Abstract Number: 9298

Abstract Title: Inventory Management of Spare Parts Subject to Deteriorationn

Jennifer K. Ryan, Zhi Zeng

Department of Industrial and Systems Engineering

Rensselaer Polytechnic Institute

Troy, NY 12180

ryanj6@rpi.edu, zengz@rpi.edu

Rong Li

Lee Kong Chian School of Business

Singapore Management University

rongli@smu.edu.sg

POMS 23rd Annual Conference

Chicago, Illinois, U.S.A.

April 20 to April 23, 2011

Abstract

Lack of coordination between machinery fault diagnosis and the inventory management of spare parts can lead to increased inventory costs and disruptions in production activity. To address this gap, we develop a framework for incorporating real-time sensor data into inventory decisions for spare parts that are subject to deterioration.

1. Introduction

The performance of complex industrial systems can be improved through the implementation of condition monitoring, which is defined as the collection of real-time sensor information from a functioning device in order to monitor the health or condition of that device. Examples of condition monitoring techniques include vibration analysis, oil analysis, and thermography. Despite rapid growth in the use of these techniques, there has been little research considering how to best use condition monitoring data to make decisions regarding the management of repair or replacement resources. In this paper, we consider the problem of using condition information to improve the inventory management of replacement parts. Specifically, we consider a system consisting of multiple machines. Each machine uses a key part or component that is subject to deterioration. We assume that the owner of these machines is able to collect real-time condition information from the operating parts and seeks to use this condition information to more effectively manage the inventory of replacement parts. Three key steps are involved in this process:

1. In order to use real-time sensor information to predict the remaining life of a device and to improve operational decisions such as part replacement, it is useful to identify and model a *degradation signal*, i.e., a quantity computed from sensor information that captures the

current condition of the device and provides information on how that condition is likely to evolve over time. The functional form of the signal is typically modeled as a continuous-time continuous-state stochastic process.

Lu and Meeker (1993) present a random coefficient degradation model that contains both common parameters shared by the entire device population and unique parameters for each individual device. In their model, the degradation signal takes the form: $S_{ij} = \eta(t_j; \Phi, \Theta_i) + \varepsilon_{ij}$, where $\eta(.)$ represents the deterministic path followed by the degradation signal, Φ represents the deterministic degradation parameter shared by all machines, Θ_i represents the degradation parameter that captures the unique characteristics of machine i, and ε_{ij} denotes the random error of machine i at time t_j.

There are two widely used approaches for modeling the random error process, independent and identically distributed random errors (e.g., Lu and Meeker 1993 and Wang 2000) and a Brownian motion error process (e.g., Doksum et al. 1992, Gebraeel et al 2005 and Whitmore 1995). As in Gebraeel et al. (2005), we assume a linear degradation model with Brownian motion error process. The model parameters are estimated in a Bayesian manner, incorporating real-time condition monitoring information.

2. Demand for spare parts inventory is generated when a part fails while operating on a machine. Hence, we need to estimate the residual life distribution (or the failure time distribution) for an operating part given an appropriately selected degradation signal model. With the remaining life distribution, we can further estimate the demand distribution for spare parts. Gebraeel et al. (2005) derive the residual life distribution of the machine parts based on a linear degradation model with i.i.d. normal error. Doksum et al. (1992) show that, for

the linear model with Brownian motion error process, the residual life of the device follows an inverse Gaussian distribution.

3. The demand distribution for spare parts is updated in real time given the observed degradation information obtained from the set of operating machines. Hence, the optimal inventory control decisions will be dynamic and will depend on the condition information obtained from the entire system of machines.

   There has been extensive research on inventory decisions considering periodically updated demand distribution, e.g., Dvoretzky et al. (1952), Scarf (1959), Karlin (1960), etc. The existing research, to the best of our knowledge, generally updates the demand distribution using previously observed demand data. However, for spare-part inventory, the demand is generally very sparse, which restricts the amount of information provided by the previous demand.

   In contrast, the research presented in this paper presents a methodology for updating the distribution of demand using the real-time sensor observations, and for incorporating this updated demand distribution into a stochastic inventory model. This approach provides improved estimation of the demand distribution, leading to improved inventory decisions, but also imposes challenges in computing the optimal inventory control policy.

The purpose of this paper is to investigate the form of the optimal inventory control policy for the spare parts inventory control problem with a periodically-updated, non-homogeneous demand distribution derived from real-time sensor information. In addition, we seek to provide effective and efficient solution approaches for this problem.

2. Condition Monitoring and Demand Estimation

In this section, we outline our model of the degradation process for the machine parts and discuss our demand estimation procedure.

2.1 Degradation Model Estimation

We assume that each operating machine part is monitored periodically and that the degradation signal captured from each part has a linear form with a Brownian motion error process. The degradation model, similar to Gebraeel et al. 2005, can be written as $Z(t) = \theta t + \varepsilon(t)$, where $Z(t)$ is the signal value as a function of time, θ is the drift or degradation parameter and $\varepsilon(t) = \sigma W(t)$, where W is a standard Brownian motion variable and σ is the associated volatility. Thus, the signal $Z(t)$ was a measure of this vibration over time. When there are multiple (m) machines in the system, we allow the degradation parameter, θ, to vary across the machines. In this case we will add the subscript i to the above model, for $i = 1, \ldots, m$.

We assume that the volatility of the Brownian motion, σ, is known, but that the degradation parameter, θ, may or may not be known. When θ is unknown, we will follow the Bayesian approach presented in Gebraeel et al. (2005) to update the probability distribution of θ in each period, i.e., each time a new degradation signal is observed. Assuming a Normal prior distribution for θ, we find that the posterior distribution of θ in period n will also be Normal, with mean μ_n and variance, ${\sigma_n}^2$, which are functions of the observed degradation signals for periods 1 through n, $z_1, \ldots, z_n$.

2.2 Demand Estimation

In our model, as well as in many practical applications, a part is considered to be failed, and is instantaneously replaced, when the observed degradation signal exceeds a predefined threshold, B. The failure threshold is typically specified by the device manufacturer. In this case, we can easily compute the remaining life distribution of an operating part given an observed degradation signal, z_n. Since the part is replaced when the degradation signal first exceeds the threshold B, the failure time T is just the first passage time for the Brownian degradation signal to reach B. Thus, when θ is known, using Cox (1965), we can write the cumulative distribution of T as follows:

$$P(T(\theta, z_n, B) \leq t) = 1 - \Phi\left(\frac{(B - z_n) - \theta t}{\sigma\sqrt{t}}\right) + \exp\left\{\frac{2\theta(B - z_n)}{\sigma^2}\right\}\Phi\left(\frac{-(B - z_n) - \theta t}{\sigma\sqrt{t}}\right), \quad (1)$$

where $T(\theta, z_n, B)$ represents the time for the degradation signal to first reach the barrier B given a current observation z_n. Further, for Brownian motion, the first passage time distribution for a single part with threshold kB is the same as the distribution of the sum of the first passage times for k successive parts with threshold B. Thus, we have the following distribution of demand for replacement parts for a single machine:

$$\begin{aligned} P(D_n = k|\theta, z_n) &= P(T(\theta, z_n, kB) \leq t_0, T(\theta, z_n, (k+1)B) > t_0) \\ &= P(T(\theta, z_n, kB) \leq t_0) - P(T(\theta, z_n, (k+1)B) \leq t_0) \end{aligned} \quad (2)$$

When there are multiple machines, the total demand for replacement parts is simply the sum of the demands from each machine, and hence we have:

$$P(D_n = k|\vec{\theta}, \vec{z}_n) = \sum_{(l_1,\ldots,l_m):\sum_{i=1}^m l_i = k} \left(\prod_{i=1}^m P(D_{in} = l_i|\theta_i, z_n^i)\right), \quad (3)$$

where D_{in}, z_n^i, θ_i are the demand, observed degradation signal and degradation parameter on machine i, respectively, while $\vec{\theta} = [\theta_1, ..., \theta_m]$ and $\vec{z_n} = [z_n^1, ..., z_n^m]$.

When the drift parameter is unknown, we must take the expectation over the posterior distribution of θ, calculated in period n, as described in the previous section. For example, for the case of a single machine, we have

$$P\{D_n = k | z_1, \ldots, z_n\} = \int_{-\infty}^{\infty} P\{D_n = k | \theta, z_n\} f_{\Theta | z_1, \ldots, z_n}(\theta) d\theta, \tag{4}$$

where $f_{\Theta | z_1, \ldots, z_n}(\theta)$ is the posterior distribution of θ calculated in period n based on the set of observed signal values.

3. Spare Parts Inventory Management

We next determine the optimal inventory control policy for this problem setting under the assumption that the manufacturer follows a periodic review inventory policy with a finite planning horizon and linear ordering, holding and penalty costs. We also assume that emergency orders are used to fill backorders.

As in the previous section, we will start by considering the case in which the degradation parameters are known. We will then consider the case in which these parameters are unknown and estimated using the observed degradation signal values.We develop a dynamic programming (DP) formulation for this problem. Let $C_n(x_n, z_n | \theta)$ denote the minimum expected cost for the remainder of the planning horizon, starting from period n, given the most recent degradation signal observation, z_n, and the current on-hand inventory level, x_n. Then, for the case of known θ, we have the following formulation:

$$C_n(x_n, z_n|\theta) = \min_{y \geq x_n} \left\{ c(y - x_n) + L_n(y|z_n, \theta) + \alpha \sum_{k=0}^{\infty} E\left[C_{n+1}(y - k, Z_{n+1}|\theta) 1_{\{D_n = k|\theta, z_n\}} \right] \right\}, \quad (5)$$

where y is the order-up-to level, c is the per unit ordering cost, $L_n(y|z_n, \theta)$ is the expected holding and penalty cost in period n, and α is the discount factor. In the final term, $1_{\{D_n=k|\theta,z_n\}}$ is an indicator function, and the expectation is taken over the random variable Z_{n+1} under the condition $D_n = k$. The state of this DP includes the on-hand inventory level, x_n, and the most recently observed degradation signal, z_n. We define $C_{N+1}(x, z) = 0, \forall x, z$. Since z_n is a continuous variable, we will need to discretize its distribution to make the DP computationally feasible. Given this DP formulation, we can prove the structure of the optimal inventory control policy with known θ:

Proposition 3.1 *When the degradation parameter, θ, is known, an order-up-to inventory policy with order-up-to level $\bar{y}_n(\theta, z_n)$ is optimal for period n, for all $n = 1, \ldots, N$. Notice that the optimal order-up-to level for this problem, $\bar{y}_n(\theta, z_n)$, depends only on θ and the most recent degradation signal observation, z_n.*

We next extend this DP formulation to the unknown parameter case.

$$C_n(x_n, z_1, \ldots, z_n) = \min_{y \geq x_n} \left\{ c(y - x_n) + L_n(y|z_1, \ldots, z_n) + \alpha \sum_{k=0}^{\infty} E\left[C_{n+1}(y - k, Z_{n+1}|z_1, \ldots, z_n) 1_{\{D_n = k|z_1, \ldots, z_n\}} \right] \right\}. \quad (6)$$

The expected cost-to-go, $C_n(x_n, z_1, ..., z_n)$ depends on all of the degradation signal observations up to period n, since these observations are required to calculate the posterior distribution of the degradation parameter. Hence, the state space of the DP includes the current inventory level, x_n, and the complete history of degradation signal observations, $z_1, ..., z_n$, which makes the computation of the optimal inventory policy infeasible. However, it is possible to reduce the state space by exploiting the conditional independence existing in the updating procedure for the posterior mean and variance of θ. In other words, given the two most recent signal observations, z_{n-1} and z_n, the observed demand in the previous period, denoted d_{n-1}, and the most recent posterior mean and variance, μ_n and σ_n^2, the posterior mean and variance for the next period, μ_{n+1} and σ_{n+1}^2, are conditionally independent of all degradation signal observations prior to period $n-1$. With this property, we can reformulate the DP as follows:

$$C_n(x_n, z_n, z_{n-1}, d_{n-1}, \mu_n, \sigma_n) =$$
$$\min_{y \geq x_n} \left\{ c(y - x_n) + L_n(y|z_n, \mu_n, \sigma_n) + \alpha \sum_{k=0}^{\infty} E\left[C_{n+1}(y-k, Z_{n+1}, z_n, k, \mu_{n+1}, \sigma_{n+1}) 1_{\{D_n = k\}}\right] \right\} \quad (7)$$

where μ_{n+1} and σ_{n+1} can be computed using the Bayesian approach described above. Given this DP formulation, we can prove the structure of the optimal policy for θ unknown:

Proposition 3.2 *When the degradation parameter, θ, is unknown, an order-up-to inventory policy with order-up-to level $\bar{y}_n(z_n, z_{n-1}, d_{n-1}, \mu_n, \sigma_n)$ is optimal for period n, for all $n = 1, \ldots, N$. Notice that the optimal order-up-to level for this problem, $\bar{y}_n(z_n, z_{n-1}, d_{n-1}, \mu_n, \sigma_n)$, depends on two recent observations (z_n, z_{n-1}), demand in previous period d_{n-1}, and posterior mean μ_n and standard deviation σ_n of the degradation parameter.*

Next, consider the problem with multiple machines and known degradation parameters. We would like to find the optimal policy for each state configuration $x_n, z_n^1, z_n^2, \ldots, z_n^m$, where z_n^i denotes the most recently observed degradation signal for machine i, and m is the number of machines. For this DP there will be m continuous state variables and one discrete state variable. Thus, the problem with multiple machines and known θ_i is likely to be computationally infeasible. When the degradation parameters are unknown, the DP becomes even more difficult to solve, since the state space will include the on-hand inventory level for the system, x_n, as well as $z_n^i, z_{n-1}^i, d_{n-1}^i, \mu_n^i, \sigma_n^i$ for each machine i, $i = 1, \ldots, m$. However, for both cases, i.e., for θ known and unknown, as above, we have shown that the optimal inventory policy is a dynamic base stock policy.

There are several approaches that can be used to solve such large-dimension DP problems. First, one can reduce the high-dimensional DP to a revised DP with a single-state variable under specific assumptions regarding the demand distribution. Alternatively, to solve this complex problem, we can resort to various approximation approaches, including Monte Carlo simulation of the optimal policy, and the use of myopic heuristic policies. In the following sections, we will briefly discuss both of these approaches to solving this complex inventory control problem.

3.2 Simulation-based Solution Approach

Since solving the DP to find the optimal inventory policy is computationally infeasible in the case of multiple machines or a single machine with unknown degradation parameter, we have developed solution approaches based on Monte Carlo simulation. The basic approach to find the optimal policy using simulation is as follows: (1) simulate the degradation processes and get an empirical estimate of the demand distribution for each period in the remainder of the planning horizon; and (2) given the estimated demand distributions, find the optimal policy by solving a one-dimensional DP, in which the current inventory level is the only state variable. Algorithm 1 outlines the basic steps for determining the optimal policy with simulation for the case of a single machine with unknown degradation parameter.

Algorithm 1 Simulation strategy for a single machine with unknown degradation parameter

For $n = 1, 2, \ldots, N$

1. **Simulate the degradation process**
   for $i = 1, 2, \ldots, S$ (number of samples)
   - Initialize the degradation signal with z_n;
   - Initialize the mean μ_θ and variance σ_θ^2 of the degradation parameter as μ_n, σ_n^2;
   - for $k = n, \ldots, N$ (simulate signal for the remainder of the planning horizon)
     - Draw drift parameter θ from distribution $N(\mu_\theta, \sigma_\theta^2)$;
     - Generate signal until time $k+1$ to obtain z_{k+1};
     - Update parameters μ_θ and σ_θ based on z_k, z_{k+1} and d_k;
   - Count the number of demands d_k^i in each period k for sample i;
2. **Find the optimal policy**
   - Estimate the demand distributions in periods n to N with the samples;
   - Solve a one-dimensional DP to find the optimal policy given the estimated demand distributions.

When the degradation parameter is known, a similar approach can be used. However, we do not need to draw a drift parameter from the posterior distribution or update the posterior distribution. For a system with multiple machines, we need to simulate the degradation processes for all machines, and then sum their demands in each period to estimate the demand distribution. In summary, with the Monte-Carlo simulation of the demand distribution, we convert a high-dimensional DP problem to a one-dimensional DP problem, which can be solved efficiently.

3.3 Myopic Critical Fractile Policy

Given the distribution of demand for replacement parts, a myopic critical fractile policy can be easily computed. For this policy, the myopic order-up-to level is just the smallest value of x satisfying $F_{D_n}(x) \geq \frac{p-c+\alpha c}{h+p}$, where F_{D_n} is the cumulative distribution of demand in period n, α is the discount factor, and c, h and p are the unit ordering cost, holding cost and penalty cost.

Lovejoy (1990) develops an upper bound on the performance gap between the optimal policy and the myopic critical fractile policy using a myopic policy with disposal. In a traditional myopic policy, we only have the choice to order or not order when making the inventory decision. But in the myopic policy with disposal, we can also choose to dispose of some on-hand inventory, with unit cost c_s. Generally, we can expect a lower inventory cost for the policy with disposal since we have the option to dispose of inventory if the holding cost is too high. Let C_o denote the expected cost over the entire planning horizon under the optimal inventory policy, let C_m denote the expected cost under the myopic policy and let C_{md} denote the expected cost under the myopic policy with disposal, with $c_s = -c$. Then, from Lovejoy (1990), we have the following upper bound on the relative cost difference between the optimal policy and myopic policy:

$$\frac{C_m - C_o}{C_o} \leq \frac{C_m - C_{md}}{C_{md}}. \tag{8}$$

4. Computational Study

We next describe the computational study that we have conducted for this spare parts inventory problem with condition monitoring information. In the experiments described in this section, we use the default parameters shown in Table 1.

Table 1: Default parameters of the inventory model

Drift parameter, θ (single machine)	0.0002	# of periods, N	40
Prior mean of the drift parameters	0.0001	Period length, t_0	24 hours
Prior variance of the drift parameter	0.00005	Ordering cost, c	3
Brownian motion variance, σ^2	$6e^{-6}$	Holding cost, h	1
Failure threshold, B	0.055	Penalty cost, p	10

4.1 Demand Estimation

In this section, we present several examples of the estimated demand distribution generated from the condition monitoring observations. Figure 1 shows the demand distribution for a single period and a single machine for several values of the observed degradation signal, z. We can see that, when z increases, the mean of the demand distribution increases. Figures 2 shows the demand distribution with different numbers of machines in the unknown parameter case. We find that the mean demand increases roughly linearly with the number of machines, while the standard deviation of demand increases more slowly. This observation is a demonstration of the risk-pooling effect.

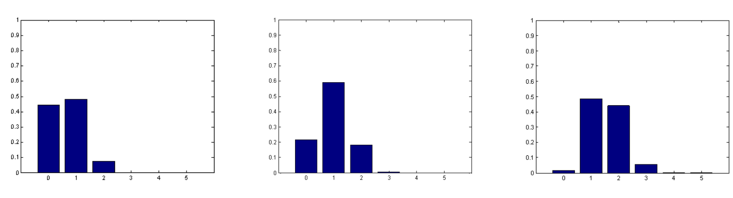

(a) z=0 (b) z=0.01 (c) z=0.05

Figure 1: Estimated demand distribution with different signal observations

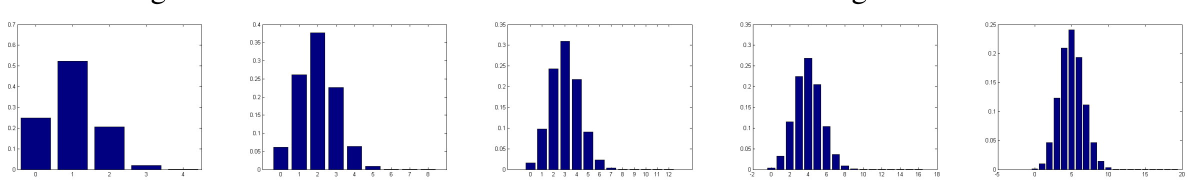

(a) m=1 (b) m=2 (c) m=3 (d) m=4 (e) m=5

Figure 2: Demand distribution for different number of machines with unknown parameters

4.2 Value of the Condition Monitoring

We next evaluate the value of condition monitoring in managing the inventory of spare parts. Figure 3(a) compares the total cost with and without condition monitoring information for a single machine. Here, the true value of θ is set as 0.0001, and we vary our prior belief about θ from 0.0001 to 0.0003. We compare two solution methods: (1) estimate the demand using only the prior belief about θ, and (2) estimate the demand using the Bayesian approach to estimate the posterior distribution of θ. From our results, we conclude that, when the prior belief about θ is further away from the true value of θ, the Bayesian estimation approach has the most value. Figure 3(b) presents a similar comparison for a system with multiple machines. In this case, we set the prior mean μ of the drift parameter to be 0.0002 and varied the standard deviation from $0, \frac{1}{8}\mu, \frac{1}{4}\mu, \frac{1}{2}\mu$ to μ. For each standard deviation value, we first sample the true drift parameters from the prior distribution, and then compare the inventory costs for two cases: (1) we update the posterior distribution of θ using the Bayesian approach and (2) we treat the prior mean 0.0002 as the known value of the parameter θ. We find that, when the prior variance of the drift parameter is large, the Bayesian updating approach has the most value.

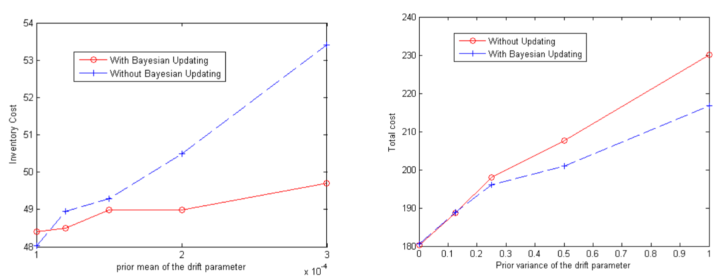

(a) Single machine (b) Multiple machines

Figure 3: Comparison of the inventory cost with and without condition monitoring information

4.3 Optimal vs. Myopic Policies

We next compare the optimal policy (using simulation) with the myopic policy for the case of unknown degradation parameters. See Tables 2 and 3 below.

Table 2: Performance gap between optimal and myopic policies with unknown parameters

# machines	Average	Min	Max
1	3.45%	2.95%	4.4%
5	0.58%	0.03%	1.2%

Since the optimal policy is computed using simulation, we evaluate the relative cost difference, r, where $r = (C_m - C_o)/C_o$, in ten different runs and report the mean, minimum and maximum, where C_o and C_m are the total inventory costs for the optimal policy and myopic policy, respectively. In the single machine case, the relative cost difference between the optimal policy and the myopic policy is only 3.45% on average, and 4.4% in the worst case, while the computational cost is significantly reduced (from 28.2s to 0.41s on average) with the myopic policy. With 5 machines, the average and maximum relative difference across the ten different runs are only 0.58% and 1.2%, respectively. It is interesting to find that the performance gap is smaller if we have multiple machines. This result shows that pooling the demand from multiple machines not only reduces the total inventory cost due to the risk pooling effect, but also reduces the performance gap between the optimal policy and the myopic policy. This property is especially important for practical applications, in which we expect to have a large number of machines and in which the optimal policy may be extremely difficult to compute due to the limited computational resources.

Table 3: Computational costs of the optimal policy and myopic policy with unknown parameters

# machines	Policy	Average(s)	Min (s)	Max (s)
1	Optimal	28.2	27.8	28.7
	Myopic	0.41	0.23	0.78
5	Optimal	61.5	60.2	64.6
	Myopic	18.8	18.1	19.2

5. Conclusion

In this paper, we have developed a framework for incorporating condition monitoring information into the inventory management decisions for spare parts. We extend the Bayesian degradation signal estimation approach developed in Gebraeel et al. 2005 to derive the demand distribution for service parts, and then demonstrate how this demand distribution can be incorporated to a classical stochastic inventory model. We then proposed and evaluated three different solution approaches for this inventory problem and demonstrated that a myopic critical fractile policy can perform quite well relative to the optimal policy, while significantly reducing the computational time. In our future research, we are interested in extending this framework to consider different condition monitoring techniques and degradation models.

6. References

Cox, D. R., H. D. Miller. 1965. The Theory of Stochastic Processes. Wiley, New York, NY.

Doksum, K. A., A. Hoyland. 1992. Models for variable-stress accelerated life testing experiments based on Wiener process and the inverse Gaussian distribution. Technometrics, 34, 74-82.

Dvoretzky, A., J. Kiefer, J. Wolfowitz. 1952. The Inventory Problem: II. Case Of Unknown Distributions Of Demand. Econometrica, 20, 450-466.

Gebraeel, N., M. Lawley, R. Li, J. K. Ryan. 2005. Residual Life Distributions from Component Degradation Signals: A Bayesian Approach. IIE Transactions on Reliability. 37, 543-557.

Karlin, S. 1960. Dynamic Inventory Policy with Varying Stochastic Demands. Management Science. 6 (3), 231-258.

Lovejoy, W. S. 1990. Myopic Policies For Some Inventory Models With Uncertain Demand Distributions. Management Science. 36 (6), 724-738.

Lu, C., W. Meeker. 1993. Using degradation measures to estimate a time-to-failure distribution. Technometrics. 35, 161-174.

Scarf, H. E. 1959. Bayes Solutions of the Statistical Inventory Problem. Annals of Mathematical Statistics. 30 (2), 490-508.

Wang, W. 2000. A model to determine the optimal critical level and the monitoring intervals in condition-based maintenance. International Journal of Production Research. 38, 1425-1436.

Whitmore, G. A. 1995. Estimating degradation by a Wiener diffusion process subject to measurement error. Lifetime Data Analysis. 1, 307-319.