate that the approximation can be generalized for systems with demand distributions that exhibit nonstationarity or correlation from period to period.

Although the Power Approximation provides excellent performance for a great diversity of inventory items, we expect that performance may degenerate if several model assumptions are allowed to vary extensively. In these instances we suggest that a new numerical analysis similar to that presented above would be successful. In particular, we caution against use of the Power Approximation when the variance of demand is very small. Notice in expression (14) that D_p vanishes when σ goes to zero, an effect which is particularly undesirable when K is large. For this situation we suggest a revised analysis for D_p, utilizing a functional form which has a finite limit as σ goes to zero.

Finally we observe that this study has shown that numerical analysis can be successfully employed to calibrate an analytically-derived approximation. It is a methodology that can provide the insights and simplicity of analytic approximations while still coming very close to the performance of exactly optimal results. We suggest that this technique can be used in other areas of applied optimization. For example, the inventory problem with discounted costs has been largely ignored in the applied literature. Analytic approximations for this problem appear in [5], and could be used as a starting point for numerical analysis.[1]

[1] This paper is based on a Ph.D. dissertation written at Yale University. The author is pleased to acknowledge the invaluable guidance and encouragement of his advisor, Professor Harvey M. Wagner. He also acknowledges George Kastner and John Klincewicz who aided in the computer calculations, and the helpful comments of the referees and the Departmental Editor. The research was supported in part by Office of Naval Research grants N00014-75-C-2041 and N00014-78-C0467, and by U. S. Army Research Office grant DAHC04-75-G-0079.

References

1. IGLEHART, D., "Optimality of (s, S) Policies in the Infinite Horizon Dynamic Inventory Problem," *Management Sci.*, Vol. 9 (1963), pp. 259–267.
2. EHRHARDT, R., "(s, S) Inventory Policies with Limited Demand Information," Ph.D. dissertation, Yale University, 1976.
3. KAUFMAN, R., "(s, S) Inventory Policies in a Nonstationary Demand Environment," School of Business Administration and Curriculum in Operations Research and Systems Analysis Technical Report #11, University of North Carolina at Chapel Hill, April 1977.
4. MACCORMICK, A., ESTEY, A. S. AND KAUFMAN, R. L., "Inventory Control with Statistical Demand Information," Operations Research and Systems Analysis Technical Report #77-8, University of North Carolina at Chapel Hill, July 1977.
5. ROBERTS, D., "Approximations to Optimal Policies in a Dynamic Inventory Model," *Studies in Applied Probability and Management Science*, edited by K. Arrow, S. Karlin and H. Scarf, Stanford Univ. Press, Stanford, Calif., 1962.
6. SCHULTZ, C. R., "(s, S) Inventory Policies for a Wholesale Warehouse Inventory System," School of Business Administration and Curriculum in Operations Research and Systems Analysis Technical Report #14, April 1979.
7. VEINOTT, A. AND WAGNER, H., "Computing Optimal (s, S) Inventory Policies," *Management Sci.*, Vol. 11 (1965), pp. 525–552.
8. WAGNER, H., O'HAGAN, M. AND LUNDH, B., "An Empirical Study of Exactly and Approximately Optimal Inventory Policies," *Management Sci.*, Vol. 11 (1965), pp. 690–723.
9. ———, *Principles of Operations Research*, 2nd Ed., Prentice-Hall, Englewood Cliffs, N.J., 1975.