Journal of the Operational Research Society (2004) 55, 647–656

www.palgrave-journals.com/jors

A periodic review production and maintenance model with random demand, deteriorating equipment, and binomial yield

TW Sloan

University of Miami, Coral Gables, USA

In many environments, product yield is heavily influenced by equipment condition. Despite this fact, previous research has either focused on the issue of maintenance, ignoring the effect of equipment condition on yield, or has focused on the issue of production, omitting the possibility of actively changing the machine state. We formulate a Markov decision process model of a single-stage production system in which demand is random. The product yield has a binomial distribution that depends on the equipment condition, which deteriorates over time. The objective is to choose simultaneously the equipment maintenance schedule as well as the quantity to produce in a way that minimizes the sum of expected production, backorder, and holding costs. After proving some results about the structural properties of the optimal policy, numerical problems are used to compare this method to the typical approach of solving the maintenance and production problems sequentially. The results show that the simultaneous solution provides substantial gains over the sequential approach. In the cases studied, the proposed method resulted in an average cost savings of approximately 18%.

Journal of the Operational Research Society (2004) **55,** 647–656. doi:10.1057/palgrave.jors.2601725

Keywords: inventory; maintenance; production; quality; scheduling

Introduction

Product quality and yield are heavily dependent on equipment condition in a wide range of manufacturing environments. For example, as drill bits in a machine shop wear out, finished product quality will deteriorate, and the number of functioning chips emerging from an integrated circuit manufacturing line will decrease as the ultra-clean production environment becomes contaminated. The goal of this paper is to answer the question: how can information about yield and equipment condition be used to improve decisions about machine maintenance and inventory management? By exploiting all the available information, firms can increase quality, reduce inventory costs, and better meet customer demand.

Despite the apparent connections, problems addressing equipment condition, product yield, and inventory management have traditionally been treated independently. Equipment maintenance research has focused on the problem of when to repair or replace deteriorating equipment in order to minimize costs, and has ignored the potential impact of equipment condition on product quality and inventory management. Most inventory models that incorporate random yield treat equipment condition as something that is beyond the decision maker's control. Models that do allow control of the process condition consider only two states, in-control and out-of-control, and all products manufactured when the process is out-of-control are defective. In short, none of these models addresses the situation in which there are intermediate machine states and the decision maker exercises some control over the machine condition.

To explore the interaction between yield, equipment condition, and inventory management, we formulate a Markov decision process (MDP) model of a single-stage production system in which demand is random and the equipment condition, which deteriorates over time, affects the product yield. The decision maker simultaneously chooses the equipment maintenance schedule as well as the production quantity in a way that minimizes the sum of expected (discounted) production, backorder, and holding costs. We refer to this as the *simultaneous* approach. After proving some results about the structural properties of the optimal policy, numerical problems are used to compare the proposed policy to the typical method, which solves the maintenance and inventory problems separately. We refer to the latter solution method as the *sequential* approach. The results of more than 26 000 problems show that the simultaneous solution provides an average cost savings of approximately 18% as compared to the sequential approach.

Correspondence: TW Sloan, Department of Management, School of Business Administration, 414 Jenkins Building, University of Miami, Coral Gables, FL 33124-9145, USA.
E-mail: tsloan@miami.edu

The paper proceeds as follows. First, we review the literature related to this problem. Next, we present the MDP model and examine the properties of the optimal policy. We then present some results of the numerical problems. The final section presents conclusions and discusses directions of future research.

Literature review

The literature related to this problem can be divided into three categories. The first category includes models of production/inventory systems in which yield is a random variable, but process condition is not under the decision maker's control. Models in the second category allow some control of the process condition. The third category includes production models with supply disruptions.

The main question addressed by papers in the first category is: how much should be ordered or produced given the uncertainty regarding yield? Yano and Lee[1] present a detailed review of the research in this category, and we mention only a few representative papers here. Gerchak *et al*[2] and Henig and Gerchak[3] examine single-stage, periodic-review systems with random demand and yield. Wang and Gerchak[4] and Parlar and Perry[5] consider systems with uncertain capacity/supply as well. All these papers primarily focus on the structural properties of the solution. Mazzola *et al*,[6] Baker and Ehrhardt,[7] and Bollapragada and Morton[8] focus on developing heuristic solution procedures to solve for near-optimal policies efficiently, making different assumptions about demand variability, cost structure, *etc*. Lee and Yano[9] and Barad and Braha[10] explore input control decisions for multi-stage systems with deterministic demand. Gerchak *et al*[11] and Gurnani *et al*[12] examine multi-stage systems with random demand. In all the models in this category, the effect of the equipment or process condition is not explicitly linked to product yield or is assumed to be beyond the decision maker's control.

The second literature category includes production/inventory models that explicitly make the link between equipment condition and yield by extending the classical economic manufacturing quantity (EMQ) model to include the possibility of imperfect process condition, and hence defective output. Rosenblatt and Lee[13] and Porteus[14] examine the case in which product quality is determined only after production. Lee and Rosenblatt[15] and Porteus[16] allow inspection during the production cycles. Lee and Park[17] treat the case where defects can be discovered before or after the item is sold, with different item costs depending on the outcome. Makis and Fung[18] allow imperfect process condition and equipment failure. In these models, the demand is deterministic, and the randomness associated with yield is reduced to one dimension: either the process is in-control, resulting in output of perfect quality, or the process is out-of-control, resulting in defective output.

The third literature category related to our problem includes lot-sizing models that incorporate supply disruptions. Groenevelt *et al*[19,20] extend the EMQ model to include the possibility of machine failure, but yield and demand are both deterministic. Arreola-Risa and DeCroix[21] study a situation in which demand is random and supply can be disrupted for random periods. In all these models, there are only two states: operational and non-operational. Venkatesan[22] models a periodic-review production/inventory system with a single machine that deteriorates over time. At the beginning of each period, one decides not only how much to produce but also whether or not to replace the machine. While the equipment condition clearly affects the *quantity* of production, the quality of output is assumed to be perfect in each of these models.

Sloan and Shanthikumar[23] examine a single-stage, multi-product system with variable yield and deteriorating equipment. Their model is related to the current problem in that the equipment condition can be reset to improve yield; however, demand does not vary and backlog and holding costs are not considered. Thus, the focus is on determining which product(s) to produce rather than how many to produce.

In summary, our model is the first of a system in which demand is stochastic, yield is variable, equipment condition is explicitly linked to product yield, and the decision maker exercises some control over the machine condition.

Model

We consider the problem of determining production quantities for a single-product, single-stage production system with uncertain demand and variable yield. Units are processed on a single machine whose condition deteriorates over time according to a Markov chain with known transition probabilities, and the level of deterioration affects the quality of the output and the operating costs. At the beginning of each period, we observe the state of the machine and the inventory level and make two decisions: whether or not to repair the machine and how many units to produce. If we choose to repair, the machine instantaneously returns to like-new condition with probability one. (Referring to our earlier examples, this would correspond to replacing the drill bit in the machine shop or cleaning the equipment in the integrated circuit plant.) The decision of how many units to release into the system must account for the fact that demand is uncertain, yield is uncertain, and there are costs associated with inventory shortages and excesses.

Model formulation

We introduce the following notation and definitions:

n	number of periods remaining
I_n	machine state; $I_n \in \{0, 1,\ldots, M\}$
X_n	inventory level; $X_n \in \{\ldots, -2, -1, 0, 1, 2,\ldots\}$
a_n	repair action taken; $a_n = 1$ if the machine is repaired, $a_n = 0$ otherwise
D_n	number of units demanded; $D_n \in \{0, 1, 2, \ldots\}$
Q	maximum production (input) quantity
q_n	input quantity; $q_n \in \{0,\ldots, Q\}$
Z_n	output quantity; $Z_n \in \{0,\ldots, Q\}$
β_i	probability that each unit produced while the machine is in state i functions properly
p_{ij}^{aq}	the probability that the machine is in state j at the beginning of the next period given that in the current period the machine is in state i and one takes repair action a and inputs production quantity q
$R(a)$	repair cost function; $R(Q) = 0$ and $R(1) = R$
h	holding cost per unit per period
b	backlog cost per unit per period
c	production cost per unit of input
α	discount factor; $0 \leqslant \alpha < 1$
$\mathscr{S}$	state space: $\{\ldots, -2, -1, 0, 1, 2, \ldots\} \times \{0, 1, \ldots, M\}$
$\mathscr{A}$	action space: $\{0, 1\} \times \{0,\ldots, Q\}$.

The machine state transitions depend only on the current state and the action taken. Specifically, $p_{ij}^{aq} = \Pr\{I_{n-1} = j | I_n = i,\ a_n = a,\ q_n = q\}$. We assume that the equipment deterioration is a result of production; any non-zero production quantity can make the machine condition worse. However, if no units are produced and the machine is not repaired, then the machine state does not change. If the machine reaches state M, it cannot leave this state unless it is repaired. For a current machine state i and repair action $a = 0$, these conditions are expressed as

$$p_{ij}^{0q} = \begin{cases} p_{ij} & \text{for } q > 0;\ i = 0, 1, \ldots, M-1;\ j = 0, 1, \ldots, M-1 \\ 1 & \text{for } q = 0;\ i = 0, 1, \ldots, M-1;\ j = i \\ 0 & \text{for } q = 0;\ i = 0, 1, \ldots, M-1;\ j \neq i \\ 1 & \text{for } q = 0,\ 1, \ldots, Q;\ i = M;\ j = M \\ 0 & \text{for } q = 0,\ 1, \ldots, Q;\ i = M;\ j \neq M \end{cases}.$$

We assume that the repair action can be taken while in any state and that repairing the machine instantaneously returns it to state 0 with probability one. Thus, when $a = 1$ we have $p_{ij}^{1q} = p_{0j}^{0q}$ for each $i = 0, 1,\ldots, M$, $j = 0, 1,\ldots, M$, and $q = 0, 1,\ldots, Q$.

We also assume that the machine deterioration rate is relatively fast as compared to the production rate: otherwise, little could be gained by considering the interaction between maintenance and production scheduling. For example, if it takes weeks of use for the machine to deteriorate from one state to the next, but it takes only hours to produce a product, then maintenance decisions would not have much of an effect on the weekly production schedule.

Inventory transitions depend on the current inventory level, the demand, and the production quantity: $X_{n-1} = X_n - D_n + Z_n$. Since the state transitions depend only on the current state and action, we can treat the problem as an MDP.

The output quantity Z_n is a random variable that depends on input quantity q_n. We assume that the yield (output) follows a binomial distribution and define $\phi_i^a(k|q)$ as the probability of having k units of output that function properly when the machine state is i, action a is taken, and the input quantity is q. When $a = 0$ we have

$$\phi_i^0(k|q) = \binom{q}{k} (\beta_i)^k (1 - \beta_i)^{q-k}$$

When $a = 1$, we have $\phi_i^1(k|q) = \phi_0^0(k|q)$ for each i, since the machine is returned to state 0 with probability one.

We assume that demand is independent and identically distributed in each period, that is, $D_n = D$ for each n, and denote the associated probability mass function as $\xi(d) \equiv \Pr\{D = d\}$. We assume that the demand probability distribution is well-behaved (eg, has finite first and second moments), but we do not assume a specific distribution or functional form.

Each period unfolds as follows. First, we observe the machine state and the inventory level. Next, we choose whether or not to repair the machine and how many units to start into the system. We then experience the demand. Finally, costs are incurred.

Define $V_n[(i, x), (a, q)]$ as the discounted expected cost when the initial machine state is i, the initial inventory level is x, repair action a is taken, q units are started into the system, and n periods remain. The objective is to determine the maintenance schedule and production quantity that minimizes the sum of expected repair, production, and backlog/holding costs. Let $V_n(i, x)$ denote the minimal discounted expected cost, that is, $V_n(i, x) = \min_{(a,q)}\{V_n[(i, x), (a, q)]\}$. The minimal cost is found by solving the following dynamic programming recursion:

$$V_n(i,x) = \min_{(a,q)} \left\{ R(a) + cq + \sum_{d=0}^{\infty} \sum_{k=0}^{q} \left[b(d-x-k)^+ + h(x+k-d)^+ + \alpha \sum_{j=0}^{M} p_{ij}^{aq} V_{n-1}(j, x+k-d) \right] \times \phi_i^a(k|q) \xi(d) \right\} \quad (1)$$

where $[y]^+ = \max[0, y]$. For the single-period problem, we have

$$V_1(i,x) = \min_{(a,q)} \left\{ R(a) + cq + \sum_{d=0}^{\infty} \sum_{k=0}^{q} [b(d-x-k)^+ + h(x+k-d)^+] \phi_i^a(k|q) \xi(d) \right\} \quad (2)$$

We assume that all costs are bounded and that $c < b$ and $b > h$. We also assume that there is sufficient capacity to meet demand on average, that is, $Q\beta_0 \geqslant E[D]$.

Note that the number of machine states is finite, the number of inventory states is countable, and the number of actions is finite. Applying standard MDP theory arguments, the value function is well-behaved, and an optimal stationary policy will exist.[24]

Structural properties

To characterize the structure of the optimal policy, we make two additional assumptions. First, we assume that the yield decreases as the machine condition gets worse. Second, we assume that as the machine deteriorates, it is more likely to go to a worse state than a better state. These assumptions can be written as

(A1) β_i is non-increasing in i.
(A2) For each l, $\sum_{j=l}^{M} p_{ij}$ is non-decreasing in i.

To further describe the structure of the optimal policy, we need the following lemma:

Lemma 3.1 *The single-period value function $V_1(i, x)$*

(a) *is non-decreasing in i,*
(b) *is convex, and*
(c) *has a unique bounded solution.*

Proof By assumption (A1), β_i is non-increasing in i. In other words, the yield decreases as the machine condition gets worse. Clearly, it is impossible to achieve lower costs when the yield is lower. Therefore, we may conclude that $V_1(i, x)$ is non-decreasing in i. The term in square brackets in Equation (2) represents the total backlog and holding costs. Since these costs are linear, this element of the cost function is convex (in a discrete sense), and the convexity is preserved under expectation.[24] Since $V_1(i, x)$ is convex and all costs are bounded, Equation (2) will have a unique bounded solution.

Define $a(i, x)$ as the smallest repair action that minimizes the right-hand side of (2) when the system is in state (i, x); that is, $a(i, x)=1$ if the equipment is repaired, and $a(i, x)=0$ otherwise. Define $q(i, x)$ as the smallest production quantity that minimizes the right-hand side of (2) when the system is in state (i, x). The next proposition states that a production threshold exists; that is, there is an inventory level below which the optimal production quantity will be greater than zero. However, the optimal production quantity is not monotone with respect to the machine state.

Proposition 3.1 *There exists an inventory level $\hat{x}$ such that if $q(i, \hat{x})>0$, then $q(i, x)>0$ for all $x \leqslant \hat{x}$ and $q(i, x)=0$ for all $x>\hat{x}$.*

Proof The result follows from the convexity of $V_1(i, x)$. By definition, $q(i, \hat{x})>0$ implies that it is preferable to raise the inventory level than to produce nothing. Then clearly it will be preferable for any $x \leqslant \hat{x}$, so $q(i, x)>0$ for all $x \leqslant \hat{x}$.

Proposition 3.2 *The optimal production quantity, $q(i, x)$ is not monotone with respect to the machine state.*

Proof Assume that $q(i, x)$ is non-decreasing in i; that is, the input quantity will increase as the machine condition gets worse, so we would have $q(i, x) \leqslant q(i+1, x)$ for $i=0, 1, \ldots, M-1$.

However, if $a(i, x)=0$ and $a(i+1, x)=1$, then $q(i+1, x)=q(0, x)$, since the machine is returned to state 0. Since $q(i+1, x)=q(0, x) \leqslant q(i, x)$, we have a contradiction.

Figure 1 illustrates the meaning of Proposition 3.2 by plotting values of $V_1(i, x)$ for different input quantities. As shown, $q(0, x)=q_0^*<q_1^*=q(1, x)$. However, $q(2, x)=q_0^*$ as well, demonstrating that the production quantity is not monotone with respect to the machine state.

The following proposition states that a repair threshold exists and that the threshold increases as the inventory level increases. In other words, we will wait to repair the machine if there is sufficient inventory on hand.

Proposition 3.3 *There exists a machine state $\hat{\imath}$ such that if $a(\hat{\imath}, x)=1$, then $a(i, x)=1$ for all $i \geqslant \hat{\imath}$ and $a(i, x)=0$ for all $i<\hat{\imath}$. Furthermore, $a(i, x)$ is non-increasing in x.*

Proof Suppose that it is optimal to repair in state (i, x), that is, $a(i, x)=1$. From (2) this means that $R+V_1(0, x) \leqslant V_1(i, x)$. Since $V_1(i, x)$ is non-decreasing in i, this must also mean that $R+V_1(0, x) \leqslant V_1(i+1, x)$. Thus, repair will be optimal in state $i+1$ and any greater state as well.

Similarly, suppose that it is optimal not to repair in state (i, x), that is, $a(i, x)=0$. This implies that $V_1(i, x) \leqslant R+V_1(0, x)$. Now suppose that it is optimal to repair in state $(i, x+1)$, that is, $a(i, x+1)=1$. This implies that $V_1(i, x+1)=R+V_1(0, x+1)$. Clearly, one would not repair unless $q(i, x+1)>0$. By Proposition 3.1, if $q(i, x+1)>0$, then $q(i, x)>0$. Furthermore, the convexity of $V_1(\cdot)$ implies that

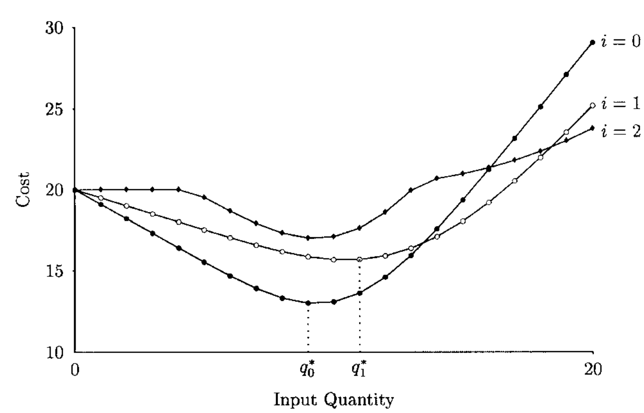

Figure 1 Sample plot of $V_1(i, x)$ as a function of input quantity.

$q(i, x) \geqslant q(i, x+1)$. But the only incentive to repair in state $(i, x+1)$ is to achieve an expected output level *greater* than that possible without repairing. In other words, one would repair the machine only when the repair cost is outweighed by the expected increase in output. But if this were the case for $x+1$, then it would also be true for x and any lower inventory level. Thus, we have a contradiction.

Infinite-horizon problem

The infinite-horizon equivalent of (1) is written as

$$V(i,x) = \inf_{(a,q)} \left\{ R(a) + cq \sum_{d=0}^{\infty} \sum_{k=0}^{q} \left[b(d-x-k)^{+} + h(x+k-d)^{+} + \alpha \sum_{j=0}^{M} p_{ij}^{aq} V(j, x+k-d) \right] \times \phi_i^a(k|q)\xi(d) \right\} \tag{3}$$

First, we must show that a solution exists.

Lemma 3.2 *$V(i, x) = \lim_{n\to\infty} V_n(i, x)$ exists for every (i, x).*

Proof Let π be any policy and let $W_\pi(i, x)$ denote the expected present value of using this policy when the initial state is (i, x). In other words, $W_\pi(i, x)$ is a return function analogous to (3) without the 'inf' function. We will apply Theorem 8–13 in Heyman and Sobel,[24] which states that if $W_\pi(i, x) < \infty$ for all (i, x), then $V(i, x) = \lim_{n\to\infty} V_n(i, x)$ exists for every (i, x). We can choose, for example, a policy that repairs the machine each period and inputs a quantity equal to the demand from the previous period; that is, $a_n = 1$ and $q_n = \min\{Q, D_{n+1}\}$ for each period n. Clearly, the single-period costs are bounded, and since $\alpha < 1$, we have $W_\pi(i, x) < \infty$ for all (i, x), and the result follows.

Now we can explore the structure of the infinite horizon problem.

Proposition 3.4 *$V(i, x)$ is non-decreasing in i.*

Proof By Lemma 3.1, $V_1(i, x)$ is non-decreasing in i. Now suppose that $V_{n-1}(i, x)$ is non-decreasing in i. By assumption (A2), $\Sigma_{j=l}^{M} p_{ij}$ is non-decreasing in i for each l. This implies that $\Sigma_{j=0}^{M} p_{ij} f(j)$ is non-decreasing for any non-decreasing function $f(\cdot)$.[25] By induction, we can conclude that $V_n(i, x)$ is also non-decreasing in i. By Lemma 3.2, $V_n(i, x) \to V(i, x)$ as $n \to \infty$, so $V(i, x)$ is also non-decreasing in i.

Proposition 3.5 *For the infinite-horizon problem, there exists a machine state $\hat{\imath}$ such that if $a(\hat{\imath}, x) = 1$, then $a(i, x) = 1$ for all $i \geqslant \hat{\imath}$ and $a(i, x) = 0$ for all $i < \hat{\imath}$.*

Proof The proof follows the same reasoning as the proof of Proposition 3.3.

Any standard MDP solution method may be used to solve (3) for the optimal policy. We use a straightforward policy improvement algorithm, the details of which are described in Appendix A.

Sequential approach

Maintenance and production problems have traditionally been treated independently. We use the phrase *sequential approach* to refer to the process of solving the two problems separately rather than simultaneously. First, the maintenance schedule is determined by finding the point (machine state) at which it is cheaper, in terms of operating costs, to repair the machine than not to repair the machine.

In the long run, we would expect to produce enough each period to meet the expected demand. Since the yield is β_i when the machine is in state i, an input quantity of $E[D]/\beta_i$ would be needed to attain an output quantity of $E[D]$, the expected demand. So the expected operating costs for state i would be $C(i) \equiv c(E[D]/\beta_i)$, and $C(i)$ is non-decreasing in i. The problem is reduced to an elementary equipment maintenance model such as that proposed by Derman.[25] The objective at this stage is to determine the repair threshold, that is, the machine state at which it becomes preferable to stop and repair the machine rather than continue producing. The threshold can be easily determined using any standard MDP solution method. More details are discussed in Appendix B.

The maintenance schedule tells us the repair action for each machine state, denoted as a_i, and thus fixes the machine state transitions. Now the production quantity can be determined by solving (3) with the modification that the repair action is specified by a_i:

$$V(i,x) = \inf_{q} \left\{ R(a_i) + cq + \sum_{d=0}^{\infty} \sum_{k=0}^{q} \left[b(d-x-k)^{+} + h(x+k-d)^{+} + \alpha \sum_{j=0}^{M} p_{ij}^{a_i q} V(j, x+k-d) \right] \times \phi_i^{a_i}(k|q)\xi(d) \right\} \tag{4}$$

This cost function is very similar to (3) and can be solved using the same policy improvement algorithm mentioned above and described in detail in Appendix A.

Clearly, the simultaneous approach will be better than the sequential approach, but the question of interest is: how much better? In the next section, we attempt to answer this question.

Numerical results

In both academic and industrial arenas, the problems of maintenance scheduling and lotsizing have traditionally been treated independently. The primary question of interest is: how much of a difference will it make to use information about equipment condition and yield to solve simultaneously these two problems? Thus, our goal is to compare the simultaneous approach outlined in the previous section with the traditional, sequential approach.

Overview of problems

We examine a single-stage, single-machine system that produces a single product. The machine has five states, $i \in \{0, 1, 2, 3, 4\}$, where state 0 indicates the best state and state 4 indicates the worst state. To get a sense of how the results are influenced by the different parameter values, we treat each parameter as a 'factor' in an experimental design and test each at different 'levels'. Each factor is tested at three levels except for the demand probability function, which is tested at four levels. A 'full factorial' design yields a total of 26 244 test problems.

Table 1 reports each factor and level tested. The equipment deterioration 'rate' listed in Table 1 refers to the likelihood of the machine state getting worse; a fast rate means that the probability of going to a worse state is high, and a low rate means that the probability of going to a worse state is low. Table 2 reports the actual machine state transition probabilities used. The yield values—low, medium, and high—are listed in Table 3. Four different demand distributions are tested: deterministic, binomial, (discrete) uniform, and geometric. The distribution parameters are set for a given value of $E[D]$ in such a way that the variance increases going from deterministic to geometric. Specifically, the coefficient of variation for deterministic demand is 0, for binomial demand is approximately 0.25, for uniform demand is approximately 0.6, and for geometric demand is approximately 1. (The specific demand parameter values used are reported in Appendix C.) For all problems, the production cost is $c = 1$ per unit of input.

The optimal policies are determined for the simultaneous and sequential approaches using a simple policy improvement algorithm to solve the optimality equations, (3) and (4), respectively. This solution procedure, the details of which are described in Appendix A, is an iterative process, and one must specify at what point to terminate it. For all the problems in our study, the search for an optimal policy is stopped when the cost for a particular state changes by less than 0.01% from one iteration to the next. In addition, the inventory level could theoretically become infinitely large or small. To address this issue, we truncate the state space at a point that is well beyond the maximum probable change in

Table 1 Summary of factors and levels for test problems

Factor	*Description*	*Factor values*
α	Discount factor	0.5, 0.7, 0.9
R	Repair cost	20, 40, 80
h	Holding cost per unit per period	0.5, 1, 2
b	Backlog cost per unit per period	5, 10, 20
Q	Maximum production quantity	12, 15, 20
$[p_{ij}]$	Equipment deterioration rate	Slow, med., fast
$[\beta_i]$	Yield values	Low, med., high
$E[D]$	Mean demand per period	6, 9, 12
$\xi\ (d)$	Demand probability function	Deterministic, binomial, uniform, geometric

Table 2 Equipment deterioration probabilities (p_{ij})

Level	*Matrix*	*Level*	*Matrix*	*Level*	*Matrix*
Slow	$\begin{bmatrix} .9 & .1 & 0 & 0 & 0 \\ 0 & .9 & .1 & 0 & 0 \\ 0 & 0 & .9 & .1 & 0 \\ 0 & 0 & 0 & .9 & .1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	Med.	$\begin{bmatrix} .5 & .5 & 0 & 0 & 0 \\ 0 & .5 & .5 & 0 & 0 \\ 0 & 0 & .5 & .5 & 0 \\ 0 & 0 & 0 & .5 & .5 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	Fast	$\begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$

Table 3 Yield values for each machine state

	Machine state				
Level	*0*	*1*	*2*	*3*	*4*
Low	1.0	0.25	0.125	0.0625	0.0
Med.	1.0	0.50	0.25	0.125	0.0
High	1.0	0.75	0.50	0.25	0.0

inventory level. Specifically, when $E[D]=6$, the state space is truncated so that $x \in [-125, 125]$ units; when $E[D]=9$, the state space is truncated so that $x \in [-200, 200]$ units: and when $E[D]=12$, the state space is truncated so that $x \in [-250, 250]$ units. It is highly unlikely that the inventory level would reach these extremes under normal conditions. This type of truncation is standard and has been shown to have a minimal impact on the solution values.[26,27]

Results

Table 4 reports the results of the numerical problems. The minimum expected discounted cost of the simultaneous approach is an average of approximately 18% less than that of the sequential approach, with a maximum of 99% and a minimum of 0%. The 75th percentile (not shown in the table) is approximately 27%; this means that for one-quarter of the cases studied, the simultaneous approach results in a cost savings of more than 27% as compared to the sequential approach.

Examining the numbers in Table 4 we see that the results are not substantially different for the three discount factors tested. However, the cost penalty for the sequential method increases dramatically as the machine repair cost increases. When the repair cost is high, the sequential method will set the repair threshold at a high level of deterioration, and this will limit the options available for the production decision and will make it difficult to raise the inventory level quickly.

Table 4 Results of test problems

		Total cost		
All problems		*Sequential**	*Simultaneous*†	*Cost penalty*‡
	Average	374.6	265.1	18.0
	Minimum	34.7	3.7	0.0
	Maximum	9396.1	3013.5	99.3
Factor	*Level*			
Discount factor	0.5	153.7	121.5	18.9
	0.7	241.6	186.2	18.4
	0.9	728.7	487.7	16.8
Repair cost	20	251.1	220.4	10.4
	40	351.4	256.0	17.5
	80	521.3	319.0	26.1
Holding cost	0.5	356.4	246.5	18.8
	1	371.6	262.1	18.1
	2	395.9	286.8	17.1
Backlog cost	5	243.0	201.3	12.1
	10	344.3	252.9	17.7
	20	536.5	341.1	24.4
Maximum production	12	560.9	366.0	20.5
	15	340.4	242.4	18.8
	20	222.5	187.0	14.9
Deterioration rate	Slow	287.5	222.0	16.5
	Med.	380.7	265.4	18.9
	Fast	455.7	308.0	18.7
Yield	Low	343.7	275.8	11.6
	Med.	403.0	267.3	21.2
	High	377.1	252.4	21.3
Mean demand	6	190.2	151.9	18.6
	9	327.6	233.5	18.1
	12	606.1	410.1	17.4
Demand distribution	Deterministic	307.3	204.2	20.0
	Binomial	342.2	228.4	19.0
	Uniform	412.0	293.8	17.5
	Geometric	437.0	334.2	15.7

*Minimum discounted cost using sequential approach averaged over all states for the given factor level.
†Minimum discounted cost using simultaneous approach averaged over all states for the given factor level.
‡Percentage cost penalty of sequential approach over simultaneous approach for the given factor level (averaged over all states).

The results are essentially the same for all three holding cost levels tested. This result makes sense given that the holding costs are small compared to other costs in the model. In contrast, the backlog costs are fairly substantial, and changes in this factor result in major differences in total cost. Intuitively, this result makes sense because we would expect a sub-optimal maintenance policy to hinder our ability to reduce inventory backlogs and thus would expect the penalty to be higher when the backlog cost is high.

The production capacity has a substantial effect on the cost differences. As the production capacity increases, the penalty decreases for the sequential method. Apparently, the additional production capacity increases the ability of the sequential approach to fine-tune production levels. The increased flexibility is particularly valuable given that the backlog costs are very high as compared to the production and holding costs.

The penalty for using the sequential method generally increases as the equipment deterioration rate increases, although not significantly. Other things being equal, we would expect the maintenance policy determined using the sequential method to repair earlier when the probability of moving to a worse state is very high. In other words, when the deterioration rate is high, the sequential method will result in repairing the machine earlier than is necessary, resulting in higher costs.

As the product yield increases, the difference between the simultaneous and sequential methods increases. At the low yield level, the change in yield from state to state is quite large. This large spread makes it easier for the sequential method to determine the 'correct' maintenance policy, that is, a policy closer to that determined by the simultaneous method. When the yield values are not as spread out (for the medium and high levels), the difference between the sequential and simultaneous methods increases.

The gains made by the simultaneous approach are relatively insensitive to the average demand level. As the mean demand increases, the penalty for using the sequential approach decreases, but only very slightly. The demand distribution, however, appears to have a much larger impact. The distributions are listed in increasing order of variance, and it is somewhat surprising at first glance to see that the biggest improvement is associated with the lowest variance: deterministic demand. This result makes more sense when one considers that a higher variance means that the demand distribution is more spread out, and hence that the maximum demand level is higher. Given that the production capacity is fixed, the opportunity for improvement by solving the problems simultaneously is decreased.

As indicated earlier, the simultaneous approach is guaranteed to be better than the sequential approach. However, it is interesting to see how much better. Analysis of more than 26 000 test problems over a wide range of parameter values indicates that the improvement provided by the simultaneous approach is substantial. This demonstrates that understanding the relationship between equipment condition and product quality, and incorporating this kind of information in maintenance scheduling and production decisions can have a significant impact on costs.

Conclusions and future directions

Improving quality and reducing costs are important goals in virtually every organization. This paper has explored the relationship between yield, equipment condition, and production management, accounting for the fact that equipment condition plays an important role in product quality and thus a firm's ability to satisfy customer demand. We formulated a Markov decision process model of a single-stage production system in which equipment condition affects the product yield. The decision maker must choose the equipment maintenance schedule as well as the quantity to produce in a way that minimizes the sum of production, backorder, and holding costs. We found that under some reasonable conditions regarding the yield and equipment deterioration, a maintenance threshold exists and a production threshold exists, but the production threshold is not monotone with respect to the machine state. That is, the optimal policy will not necessarily input more units in a worse state, even though the expected yield is decreasing.

Numerical problems were used to compare the proposed policy to a sequential approach, which solves the maintenance and production problems independently. The results of more than 26 000 test problems show that using the sequential approach results in an average cost that is approximately 18% greater than that of the simultaneous approach.

Until recently, researchers and practitioners alike have discounted or ignored the interaction between equipment condition, product quality, and inventory control. This paper demonstrates that substantial improvements can be realized by using equipment condition and yield information for maintenance and production decisions. While this work represents an important first step toward addressing an issue that has received relatively little attention, several enhancements are possible. First, treating the equipment repair time as a random variable rather than a fixed value would be of interest. A second enhancement would be to include multiple product types, each with a different yield distribution. Finally, extending the model to a multi-stage system would be of interest. Future research will be directed towards these extensions.

Acknowledgements—I thank two anonymous referees for their helpful comments and suggestions.

References

1 Yano CA and Lee HL (1995). Lot sizing with random yields: a review. *Opns Res* **43**: 311–334.

2 Gerchak Y, Vickson RG and Parlar M (1988). Periodic review production models with variable yield and uncertain demand. *IIE Trans* **20**: 144–150.

3 Henig M and Gerchak Y (1990). The structure of periodic review policies in the presence of random yield. *Opns Res* **38**: 634–643.
4 Wang Y and Gerchak Y (1996). Periodic review production models with variable capacity, random yield and uncertain demand. *Mngt Sci* **42**: 130–137.
5 Parlar M and Perry D (1995). Analysis of a (Q, r, T) inventory policy with deterministic and random yields when future supply is uncertain. *Eur J Opl Res* **84**: 431–443.
6 Mazzola JB, McCoy WF and Wagner HM (1987). Algorithms and heuristics for variable-yield lot sizing. *Naval Res Logist* **34**: 67–86.
7 Baker H and Ehrhardt R (1995). A dynamic inventory model with random replenishment quantities. *Omega* **23**: 109–116.
8 Bollapragada S and Morton TE (1999). Myopic heuristics for the random yield problem. *Opns Res* **47**: 713–722.
9 Lee HL and Yano CA (1988). Production control in multistage systems with variable yield losses. *Opns Res* **36**: 269–278.
10 Barad M and Braha D (1996). Control limits for multi-stage manufacturing processes with binomial yield (single and multiple production runs). *J Opl Res Soc* **47**: 98–112.
11 Gerchak Y, Wang Y and Yano CA (1994). Lot sizing in assembly systems with random component yields. *IIE Trans* **26**: 19–24.
12 Gurnani H, Akella R and Lehoczky J (2000). Supply management in assembly systems with random yield and random demand. *IIE Trans* **32**: 701–714.
13 Rosenblatt MJ and Lee HL (1986). Economic production cycles with imperfect production processes. *IIE Trans* **18**: 48–54.
14 Porteus EL (1986). Optimal lot sizing, process quality improvement and setup cost reduction. *Opns Res* **34**: 137–144.
15 Lee HL and Rosenblatt MJ (1987). Simultaneous determination of production cycle and inspection schedules in a production system. *Mngt Sci* **33**: 1125–1136.
16 Porteus EL (1990). The impact of inspection delay on process arid inspection lot sizing. *Mngt Sci* **36**: 999–1007.
17 Lee JS and Park KS (1991). Joint determination of production cycle and inspection intervals in a deteriorating production system. *J Opl Res Soc* **42**: 775–783.
18 Makis V and Fung J (1998). An EMQ model with inpsections and random machine failures. *J Opl Res Soc* **49**: 66–76.
19 Groenevelt H, Pintelon L and Seidmann A (1992). Production batching with machine breakdowns and safety stocks. *Opns Res* **40**: 959–971.
20 Groenevelt H, Pintelon L and Seidmann A (1992). Production lot sizing with machine breakdowns. *Mngt Sci* **38**: 104–123.
21 Arreola-Risa A and DeCroix GA (1998). Inventory management under random supply disruptions and partial backorders. *Naval Res Logist* **45**: 687–703.
22 Venkatesan M (1984). Production-inventory with equipment replacement—PIER. *Opns Res* **32**: 1286–1295.
23 Sloan TW and Shanthikumar JG (2000). Combined production and maintenance scheduling for a multiple-product, single-machine production system. *Product Opns Mngt* **9**: 379–399.
24 Heyman DP and Sobel MJ (1984). *Stochastic Models in Operations Research. Volume II: Stochastic Optimization*. McGraw-Hill: New York.
25 Derman C (1963). On optimal replacement rules when changes of state are Markovian. In: Bellman R (ed), *Mathematical Optimization Techniques*. University of California Press: Berkeley, CA, pp 201–210.
26 Puterman ML (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley: New York.
27 Sennott LI (1999). *Stochastic Dynamic Programming and the Control of Queueing Systems*. Wiley: New York.

Appendix A: policy improvement algorithm

Several methods are available to solve dynamic programming recursions such as (3). The books by Heyman and Sobel[24] and Puterman[26] provide detailed explanations of different approaches and explain the advantages and disadvantages of each. We use a simple policy improvement algorithm: this technique is not the most efficient computationally, but it is very robust, easy to understand, and easy to implement. In a nutshell, the algorithm tries all possible actions for all possible states and picks the best action. For a single iteration of the algorithm, one steps through each state (i, x) and tests each action (a, q), retaining the action that minimizes the costs. This process is repeated until the value function converges, that is, does not change significantly from one iteration to the next.

A formal statement of the policy improvement algorithm is presented below. Define $\pi(i, x)$ as a decision rule that specifies the action (a, q) to be taken when the process is in state (i, x). Let π_n refer to the policy tested on the nth iteration, and let W_n refer to the cost incurred from using policy π_n. The value of W_n is computed via (3) using the action specified by π_n (without the 'inf' operator).

Policy improvement algorithm

0: Choose any stationary policy and label it π_1.

1: For each state $(i, x) \in \mathscr{S}$, compute the cost function vector for policy π_n: $[W_n(i, x)]$.

2: Define $\mathscr{A}_n$ as the set of all actions *except* the action specified by the current policy: $\mathscr{A}_n = \mathscr{A} \backslash \{\pi_n(i, x)\}$. For each state $(i, x) \in \mathscr{S}$, compute the *difference* between the cost for the action specified by policy π_n and the minimal cost for actions in $\mathscr{A}_n$:

$$\Delta_n(i,x) = \min_{(a,q)\in\mathcal{A}_n} \Bigg\{ R(a) + cq + \sum_{d=0}^{\infty}\sum_{k=0}^{q} \Bigg[b(d-x-k)^+ + h(x+k-d)^+ + \alpha \sum_{j=0}^{M} p_{ij}^{aq} W_n(j, x+k-d) \Bigg] \times \phi_i^a(k|q)\xi(d) \Bigg\} - W_n(i,x) \tag{5}$$

If $\Delta_n(i, x) \geqslant 0$, then $\pi_n(i, x)$ is the cost-minimizing action, so let $\pi_{n+1}(i, x) = \pi_n(i, x)$. If $\Delta_n(i, x) < 0$, then let $\pi_{n+1}(i, x)$ be the smallest action that minimizes (5).

3: If $\pi_{n+1}(i, x) = \pi_n(i, x)$ for all (i, x), then stop: π_n is optimal. Otherwise, replace n with $n+1$ and return to Step 1.

Appendix B: maintenance policy for the sequential approach

The first step in the sequential approach is to determine a maintenance policy, that is, to determine the machine states in which it is cheaper to stop production and repair the machine than to continue producing. This is accomplished using an elementary maintenance model such as that proposed by Derman.[25]

As stated in the Sequential approach section, $C(i)$ denotes the cost to operate the machine when it is in state i, and $C(i) \equiv c(E[D]/\beta_i)$. The objective for this sub-problem is to choose the action—'repair' or 'do not repair'—which minimizes the sum of discounted costs. Define $V'(i)$ as the minimal expected discounted cost function when the initial machine state is i. Since repairing the machine returns it to state 0 with probability one, the cost function can be written as

$$V'(i) = \min\left\{R + C(0) + \alpha \sum_{j=0}^{M} p_{0j} V'(0), \; C(i) + \alpha \sum_{j=0}^{M} p_{ij} V'(j)\right\}, \tag{B1}$$

where the terms within the braces refer to the costs associated with repairing and not repairing, respectively.

This dynamic programming recursion is solved using the same policy improvement method described in Appendix A with the appropriate modifications of cost functions, states, and actions. The optimal policy tells us whether or not to repair in each state and fixes the machine state transitions for the second step in the sequential approach in which the production quantity is determined.

Appendix C: demand distribution parameters

Four demand probability functions are tested in the Numerical results section: deterministic, binomial, uniform, and geometric. Figure 2 illustrates what the latter three probability functions look like for $E[D] = 6$. (Note that the actual functions are discrete, but they have been sketched as continuous functions here for illustrative purposes.) The specific parameter values used for each distribution are reported in Table 5.

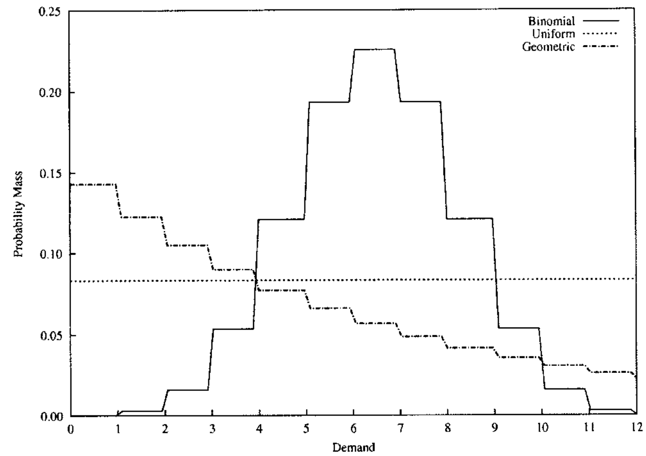

Figure 2 Illustration of different demand probability functions ($E[D] = 6$).

Table 5 Parameter values for demand distributions tested

Distribution (parameters)	*Parameter values*	*Mean*	*CV**
Deterministic (none)	n/a	6	0
	n/a	9	0
	n/a	12	0
Binomial *(n, p)*	$n = 12, p = 0.5$	6	0.29
	$n = 18, p = 0.5$	9	0.24
	$n = 24, p = 0.5$	12	0.20
Uniform *(a, b)*	$a = 0, b = 12$	6	0.62
	$a = 0, b = 18$	9	0.61
	$a = 0, b = 24$	12	0.60
Geometric *(p)*†	$p = 1/7$	6	1.08
	$p = 1/10$	9	1.05
	$p = 1/13$	12	1.04

*Coefficient of variation, which is equal to the standard deviation divided by the mean.
†There is no upper bound on the range of the geometric distribution, and thus it must be truncated and normalized for numerical solutions. We truncate the distribution at $4E[D]$, which represents about 98% of the distribution.

Received 1 April 2002; accepted August 2003 after two revisions