Bayesian design of stochastic inventory systems

Tapan P Bagchi
tapan.bagchi@nmims.edu
NMIMS Shirpur Campus, India 425405

Abstract
This paper re-visits the Bayesian approach to test its efficacy in optimally designing the classical (*s*, *Q*) inventory model. A heuristic search of the unstructured (*s*, *Q*) decision space finds that one can indeed make a decent start by the Bayesian approach—and keep total cost nearly optimally low throughout.

Keywords Stochastic Demand, Bayesian Design, (s, Q) Inventory Policy, Optimization

Inventory Control: Theory and Practice
Inventories are a common insurance against uncertainties impacting most production or service operations. However, inventories do not producing any "return". Yet inventories help out in breakdowns or crisis and also improve customer service. So a balance is needed here to optimally determining (a) "when to order", and (b) "how much to order". Uncertainties are particularly high for a new business. One wants here two answers—how much should the organization stock initially, and should it adjust decisions (a) and (b) as time progresses? This paper re-visits the Bayesian approach to test its efficacy in answering these two questions.

Stochastic demand condition
Prominent here are the single period stochastic model, the (*s*, *S*) model, and the (*s*, *Q*) model, descriptions being given in Silver et al. (1998), and Jensen and Bard (2003). Murray and Silver (1966) state that initially one would have great uncertainty concerning the sales potential of an item. *Learning*, in the decision theoretic sense, is the process of basing one's initial decisions on an informed or inspired guess to start one's business, and subsequently updating that initial guess by some rational logic—to assure optimality of future decisions.

Bayesian adaptation to help better sales forecasting was used first by Murray and Silver (1966). They adaptively changed the distribution of sales to show how to improve decisions. Others have also used adaptation-aimed learning but few have assessed the potential benefits of Bayesian adaptation/learning with inventories with the exception of Azoury and Miller (1994), who show that the quantity ordered under the non-Bayesian policy would be greater than that under a Bayesian policy. Bayesian methods were used also to design queuing systems (see Bagchi and Cunningham (1972) and Morales et al. (2005). In this study we inquire:

Can the Bayesian learning logic (prior → prior + data → posterior) of observing and updating stochastic information help reduce total inventory operating costs?

As shown by Azoury and Miller (1994), we hypothesize that for the popular (*s, Q*) policy, Bayesian learning would let one see how the successive incorporation of new data would improve decisions (reduce the total expect cost (*EC*) and/or improve service level) as $Q*$ and $s*$ are continually updated with accumulating demand data. But, up to what point such updating would be meaningful? We expect that that answer may depend on the estimated unknown but stationary demand rate for that might require meeting certain minimum sample size. How long one should sample demand (*X*) might depend on the cost of using non-optimal rather than a near-optimal *EC*? We do not probe this. Here we use a well-developed stochastic inventory model from the literature—the (*s, Q*) model. It incorporates safety stock into the reorder stock level (*s*) and uses an optimum constant quantity *Q* of an order every time the current stock level touches or falls below *s*. Optimum *s* and *Q* minimize the total expected cost/unit time (Silver 2007).

The (*s, Q*) Inventory Model

This study uses the results of Jensen and Bard (2003). We assume that only a single item is stocked and sold whose inventory is managed to keep the expected total cost minimum, comprising holding, replenishment and stock out costs.

Ordering too much or too little or at the wrong time can disrupt the optimal control of inventory, an event easily caused by uncertainty in demand. Here the deterministic approach clearly does not minimize the expected total cost. At some instant of time if inventory level is *z*, then the probability of shortage (P_s), the probability of excess (P_e), the expected shortage (E_s) and expected excess (E_e) are, respectively,

$$P_s = P[x > z] = 1 - F(z)$$
$$P_e = P[x \le z] = F(z)$$
$$E_s = \int_z^{\infty} (z - x) f(x) dx \quad (\textit{for continuous demand } x)$$
$$E_s = \sum_z^{\infty} (x - z) P(x) \quad (\textit{for discrete demand } x), \quad \textit{and}$$
$$E_e = z - \mu + E_s$$

This study used the (*s, Q*) inventory management policy to serve as the test bed to probe our query. Here demand is stochastic. (*s, Q*) first determines the optimum values (s^* and Q^*) for the reorder point (*s*) and the order quantity (*Q*). It then monitors inventory continuously through the repeated execution of order cycles. An order of size Q^* is placed when the current inventory level *z* touches s^*. The order (quantity = Q^*) is received after lead time *L* to replenishes stock.

Optimum s^* and Q^* are found as follows. When *L* is small compared to the expected time required to exhaust *Q*, only 1 order would be outstanding. (In practice a plant may place multiple orders on a vendor when expediting becomes ineffective, but we do not consider this case here.) An order cycle is the time between two successive receipts. If *a* and *L* respectively represent the average demand rate and lead time, then the *mean demand during lead time* is $\mu = aL$. The reorder point being *s*, P_s is $1 - F(s)$, and the system's *service level* (fraction of demand during lead time that is met) is $1 - P_s = F(s)$. The safety stock (excess stock beyond μ) is $SS = s - \mu$. The general solution for this situation has been given by Jensen and Bard (2003) as follows.

If the per SKU unit holding cost is h per unit time, then

Expected holding cost/unit time = $h\left(\frac{Q}{a}+s-\mu\right)$

The time between orders is random with mean of Q/a. If the cost of replenishment/order is K then the expected replenishment cost/unit time is (Ka/Q). If the expected shortage cost/order cycle is C_s, then the expected shortage cost/unit time will be $C_s/(Q/a) = C_s a/Q$. The general model for the expected total cost/unit time for the (s, Q) policy will thus be

$$EC(s,Q) = h\left(\frac{Q}{a}+s-\mu\right)+\frac{Ka}{Q}+\frac{a}{Q}C_s \tag{1}$$

Equation (1) is the expression for the expected total cost/unit time for an inventory system being operated by the (s, Q) policy. To optimize it one uses decision variables s, the reorder point, and Q, the quantity ordered in each order cycle. Analytically, (1) may be partially differentiated with respect to s and Q and the derivatives equated to zero. Doing this yields two conditions that simultaneously characterize the two optimal values Q^* and s^*. These are

$$Q^* = \sqrt{\frac{2a(K+C_s}{h}} \tag{2}$$

and

$$\frac{\partial C_s}{\partial s} = -\frac{hQ}{a} \tag{3}$$

Peterson and Silver (1979) employ special cases for obtaining the optimal values Q^* and s^*. The first case assumes that a constant cost π_1 is expended whenever a stock out event occurs. This gives us a quick way to evaluate C_s—the expected shortage cost/order cycle. This is

$$C_s = \pi_1 P[x>s] = \pi_1\int_s^\infty f(x)dx = \pi_1[1-F(s)] \tag{4}$$

Equation (3) may be now utilized since we have C_s expressed in (4) as a function of s. Thus

$$\frac{\partial C_s}{\partial s} = -\pi_1 f(s^*) = -\frac{hQ}{a} \tag{5}$$

which gives $f(s^*) = \dfrac{hQ}{\pi_1 a}$ (6)

with $C_s = \pi_1(1-F(s^*)$ (7)

Equation (6) helps link s^* with Q^* via (2). Note that seeking a solution to the (s, Q) policy problem by simultaneously solving (2), (6) and (7) for arbitrary demand distribution $F(s)$ is not trivial.

A variant of the constant cost π_1 *per stock out event* is a cost π_2 incurred for *every unit short* in a stock out. Then the expected shortage cost/order cycle will be dependent on how many units are expected to be shorted in each order cycle (E_s). Here,

$$E_s = \int_s^\infty (x - s) f(x) dx$$

and therefore $C_s = \pi_2 E_s$. This gives

$$\frac{\partial C_s}{\partial s} = -\pi_2 \int_s^\infty f(x) dx = -\pi_2 (1 - F(s)) \qquad (8)$$

Combining (3) and (8) one obtains

$$\frac{\partial C_s}{\partial s} = -\pi_2 (1 - F(s)) = -\frac{hQ}{a}$$

which for a specified Q gives the condition for the optimal reorder point s^* as

$$F(s^*) = 1 - \frac{hQ}{\pi_2 a} \qquad (9)$$

The optimum decision (s^*, Q^*) is the combination of s and Q that minimizes EC given by (1).

Owing to the non-linear and complex nature of (1) through (9) we used an orthogonal array experimental computational framework to first determine the sensitivity of the two performances (responses)--% service level and total expected cost. For this we selected two working levels for each of the factors—monthly demand (a), holding cost (h), shortage cost ($\pi 1$) and order cost (K) as shown in Table 1, and a L_8 array (Montgomery 2008). Table 1 and Figures 1 and 2 show the results. The following inferences may be drawn:

- Optimal Order Quantity Q^* is reactive to a, h and K, but only mildly to shortage cost $\pi 1$.
- Service level seems to be robust relative to most factors considered in the region of the cost parameters studied. It is closely related to the setting of the reorder point s^*.
- Total Expected Cost/unit time is relatively robust with respect to shortage cost per stock out event $\pi 1$, *but* sensitive to a, h and K.

These deductions—typically unavailable to the inventory manager—suggest that it would be wise to spend effort in *optimally setting* the reorder point s^* and order quantity Q^* before one declares the operational policies of a stochastic inventory system. In one sense, this information is analogous to the *relative robustness* of the total operating cost/unit time to EOQ for a deterministic inventory system. Limited generalization of such deductions may be attempted in a given cost-demand scenario to assess how accurately the parameters a, h, K and $\pi 1$ need to be estimated, to assure minimum cost operation of the system. In this study we focus on the factor with perhaps the highest uncertainty—the stochastic nature of X, the demand per unit time.

***Table 1**: An L_8 Computational Experiment to uncover the Sensitivity of Service Level and Total Expected Cost of an (s, Q) inventory System*

Orthogonal Experiment #	*a* Units /month	h	*π*1	*K*	s^*	Q^*	%Service Level	Total Expected Cost/unit time
1	50	5	500	400	19.25	91	97.2	488.76
2	50	5	1000	800	19.88	127.9	98.2	676.6
3	50	10	500	400	18.56	64.9	95.7	709.93
4	50	10	1000	800	19.25	91	97.2	977.51
5	100	5	500	800	33.61	181.2	95.74	949.29
6	100	5	1000	400	36.22	128.4	98.76	698.3
7	100	10	500	800	32.51	129.1	93.35	1366.1
8	100	10	1000	400	35.41	91.5	98.13	1019.2

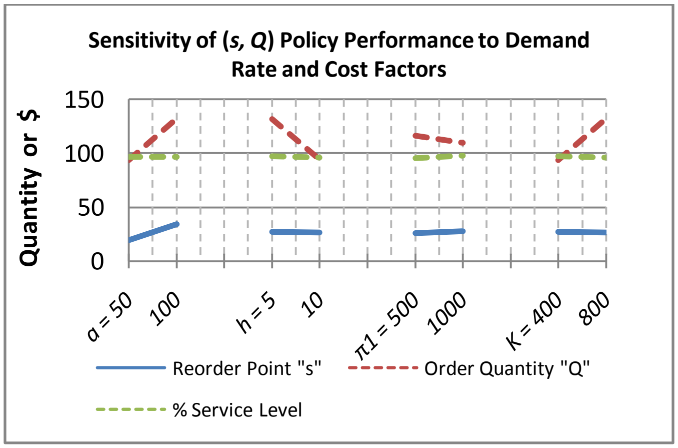

***Figure 1** Sensitivity of Optimal Reorder Point (s^*), Order Quantity (Q^*) to Monthly Demand Rate and Costs*

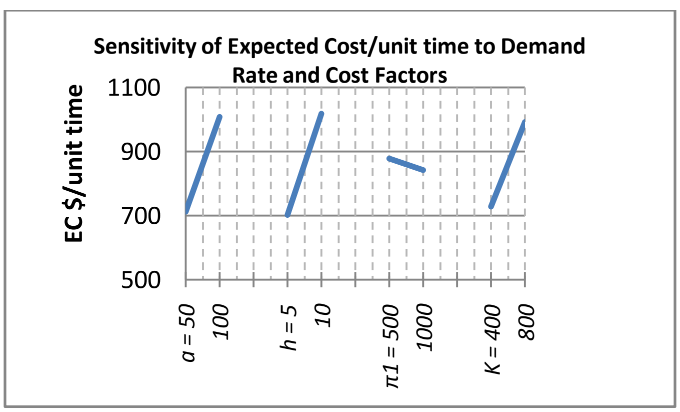

***Figure 2** Sensitivity of Expected Total Cost/unit time to Monthly Demand Rate and Costs*

Bayesian learning and Inferencing

The assumption of stationarity is generally made to keep the analysis straightforward. In the present case also we assume stationarity; specifically, the parameters that control the distribution of demand are assumed to be unknown, but stationary.

Since the Bayesian learning logic (*prior* → *prior* + *data* → *posterior*) follows the path of pre-supposing information, observing the phenomena and then repeatedly updating stochastic information, one is advised to use an appropriate subjective probability function for the "prior". In theory this is done by deriving the posterior density from the likelihood function and the prior density, and deriving the distribution of the reduced-form parameters from the initial information on the unknown parameters controlling the stochastic process (here the random demand).

For the present case, demand is assumed to be random, Poisson distributed with an unknown parameter (average rate) λ/unit time. This extends two advantages: First, the Poisson distribution is often quite realistic when demand is random and independent of earlier and future demands. Secondly, from analytical point of view, the Bayesian prior-posterior *conjugate family* (Raiffa and Schlaifer 1961) of the distribution of the possible values of λ is *Gamma*, a two-parameter distribution convenient to update. However, this is a minor restriction—Bayesian inference may be performed using stochastic simulation of the process also (Morales et al. 2005).

The Bayesian Learning Framework

The conjugate prior distribution for the Poisson rate parameter is the Gamma distribution with two parameters α and β, the density function being

$$Gamma\,(x) = \frac{\beta^{\alpha}}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x}, x \geq 0, \alpha > 0, \beta > 0$$

where $\Gamma(\alpha) = \int_0^{\infty} t^{\alpha-1} e^{-t} dt$.

Here the Gamma distribution is used as a conjugate prior (Raiffa and Schlaifer 1961) for the rate parameter λ Gamma has a mean of α/β, variance α/β^2 and it is known to be flexible in shape. The interpretation is as follows: There are α total occurrences in β time intervals. Updating the Gamma prior is straightforward. For instance, if *r* quantities are demanded from the inventory in a time period of length *t*, the posterior density of λ will be a Gamma distribution with α' = α + *r* and β' = β + *t*. The maximum likelihood estimated of λ is obtained from the posterior mean $E(\lambda) = (\alpha + r)/(\beta + t)$. This posterior mean $E(\lambda)$ approaches λ_{MLE} in the limit as α → 0 and β → 0. In the present application our intention will be to start with some reasonably assumed prior value of Poisson demand rate λ (= α/β) and then continually update it using *r* (the demand observed in the time span *t*) as time *t* advances. Bagchi and Cunningham (1972) provide the decision theoretic rationale for adopting this procedure.

Table 2 Sample Poisson Random Demand Values simulated with λ = 100. Assuming a prior as a Gamma(α = 5, β = 1) distribution, we compute Mean of the Gamma Posterior distributions as monthly demand is successively observed.

Month #	**Observed Monthly Demand**		**For Posterior Gamma ($\underline{\alpha}$, $\underline{\beta}$):**		**λ = Mean of Gamma Posterior**	**SD = SQRT($\underline{\alpha}$/$\underline{\beta}$^2)**	**Upper/Lower estimates of λ**	
i	**ki**	**Sum ki**	**$\underline{\alpha}$ = α + sum ki**	**$\underline{\beta}$ = β + n**	**λ estimate = $\underline{\alpha}$/$\underline{\beta}$**	**SD (λ est)**	**+95%**	**-95%**
Month 1	111	111	116	2	58.00	5.39	68.77	47.23
2	111	222	227	3	75.67	5.02	85.71	65.62
3	92	314	319	4	79.75	4.47	88.68	70.82
4	104	418	423	5	84.60	4.11	92.83	76.37
5	102	520	525	6	87.50	3.82	95.14	79.86
6	98	618	623	7	89.00	3.57	96.13	81.87
7	etc.							

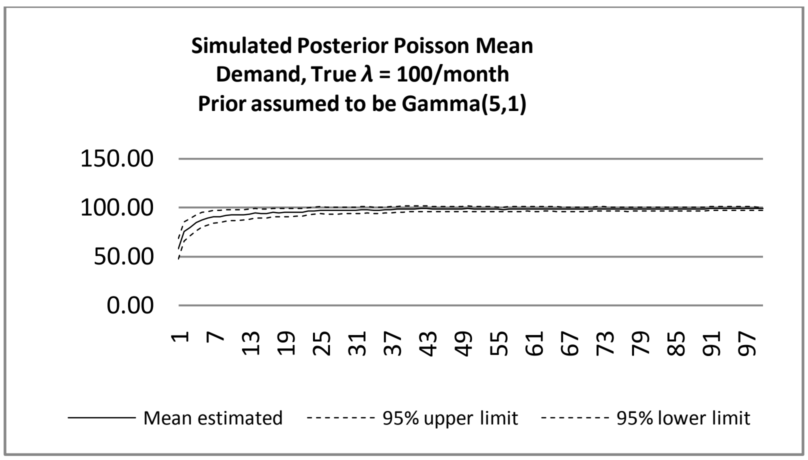

Figure 3 *Trace of Continually Updated Posterior Mean Estimates produced by Succession of Samples drawn from a Simulated Poisson Distribution with true λ = 100/month*

The Bayesian approach gives us a way around a special decision problem. Sometimes one makes the best decision on the basis of a given set of data available from past history, or produced by conducting some special statistical experiments. What if there is no such history available, or there is no opportunity to conduct experiments? As said above, the Bayesian approach begins with an assumed prior about the decision environment, and then by using a learning logic (prior → prior + data → posterior) follows the path of pre-supposing information, observes the phenomena and then repeatedly updates information.

Table 2 and Figure 3 illustrate the effect of updating the maximum likelihood estimate of the mean demand of a Poisson distribution. Here a Poisson demand process with true mean demand (λ) of 100 units/month was observed (simulated) in succession of months 1, 2, 3, … etc. and the estimates of the posterior along with the ±95% limits of this estimate was calculated. The prior of λ was assumed to be a Gamma(α = 5, β = 1) distribution with an average of 5 units/month. The simulated monthly demands are shown in the second column (ki) of Table 2.

Continual updating of the estimated value of λ produced a 95% confidence band for λ as (96.81, 100.77) after 100 updates. Clearly, another assumed value of the prior would produce another trace of estimates. But it can be assured based on the standard deviation of the estimated λ that they would all eventually converge toward the true demand as time t increases.

Regular Bayesian updates of Demand Rate help keep Total Expected Cost near Minimum

One way to operate the (s, Q) inventory policy is to be myopic (Levy et al. 2007) in setting the operating values of s and Q at some initial (prior) guess for the demand rate λ, or use only a limited amount of data to estimate λ. Such trust on an initial guess for λ and not changing it later may even be favored, for it saves the effort needed to incorporate any emergent evidence about true demand as the business moves forward. Many practitioners indeed do not change the initial assumption about a or λ or even costs, though Azoury and Miller (1994) deplore this.

Plainly, the s and Q derived using the initial guess for λ would almost surely be suboptimal, except by accident (Figure 2 indicates the strong dependence of EC on a (hence on λ)). To test any possible merit of the myopic approach we studied it computationally. Table 3 displays the effect of setting the unknown demand rate λ at some unsubstantiated value (assuming mistakenly this to be the true demand), deriving the flawed s^* and Q^* from it, and incurring the consequent expected total cost EC when the (s, Q) policy is used. It is straightforward to compare these higher values of EC with the near optimal EC achievable by getting close to the *true demand rate* by using Bayesian updates. Table 3 used 100 units/month as the true demand rate whereas the ("wrongly") presumed values of λ were set respectively at 25, 50, 75, 150, 200, and 300 units/month. The costs used were h = \$10/unit-month, $\pi 1$ = \$500 per backorder event and K = \$800/order placed. Figures 3 and 4 show the effect of *flawed* (differing from the true) demand conjecture ("prior") on EC, and the service level projected.

***Table 3** Costs and Service Levels Experienced when s^* and Q^* are set based on a prior estimated demand, but true demand (aT) is different from the prior estimate (a)*

a = Estimated demand/month based on a *flawed* prior guess	25	50	75	100	150	200	300
aT = True demand/month	**100**	**100**	**100**	**100**	**100**	**100**	**100**
s^* = Optimum Reorder Point based on prior	10	18	25	33	47	61	88
Q^* = Optimum Order Quantity based on prior	**65**	**91**	**111**	**129**	**158**	**183**	**224**
Holding Cost based on *a and (s*, Q*)*	360	510	624	720	883	1019	1248
Holding Cost based on *(s*, Q*) but facing True demand*	**173**	**385**	**569**	**721**	**989**	**1269**	**1748**
Replenishment Cost by *a and (s*, Q*)*	310	438	537	620	759	876	1073
Replenishment Cost based on *(s*, Q*) but facing True demand*	**1239**	**876**	**715**	**620**	**506**	**438**	**358**
Shortage Cost by *a and (s*, Q*)*	13	18	22	26	32	36	45
Shortage Cost based on *(s*, Q*) but facing True demand*	**52**	**36**	**21**	**26**	**36**	**18**	**15**
Expected Total Cost by *a and (s*, Q*)*	683	966	1183	1366	1673	1932	2366
Expected Total Cost based on *(s*, Q*) but facing True demand*	**1464**	**1297**	**1305**	**1366**	**1531**	**1725**	**2121**
Service level experienced at True Demand but operating at (s^*, Q^*)	98.8 %	97.2 %	95.4 %	93.4 %	88.5 %	82.5 %	59.6 %

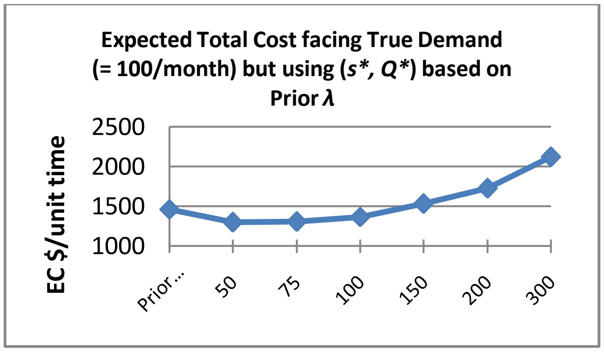

Figure 4 *The Expected Total Cost for a (s, Q) System that uses a flawed prior estimate that is far from the true demand value*

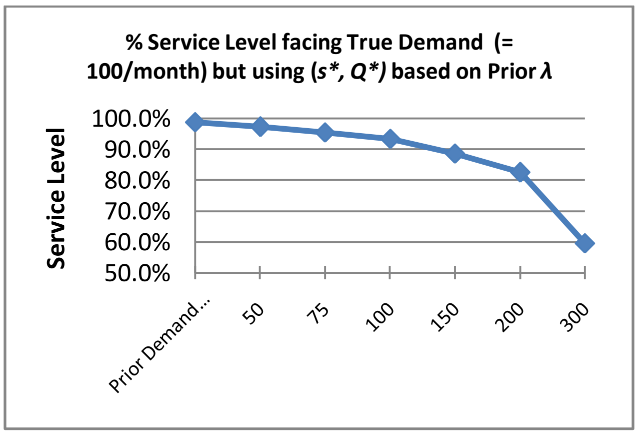

Figure 5 *Service Level provided by a (s, Q) System that uses a flawed prior estimate that is far from the true demand value*

A review of Table 3 and Figures 4 and 5 would suggest that the results of using Bayesian learning to keep continually updating the demand estimate provide a mixed message. But a closer look reveals that the minimum total cost (*s*, *Q*) policy should indeed be based on a demand estimate *as close to* the true demand as is possible. But Figure 4 appears to suggest that a lower guess for demand actually improves customer service! Prima facie, therefore, guessing a low value of demand appears to be doing something good. But such inference is shortsighted and most misleading. More seriously, this is not a defect in the *model* or its analysis.

Recall that our objective of setting up the (*s*, *Q*) model to help find a rational way to manage inventories when demand is stochastic included spelling out the objective first—that of minimizing (1), the expected total cost/unit time. This total cost included three components—the holding cost, the replenishment cost, and the shortage cost. At least for this model, therefore, maximizing customer service per se was not the objective. Customer service ($F(s)$) enters into (1) via Cs (= $\pi_1(1 - F(s))$), the cost of short shipment. If one requires the final (*s*, *Q*) solution to assure a high level of customer service, one would need to use a large value for π_1. Thus, like in any optimization, one must be clear about the objective—where does he want to put priority?

Conclusions

This study has investigated the value of incorporating Bayesian learning into the popular (*s*, *Q*) model for managing inventories when demand is stochastic. The study finds that one can indeed make a decent start by the Bayesian approach—*and stay the course nearly optimally*—while upholding a target service level by suitably selecting costs *and also* keeping the expected operating cost/unit time minimum. Specifically, this study uncovers the high value in continually updating the decisions (a) "when to order", and (b) "how much to order", rather than sticking to the initial guess for the demand average, as is frequently practiced.

This work utilized an orthogonal array experimental framework to determine the sensitivity of the two performances (responses)--% service level and total expected cost. For this two working levels for each of the factors—monthly demand (*a*), holding cost (*h*), shortage cost (π_1) and order cost (*K*) were selected and a L_8 array was adopted to guide the computations. Furthermore, importantly, myopic estimate of true demand would surely yield suboptimal *EC*.

These deductions—typically unavailable to the inventory manager—suggest that it would be wise to spend effort in *optimally setting* the reorder point s^* and order quantity Q^* before one sets out to declare the operational policies to cope with stochastic demand.

References

Azoury Katy S and Bruce L Miller (1994). A Comparison of the Optimal Ordering Levels of Bayesian and Non-Bayesian Inventory Models, *Management Science*,

Bagchi, Tapan P and A A Cunnigham (1972). Bayesian Approach to the Design of Queuing Systems, *INFOR Journal*, Vol 10(1).

Jensen, P A and J F Bard (2003). *Operations research: Models and Methods*, John Wiley & Sons Inc.

Levi, T, R O Roundy and D B Shmoys (2007). Provably near-optimal sampling-based policies for stochastic inventory control models, *Operations Research*, Vol 32, 821-839.

Montgomery, D C (2008). *Design and Analysis of Experiments*, 5th ed., John Wiley.

Morales, J, M E Castellanos, A M Mayoral and C Armero (2005). Bayesian Design in Queues: An Application to Aeronautic Maintenance, Centro de Investigacion Operative, Elche

Murray, George R and Edward A Silver (1966). A Bayesian Analysis of the Style Goods Inventory Problem, *Management Science*, 1966, vol. 12 (11), pages 785-797

Peterson R and E A Silver (1979). *Decision Systems for Inventory Management and Production Planning*, John Wiley.

Raiffa, Howard and Robert Schlaifer (1961). *Applied Statistical Decision Theory*. Division of Research, Graduate School of Business Administration, Harvard University.

Silver, E A (1981). Operations Research in Inventory Management: A Review and Critique, *Operations Research,* vol. 29 no. 4 628-645.

Silver, E A, D F Pyke, R Peterson (1998). *Inventory Management and Production Planning and Scheduling*, Wiley

Silver, E A (2007). *Inventory Management: A Tutorial 2007-03,* University of Calgary.