ARMY MATERIEL SYSTEMS ANALYSIS ACTIVITY
AMSAA

AMSAA

AD-A150 812

3

AD-
TECHNICAL REPORT
USAIRO-TR-84/5

THE BAYESIAN INVENTORY PROBLEM

INVENTORY RESEARCH OFFICE

May 1984

Alan J. Kaplan

US Army Inventory Research Office
US Army Materiel Systems Analysis Activity
800 Custom House, 2d & Chestnut Sts.
Philadelphia, PA 19106

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER	2. GOVT ACCESSION NO.
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle) THE BAYESIAN INVENTORY PROBLEM	5. TYPE OF REPORT & PERIOD COVERED Technical Report
	6. PERFORMING ORG. REPORT NUMBER TR-84/5
7. AUTHOR(s) ALAN J. KAPLAN	8. CONTRACT OR GRANT NUMBER(s)
9. PERFORMING ORGANIZATION NAME AND ADDRESS US Army Inventory Research Office Room 800, US Custom House, 2nd & Chestnut Sts. Philadelphia, PA 19106-2976	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS US Army Materiel Development & Readiness Command 5001 Eisenhower Avenue Alexandria, VA 22333	12. REPORT DATE May 1984
	13. NUMBER OF PAGES 44
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) US Army Materiel Systems Analysis Activity Aberdeen Proving Ground, MD 21005-5071	15. SECURITY CLASS. (of this report) UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for Public Release; Distribution Unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Dynamic Programming,
Bayes
Inventory Models,

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Our objective is to improve the management of repair parts for newly fielded weapon systems. The initial procurement of each part is made a lead time before the system is fielded, and is based on an engineering estimate of the mean part demand rate per fielded system. There may be additional procurements before the fielding date. Once the system is fielded, demand experience accrues and is used to update the forecast of the demand rate, improving its accuracy.

DD FORM 1473, 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

lower case s

ABSTRACT (CONT)

Inventories are managed under a periodic review, (s,S) policy: when assets fall to s, order up to S. The period is as small as one week. The issue of concern is how the expected improvement in accuracy of the demand forecast should affect the values of the inventory control parameters.

Formally, we are seeking to determine optimum (s,S) parameters - they change each period - when there is Bayesian updating, periodic review and a dynamic mean with demand randomly distributed about the mean.

An algorithm is programmed, sample results obtained, and conclusions drawn. Originator-supplied key words included. I front p.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

ACKNOWLEDGEMENT

The US Army Materiel Systems Analysis Activity (AMSAA) recognizes the following individual for contributing to this report:

PEER REVIEWER: Dr. Gerald Garfinkel.

Dr. Garfinkel's review significantly improved the readability of this manuscript.

CONTENTS

Next page is blank

THE BAYESIAN INVENTORY PROBLEM

Chapter 1. INTRODUCTION

1.1 Background.

Our objective is to improve the management of repair parts for newly fielded weapon systems. The initial procurement of each part is made a lead time before the system is fielded, and is based on an engineering estimate of the mean part demand rate per fielded system. There may be additional procurements before the fielding date. Once the system is fielded, demand experience accrues and is used to update the forecast of the demand rate, improving its accuracy.

Inventories are managed under a periodic review, (s,S) policy: when assets fall to s, order up to S. The period is as small as one week. The issue of concern is how the expected improvement in accuracy of the demand forecast should affect the values of the inventory control parameters.

Formally, we are seeking to determine optimum (s,S) parameters - they change each period - when there is Bayesian updating, periodic review and a dynamic mean with demand randomly distributed about the mean. It has long been known that this problem can be solved by dynamic programming, Cf. [11]. However, in the general case a multi-dimensional state vector is required, and the dynamic programming formulation has not been pursued. Thus Kaplan and Kruse unsuccessfully solicited interest in designing a computationally feasible algorithm [9]. In fact we have not even found useful qualitative conclusions about the Bayesian solution in the general case.

Instead, the literature has focused on the specific situation in which there is no cost to order so that the purchase price is simply (Unit Price) x (amount ordered). In this case several interesting results have been obtained:

a. Optimum solutions are well behaved in that if "t" is a sufficient statistic for the demand experienced to date, the optimum policy is to order each period up to the optimum S value for that period, and for a given period S is a non-decreasing function of "t".

b. If the distribution of demand, conditional on "t", has certain properties, the optimum solution may be found by solving a dynamic program with only a single dimensional state vector. For example, if the conditional distribution is Weibull:

$$f(x;w) = w \cdot k \cdot (xw)^{k-1} \exp[-xw]^k \quad x > 0$$

and the prior on w given "t" is Gamma, then only a single dimensional state vector is required.

These results were first obtained by Scarf [11,12] and are reviewed and extended by Azoury [1].

Efforts have also been made to determine how Bayesian solutions relate to non-Bayesian solutions. Thus, Scarf looked at how the Bayes solution asymptotically approached the non-Bayes solution with increasing experience, (again with no cost to order), and found that the Bayes S values could approach from above or below the non-Bayes S values. However, Azoury and Miller do show, for an n period model for repairable items, that the Bayes S is always less than or equal to the non-Bayes S [2]. Again, they assume no cost to order.

The purpose of this research is to explore the feasibility of the dynamic programming formulation as a basis for a computer algorithm, and to obtain some insight into the impact and significance of uncertainty about the demand rate.

1.2 Model.

A difficulty arose in formulating a model for newly fielded weapon systems,i.e. for the "provisioning" phase of repair part support, in that no rigorous model for managing inventories after provisioning, when the demand forecast stops improving, really exists! Few if any real world inventory systems forecast demand as the average of all the historical demand ever received; rather, some form of moving average, exponential smoothing or related technique is used. Thus, the forecast cannot even asymptotically approach the true mean so long as demand variability does not diminish; yet, there is no practical model we are familiar with which rigorously deals with the fact that the demand forecast is neither the true mean nor approaches it. Students at the University of North Carolina under Harvey Wagner have been interested in the problem and Ehrhardt has developed a "power" approximation which is more robust to the use of estimated means and variances than the more rigorous computations which assume the mean and variance are known [5]. Some promising recent work of Miller [10] attacks a related problem theoretically - he requires the forecast, derived by exponential smoothing, to be the true mean, but allows this mean to change stochastically each period.

Developing a rigorous post-provisioning model would be a major undertaking in its own right. As an alternative, we assume that at some future period, period "M", updating of the inventory control parameters ceases. While unrealistic, this assumption captures the reality that even after provisioning, the (s,S) parameters will not reflect knowledge of the true mean.

The dynamic programming formulation can quite easily handle non-stationarities; these may be due to changes in cost and related parameters such as an obsolescence rate, which is high only until weapon design stabilizes, or non-stationarity may reflect changes in demand per period as the number of fielded weapon systems increases. The heuristic used to model expected costs after period "M" cannot handle these dynamic changes and is therefore limiting. We will discuss later how this might be overcome under the "alternative formulation" of the dynamic programming approach.

The standard dynamic programming formulation limits the choice of demand distributions to those for which a sufficient statistic exists for the mean. Moreover, practicality pretty much limits choice of the prior to the conjugate distribution (Appendix A reviews sufficient statistics and conjugate distributions). We discuss an alternative formulation which allows any demand distribution but has these disadvantages: the prior must be expressed as a discrete histogram[1] and it is more time consuming to run on a computer.

The cost structure allowed in the dynamic programming formulation is quite general, but the heuristic used to model costs after period "M" does require a discount rate. For readers unfamiliar with the discount rate concept and its application to treatment of interest rates and obsolescence the paper on multi-year holding costs is recommended [6]. In brief, use of a discount rate is crucial to cost analysis when inventories may be held for many years, as is possible when buys are made based on forecasts which may be much too high. If i is the interest rate, per period, and θ is the probability of obsolescence in a period, the one period discount is $(1-\theta)/(1+i)$, where $1/(1+i)$ is the present value and $(1-\theta)$ is the probability the item can incur demand (it is not obsolete).

---

[1] By this we mean a discrete distribution with a limited number of points with positive probability, and in which it may not be possible to express the probabilities other than by enumeration.

1.3 Organization of Report.

Both the "standard" dynamic programming formulation based on a sufficient statistic and the alternative formulation based on a histogram prior are presented. Issues in moving from formulation to algorithm are discussed. Starting the dynamic programming recursion is the most difficult issue and receives an extensive treatment.

An algorithm based on the standard formulation was programmed and applied to a simple experimental design. Results are presented and conclusions are drawn.

Chapter 2. DYNAMIC PROGRAMMING FORMULATION

2.1 Recursive Equation.

A periodic review inventory policy is assumed. The sequence of events in a period is assumed to be: receive demand, debiting assets accordingly, record holding or backorder costs based on current on hand; receive delivery of inventory ordered a lead time (L) ago; order new inventory if warranted. While the assumed sequence of events within a period is arbitrary, the formulation could be adjusted to accommodate any other assumed sequence.

Inventory ordered at period n does not affect on hand, and therefore backorder and holding costs in the interval $[n,n+L]$. Therefore, in developing an inventory policy for period n, we do not need to include those costs in our cost expression.

Define $C_n(x,t_n)$ as the expected value of the present value of all future costs at period n, omitting backorder and holding costs in $[n,n+L]$. "x" are assets just before ordering, and demand experience is summarized by the sufficient statistic "t". Implicit is the use of a particular inventory policy for deciding what to order in period n and thereafter; e.g. the policy may be to minimize costs.

Then,

$$C_n(x,t_n) = f_1(y-x) + f_2(y|t_n) + a\sum_d C_{n+1}[y-d,f_3(t_n,d)] \cdot f_4(d|t_n) \tag{2.1}$$

$$y \equiv f_5(x,t_n)$$

where:

y: assets after ordering in period n.

$f_1(z)$: cost of ordering z units of inventory.

$f_2(y|t_n)$: expected value of the present value of backorders and holding costs in period (n+L+1) given inventory at the end of the current period of y, and future demand whose distribution is conditional on t_n, as well as on the original prior on the demand distribution.

a: the discount rate which accounts for the time value of money and the possibility of obsolescence.

d: demand in period n+1.

$f_3(t_n,d)$: the function for recomputing the sufficient statistic after one additional period has elapsed and demand of d has occurred. Typically, but not always, t_n is the 2 dimensional vector (n, sum of demands to date) so $f_3(t_n,d) = t_n$ + the vector (1,d). If demand is uniform, t_n would be (n, maximum demand in any period experienced)[4].

$f_4(d|t_n)$: probability function for demand given t_n

$f_5(x;t_n)$: A mapping from (x,t_n) to y reflecting the inventory policy, where y may equal x. If the objective is solely to minimize costs, then y is found as the value which when substituted into the RHS (right hand side) of equation (2.1) minimizes $C_n(x,t_n)$.

Note that to handle non-stationarities, a subscript of n would be appended to f_1, f_2, a, f_3, f_4 where f_{1n} would mean ordering cost parameters are time dependent, f_{2n} would reflect time dependency of backorder, holding or discount parameters, and f_{3n} and f_{4n} would reflect the need to incorporate the number of systems fielded in every period. For the first L periods before any system is fielded, $f_4(d|t_n)$ is always 0 for d > 0.

2.2 Use of Recursive Equation in an Algorithm.

The major problem is how to get $C_{n+1}(\cdot,\cdot)$ for some n in order to start the iterative process by which $C_n(\cdot,\cdot)$ can eventually be determined. Period zero represents the present, at which a decision must be made. As discussed in the introduction, our efforts are devoted to determining $C_M(\cdot,\cdot)$, M being when updating of the inventory control parameters ceases. An entire section will be devoted to this.

The next problem is that there is no natural upper bound on the demand which may be received, so some kind of truncation is necessary. Truncation arises in three distinct contexts:

a. Truncation on demand received in period n + 1 (the d values in the recursive equation).

b. Truncation on total demand in the next (L+1) periods, used in getting $f_2(y|t_n)$.

c. Truncation on the values of t_n. In particular, where t_n is the vector (n,D_n), D_n being total experienced demand in the first n periods, we are referring to truncation on possible values of D_n.

The single period (one week) demand values were truncated at a cumulative probability of $1 - 10^{-4}$; i.e., the value k was found such that:

$$\sum_{d=0}^{k} f_4(d|t_n) \geq 1 - 10^{-4}$$

and then $f_4(k|t_n)$ was reset:

$$f_4(k|t_n) = 1 - \sum_{d=0}^{k} f_4(d|t_n)$$

Lead time demand was truncated at a cumulative probability of $1 - 10^{-5}$. These truncation values were subjectively chosen.

Now let t_n^{max} be the maximum value of the sufficient statistic "t" considered in period n and let d_{n+1}^{max} be the single period demand maximum in period n+1, based on $f_4(d|t_n^{max})$. It would be natural to set $t_{n+1}^{max} = t_n^{max}$ plus the vector $(1, d_{n+1}^{max})$. However, based on a kind of bootstrap mechanism at work, as n increased, t_n^{max} would tend to increase out of control, with the state space, the values of (x, t_n) to be evaluated, getting out of hand. Moreover, the probability of values as large as such a t_n^{max} actually occurring on a real item would get increasingly small as n increased.

Therefore, let $PMEAN_n$ and $PSDEV_n$ be the mean and standard deviation of demand over n periods calculated using the original prior, before any updating. Set:

$$t_n^{max} = [n,\ PMEAN_n + 4 \cdot PSDEV_n]$$

where the value 4 was chosen after some experimentation on the sensitivity of results to this value (see Appendix C). Truncate single period demand as necessary so as not to violate the upper bound on t_{n+1}^{max} (the truncation of lead time demand is not affected.) We are excluding values of t_n^{max} which can occur only with very small probability, although if our initial knowledge about the true mean is poor, $PSDEV_n$ will be large and t_n^{max} will be large. Also, t_n^{max} does increase with n.

One might ask how we limit the values of assets, x, to consider. The maximum x would be the largest order up to value, S, for any t_n^{max}. The current computer program makes a very conservative guess on what this will be, and issues a warning if this guess proves to be too low.

The remaining issue relating to use of the recursion equation is whether to assume any particular form for the inventory policy. It is perfectly feasible to find the optimum y for each (x, t_n) vector, which we will label $y^*(x, t_n)$. Such answers, however, are not nearly as convenient to work with as the optimum values for an (s,S) policy since potentially we need to record a different y^* for each (x, t_n) whereas the (s,S) parameters, by definition, are independent of x.

An interesting question is just what can we say theoretically about the $y^*(x,t_n)$. We assume order cost is of the form $C_P + (UP)(y-x)$ and is zero for $y = x$.

Claim

$$\text{If} \quad x < y^*(o,t_n)$$
$$x < y^*(x,t_n)$$
$$\text{Then} \quad y^*(x,t_n) = y^*(o,t_n)$$

Interpretation

The optimum order-up-to level when assets are 0, $y^*(o,t_n)$, will also be the optimum order-up-to level when assets are x, for all $x < y^*(o,t_n)$ at which it is worthwhile to place an order.

Proof

Let $V_n(x;y,t_n)$ be the value of $C_n(x;t_n)$ if the amount y-x is ordered in period n. For convenience drop the t_n argument.

By inspection of the recursion equation, (2.1), $V_n(x_1;y) = V_n(x_2;y) + (UP)(x_2-x_1)$ for all values of x_1, x_2 such that $y > x_1, x_2$. Hence

$$V_n[o,y^*(x)] - V_n[x, y^*(x)] = (UP)(x)$$
$$V_n[o,y^*(o)] - V_n[x, y^*(o)] = (UP)(x)$$

Also, by definition of $y^*(o)$

$$V_n[o,y^*(x)] - V_n[o,y^*(o)] \geq 0$$

Adding the second and third equation,

$$V_n[o,y^*(x)] - V_n[x,y^*(o)] \geq (UP)(x)$$

Subtracting this from the first equation,

$$V_n[x,y^*(o)] - V_n[x,y^*(x)] \leq 0$$

Q.E.D.

The methodology for determining (s,S) values in the algorithm is therefore:

a. Find $y^*(o,t_n)$ and set $S = y^*(o,t_n)$

b. Find the smallest x such that $y^*(x,t_n) = x$ and set $s = x - 1$. It is easy to check that $y^*(S,t_n) = S$ and thus $s < S$.

Since it has never been shown an (s,S) policy is optimal, there is no assurance there is not some x greater than the s identified for which it would decrease costs to order.

2.3 Alternative Formulation.

If the prior is expressed in the form of a histogram it is always possible to use dynamic programming as an evaluator of some given policy. If, in addition, we restrict the set of policies to those which are based on the total demand experienced to date, and do not consider the period to period pattern of demand, then dynamic programming can be used to find the optimum policy or optimum (s,S) policy. The technique does not depend on any assumption about what the demand distribution is, but has computational disadvantages compared to the sufficient statistic approach.

We will first show how to write the recursion equation, (2.1), in a different way, one which is more readily generalized to our current assumptions. Let $C'_n(y;t_n)$ - note the prime - represent expected future discounted costs when assets are y after the order is placed in period n. From the definition of $f_5(x,t_n)$ if assets after ordering in period n are y, then assets after ordering in period n + 1, given that demand of d occurs, are:

$$f_5[y-d,\ f_3(t_n,d)]$$

Thus,

$$C'_n(y;t_n) = f_2(y|t_n) + a\sum_d f_4(d|t_n)(\text{Expression}) \tag{2.2}$$

$$\text{Expression} = f_1\{f_5[y-d,f_3(t_n,d)] - (y-d)\} + C'_{n+1}\{f_5[y-d,\ f_3[t_n,d)];\ f_3(t_n,d)]\}$$

To generalize this to a histogram prior, let i be the set of possible mean demand rates; let $f_6(i;t_n)$ be the probability of i after observing t_n where t_n is the 2-dimensional vector (n, total demand); and let $C'(y;t_n,i)$ be the cost conditional on what the true mean demand rate is. It is $C'(y;t_n,i)$ we calculate recursively:

$$C'_n(y;t_n,i) = f_2(y;i) + a\sum_d f_4(d;i)(\text{Expression}) \tag{2.3}$$

$$\text{Expression} = f_1\{f_5[y-d,f_3(t_n,d)] - (y-d)\} + C'_{n+1}\{f_5[y-d,\ f_3(t_n,d)];\ f_3(t_n,d),\ i\}$$

For purposes of finding (s,S) values, we set:

$$C'_n(y;t_n) = \sum_i f_6(i;t_n)\ C'_n(y;t_n,i)$$

and use the $C'_n(y;t_n)$;e.g. we set S to y which minimizes $(UP)(y) + C'_n(y;t_n)$

The histogram form of the recursion has the obvious computational disadvantage of requiring values of C'_n for each i. An important advantage is that the number of probability distributions which must be calculated depends on the number of histogram points rather than on the number of values of t_n.

Chapter 3. STARTING THE RECURSION

We wish to determine values of $C_M(x,t_M)$ over the range of (x,t_M), assuming that the (s,S) values are not updated after period M.

Our approach relies on being able to approximate the priors at period M (there is a different prior for each value of t_M) by histograms. Other simplifications were made which while not essential, simplified coding and reduced running times. Before reviewing these, let us make several observations in support of the belief that the details of the approach are not critical.

a. The primary purpose of this research or any implemented algorithm would be to improve the decisions made before or around the time of fielding when uncertainty is greatest. Intuitively, because of discounting, the impact of the values $C_M(\cdot,\cdot)$ on the values $C_n(\cdot,\cdot)$, $M > n$, must decrease as the time between M and n increases.

b. The decisions made at period n do not depend on the absolute values of $C_n(\cdot,\cdot)$, but the relative values, i.e., $C_n(0,t_n) - C_n(x,t_n)$, and therefore depend only on the relative costs at period M. Inspection of the recursion equation shows that the decrease in cost for x assets, relative to 0 assets, for $x < s$, is always (up)(x).

c. The impact of non-optimal (s,S) values at period M on the relative costs associated with different asset levels at period M is likely to be a second order effect. The primary effect would be to overstate somewhat the costs which must be incurred for all values of starting assets, and even then it is frequently found in inventory modelling that costs are flat in the region of optimality.

3.1 Simplification.

We calculate the reorder quantity (S minus s) without regard to its impact on backorders, then calculate s, rather than jointly optimizing s,S. We also assume that when an order is placed, assets are exactly "s"; and we also assume when it is convenient to do so, that backorder and holding costs are recorded whenever the asset level changes, rather than only at the end of the period. With a review period of only one week, these simplifications are not unreasonable, but the assumption that we can order at exactly "s" would be inadequate for items characterized by low frequency of demand, but large requisition sizes [7].

3.2 Notation.

$L(s,S;i)$: expected lifetime costs including backorder costs when the mean demand rate is i, an (s,S) policy is followed and assets of s are initially on hand and zero are due-in.

$B(x;p,i)$:, $H(x;p,i)$: expected backorder and holding costs, respectively over the next p periods when the mean demand rate is i, initial on hand is x, and there will be no additional stock arriving during the p periods.

$f_6(i,t_M)$: as earlier, probability the true demand rate is i; the t_M subscript will be dropped in this section.

$f_7(p;x,i)$: probability mass function on the number of periods to accumulate demand $\geq$ x units, given a mean demand rate of i.

$G(x;i)$: $\sum_p a^p f_7(p;x,i)$ where a is the discount rate

C_p: fixed cost to order

Q: the order quantity, S minus s

Cost(Q;i): expected lifetime ordering and holding costs, but not backorder costs, for initial assets of 0, given i, and assuming Q is ordered whenever assets are zero.

3.3 Determination of (S-s):

$$\underset{Q}{\text{Minimize}} \quad \Sigma \quad f_6(i) \text{ Cost } (Q;i)$$

$$\begin{aligned}\text{Cost } (Q;i) &= C_p + (up)(Q) + [H(Q;\infty,i) - H(Q;L,i)] \\ &\qquad + G(Q;i) \text{ Cost } (Q;i) \\ &= [C_p + (up)\ Q + H\ (Q;\infty,i) - H(Q;L,i)]/[1-G(Q;i)]\end{aligned}$$

In the expression for holding cost, $H(Q;\infty,i)$ is total lifetime holding cost on the Q units assuming instantaneous delivery, from which costs over the first L periods, $H(Q;L,i)$ are subtracted, since the Q units are still on order then.

$G(Q;i)$ is the expected value of the present value of spending one dollar, given that the dollar will be spent after Q units are demanded. If the item survives that long, i.e., it does not become obsolete before then, expected costs from that point on are Cost(Q;i) again.

3.4 Determination of s.

The cost of a policy cannot be defined independently of the current asset position. The s chosen is the smallest s such that it is as good or better to use an (s,s+Q) policy rather than an (s+1, s+Q+1) policy, given assets of s+1. It is assumed for convenience all s+1 assets are on hand, but the cost difference between any two policies is independent of the status of assets - on hand or on order - since policy decisions cannot affect backorder or holding costs for a lead time.

We wish to find the lowest s for which $\Delta(s)$ is > o, where:

$$\Delta(s) = \sum_i f_6(i)\ [L_1(i) - L_2(i)]$$

$$L_1(i) \equiv L(s+1,\ s+Q+1;i) - H(s+1;\infty,i)$$

$$L_2(i) \equiv G(1;i)\ [L(s,s+Q;i) - H(s;\infty,i)]$$

Both $L_1(i)$ and $L_2(i)$ exclude holding cost on the initial s+1 assets. In the case of $L_2(i)$, holding cost until the first demand - at which time assets fall to s - are never charged, so only the remaining costs need to be netted out. The use of G(1;i) discounts back to the initial time when assets were (s+1).

Computation of L(s,s+Q,i).

In the computation of the Cost(Q;i) we saw a renewal process at work in that if the item survived, free of obsolescence, until Q demands were incurred, expected costs from that point equalled expected costs from the original starting point. We can take advantage of the same logic if we exclude holding cost on the initial s assets, defining:

$$L'(s,s+Q;i) = L(s,s+Q;i) - H(s;\infty,i)$$

Just as we will associate with each order of Q the incremental holding costs caused by ordering those Q units, we can associate the incremental backorder cost caused by backorders which will be eliminated by these Q units; every backorder is eventually eliminated by a due-in, so if we associate the backorder and the due-in which eliminates it, we will account for every backorder, but never double count.

Thus:

$$L'(s,s+Q;i) = Cp + (up)\,Q + H\,(Q) + B\,(Q) + G\,(Q;i)\,L'(s,s+Q;i)$$

where $H(Q)$ is the incremental holding cost caused by ordering the Q units and $B(Q)$ is a backorder computation to be described.

Note first that

$$H(Q) = [H(s+Q;\infty,i) - H(s;\infty,i)] - [H(s+Q;L,i) - H(s;L,i)]$$

where the first term in brackets is the added cost assuming instantaneous delivery, and the second term corrects for the lead time.

Also:

$$B(Q) = B(s;L,i) - B(s+Q;L,i)$$

We are netting out those backorders, the $B(s+Q;L,i)$, which would occur even if all Q units were delivered instantaneously. These backorders will be eliminated by some subsequent orders and will be associated with them.

Summarizing

$$L(s,s+Q;i) = L'(s,s+Q;i) + H(s;\infty,i)$$

$$L'(s,s+Q;i) = [C_p + (up)\,Q + H\,(Q) + B\,(Q)]/[1 - G(Q;i)]$$

$$H(Q) = [H(s+Q;\infty,i) - H(s;\infty,i)] - [H(s+Q;L,i) - H(s;L,i)]$$

$$B(Q) = B(s;L,i) - B(s+Q;L,i)$$

3.5 Computation of $C_M(x,t_M)$.

Given the form of the recursion equation, (2.1), all backorder and holding costs in [M, M+L] are accounted for by use of the $f_2(y_1 t_n)$ in periods (M-L) thru (M-1). Therefore, the $C_M(x,t_M)$ represent the expected value of all costs after period M, net of backorder and holding costs in [M+1, M+L]. The computation makes uses of the L(s,S ;i) with appropriate netting out. It is assumed the x assets referenced by $C_M(x,t_M)$ are all on hand, but this is for convenience only since it is only holding and backorder costs in [M,M+L] which are affected by the status of the x units, and the final answer does not reflect these costs.

Claim.

$$C_M(s;t_M) = C_M(s;t_M) + (up)(s-x) \qquad \text{for } x < s \tag{3.5a}$$

$$C_M(s;t_M) = \sum_i f_6(i;t_M) \{L(s,S;i) - B(s;L,i) - H(s;L,i)\} \tag{3.5b}$$

$$C_M(x;t_M) = \sum_i f_6(i;t_M) [T_1(i) + T_2(i)] \text{ for } x > s \tag{3.5c}$$

$$T_1(i) = G(x-s;i) [L(s,S;i) - H(s;\infty,i)]$$

$$T_2(i) = H(x;\infty,i) - H(x;L,i) - B(x;L,i)$$

Explanation. The equations follow almost entirely from the same arguments we have been making throughout this and earlier sections. The logic underlying (a) was used in the proof of the lemma on the properties of the optimum inventory control parameters. Equation (b) is a special case of (c).

T_1 (i) represents the expected value of the present value of all future costs, net of all holding costs on the initial x units. The approach to holding cost is supported by the same logic underlying the derivation of $L_2(i)$ in the sub-section on determination of s. Since the x units are initially on hand, no backorders can occur in the time interval until assets fall to s, so $T_1(i)$ includes all backorder costs. L(s,S;i) discounts them back to the time when assets first fall to s, and G(x-s;i) discounts them further back to period M.

$T_2(i)$ adds back in the holding cost on the initial assets which is incurred after period M+L, and subtracts out backorders in [M,M+L].

3.6 More General Formulations.

Suppose that updating of the demand forecast ceases after period M, but non-stationarity persists in that additional weapon systems are fielded.

We may identify a period M_2, $M_2 > M$, after which we are willing to assume conditions stabilize. If we choose M_2 large enough, it will make little difference what we assume, because discounting will minimize the impact on our current decisions, and we may use a very crude scheme for evaluating the worth of assets at M_2 relative to 0 assets.

Under the histogram approach for periods $M_2 \geq n \geq M$, referring to equation (2.3),

$$C'_n\ (y;t_M,i) = f_2(y;i) + a \sum_d f_4(d;i) \quad \text{(Expression)}$$

$$\text{Expression} = f_1[f_5(y-d,t_M) - (y-d)] + C'_{n+1}[f_5(y-d,t_M);t_M,i]$$

Note that in accordance with our assumption the sufficient statistic is not updated. Because of this it is more feasible to extend the dynamic programming to the additional interval $[M,M_2]$. If t_M must be updated, t_n^{max} will continue to grow as n increases past M (recall discussions on truncation) and the number of states which must be evaluated will continue to grow, although interpolation may be possible (see Appendix C).

Under the sufficient statistic approach (equation 2.2), we must write:

$$C'_n\ (y;t_M,t_n) = f_2(y;t_n) + a \sum_d f_4(d|t_n) \quad \text{(Expression)}$$

$$\text{Expression} = f_1[f_5(y-d,t_M) - y-d] + C'_{n+1}\ [f_5(y-d,t_M);t_M,t_n]$$

In particular, the sufficient statistic must be updated since it determines the distribution of future demands. If we set $f_4(d|t_n) = f_4(d|t_M)$ for all $n > M$, we are ignoring the correlation between demand in successive periods, i.e., if demand in M+1 is large, this makes it more likely demand in M+2 will be large since it is more likely the true mean is higher than the mean implied by $f_4(d|t_M)$.

Chapter 4. MODEL RESULTS

4.1 Inputs and Assumptions.

The model was run for a number of different examples to get a feel for how uncertainty affects the optimum (s,S) values, and what the cost implications of ignoring uncertainty are.

All examples were set in the provisioning context in that there was no demand for the first lead time before fielding. To simplify the interpretation, it was assumed that full deployment occurred at the fielding date, so that once demand began it occurred at a constant mean rate per period. The prior on the mean was assumed to be Gamma, and the distribution of demand was assumed to be Poisson. The mathematics for these specific distributions, including how to update the prior when deployment is increasing, is given in Appendix B.

Other assumptions were that the review period was one week and that updating of the (s,S) values ceased after one year. In all cases the cost to procure was $450 and the discount rate was 80 percent for one year and therefore $.80^{1/52}$ for one week. No holding cost was charged since estimates for US Army wholesale inventories are only 3%, and ignoring holding costs simplified the coding of the algorithm. The discount rate reflects interest and obsolescence costs which are significant.

A simple experimental design was used. A base case was defined as:

Procurement Lead Time in Weeks (L)	52.
Unit Price (UP):	$1400.
Average Yearly Demand, i.e. Mean of the Prior (AYD)	10.

Each of these variables was then varied, one at a time, to two other settings: L of 26 or 78, UP of $350 or $5600, AYD of 5 or 20. For example, one case run was for L = 26, UP = 1400, AYD = 10.

Each case was run for three uncertainty levels as measured by the coefficient of variation (γ): the ratio of the standard deviation of the prior to its mean. This ratio is independent of the period of time to which the prior is applied, whether it be the prior on mean demand for one week or one year (Appendix B). Uncertainty settings used were 10^{-6} (essentially no uncertainty), 0.5 and 1.0.

Each case was also run for at least two different levels of backorder cost. For steady state models without discounting, for which the holding cost includes interest and obsolescence costs, there is a relationship between the degree of protection afforded by the optimum reorder point and the ratio of backorder cost to holding costs, so that it is more informative to report this ratio than the specific value of the backorder cost. In our model holding cost in a sense is the complement of the discount rate, i.e., 20 percent, and all backorder costs are discounted back to the period for which the (s,S) values are being determined. This motivated us to define a backorder ratio (BR) and report this value, where for a discount of 80 percent,

$$[\text{Backorder Cost}]^{.8L/52} = (BR)(UP)(.20)$$

The base case was run for BR values of 4, 10, 100, 500 and all other cases were run with BR values of 10 and 100.

4.2 Findings.

Table 1 shows the optimum (s,Q) values at time of fielding where Q is defined as S minus s. Thus, for the case where lead time differs from the base case and equals 78, BR = 100 and $\gamma = 1/2$, the optimum reorder point (s) is 26 and the optimum reorder quantity (Q) is 7.

Let us first say something about the results for when there is no uncertainty, $\gamma = 0$. There is some tendency for Q values to decrease as BR increases; while not very intuitive, this is characteristic of stationary, undiscounted models with certainty [8]. Otherwise Q's behavior is intuitive, and so is that of s, which increases with BR, L and UP, the last because higher UP results in lower Q's.

Uncertainty had relatively little impact on Q for $\gamma = 0.5$; and when Q values were lower, s values, with one exception, were higher (the exception was for AYD of 5 and BR of 10). For $\gamma = 1$, Q declined in all cases, but showed less sensitivity to BR than under certainty.

The impact of uncertainty on s depended on BR and γ. For lower values of BR, s decreased, while for higher values s increased (for $\gamma = .5$) or the discrepancy between certain and uncertain values narrowed (for $\gamma = 1.0$). For a very high BR of 500, s increased significantly even for $\gamma = 1$.

Intuitively, uncertainty suggests lower values of s in order to hedge until more is known, but at the same time with uncertainty higher s values are generally needed to get the same protection. Backup Table 1 offers insight into how this tradeoff is apparently working to determine the s values.

Reported for each case are the probability that demand in the first (L+1) weeks will exceed the s value found to be optimal at time of fielding, with the probability accounting for uncertainty as to the true mean as well as variance of demand around the mean. With two minor exceptions (L:26, BR:100, γ = .5 and AYD:5, BR:10, γ = .5), less protection was provided by the optimum s as uncertainty increased, even in cases where the optimum s upon which the calculations were based was higher for uncertainty than for certainty.

Table 2 shows the pre-fielding pattern of buys for each case. For example, for the case (L:78, BR:10, γ:0) the first buy was for 6, a second buy was made at period (-49), or 49 weeks before fielding, increasing on-order to 14, and a third buy at period (-13) increased on order to 22. The first buy is always made L weeks before fielding.

In every case for which BR was 100, the optimum amount on order increased monotonically as a function of γ, for all pre-fielding periods. For example, comparing γ = 0 and γ = 1.0, for the base case, the optimum initial order is 12 versus 9, and on-order is increased to 29 at (-34) versus to only 17 at (-25); (-25) is 9 periods after (-34). Results are more mixed at BR = 10, but clearly there is a greater tendency to react to uncertainty by increasing buys before fielding than there is after fielding.

The explanation is that hedging - buying less and waiting for more information - is a more viable strategy after fielding than pre-fielding. After fielding new information becomes available each week, while during pre-fielding nothing is learned until demand begins. To support this theory we show below the optimum s values from period (-52) thru period (-1) for the case (Base, BR:100, γ = 1.0) and we observe that the largest s value of 25 actually occurs before fielding, and then the optimum s start to decline as hedging becomes more viable:

Periods (-52) thru (-24): optimum s increase from 0 to 25

Periods (-24) thru (-14): optimum s stays at 25

Periods (-14) thru (-1): optimum s declines from 25 to 17

Incidentally, the pre-fielding buys for this case were inferred as follows: at (-52) optimum S was 12; s first increased to 12 at period (-34), requiring another buy then.

As further corroboration of the importance of hedging a special run was made of the same (Base, BR:100, γ:1.0) case we have been examining. In this run (s,S) values were fixed at (15,20) for periods 1 thru 13. In the original run for this case (15,20) was optimum for period 0 and subsequent (s,S) values depended on the demand experienced. Now, without the flexibility to revise (s,S) for periods 1 thru 13, the optimum (s,S) values found for period 0 were (28,32).

We have observed that the impact of uncertainty on the optimum (s,S) values depends on a number of factors including the backorder cost, amount of uncertainty, time until fielding. It would be difficult to imagine a simple heuristic which takes all these factors into account, so it is of some interest to determine the cost of using a heuristic which ignores the impact of uncertainty. Such a policy is described in Appendix D.

Table 3 shows the cost increase when this heuristic is used as the pre-fielding policy. Cost increases are expressed as a percent of the total expected lifetime cost under an optimum policy. Since the heuristic is used only through period 0, but the costs reported are lifetime, the costs increases reported tend to minimize the adverse impact of the heuristic. Nevertheless it is clear that ignoring uncertainty may be feasible for lower values of backorder cost but is costly for higher values.

4.3 Summary.

The demand forecasting techniques used for real world inventory problems suggest that the mean demand is not known with certainty, and in fact is subject to change. This paper addresses uncertainty without ever really coming to grips with the possibility of mean changes. Nevertheless, the results should contribute to the formulation of improved policies for the early stages of an item's life, when uncertainty diminishes as the basis of forecast changes from pre-fielding estimates to actual demand experience.

Two variants of a dynamic programming approach of the problem were considered; one requires the existence of a sufficient statistic for the mean, and the other requires that the prior be specified as a histogram. The sufficient statistic approach was explored in depth and results were obtained which provide insight into how uncertainty and learning impact on the optimum inventory control parameters.

Even from the limited range of cases examined, it emerges that the question of whether the proper response to uncertainty is to raise or lower inventory levels has no simple answer. Once learning began, it was found to be worthwhile to tolerate higher probabilities of stockout in the lead time, but this does not always translate into lower "s" values. For higher backorders cost parameters, it became cost effective to raise s to partially offset the greater stockout risk created by uncertainty. Optimum reorder quantities were insensitive to uncertainty at the lower uncertainty level examined, but did drop at the higher level.

During pre-fielding there is uncertainty, but no immediate learning, and the proper response was more likely to be to raise inventory levels, although this also depended on the level of the backorder cost parameters, and how close to fielding the item was. A heuristic based on ignoring uncertainty performed satisfactorily at lower backorder cost parameters, but under-bought seriously for higher settings.

Computer processing times would permit implementation of the algorithm developed, at least for selected items when high speed computers were used for inventory control. The current algorithm is somewhat limiting in its assumptions, and we discussed why a more general model may be more feasibly developed under the alternative dynamic programming formulation based on a histogram prior. Our belief, however, is that the most promising path for the future is to develop heuristics based on the dynamic programming results and validated by sensitivity testing. Hopefully, such heuristics would be less sensitive to correct assumptions about the exact distribution of demand, the exact distribution of the prior, and the exact dynamics of changes in the mean than a more exact model. The success of the Power Approximation provides some precedent for such a hope [5].

TABLE 1. OPTIMUM (s,Q) AT FIELDING

	BR	$\gamma = 0$	$\gamma = 1/2$	$\gamma = 1$
BASE:	4	8/8	5/8	0/5
	10	11/7	9/8	3/5
	100	15/7	18/6	15/5
	500	18/6	24/5	26/5
L-26:	10	5/7	4/7	1/4
	100	8/7	10/5	8/4
L-78:	10	16/8	14/8	4/6
	100	22/7	26/7	22/6
UP-350:	10	9/13	7/13	1/10
	100	14/12	16/11	12/10
UP-5600:	10	12/4	11/4	4/3
	100	16/4	19/3	17/2
AYD-5:	10	5/6	5/5	1/4
	100	9/5	10/4	9/4
AYD-30:	10	22/10	18/10	4/8
	100	28/9	33/9	25/7

TABLE 1 BACKUP. PROBABILITY OF STOCKOUT

	BR	$\gamma = 0$	$\gamma = 1/2$	$\gamma = 1$
BASE:	4	68.8%	77.0%	91.1%
	10	32.5%	47.1%	68.8%
	100	5.6%	9.5%	22.4%
	500	0.9%	2.5%	7.3%
L-26:	10	41.8%	51.2%	70.3%
	100	8.1%	7.8%	20.5%
L-78:	10	35.4%	46.5%	72.7%
	100	3.7%	10.2%	23.1%
UP-350:	10	56.6%	62.3%	82.9%
	100	9.4%	14.3%	29.6%
UP-5600:	10	22.7%	35.2%	62.6%
	100	3.1%	7.7%	18.6%
AYD-5:	10	40.1%	38.5%	69.9%
	100	3.5%	7.3%	16.7%
AYD-30:	10	31.0%	50.2%	78.7%
	100	4.2%	12.2%	28.8%

TABLE 2. OPTIMUM PRE-FIELDING ASSETS

	BR	$\gamma = 0$	$\gamma = 1/2$	$\gamma = 1$
BASE:	4	5; 13(-17);	4; 12(-24);	6;
	10	6; 14(-23);	6; 15(-26);	5; 13(-35);
	100	9; 17(-25);	11; 24(-27);	12; 29(-34);
	500	9; 16(-29); 23(-7)	10; 20(-36);30(-18);	12; 26(-41); 42(-28);
L-26:	10	7;	8;	6;
	100	7; 14(-6);	7; 14(-12);	14;
L-78:	10	6; 14(-49); 22(-13)	8; 21(-42)	7; 19(-53);
	100	9; 17(-51); 25(-20)	10;21(-56);34(-31);	11; 26(-62); 43(-43);
UP-350:	10	12;	14;	12;
	100	16;	23;	28;
UP-5600:	10	4; 9(-34); 14(-13)	4; 9(-37); 15(-20);	4; 12(-37);
	100	5;9(-41);13(-28); 18(-12)	6; 11(-41); 17(-30) 24(-17);	7; 16(-43); 28(-31);
AYD-5:	10	4; 9(-16);	4; 9(-21);	6;
	100	6; 12(-20);	7; 14(-23);	7; 16(-33);
AYD-30:	10	8; 18(-32); 29(-9)	11; 27(-28)	8; 23(-37)
	100	11; 21(-34);32(-14)	13; 27(-33); 43(-21)	14;32(-42);54(-30);

TABLE 3. INCREASE IN COST

	BR	$\gamma = 0$	$\gamma = 1/2$	$\gamma = 1$
BASE:	4		0.1%	1.3%
	10		0.2%	0.3%
	100		1.9%	8.4%
	500		6.3%	48.8%
L-26:	10		0.0%	0.0%
	100		0.4%	1.4%
L-78:	10		0.6%	1.8%
	100		4.8%	21.1%
UP-350:	10		0.1%	0.0%
	100		2.2%	8.1%
UP-5600:	10		0.2%	0.3%
	100		2.5%	10.6%
AYD-5:	10		0.2%	0.0%
	100		0.8%	4.5%
AYD-30:	10		0.4%	1.4%
	100		3.6%	13.8%

Next page is blank

BIBLIOGRAPHY

1 Azoury, Katy S., Bayes Solution to Dynamic Inventory Models Under Unknown Demand Distribution, College of Business Administration, University of Illinois, Chicago, 1981.

2 Azoury, Katy S. and Bruce L. Miller, A Comparison of the Optimal Ordering Levels of Bayesian and Non-Bayesian Inventory Models, School of Business Administration and Economics, California State University, Northridge, 1982.

3 Brown, G. F. and W. F. Rogers, "A Bayesian Approach to Demand Estimation and Inventory Provisioning," Naval Research Logistics Quarterly, Vol 20, 1973, (p607-624).

4 DeGroot, Morris H., Optimal Statistical Decisions, McGraw-Hill, New York, 1970.

5 Ehrhardt, Richard, (s,S) Inventory Policies With Limited Demand Information, PhD Dissertation, Yale University, 1976.

6 Kaplan, Alan, Multi Year Holding Costs for Secondary Items, Inventory Research Office, 1969, AD712315, LD25575.

7 Kaplan, Alan "A Note on Initial Fill Rate," Canadian Journal of Operations Research, Vol 19, February 1981, LD45729MA.

8 Kaplan, Alan, "The Relationship Between Decision Variables and Penalty Cost Parameters in Q,R Inventory Models," Naval Research Logistics Quarterly, Vol 17, 1970, (p253-258).

9 Kaplan, Alan and Karl Kruse, "Dynamic Model of a Continuous Review Inventory System," presented at Fall 1978 TIMS/ORSA National Meeting.

10 Miller, Bruce L., Scarf's State Reduction Method, Flexibility, and a Dependent Demand Inventory Model, Air Force Institute of Scientific Research, 1983, ADA127483, LD55771K.

11 Scarf, Herbert E, "Bayes Solution of the Statistical Inventory Problem," Annals of Mathematical Statistics, Vol 30, 1959, (p490-508), LD08436.

12 Scarf, Herbert E, "Some Remarks on Bayes Solutions to the Inventory Problem," Naval Research Logistics Quarterly, Vol 7, (p591-596).

Next page is blank

APPENDIX A

SUFFICIENT STATISTICS

A convenient reference for most of this material is DeGroot [4].

Suppose that we are sampling from a random variable X with probability function (or probability density function if X is continuous) $f(X;w)$, where w is an unknown parameter. All the sample values we see are governed by the same w, but w is the realization of a random variable W, with distribution $g(W)$, termed the prior.

Let $\tilde{x}$ be the sample values, where $\sim$ is used to denote a vector. In accordance with the Bayesian approach to statistics, it is meaningful to reference the updated prior $g(W|\tilde{x})$, which reflects the changed probabilities of what W might be after taking into account the observed $\tilde{x}$, as well as the original prior $g(W)$.

A statistic $T(\tilde{x})$ is any function defined over all possible values of $\tilde{x}$. A sufficient statistic is defined by the fact that $g(W|T(\tilde{x})) = g(W)|\tilde{x})$ for all possible values of x, where the equality must pertain only to values of W with positive probability. There is always at least one sufficient statistic, namely $T(\tilde{x}) = \tilde{x}$, but we are interested in sufficient statistics of fixed dimensionality; i.e., the dimensionality of $T(\tilde{x})$ does not increase as sample size increases. An example of $T(\tilde{x})$ is the 2-dimensional vector $(n, \Sigma x_i$ where n is the sample size and x_i the sample values.

In the definition of a sufficient statistic the function g denotes a mapping from the sample space to a probability space. Now suppose g can be written as a mathematical formula with $T(\tilde{x})$ and the sample size as parameters; also, suppose that if $T(\tilde{x},\tilde{y})$ is the sufficient statistic after observing a total sample $(\tilde{x},\tilde{y})$ where x is n-dimensional and y is m-dimensional, then

$$g(W|T(x,y),\ n+m) = K\ g\ (W|T(\tilde{x}),\ n)\ g\ (W|T(\tilde{y}),\ m)$$

where K is a normalizing constant chosen so that the right hand side integrated (or summed) over W equals one. The function g is then called a conjugate prior to the distribution $f(X;w)$.

It turns out that whenever there is a fixed dimensional sufficient statistic, there is a conjugate prior. In Appendix B we treat explicitly the Poisson conditional, $f(X;w)$, and the Gamma conjugate.

$f(X;w)$ is said to belong to the exponential family if for any possible values of X,w,

$$f(x;w) = a(w)\ b(x) \exp \sum_{i=1}^{k} g_i(w)\ h_i(x)$$

For any member of the exponential family there is a sufficient statistic of dimension k.

However, while the exponential family is very rich in that it encompasses many distributions, it is possible for the same distribution both to be and not to be part of the exponential family! Let us illustrate with the Negative Binomial which is much used in Army inventory models.

The Negative Binomial may be written:

$$f(x;r,p) = \frac{\Gamma(r+x}{\Gamma(r)\Gamma(x+1)} p^r(1-p)^x = \frac{\Gamma(r+x)}{\Gamma(r)\Gamma(x+1)} p^r \exp[\log(1-p)x]$$

Now if r is known, and p is the unknown parameter, the Negative Binomial satisfies the condition for membership in the exponential family. However, in an inventory context this would be an unlikely situation since r equals the square of the mean divided by the difference between the variance and mean. If, as is more likely, the p parameter is known, p being the reciprocal of the variance to mean ratio, and r is unknown, $f(X;r)$ does not satisfy the condition for membership in the exponential family since $\Gamma(r+x)$ cannot be written as $\exp[\Sigma g_i(r)h_i(x)]$.

APPENDIX B

MATHEMATICS FOR GAMMA-POISSON

<u>Bayesian Updating</u>

A convenient reference for most of this material is Brown and Rogers [3].

Write the Poisson, Gamma and Negative Binomial distributions as:

Poisson: $f(d;\lambda) = \frac{\lambda^{d} e^{-\lambda}}{d!}$

Gamma: $g(\lambda;a,b) = \frac{b^{a}}{\Gamma(a)} \lambda^{a-1} e^{-b\lambda}$

Negative Binomial: $h[d;a,b/(b+1)] = \frac{\Gamma(a+d-1)}{\Gamma(d)\Gamma(a-1)} \left(\frac{b}{b+1}\right)^{a} \left(\frac{1}{b+1}\right)^{d}$

Then the following are the means and variances:

	Mean	Variance
Poisson	λ	λ
Gamma	a/b	a/b^2
Negative Binomial	a/b	$a/b^2 + a/b$

Note the variance of the Negative Binomial is the variance of a Poisson with mean a/b plus uncertainty a/b^2 around the mean.

If demand in a period is Poisson, with unknown λ sampled from a Gamma prior, the distribution of demand allowing for uncertainty about λ is Negative Binomial:

$$h[d;a,b/(b+1)] = \int_{\lambda} g(\lambda;a,b)\ f(d;\lambda)$$

If after n periods demands of $d_1, d_2 \ldots d_n$ are observed, then $(n, \Sigma d_i)$ is a sufficient statistic and the updated prior is a Gamma with revised parameters $a = a_o + \Sigma d_i$, $b_o = b + n$; the distribution of demand in one period, given $(n, \Sigma d_i)$ is Negative Binomial with mean a/b and variance $a/b^2 + a/b$. The distribution over p periods is Negative Binomial with mean ap/b and variance $ap^2/b^2 + ap/b$.

We can redefine λ as the expected demand per weapon system per period. Then if W_i are the number of systems deployed in period i we could use the above results with these substitutions:

$\sum_{i=1}^{n} W_i$ is substituted for n

$\sum_{i=n+1}^{n+p} W_i$ is substituted for p

Computation of Functional Values

Let B_c be backorder cost per backorder per period and H_c be holding cost. Let f_2 be as defined in Section 2.1. Capital letter functions are from Chapter 3 and the f, g and h functions are as defined in this Appendix.

$$f_2(y;t_n) = a^L H_c \sum_{d=o}^{y} h(d|t_n) + a^L B_c \sum_{d=y+1}^{\infty} h(d|t_n)(d-y)$$

$$\sum_{d=y+1}^{\infty} h(d|t_n)(y-d) = \sum_{d=o}^{\infty} h(d|t_n)(y-d) - \sum_{d=o}^{y} h(d|t_n)(y-d)$$

$$= y - \text{Expected Value } (d|t_n) - \sum_{d=o}^{y} h(d|t_n)(y-d)$$

$$B(x;P,i) = B_c \sum_{p=1}^{P} \sum_{d=x+1}^{\infty} a^p \; f(d;i_p)(d-x) \qquad i_p = (i)(p)$$

$$H(x;P,i) = H_c \sum_{p=1}^{P} \sum_{d=o}^{x} a^p f(d;i_p)(x-d)$$

For $G(x;i)$ and $H(x;\infty,i)$ we assume costs are recorded continuously with a discount of $e^{-\alpha t}$ where t is the elapsed time in periods and $e^{-\alpha} = a$. Thus, for p periods, the discount is $e^{-\alpha p} = (e^{-\alpha})^p = a^p$.

Now if demand is Poisson with mean of i per period, the time to order x units is Gamma with mean x/i and variance[2] x/i^2 from which we may infer that the Gamma parameters are a = x and b = i. Therefore,[3]

$$G(x;i) = \int_o^{\infty} e^{-\alpha t} \; g(t;x,i)dt = \int_o^{\infty} \frac{i^x}{\Gamma(x)} \; t^{x-1} \; e^{-(\alpha+i)t} \; dt$$

$$= \left(\frac{i}{\alpha+i}\right)^x \int_o^{\infty} \frac{(\alpha+i)^x}{\Gamma(x)} \; t^{x-1} \; e^{-(\alpha+i)t} \; dt = \left(\frac{i}{\alpha+i}\right)^x$$

since the last integral is the integral of a Gamma density over its range.

Now suppose we numbered (conceptually) x assets from 1 to x with

---

[2]This well known result follows from the fact that Poisson demand implies exponential time between demand, and the sum of independent exponentials is Gamma.

[3]The use of the Gamma as shown here was suggested by Karl Kruse [9]

asset 1 issued first and so on. Suppose the k^{th} asset is issued at time t. Total holding costs for the k^{th} asset would be:

$$H_c \int_o^t e^{-\alpha j}\, dj = \frac{H_c}{\alpha} \, [1 - e^{-\alpha t}]$$

Since t is a random variable with distribution g(t;k,i),

$$H(x;\infty,i) = \frac{H_c}{\alpha} \sum_{k=1}^{x} \int_o^{\infty} g(t;k,i)\, [1 - e^{-\alpha t}]$$

$$= \frac{H_c}{\alpha} \sum_{k=1}^{x} [1 - (\frac{i}{\alpha+i})^k]$$

$$= \frac{H_c}{\alpha} [x - \frac{(r)(r^x-1)}{r-1}], \; r = i/(\alpha+i)$$

A continuous analogue of B(Q) was used in determining the optimum "s" and developing L(s,s+Q;i). Recall

$$B(Q) = B(s;L,i) - B(s+Q;L,i)$$

Consider each of the Q units separately and index them by k. Let t_k be the time until the k^{th} of the Q assets is demanded measured from the time the order is placed. For the k^{th} of the Q assets to be demanded, there must be (s+k) demands, so the density on t is $g(t_k;\, s+k,\, i)$. If t occurs before L there is a backorder and total backorder cost for that backorder is:

$$B_c \int_{t_k}^{L} e^{-\alpha t} dt = B_c [e^{-\alpha t_k} - e^{-\alpha L}]$$

Thus,

$$\frac{B(Q)}{B_c} = \sum_{k=1}^{Q} \int_o^L g(t;s+k,i)\, [e^{-\alpha t} - e^{-\alpha L}]$$

Now,

$$\int_o^L g(t;s+k,i)\, e^{-\alpha t} = (\frac{i}{i+\alpha})^{s+k} \int_o^L g(t;s+k,i+\alpha)$$

$$= (\frac{i}{i+\alpha})^{s+k} \, [1 - F(s+k-1;\, (i+\alpha)L)]$$

where F(s+k-1;(i+α)L) is the probability a Poisson random variable with mean (i+α)L is $\leq$ s + k - 1; i.e. t $\leq$ L is equivalent to s + k demands by time L.

Putting all these results together,

$$\frac{B(Q)}{B_c} = \sum_{k=1}^{Q} \left(\frac{i}{i+\alpha}\right)^{s+k} [1 - F(s+k-1;\ (i+\alpha)L)]$$

$$- e^{-\alpha L} \sum_{k=1}^{Q} [1 - F(s+k-1; i\ L)]$$

Application of Functional Values.

The various functional values H( ) and B( ) are used to derive the L(s,s+Q,i) and then the C_M (x,t) as discussed in the section of the report on Starting the Recursion. Then the recursive equation is used.

APPENDIX C

ALGORITHM PRECISION AND PROCESSING TIMES

Precision. The issues investigated briefly were:

(1) In the formula from Section 2.2 for t_n^{max}, what should PSDEV be multiplied by?

(2) In approximating the Gamma prior by a histogram, how many points should be used (Section 3.0)?

(3) How much accuracy would be lost if we limited the number of values of the sufficient statistic for which costs were calculated, and found the rest by linear interpolation?

Four runs were made for the base case with γ = 1.0 and BR = 100 (see model results for a description of the base case). The first and second runs multiplied PSDEV by 4 and used 10 histogram points, but the second run calculated costs for a maximum of 30 values of the sufficient statistic in any given period. At period 52, there were (coincidentally) 52 values so that 22 were found by interpolation. Run 3 differed from run 1 in that PSDEV was multiplied by 5, while run 4 differed from run 1 in that only 5 points were used for the histogram in starting the recursion.

Our conclusions were that Run 1 settings were adequate and that interpolation might be practical. We recorded for each run the (s,S) values at fielding and a lead time before fielding. They were the same for each run, (15-20) and (0,12) respectively. At fielding we recorded the expected lifetime costs (in thousands) if assets were 1 unit and you ordered up to either 13, 14, 15,....21, or 22 (Table C.1). A lead time before fielding we recorded expected cost if assets were 0 and you ordered up to 8, 9, 10...15, 16 units. We focused on how the costs changed as a function of assets bought, looking at asset values around the inventory control parameters selected.

Processing Times

The algorithm was run on a CDC CYBER 76 computer under the SCOPE 2.1.5 operating system. For the base case, with BR = 100 and γ = 1.0, 17.5 CPU seconds were used. For Run 2, using interpolation, CPU seconds dropped to 13.4 seconds. Of the 17.5 CPU seconds for Run 1, approximately 4.5 were used to get through the "Starting the Recursion" step.

TABLE C.1. COMPARISON OF EXPECTED COSTS (AT FIELDING)

RUN	ORDER TO:									
	13	14	15	16	17	18	19	20	21	22
1	\$89.7	\$88.1	\$86.4	\$84.8	\$83.2	\$81.7	\$80.2	\$78.8	\$77.4	\$76.1
2	89.9	88.3	86.6	84,9	83.4	81.9	80.4	79.0	71.6	76.3
3	90.3	88.7	87.0	85.3	83.8	82.3	80.8	79.4	78.0	76.7
4	89.3	87.7	86.0	84.4	82.8	81.3	79.8	78.4	77.1	75.7

TABLE C.2. COMPARISON OF EXPECTED COST (AT FIRST BUY)

RUN	ORDER TO:								
	8	7	10	11	12	13	14	15	16
1	\$87.4	\$86.0	\$84.5	\$83.1	\$81.7	\$80.3	\$78.9	\$77.6	\$76.2
2	87.6	86.2	84.8	83.3	81.9	80.5	79.2	77.8	76.4
3	87.9	86.5	85.0	83.6	82.2	80.8	79.4	79.0	76.7
4	Not Recorded								

APPENDIX D

PRE-FIELDING HEURISTIC

Suppose we wish to order every QT (for Q Time) weeks, and for a moment assume we order whenever assets fall below lead time demand.

Then, we draw a time line as follows:

QT WEEKS | (L+1) WEEKS
PRESENT

If we order up to our expected demand in the next (QT+L+1) weeks, and demand occurs as expected, we will order next in QT-weeks. If "present" is pre-fielding, some of the weeks on the time line will occur before fielding and we add in zeroes for those weeks.

Now suppose the reorder point (the s value) is based on lead time demand plus "n" standard deviations, and that based on the Poisson assumption the standard deviation is the square root of the mean. Then we must add to our order-up-to level "n" multiplied by the square root of demand in the (L+1) weeks shown on the time line.

The heuristic takes as input the (s,Q) values found to be optimum at time 0 (just before fielding) for the case with no uncertainty. From these it computes "n" and QT and applies them as discussed to get the order-up-to level at any time prior to fielding.

For the purpose of getting the reorder point (s), a different time line is developed:

(L+1) WEEKS QT WEEKS
PRESENT

Now demand in the first (L+1) weeks is used to get s, since this determines exposure to backorders in the lead time. This is still consistent with the derivation of the order-up-to level, since that derivation assumed that in QT weeks the s value would be based on demand in the next (L+1) weeks counting from that point.

Next page is blank

DISTRIBUTION

COPIES

1	Deputy Under Sec'y of the Army, ATTN: Office of Op Resch
	Headquarters, US Army Materiel Command
1	AMCSM
1	AMCSM-E
1	AMCSM-W
1	AMCSM-WRS
2	Defense Technical Info Center, Cameron Station, Alexandria, VA 22314
	Commander, USA Communications-Electronics Cmd, Ft. Monmouth, NJ 07703
1	DRSEL-PL-SA
	Director, Army Materiel Systems Analysis Activity, APG, MD 21005
1	AMXSY-D
1	AMXSY-G
1	AMXSY-A
1	AMXSY-C
1	AMXSY-R
1	AMXSY-RE
1	AMXSY-RM
1	AMXSY-RV
1	AMXSY-L
1	AMXSY-LX
1	AMXSY-LA
1	AMXSY-LM
1	AMXSY-LR
1	AMXSY-LO
1	AMXSY-DA
1	AMXSY-PM, ATTN: Pat Andrews
1	Commander, US Army Logistics Evaluation Agency, New Cumberland Army Depot New Cumberland, PA 17070
1	Commander, US Air Force Logistics Cmd, WPAFB, ATTN: AFLC/XRS, Dayton, OH 45433
1	US Navy Fleet Materiel Support Office, Naval Support Depot, Operations Analysis Dept, Code 93, Mechanicsburg, PA 17055
1	George Washington University, Inst. of Management Science & Engr., 707 22nd St., N.W., Wash., DC 20006
1	Naval Postgraduate School, ATTN; Dept of Opns Anal, Monterey, CA 93940
1	Air Force Institute of Technology, ATTN: SLGQ Head Quantitative Studies Dept., Dayton, OH 43433
1	Commander, US Army Logistics Center, ATTN: Concepts & Doctrine Directorate, Ft. Lee, VA 23801
1	Librarian, Logistics Mgt Inst., 4701 Sangamore Rd., Wash., DC 20016
1	University of Florida, ATTN; Dept of Industrial Systems Engr., Gainesville, FL 32601
1	RAND Corp., ATTN: S.M. Drezner, 1700 Main St., Santa Monica, CA 90406
1	Director, Army Management Engineering Training Agency, Rock Island Arsenal, Rock Island, IL 61299
1	Defense Logistics Agcy, ATTN: DLA-LO, Cameron Sta, Alexandria, VA 22314
1	Operations Research Center, 3115 Etcheverry Hall, University of California, Berkeley, CA 94720
1	Dr. Jack Muckstadt, Dept of Industrial Engr & Operations Research, Upson Hall, Cornell University, Ithaca, NY 14890

1 Dept of Industrial Engr, University of Michigan, Ann Arbor, MI 48104

1 Prof Robert M. Stark, Dept of Stat & Computer Sciences, University of Delaware, Newark, DE 19711

1 Prof E. Gerald Hurst, Jr., Dept of Decision Science, The Wharton School, University of Penna., Phila., PA 19104

1 Director, AMSAA, ATTN: AMXSY-LISO, Ft. Lee, VA 23801

1 Director, AMSAA, ATTN: Paul Arvis, AMXSY-PRO, Ft. Lee, VA 23801

1 Dept of Industrial Engr & Engr Management, Stanford University, Stanford, CA 94305

1 Prof Harvey M. Wagner, Dean, School of Business Adm, University of North Carolina, Chapel Hill, NC 27514

1 Pentagon Library, ATTN: ANR-AL-RS (Army Studies), Pentagon, Rm 1A518. Wash., DC 20310

1 DARCOM Intern Training Center, Red River Army Depot, Texarkana, TX 75501

1 Prof Leroy B. Schwarz, Dept of Management, Purdue University, Krannert Bldg., West Lafayette, IN 47904

1 Operations & Inventory Analysis Ofc, NAVSUP (Code 04A), Dept of Navy, Wash., DC 20376

1 US Army Research Office, ATTN: Robert Launer, Math. Div., P.O. Box 12211, Research Triangle Park, NC 27709

1 Prof William P. Pierskalla, 3641 Locust Walk CE, Phila., PA 19104

Air Force Logistics Mgt Center, Gunter Air Force Station, AL 36144

1 AFLMC/LGY

1 AFLMC/XRP, Bldg. 205

1 Dr. Carl Weisman, C.A.C.I., 1815 N. Fort Myer Drive, Arlington, VA 22209

1 Commander, US Army Missile Cmd, ATTN: Ray Dotson, DRSMI-DS, Redstone Arsenal, AL 35898

1 Prof Peter Kubat, Graduate School of Mgmt, University of Rochester, Rochester, NY 14627

1 Prof Barney Bissinger, The Pennsylvania State University, Middletown, PA 17057

1 Prof Garry D. Scudder, Mgmt Sciences Dept, College of Business Adm, University of Minnesota, 271 19th Ave., South, Minneapolis, MN 55455

1 Center for Naval Analyses, ATTN: Stan Horowitz, 2000 N. Beauregard St., Alexandria, VA 22311

1 Dr. Warren H. Hausman, Industrial Engr. & Engr. Management, Stanford University, Stanford, CA 94305

1 Dr. Paul Zipkin, Graduate School of Business, Columbia University, 416 Uris Hall, New York, NY 10027

1 Mr. Craig Sherbrooke, Logistics Management Institute, 6400 Goldsboro Rd., Bethesda, MD 20827-5886