OPTIMAL CONTROL OF A SINGLE STAGE PRODUCTION SYSTEM SUBJECT TO RANDOM PROCESS SHIFTS

RAMAN K. NURANI

Carnegie Mellon University, Pittsburgh, Pennsylvania

SRIDHAR SESHADRI

New York University, New York, New York

J. GEORGE SHANTHIKUMAR

University of California, Berkeley, California

(Received February 1994; revisions received March 1995; accepted June 1995)

We consider a single stage production system with Poisson demand and exponential processing times. After producing a good item, the production process can shift to an "out-of-control" state with a given probability and start producing bad items. The state of the process is known only when the next stage (or customer) receives the item. Once an out-of-control process is detected, process correction is instantaneous. Customers arriving to an empty system get backlogged. In this framework, we examine FIFO (First In First Out) and LIFO (Last In First Out) issuing policies. The objective is to minimize the total expected discounted or average costs over an infinite time horizon. We characterize the structure of the optimal production policy for FIFO and LIFO, show that LIFO is better than FIFO and, in general, better than a large class of issuing policies. A numerical example illustrates that savings up to 20 percent can be obtained from using LIFO over FIFO. We also derive conditions under which maintaining *zero inventory* is optimal, and show that zero inventory is more likely to be optimal when either the backlogging cost or arrival rate of customers is small, and when the inventory carrying cost or the processing rate or the probability of getting a good item is large.

The goal of an ideal production system is producing right the first time. This goal is often not achieved in practice in many situations because of process uncertainties, leading to loss of yield, rework, and starvation of downstream stages. Our focus in this paper is to determine how to control a production system with an imperfect production process when there is delay in getting information about its state. In this framework, we bring out the relationships between process quality, the optimal production policy, and the method of issuing items to the customer.

We model a single stage production system with Poisson demand. The demand could be from external customers or the downstream stages. The items require exponential processing times, and are stored in an output buffer and issued when demanded. If there is no item in the buffer, the demand gets backlogged. The production process is assumed to be imperfect, i.e., after producing a good item, the process can shift from the in-control-state to the out-of-control state with a given probability, as in Porteus (1986). Once the process shifts to the out-of-control state, it remains in that state producing defective items until discovered. Due to technological constraints, the detection of process shifts cannot be performed on-line, and the quality of the items is known only when the subsequent stage receives the item. Once the discovery is made, the process is corrected immediately.

This problem is motivated by practices in the semiconductor, as well as other industries where there is delay in transmitting the information about the quality of the items. For example, in semiconductor manufacturing a wafer cannot be tested till subsequent operations are performed. Also see Porteus (1986) for other examples where the detection of process shifts cannot be done on-line. (A different modeling approach is to assume that the manager has to live with a *random yield*, i.e., the process quality cannot be controlled; see Yano and Lee (1989) for a survey of the random yield problem.) Our model will also fit (though not directly) into Taguchi's framework (Taguchi et al. 1989), where the items, even if inspected and found to be within the process control limits, might result in increased cost in the downstream stages due to deviation from the target quality characteristics.

In systems such as ours, there is an interaction between production and process control. We do not want to keep too much stock, as many of the items in the buffer could be defective, nor do we wish to keep too little stock because that may result in excessive backlog. Apart from this trade-off, it also becomes important to ask which item from the output buffer should be issued to the customer. We discuss two methods of issuing items, namely FIFO (First In first Out) and LIFO (Last In First Out). In FIFO, the item produced first is issued to the customer, whereas in LIFO the most recently produced item is issued. Due to the difference in issuing policies, LIFO provides information about the state of the system faster than FIFO, and in

Subject classifications: Inventory/Production: Optimal inventory, zero inventory, quality, process shifts. Dynamic Programming/Control: Markov.
Area of review: Manufacturing, Operations and Scheduling.

0030-364X/97/4505-0713 $05.00

general any other issuing policy (and thus LIFO is expected to dominate any issuing policy).

The objective is to minimize the long run (discounted or average) expected total costs, where the costs included are processing cost, inventory carrying cost, scrap cost arising due to production of defective items, and backlogging cost. Our interest is in obtaining the structure of the optimal policies for the FIFO and LIFO issuing methods and in comparing the benefits of using LIFO. For FIFO, we show that there exists an *upper threshold level* on inventory, up to which it is optimal to produce and above which it is optimal to stop producing. For LIFO such an upper threshold level is characterized by a pair of numbers. We also numerically investigate the benefits of using LIFO.

With the advent of new manufacturing ideologies such as Single Minute Exchange of Dies (SMED) and Zero Defects, much attention has been focussed on *zero inventory*. Bielecki and Kumar (1988) show that an optimal zero inventory policy may exist even when there is uncertainty and derive conditions in which zero inventory is optimal for a failure prone manufacturing system. For our model we derive conditions in which maintaining zero inventory is optimal. We show that zero inventory is more likely to be optimal when either the backlogging cost or arrival rate of customers is small, and when the inventory carrying cost or the processing rate or the probability of getting a good item is large.

Our modeling assumptions (e.g., Poisson demand, exponential processing time, and geometric process shift) permit us to use Markov decision theory along with a novel use of sample path analysis to infer the structure of optimal control policies. These assumptions also allow us to use dynamic programming to compare the FIFO and LIFO issuing policies. While the modeling assumptions are restrictive, the insights provided by the model are general: (i) LIFO is better than FIFO, and numerical examples (see Section 6) indicate that the long run average cost of using LIFO can be about 20% lower compared to using FIFO in some cases. (ii) When LIFO is used the structure of the control policy has to be thought through as explained in Section 1. (iii) Insights obtained from the Markovian model can then be used for the design and control of manufacturing systems as done in Buzacott and Shanthikumar (1993) under most general modeling assumptions. (iv) It appears that LIFO permits the use of much larger in-process inventories than FIFO. The implications seem to be that with process uncertainties reduced, the focus shifts to reducing backlog.

FIFO and LIFO issuing policies have been analyzed in other contexts such as perishable inventory management, accounting, and tax valuation. For perishable goods, FIFO is the optimal issuing policy over a wide range of assumptions, particularly where the issuing organization has complete control over issuing actions; see Nahmias (1982) and Silver and Peterson (1985). (However, where the customer makes the selection, for example in retail food distribution where the expiration date is shown on each unit, the LIFO policy is likely to be observed.) In the context of tax valuation, Cohen and Pekelman (1979) examine the effect of LIFO and FIFO accounting systems on inventory control policies. They observe that the optimal inventory policy does not vary greatly with the valuation scheme but does depend on the explicit inclusion of taxation. Based on our results in this paper, apparently LIFO is superior in several instances due to its impact on process quality. The ramifications of this finding on accounting policies are discussed in Section 4.

Porteus (1986) models a single stage production system that is similar to ours except that the demand rate is constant and there is a setup cost for production such as in the EOQ model. The main decision variable in his setting is the lot size. The process can be inspected after finishing the production of a lot. Porteus shows that reducing lot sizes can improve quality levels. In contrast, we look at the structure of the optimal control policy under stochastic demand and in addition examine the effect of issuing policies. In a single stage ($M/M/1$) queue setting, Tapiero and Hsu (1987) present numerical results for different inspection schemes such as full or 100-percent inspection and barrier inspection (i.e., inspect if the number of items in the system is less than a critical number). Hsu and Tapiero (1992) model the production system as an $M/G/1$ queue, which may shift to a substandard processing state in random time. They evaluate various sampling schemes such as no quality control, full sampling, and random sampling, and examine the impact and sensitivity of system parameters numerically. These papers do not address the production control and issuing problems. To the best of our knowledge, the optimal policy structure and the benefits of using LIFO have not previously been addressed in the quality management literature.

The rest of our paper is organized as follows. Section 1 discusses the basic model and the assumptions. Section 2 analyzes the FIFO issuing method and presents optimal control policies for the discounted and the average cost problems. Section 3 analyzes LIFO. Section 4 compares FIFO and LIFO and shows that LIFO is optimal over all class of issuing policies, in which an arriving customer does not wait for the item under process, if any, to be completed if there is already inventory in the output buffer. Section 5 examines the conditions for zero inventory. Section 6 presents a numerical example comparing FIFO to LIFO. Section 7 discusses the conclusions and some extensions.

1. MODEL DESCRIPTION

Consider a single stage manufacturing system producing items with abundant supply of raw materials. We identify the random demand from the subsequent stage as a random "customer" arrival to the system. We model the demand process as Poisson with rate λ. Processing times are assumed to be i.i.d. and exponentially distributed with mean $1/\mu$. We assume that there is a processing cost that is

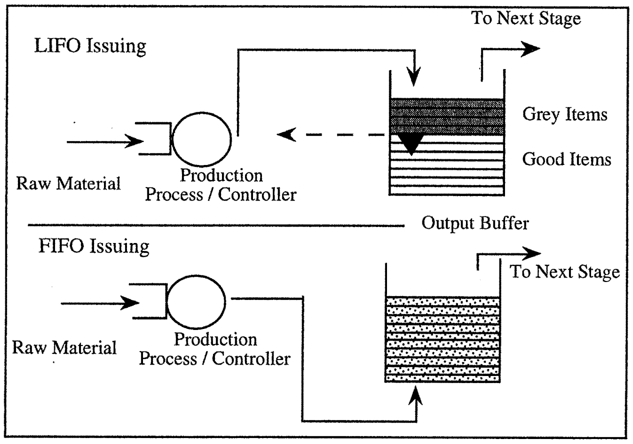

Figure 1. The production system under FIFO and LIFO issuing.

directly proportional to the time taken to produce the item, and that after producing a good item, the process may shift with probability p from being in the in-control state to the out-of-control state. This is termed a *geometric process shift*. After the process reaches the out-of-control state it remains in that state until detected and corrective action has been taken. Items produced in the out-of-control state are of poor quality and scrapped at a cost. The out-of-control state is detected when the subsequent processing stage (or customer) finds a defective item. Once the process shift is detected the correction is assumed to be instantaneous and the machine is as good as new after correction. We assume that the item under process, if any, can be reprocessed at no extra cost whenever an arrival detects the out-of-control state of the machine.

This demand model is applicable in PULL manufacturing systems, or when supplying to a bottleneck station or in some instances even supplying to customers. Examples include preparation of molds prior to pouring of metal in a foundry, supply of completed wafers from wafer fabrication to testing/probing stage in semiconductor manufacturing and supply of components from one plant to another.

FIFO and LIFO methods of issuing the items from the output buffer are operationalized by stacking differently in the output buffer and letting the customer pick from the top of the stack (Figure 1). In FIFO the item produced first will be on top of the stack. An example where FIFO is used could be a two-stage system with stages interconnected by an automatic conveyor. In LIFO, the most recent item is placed on top of the buffer stack.

If there are no items in the output buffer when a customer arrives, the customer waits (gets backlogged). If there is an item in the buffer, she picks up the first item in the stack for use. If the item is found to be bad, she informs the system instantaneously. The customer then picks up the next item and so on until she either finds a good item or gets backlogged. We assume that all these actions occur instantaneously. In the case of LIFO, we also assume that if the machine is processing an item at the instant an arrival takes place and if there are items in the output buffer stack, then the customer does not wait but takes the topmost item (last produced) in the stack.

As stated earlier, our objective is to find the optimal production control policy for minimizing the total expected discounted costs over the infinite time horizon for FIFO and LIFO. Define $\Lambda = \lambda + \mu$ and let

$h\Lambda$ = inventory carrying cost per unit per unit time,
$r\Lambda$ = processing cost per unit time,
$b\Lambda$ = backlogging cost per unit per unit time,
s = scrap cost per unit, and
α = rate of discount.

Consider the FIFO stacking policy. If a customer picks up a bad item, then *all* the items in the output buffer are defective and the process is in the out-of-control state. If the item is good, then the inventory in the output buffer reduces by one. We assume that the processing of an item can be switched on or off at any stage of processing (and later show that this is not a restriction in practice). The state of the system will be represented by a single number n, with positive n representing the inventory of items in the output buffer and negative n the number of backlogged customers. The state space is $\{\cdots, -1, 0, 1, 2, \cdots\}$. With this state representation and an admissible control policy (described below), the evolution of the system under FIFO is a Markov process.

On the other hand, the LIFO stacking policy needs a different state representation, because even if the item on top of the stack is bad it is possible to have some good items below. Therefore, we represent the state of the system as a pair of numbers (n, m), where n denotes the number of *good* items in the stack and m denotes the number of *grey* items, i.e., items that may be either good or bad. The n good items are at the bottom of the output buffer stack, whereas the m grey items are on top of the n good items. The controller, based on the past history, knows which n are good and which m are grey (see Figure 1), whereas the arrivals do not have this information. (For example, the controller can keep a *signal* kanban in the buffer stack to separate the *good* from the *grey* items as in Monden 1981).

As in FIFO, we assume that the processing of an item can be switched on or off at any stage of processing, and later show that the optimal policy corresponds to the case when the machine is never switched off while processing an item. Recall the assumption that the customer will always take an item from the output buffer stack if there are any, and not wait until the processing of the item (if any) on the machine is completed. On customer arrival, if k out of m grey items are good with $k < m$, then the new state $(n + k - 1, 0)$ is reached, and $(m - k)$ bad items are scrapped. This event has probability $q^k p$, where $q = 1 - p$. If all the m grey items are good then the new state $(n + m - 1, 0)$ is reached, and the probability of this event is q^m. The state space under the LIFO stacking policy is $\{(n, m), n \in \{\cdots, -1, 0, 1, \cdots\}; m \in \{0, 1, 2, \cdots\}$ *and* $m = 0$ *when*

$n < 0\}$, where negative n represents the number of backlogged customers.

The control decision available in state i (for any issuing policy) is simply whether to produce or not. We denote this decision by:

$$u(i) = \begin{cases} 1 & \text{if we decide to produce when the state of the system is } i, \\ 0 & \text{otherwise.} \end{cases}$$

We say that the control $u(i)$ is admissible if it does not use information about the future evolution of the system. We have a continuous time Markov process under both the FIFO and the LIFO stacking policies when admissible controls are used. Using the technique of Lippman (1975), we shall uniformize the transitions in this continuous process using a fast Poisson process of rate Λ and then pose the problem of optimally controlling a discrete time Markov chain.

2. FIFO

In this section we characterize the optimal control policy for the FIFO stacking policy for minimizing the discounted cost over the infinite time horizon. We assume that $u(n) = 1$ for all $n < 0$, for otherwise running such a production system is meaningless. Likewise, asking a customer to wait when there is inventory in the buffer is not desirable. These cases are considered in Nurani et al. (1995), where we provide sufficient conditions for the validity of these assumptions.

Define $g(n, u(n))$ to be the one step cost of being in state n and applying the controls $u(n)$. Here, one step is equal to one transition in the fast Poisson process with rate Λ. Using the indicator function $I(.)$, $g(n, u(n))$ can be written as

$$g(n, u(n)) = \frac{\Lambda}{\Lambda + \alpha} n^{-} b + \frac{\Lambda}{\Lambda + \alpha} n^{+} h + \frac{\mu}{\Lambda + \alpha} spI(n < 0) + \frac{u(n)\Lambda}{\Lambda + \alpha} r + \frac{\lambda}{\Lambda + \alpha} nspI(n \geq 0) \tag{1}$$

where $n^{+} = \max(n, 0)$ and $n^{-} = \max(-n, 0)$. The first term on the right-hand side (rhs) is the cost of customers backlogged at the beginning of the transition, the second term is the carrying cost of items in the buffer that were present at the beginning of the transition, the third is the cost of scrapping if a defective item is produced when there is backlogged demand, the fourth term is the processing cost incurred if the decision is to produce in the state n, and the fifth term on the rhs is the cost of scrapping all items if an arrival finds a defective item (recall that the customer gets backlogged).

Let $V(n)$ denote the minimum expected total discounted cost over the infinite time horizon starting from the state n. Then we have

$$V(n) = \min_{u(n)\in[0,1]} \Big\{ g(n, u(n)) + \frac{u(n)\mu}{\Lambda + \alpha} [I(n \geq 0)V(n+1) + I(n < 0)\{qV(n+1) + pV(n)\}] + \frac{\lambda}{\Lambda + \alpha} [I(n > 0)\{qV(n-1) + pV(-1)\} + I(n \leq 0)V(n-1)] + \frac{[\Lambda - u(n)\mu - \lambda]}{\Lambda + \alpha} V(n) \Big\} \tag{2}$$

$$= A(n) + \frac{I(n \geq 0)}{\Lambda + \alpha} \min[\mu V(n), \Lambda r + \mu V(n+1)], \tag{3}$$

where

$$A(n) = \frac{\Lambda}{\Lambda + \alpha} n^{-} b + \frac{\Lambda}{\Lambda + \alpha} n^{+} h + \frac{\mu}{\Lambda + \alpha} spI(n < 0) + \frac{\lambda}{\Lambda + \alpha} nspI(n \geq 0) + \frac{\Lambda}{\Lambda + \alpha} I(n < 0) + \frac{\mu}{\Lambda + \alpha} [I(n < 0)\{qV(n+1) + pV(n)\}] + \frac{\lambda}{\Lambda + \alpha} [I(n > 0)\{qV(n-1) + pV(-1)\} + I(n \leq 0)V(n-1)].$$

Consider the states $n \geq 0$. From (3) it is optimal to produce in state n if

$$V(n+1) - V(n) \leq -\frac{\Lambda r}{\mu}. \tag{4}$$

We prove below that there exists a number $z_1(\geq 0)$ such that it is optimal to produce at all levels below z_1 · z_1 is called the *upper threshold level*. It can be either a finite positive number or positive infinity. Such a policy is called a threshold type policy.

Theorem 1. *Consider a single stage system with Poisson demand, exponential service times and geometric process shifts with FIFO stacking in the buffer and where the objective is to minimize the expected discounted costs over the infinite time horizon. Then there exists an upper threshold level z_1 (≥ 0) such that it is optimal to always produce when the inventory is below z_1 and not to produce when the inventory exceeds this level.*

Proof. Consider two systems, labeled I and II. Let $n \geq 0$ and let systems I and II start with initial inventory in the output buffer of $(n + 1)$ and $(n + 2)$ items, respectively. Label the items in the order in which they were produced. Assume that the machine states were and will be the same in both systems when producing the ith item. This ensures that the marginal distributions of the first $(n + 1)$ items in both the systems are identical. If initially there is even one bad item in system I, then the number of bad items in system II will be one more than the number of bad items in system I. Arrival epochs are assumed to be the same for both the systems. It is assumed that the service time for producing the ith item in system I and the service time for the $(i+1)$st item in system II are identical. Assume that the optimal control policy is used in system II. For system I we use the control policy used for controlling system II. The two systems are now said to be coupled. We stop the coupling either when systems I and II hit the states n and $(n+1)$ respectively or when both of them hit the state -1. Under this coupling arrangement, we pass information about the future evolution of system I only when there is one item in system II and none in system I and an arrival occurs. This cannot happen because we would have already stopped the coupling before such a state is reached (as $n \geq 0$). Therefore, we do not pass on information about the future to the controller of system I, and the control policy of II (till the coupling ends) is admissible for system I.

Let $P(y)$ where $y \in \{I, II\}$ be the expected cost obtained using these control policies over the infinite time horizon in system y. Then

$$V(n+2) - V(n+1) \geq P(II) - P(I), \tag{5}$$

as we are using a suboptimal policy for system I, i.e., $P(II) = V(n+2)$, $P(I) \geq V(n+1)$. Define $P(y, [t_1, t_2], B)$ to be the expected cost obtained in system y over time $[t_1, t_2]$ on the set of events $\{B\}$. Consider the following possibilities.

1. System II never hits the state $(n+1)$, and system I never hits the state n. Call the set of events on which this happens as $\{B_1\}$. Since system II started with one more extra item:

$$P(II, [0, \infty], B_1) - P(I, [0, \infty], B_1) \geq 0. \tag{6}$$

2. There are defective items in system I and system II, and both of them hit the state -1 before hitting the states n and $(n+1)$, respectively. Call the set of events on which this happens as $\{B_2\}$ and the time at which this happens as τ_{-1}. Then as above:

$$P(II, [0, \tau_{-1}], B_2) - P(I, [0, \tau_{-1}], B_2) \geq 0. \tag{7}$$

3. On the set $(B_1 U B_2)^c$ system I hits the state n and system II hits the state $(n+1)$, say at time τ_n. Let the number of transitions of the fast Poisson process in time τ_n be τ and $E(\rho^\tau)$ be the expected discount factor, where $\rho = \Lambda/(\Lambda + \alpha)$. We have:

$$P(II, [0, \tau_n], (B_1 U B_2)^c) - P(I, [0, \tau_n], (B_1 U B_2)^c) \geq 0. \tag{8}$$

The coupling stops at time ∞ on $\{B_1\}$, at time τ_{-1} on $\{B_2\}$ and at time τ_n on $\{(B_1 U B_2)^c\}$. Define $\tau = \infty$ on $(B_1 U B_2)$. Using this definition the difference in cost obtained *after* time τ_n can be concisely written as $E(\rho^\tau)[V(n+1) - V(n)]$ (because $E(\rho^\tau I(B_1 U B_2))$ is zero). Since system I follows a suboptimal policy and system II follows the optimal policy and using (6) and (7) we have:

$$\begin{aligned} V(n+2) - V(n+1) \geq \{&P(II, [0, \tau_n], (B_1 U B_2)^c) \\ &- P(I, [0, \tau_n], (B_1 U B_2)^c) \\ &+ E(\rho^\tau)[V(n+1) - V(n)]\}. \end{aligned} \tag{9}$$

The first two terms on the rhs of this equation represent the cost of carrying one extra item of inventory for τ transitions. To obtain an expression for this cost, define t_k to be time the $(k-1)$th transition took place. The expected cost of carrying one item from t_k to t_{k+1} discounted continuously is

$$h\Lambda E\left[\int_{t_k}^{t_{k+1}} e^{-\alpha t}\, dt\right] = h\left(\frac{\Lambda}{\Lambda + \alpha}\right)^{k+1}.$$

We can now express the expected carrying cost of one item for τ transitions on the set $(B_1 U B_2)^c$ as:

$$\begin{aligned} &P(II, [0, \tau_n], (B_1 U B_2)^c) - P(I, [0, \tau_n], (B_1 U B_2)^c) \\ &= \frac{h\Lambda}{\alpha + \Lambda} E\left[I((B_1 U B_2)^c) \sum_{k=0}^{\tau-1} \left(\frac{\Lambda}{\alpha + \Lambda}\right)^k \right] \\ &= \frac{h\Lambda}{\alpha} [\Pr((B_1 U B_2)^c) - E(\rho^\tau I((B_1 U B_2)^c))]. \end{aligned} \tag{10}$$

Using (9) and (10):

$$\begin{aligned} &V(n+2) - V(n+1) \\ &\geq \left\{\frac{h\Lambda}{\alpha} [\Pr((B_1 U B_2)^c) - E(\rho^\tau)] \right. \\ &\quad \left. + E(\rho^\tau)[V(n+1) - V(n)]\right\}. \end{aligned} \tag{11}$$

When in state $(n+1)$, by considering the suboptimal policy of setting aside one item by incurring the carrying cost for that one item forever, we have:

$$V(n+1) - V(n) \leq \frac{h\Lambda}{\alpha}, \tag{12}$$

where $(h\Lambda/\alpha)$ is the cost of carrying one item forever. Substituting this in (11) we get

$$\begin{aligned} &V(n+2) - V(n+1) \\ &\geq [V(n+1) - V(n)][\Pr((B_1 U B_2)^c)]. \end{aligned} \tag{13}$$

From Equations (4) and (13) it is clear that if it is optimal to produce in the state $(n+1)$ then it is optimal to produce in the state n also. Thus the existence of the threshold level $z_1 (\geq 0)$ is proved. □

Remarks. (i) As we never stop producing when the inventory in the system is below z_1, the assumption that we can switch on or switch off the production at any stage of processing an item becomes equivalent to the assumption that we never switch off the processing of an item till the production of the item is complete. (ii) By adding an inventory carrying cost to the processing cost, we can model the carrying cost of the item under process.

As a result of the above theorem, production control decisions are made simple if we know the threshold level

z_1. Successive approximation can be used to solve the cost functions $V(n)$ and the optimal policy. An efficient computational procedure for the average cost case (based on determining the stationary probability distribution of the controlled Markov chain) is given in Nurani et al. (1995).

3. LIFO

The state of the system is given by a pair of numbers, (n, m). The number of good items in the output buffer is denoted by n. The status of the other items in the output buffer is not known until an item is withdrawn by the customer. These items are called *grey* items and are denoted as m. These m grey items are on top of the n good items in the output buffer stack. We assume that: (i) if there are items in the output buffer the customer will select the top most item in the output buffer stack (and not wait for the machine to finish processing an item, if any), and (ii) production can be switched on or off at any stage of processing an item. We justify (i) in Nurani et al. (1995) and (ii) can be justified based on the structure of the optimal policy (see Remark (i) in Section 2).

The dynamic programming recursion equations are formulated below for LIFO. The one step cost $g(n, m, u(n, m))$ incurred in state (n, m) by using the control $u(n, m)$ is given by:

$$g(n, m, u(n, m)) = \frac{\Lambda b}{\Lambda + \alpha}(n + m)^{-} + \frac{\Lambda h}{\Lambda + \alpha}(n + m)^{+} + \frac{\mu}{\Lambda + \alpha} spI(n < 0, m = 0) + \frac{u(n, m)\Lambda}{\Lambda + \alpha} r + \frac{\lambda}{\Lambda + \alpha}\left[\left(\sum_{k=0}^{m-1} q^k p(m - k)s\right) I(n \geq 0, m > 0)\right]. \quad (14)$$

In this expression, the first term accounts for the cost of backlogged items present at the beginning of the transition, the second term is the cost of carrying the inventory of items present in the output buffer at the beginning of the transition, the third term is the cost of scrapping if a defective item is produced in a backlogged state, the fourth term is the processing cost, and the fifth is the cost of scrapping items when a grey item is drawn from the output buffer and a defective item detected.

The cost function $V(n, m)$ can then be written as

$$V(n, m) = \min_{u(n,m)\in[0,1]} \Bigg\{ g(n, m, u(n, m)) + \frac{u(n, m)\mu}{\Lambda + \alpha} \times [V(n, m + 1)I(n \geq 0) + \{qV(n + 1, 0) + pV(n, 0)\} \times I(n < 0, m = 0)] + \frac{\lambda}{\Lambda + \alpha}\left[\sum_{k=0}^{m-1} q^k pV(n + k - 1, 0) + q^m V(n + m - 1, 0)\right] I(n \geq 0, m > 0) + \frac{\lambda}{\Lambda + \alpha} \times [V(n - 1, 0) \times I(n > 0, m = 0) + V(n - 1, 0) \times I(n \leq 0, m = 0)] + \frac{(\Lambda - u(n, m)\mu - \lambda)}{\Lambda + \alpha} V(n, m)\Bigg\}. \quad (15)$$

Rewriting this equation we get

$$V(n, m) = B(n, m) + \frac{I(n \geq 0)}{\Lambda + \alpha} \min [\mu V(n, m), \Lambda r + \mu V(n, m + 1)], \quad (16)$$

where

$$B(n, m) = \frac{\Lambda b}{\Lambda + \alpha}(n + m)^{-} + \frac{\Lambda h}{\Lambda + \alpha}(n + m)^{+} + \frac{\mu}{\Lambda + \alpha} spI(n < 0, m = 0) + \frac{\Lambda}{\Lambda + \alpha} I(n < 0) + \frac{\lambda}{\Lambda + \alpha}\left[\left(\sum_{k=0}^{m-1} q^k p(m - k)s\right) I(n \geq 0, m > 0)\right] + \frac{\mu}{\Lambda + \alpha}[\{qV(n + 1, 0) + pV(n, 0)\}I(n < 0, m = 0] + \frac{\lambda}{\Lambda + \alpha}\left[\sum_{k=0}^{m-1} q^k pV(n + k - 1, 0) + q^m V(n + m - 1, 0)\right] I(n \geq 0, m > 0) + \frac{\lambda}{\Lambda + \alpha}[V(n - 1, 0)I(n > 0, m = 0) + V(n - 1, 0)I(n \leq 0, m = 0)].$$

The next step is to obtain the structure of the optimal policy for LIFO. The results of the following three lemmas lead to the structure of the optimal policy given in Theorem 2. The proofs of the lemmas are given in the appendix.

Lemma 1. *$V(n, 0)$ is convex in n, i.e., $V(n + 2, 0) - V(n + 1, 0) \geq V(n + 1, 0) - V(n, 0)$ for $n = -1, 0, 1, 2, \ldots$*

Lemma 2. *If it is optimal to produce in the state $(n, m + 1)$ then it is optimal to produce in the state (n, m) for $n \geq 0$.*

Lemma 3. *If it is optimal not to produce in the state (n, m) then not producing in the state $(n + 1, m)$ is optimal.*

Theorem 2. *Consider a single stage system with Poisson demand, exponential service times and geometric process shifts with LIFO stacking and where the objective is to minimize the expected discounted costs over the infinite time horizon. The structure of the optimal policy is such that:*

(i) *There exists upper threshold levels of the form $z_1 = (n, m) = (n, f(n))$, such that for a given number, n, of good items, it is optimal to produce when the number m of grey items is below $m = f(n)$. Moreover, the function f is nonincreasing in n.*

(ii) *$n + f(n)$ is nondecreasing in n.*

Proof. (i) From Lemma 2, as it is optimal to produce in the state (n, m) when it is optimal to produce in the state

$(n, m + 1)$, the existence of a function f such that we produce in all states $\in \{(n, m) : m \leq f(n)\}$ follows. In Lemma 3 we showed that if it is not optimal to produce in the state (n, m) then it is not optimal to produce in the state $(n + 1, m)$. Therefore the function f is seen to be nonincreasing in n.

(ii) We need to prove that if it is optimal to produce in the state (n, m) then it is also optimal to produce in the state $(n + 1, m - 1)$. The construction and the proof are similar to Lemmas 2 and 3 and are therefore omitted (see Nurani et al. 1995). □

Remarks. (i) Like in FIFO, the optimal policy corresponds to never halting production during the processing of an item. (ii) The carrying cost of an item under process can be added to the processing cost.

4. LIFO VS. FIFO

It is generally a good practice to use the FIFO policy when using the item at the subsequent stage provides no information about the current state of the process. However, intuitively speaking, when the state of the process can be inferred (and this information utilized) it is better to issue items using LIFO. In the following theorem we prove this intuitive result that LIFO is better than FIFO for our model.

Theorem 3. *Under the same operating conditions, LIFO is better than FIFO for minimizing the expected average or discounted costs over the finite or infinite time horizon.*

Proof. Consider two systems I and II. For system I use the optimal FIFO production policy. For system II use the production policy of system I, but use the LIFO issuing policy. Assume that the machine states are the same in both systems, and if the failure of the machine is detected in system II, we do not produce in system II until failure is detected in I and corrected (failure in system I will be detected later because the bad items will be at the bottom, whereas in the LIFO stacking policy they will be on top of the stack). Then the following are true.

1. On every sample path system II has lower inventory than system I and the same level of backlogged demand.
2. The evolution of the state of system I can be deduced from that of system II.

To prove (2), assume that there are x good items and y grey items in LIFO and $x + y$ items in FIFO. On customer arrival, the state of LIFO becomes $(x + y - 1, 0)$ if all the grey items are good. If only $y_1(\geq 0)$ of the y grey items are good then the bad item is detected *immediately* and the state reached on customer arrival is $(x + y_1 - 1, 0)$. On the other hand, the state of FIFO system will be $(x + y - 1)$. If $y_1(\geq 0)$ of the y grey items are good then the bad state of the machine will be detected only after $(x + y_1 - 1)$ more arrivals. During the time when system II is idle waiting for system I to detect the bad item, system I is producing bad items. This argument can be extended to show that the state of system I can be traced using the past history of production, detection of failures and customer arrivals. This also shows that system I will have higher inventory than system II. The policy used for controlling system II is admissible, as it uses information only about its past. Conclusions 1 and 2 imply that FIFO is inferior to LIFO. □

Theorem 4. *LIFO is optimal over all issuing policies.*

Proof. In the construction given in Theorem 3 just substitute any nonanticipative policy for FIFO. Let system I be controlled using any nonanticipative policy. In system II use the LIFO policy. Whenever system I initiates production so does LIFO. Whenever a defective item is found in LIFO, all items are inspected and bad items discarded. The system controlled using LIFO does nothing till the defect in the other system gets corrected. Note that production and arrival events are the same in both systems. So the only uncertainty is regarding when a defective item gets issued to a customer. In LIFO this happens earlier than in any nonanticipative policy. So the LIFO controller has more information than the other controller at all points of time. Therefore the policy used in system I is admissible for system II. Finally system II always has less or equal inventory compared to system I. The conclusion follows. □

Remark. Theorem (4) holds under the assumption that if there is a cost for scrapping defective items and if there is inventory then when a customer arrives, she does not wait for the machine to finish processing the item (if any) but takes the item from the output buffer stack as per the issuing policy.

The above results can be extended beyond the setting of our model, provided the items produced do not deteriorate with time. In the case of perishable inventory (Nahmias 1982), due to perishability, it is natural to use the FIFO scheme for the issuing organization whereas if the customer has control over issuing, LIFO will be followed. In the case of tax valuation schemes, LIFO has marginal benefits over FIFO (Cohen et al. 1978). It is possible that good accounting practices dictate the use of FIFO or weighted average costing methods for inventory valuations (Fogarty et al. 1991). In such situations the actual valuation of the inventory can indeed be done on a basis that is different from the physical issuing policy. For example, the inventory can be valued on FIFO basis and the material issued on LIFO basis.

5. ZERO INVENTORY

In this section, we find conditions when *zero inventory* is optimal. Let $J(0)$ and $J(1)$ be the long run *average* costs (per transition) obtained using the zero threshold policy. If it turns out that $J(1) - J(0) + \Lambda r/\mu \geq 0$, then the zero inventory policy is optimal. We compute $J(1) - J(0)$ using a coupling argument and derive conditions for the optimality of zero inventory.

Consider two systems I and II. Initially system I and system II are assumed to be in the states 0 and 1, respectively. System I uses zero as the threshold level. System II starts from state 1 and also follows the zero threshold production policy. The states of system I and II will be referred as a pair with the first element of the pair representing the state of system I and the second element representing the state of system II. We define $(-1, -1)$ and $(0, 0)$ as the absorbing states when the two systems couple. From the initial states, on customer arrival the systems couple with probability p as they reach the absorbing state $(-1, -1)$ from $(0, 1)$ and the mean time for this transition is $1/\lambda$. With probability q they reach the state $(-1, 0)$. The next coupling can take place only when they reach $(0, 0)$. Till then the pair of states evolve as a Markov chain with transition probability matrix

$$\begin{array}{c} \\ \text{abs} \\ (-1,0) \\ (-2,-1) \\ (-3,-2) \\ \cdots \end{array} \begin{array}{cccccc} \text{abs} & (-1,0) & (-2,-1) & (-3,-2) & (-4,-3) & \cdots \end{array} \left|\begin{array}{cccccc} 1 & 0 & 0 & 0 & 0 & \cdots \\ (\mu q/\Lambda) & (\mu p/\Lambda) & (\lambda/\Lambda) & 0 & 0 & \cdots \\ 0 & (\mu q/\Lambda) & (\mu p/\Lambda) & (\lambda/\Lambda) & 0 & \cdots \\ 0 & 0 & (\mu q/\Lambda) & (\mu p/\Lambda) & (\lambda/\Lambda) & \cdots \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \end{array}\right|,$$

where abs denotes the coupling or absorbing state.

Let θ_1 denote the mean time for the systems to couple from the state $(-1, 0)$. Note that system II has one less customer backlogged than system I until the systems couple. Define N to be the mean number of visits to state $(-1, 0)$ starting from $(-1, 0)$ before absorption (coupling). Every time the state $(-1, 0)$ is visited, the conditional probability of incurring a loss due to scrap, given that the systems do not couple, is $\mu p/(\mu p + \lambda)$, and the mean number of visits in which scrap cost might be incurred is $(N - 1)$ because the systems couple on the final visit to $(-1, 0)$. Thus we can write

$$J(1) - J(0) = \frac{h\Lambda}{\lambda} + ps - qb\Lambda\theta_1 - qNr - q(N-1) \times \left(\frac{\mu p}{\mu p + \lambda}\right)s. \tag{17}$$

In this equation, the first term is the carrying cost of one item for mean length of time $1/\lambda$, the second term is the scrap cost if that item is defective, the third term is the cost of backlog for mean time θ_1, the fourth term is the processing cost incurred while processing in state $(-1, 0)$, and the last term is the cost of scrapping if the item produced in state $(-1, 0)$ is defective. It can be verified that $N = \Lambda/(\mu q)$.

Now we need to know θ_1. Define θ_i to be the mean time for the two systems to couple from the state $(-i, -i + 1)$ for all $i \in \{1, 2, 3, \cdots\}$. Then we can write

$$\theta_i = \frac{1}{\lambda+\mu} + \frac{\mu q}{\lambda+\mu}\theta_{i-1} + \frac{\mu p}{\lambda+\mu}\theta_i + \frac{\lambda}{\lambda+\mu}\theta_{i+1} \quad \text{for } i \geq 1. \tag{18}$$

This set of equations can be solved by setting up a set of difference equations in $\omega_i = (\theta_i - \theta_{i+1})$ (see Karlin 1975 for details). Using the assumption that $\mu q > \lambda$ we obtain $\theta_1 = 1/(\mu q - \lambda)$. Let $\beta = (\lambda/\mu q)$ denote the load on the system. Then θ_1 can be written as $\beta/(\lambda(1 - \beta))$.

Lemma 4. *Zero inventory is optimal when* $h \geq (qb\beta/(1 - \beta))$.

Proof. Substituting $\theta_1 = \beta/(\lambda(1 - \beta))$ in the Equation (17) and using the condition $J(1) - J(0) + \Lambda r/\mu \geq 0$ gives the desired result. □

Rewriting the condition for zero inventory as $h \geq b/(\mu/\lambda - 1/q)$ it can be noted that having zero inventory becomes more attractive when the backlogging cost b or arrival rate of customers λ, is small, and when the inventory carrying cost h or the processing rate, μ or the probability of getting a good item, q, is large. The condition $h \geq (qb\beta/(1 - \beta))$ does not involve processing and scrap costs, because in the states below zero, processing and scrap costs do not play a role in the control. They have to be incurred to satisfy demand and the trade off is only between the carrying and backlog costs.

6. NUMERICAL EXAMPLE

An example is presented in this section to provide some insights into the relative advantage of LIFO. Table I gives the long run average cost (per transition $\Lambda = \lambda + \mu$) for both FIFO and LIFO for different system parameters. The column under % Imp. gives the % improvement of the LIFO issuing policy over FIFO. As p increases, the benefit of using LIFO initially increases and then decreases. This phenomenon is due to two conflicting effects discussed below.

Tables II and III present the threshold levels for FIFO and LIFO. The upper threshold level z_1 represents states up to which it is optimal to produce. Note that the upper

Table I
Comparison of FIFO and LIFO Average Costs

Parameters: h = 0.02; r = 0.05; s = 0.025; b = 0.1

	$\lambda = 10$, $\mu = 15$			$\lambda = 20$, $\mu = 35$		
p	FIFO	LIFO	% Imp.	FIFO	LIFO	% Imp.
0.025	0.1426	0.1313	7.9	0.1025	0.0975	4.9
0.05	0.1666	0.1436	13.8	0.1141	0.1044	8.5
0.1	0.2245	0.1769	21.2	0.1408	0.1221	13.3
0.2	0.4487	0.3427	23.6	0.2225	0.1820	18.2
0.3	1.7640	1.5305	13.2	0.4223	0.3589	15.0

threshold levels z_1 for LIFO are represented by pairs with the first number in the pair denoting the number of good items and the second denoting the number of grey items. It can be verified for LIFO that the optimal threshold of the number of grey items is nonincreasing in the number of good items, and if it is optimal to produce in the state (n, m), then it is optimal to do so in the state $(n + m, 0)$ (but not necessarily vice-versa). Also, LIFO gives higher optimal positive inventory levels than FIFO. This is because the danger of items getting scrapped is more in FIFO than in LIFO. This is rather surprising because *reducing* the level of process uncertainty (by using LIFO) leads to *larger* inventories. The process quality, p, seems to have two opposing effects on inventory. (i) Suppose p is large, i.e., the process has poor quality. This means that the expected scrap cost due to producing a grey item increases rapidly with the level of grey items in the inventory. This effect favors less inventory of grey items. (ii) On the other hand, the utilization of the process $\beta = \lambda/\mu q$ becomes larger which favors keeping more inventory (including grey items). These opposing effects result in an increase in benefits from LIFO with increasing p, due to lower scrap costs. Later when the load (utilization of the process) increases with p, LIFO is forced to maintain high inventories leading to smaller benefits. When $m = 0$, the threshold level of n for stopping production is increasing in p, since only the utilization factor is relevant in this case, and the scrap cost is unimportant.

Table II
Upper Threshold (z_1) for FIFO and LIFO

Parameters: h = 0.02; r = 0.05; s = 0.025; b = 0.1; $\lambda = 10$; $\mu = 15$

p	FIFO z_1	LIFO $z_1 = \{n, f(n)\} = \{good, grey\}$
0.025	3	(0,3), (1,2), (2,1), (3,0)
0.05	3	(0,4), (1,3), (2,2), (3,1), (4,0)
0.1	2	(0,4), (1,3), (2,3), (3,2), (4,1), (5,0)
0.2	2	(0,7), (1,7), (2,6), (3,6), (4,5), (5,5), (6,4), (7,4), (8,3), (9,2), (10,2), (11,1), (12,1), (13,0)
0.3	3	(0,10), (1,10), (2,10), (3,10), (4,10), (5,10), (6,10), (7,10), (8,10), (9,10) · · ·, (49,5), (50,5), (51,5), (52,5), (53,4), · · ·

Table III
Upper Threshold (z_1) Level for FIFO and LIFO

Parameters: h = 0.02; r = 0.05; s = 0.025; b = 0.1; $\lambda = 20$; $\mu = 35$

p	FIFO z_1	LIFO $z_1 = \{n, f(n)\} = \{good, grey\}$
0.025	2	(0,2), (1,1), (2,0)
0.05	2	(0,2), (1,1), (2,0)
0.1	1	(0,2), (1,2), (2,1), (3,0)
0.2	1	(0,3), (1,3), (2,2), (3,1), (4,1), (5,0)
0.3	1	(0,5), (1,5), (2,4), (3,4), (4,4), (5,3), (6,3), (7,3), (8,2), (9,2) (10,2), (11,1), (12,1), (13,0), (14,0)

7. CONCLUSIONS AND EXTENSIONS

In this paper we considered a single stage manufacturing system subject to random process shifts. There is delay in getting information about the state of the process because the process quality cannot be measured on-line due to technological limitations. In this setting we analyzed optimal production policies for FIFO and LIFO methods of issuing items to the customer. We proved that LIFO is optimal over all issuing policies when an arriving customer does not wait for the processing of an item, if any, to be completed if there is inventory in the output buffer. A numerical example suggests that up to 20% cost savings can be obtained from LIFO over FIFO. We also derived conditions in which maintaining zero-inventory is optimal. The model and the results presented in this paper can be extended to the following cases (Nurani et al. 1995): (i) When customers can be turned away at a cost (instead of incurring the cost of backlogging), the optimal policy is characterized by an additional *lower threshold level* (z_2). Below this level of backlogged customers, it is optimal to turn away customers. (ii) When a customer arrives, if there is no inventory in the buffer then the subsequent customer arrivals could get stopped. This is equivalent to a pull manufacturing system where the downstream line is stopped if there are no kanbans.

APPENDIX

Proof of Lemma 1. We start with two systems, namely system I and system II, having initial states of $(n + 1, 0)$ and $(n + 2, 0)$, respectively. System II is assumed to be optimally controlled. System I follows the control policy of system II. The machine states and the occurrence of production and arrival events are assumed to be the same in both systems until time τ_n defined below. It may be verified that the control policy is admissible for system I as random events occur together on these systems. Let τ_n be the first time system II hits the state $(n + 1, 0)$ and $\{A\}$ be the set of events where $\{\tau_n < \infty\}$. Denote the number of transitions during the time $[0, \tau_n]$ as τ. We have

$$\begin{aligned} &V(n+2, 0) - V(n+1, 0) \\ &\quad \geq P(II, [0, \infty]) - P(I, [0, \infty]) \\ &\quad = \frac{h\Lambda}{\alpha}\{[1 - \Pr(A)] + [\Pr(A) - E(\rho^{\tau})]\} \\ &\qquad + E(\rho^{\tau})[V(n+1, 0) - V(n, 0)] \\ &\quad \geq [1 - E(\rho^{\tau})][V(n+1, 0) - V(n, 0)] \\ &\qquad + E(\rho^{\tau})[V(n+1, 0) - V(n, 0)] \\ &\quad = V(n+1, 0) - V(n, 0). \end{aligned}$$

The first inequality follows from the fact that $h\Lambda/\alpha \geq V(n+1, 0) - V(n, 0)$. Note that the proof will go through for $n = -1$ because of our assumption that we always produce in a backlogged state, i.e., $V(n+1, 0) - V(n, 0) \leq -\Lambda r/\mu$ for $n < 0$. □

Proof of Lemma 2. Define $C(\tau, A)$ to be the expected discounted cost of carrying one unit for τ transitions on the set of events $\{A\}$. Define $\tau = \infty$ on the set A^c. From Equation (10) we have $C(\tau, A) = (h\Lambda/\alpha)[Pr(A) - E(\rho^{\tau}I(A))]$. This definition is needed for simplifying the equations. We will prove this lemma by contradiction. Assume that the optimal policy is to produce in the state $(n, m+1)$ and not to produce in the state (n, m). These assumptions and the equation for the optimal cost function give

$$V(n, m+2) - V(n, m+1) \leq V(n, m+1) - V(n, m), \tag{19}$$

$$V(n, m+2) - V(n, m+1) \leq -\frac{\Lambda r}{\mu}. \tag{20}$$

As before, we start with two systems I and II, having initial inventory of $(n, m+1)$ and $(n, m+2)$ respectively. System I follows the control policy of system II. System II is assumed to be operated using the optimal policy. If there is a bad item in system I then we assume that the number of bad items in system II is one more than that of system I.

To formalize and justify this assumption assume that the distributions of the first $(m+1)$ items (from bottom of the output buffer stack) in both systems are identical. If any of these items is bad then let the $(m+2)$nd item to be bad and the machines to be in the out-of-control state in both systems. The two systems will hit the same state, i.e., $(n-1, 0)$ when an arrival occurs. If none of these items is bad, then generate a random number, say z, between [0, 1]. If this number is less than p then set the $(m+2)$nd item to be bad in system II. Now wait for the next transition in the uniformizing Poisson process to occur. When it occurs, generate a random number to decide whether the next event will be a production event or an arrival. If it is a production event, add one bad item to both systems and set the machines to be in the out-of-control state for both systems. (This does not benefit system I.) If the next event is an arrival, then the two systems will couple. If z is greater than p, set the $(m+2)$nd item to be good, and assume that the machine states are the same in both systems till a random time to be specified next. The coupling will stop when the states of the systems are the same (an event that occurs if there is a bad item in system II to begin with) or when systems I and II hit the states $(n+m, 0)$ and $(n+m+1, 0)$, respectively (which can happen only if all the $(n+m+2)$ items are good in system II). It can be verified that the control policy for system I is admissible and the only time there is gain in producing in the state $(n, m+2)$ is when all the $(n+m+2)$ items are good.

If there is at least one bad in system II, then the two systems hit the same state at the instant of the next arrival. The time for hitting the same state is exponentially distributed with mean $1/\lambda$. Let η denote the number of transitions needed for the systems to hit the same state through the detection of a defective item and $\{D_1\}$ be the set of events on which it happens. In other words, $\Pr\{D_1\}$ denotes the probability of having at least one bad item in system II. Until the systems couple system II incurs additional cost due to the extra item. Next, consider the case when all the initial items are good in both systems. Let $\{G_1\}$ be the set of events on which this happens. Let τ_n be the time for the systems I and II to hit the states $(n+m, 0)$ and $(n+m+1, 0)$, respectively, with $\tau_n = \infty$ if that event never occurs, and let τ be the number of transitions during τ_n. We have

$$\begin{aligned} &V(n, m+2) - V(n, m+1) \\ &\quad \geq P(II, [0, \infty]) - P(I, [0, \infty]) \\ &\quad = C(\eta, D_1) + C(\tau, G_1) + E(\rho^{\tau}I(G_1)) \\ &\quad \times [V(n+m+1, 0) - V(n+m, 0)]. \end{aligned} \tag{21}$$

The next step is to get a bound for $V(n, m+1) - V(n, m)$. To do this, we consider two more systems, called III and IV, starting with the initial inventory (n, m) and $(n, m+1)$, respectively. Assume that system III is optimally controlled and hence we do not produce in the state (n, m) according to our assumption. Let system IV follow the control policy of system III. As done before, if there are bad items in system III assume that the system IV has one more bad item than system III, and that the states of the machines in the two systems are the same till the random time specified below. To formalize this, let the distribution of the first m items be the same in both systems, and also be the same as that in systems I and II. If any of these items is bad assume that the machines are in the out-of-control state in both systems and let the $(m+1)$st item be bad in system IV. (This does not benefit system IV.) Else if all m items are good generate a random number z between [0, 1]. If this is less than p let the machines be in out-of-control states in both systems and the $(m+1)$st item be bad in system IV. Else let the $(m+1)$st item be good and the machine states in both systems be the same. System IV has more information than system III, and therefore the control policy of system III is admissible for system IV.

As before, if a bad item is found in one of the systems then they hit the same state at the instant of the next arrival, and this time is exponentially distributed with mean $1/\lambda$. Thus η is the number of transitions needed for the systems to hit the same state and let $\{D_2\}$ be the set of events in which this happens. However, until such a time system IV incurs additional carrying cost due to the extra item. Next, consider the case when all the items are good. Let $\{G_2\}$ be the set of events on which this happens. Let ν_n be the time for systems III and IV to hit the states $(n + m - 1, 0)$ and $(n + m, 0)$, respectively, with $\nu_n = \infty$ if that event never happens. Note that since we do not produce in the state (n, m) in system III, ν_n is also exponentially distributed with mean $1/\lambda$. Let η (once again) be the number of transitions. (Note η is the same on D_1, D_2 and G_2 because the coupling time is the time for the next arrival.) Then we have

$$\begin{aligned} V(n, m+1) - V(n, m) &\leq P(IV, [0, \infty]) - P(III, [0, \infty]) \\ &= C(\eta, D_2) + C(\eta, G_2) + E(\rho^{\eta} I(G_2)) \\ &\times [V(n+m, 0) - V(n+m-1, 0)]. \end{aligned} \tag{22}$$

Using Equations (19), (21), and (22) we get

$$\begin{aligned} &C(\eta, D_2) + C(\eta, G_2) + E(\rho^{\eta} I(G_2)) \\ &\times [V(n+m, 0) - V(n+m-1, 0)] \\ &\geq C(\eta, D_1) + C(\tau, G_1) + E(\rho^{\tau} I(G_1)) \\ &\times [V(n+m+1, 0) - V(n+m, 0)]. \end{aligned} \tag{23}$$

The probability of there being one more bad item in system II compared to system I is greater than the probability of having one more bad item in system IV compared to system III, because even when all the items are good in systems III and IV, the probability that there is a bad item in system II and none in system I is positive. Thus we have:

$$\Pr\{D_1\} + \Pr\{G_1\} = \Pr\{D_2\} + \Pr\{G_2\}, \tag{24}$$

$$\Pr\{D_1\} \geq \Pr\{D_2\}, \tag{25}$$

$$\Pr\{G_1\} \leq \Pr\{G_2\}. \tag{26}$$

The time taken to hit the states $(n + m + 1, 0)$ and $(n + m, 0)$ by systems I and II cannot be less than the time for the next arrival in systems III and IV, because we may produce in the state $(n, m + 1)$ in system I and do not produce in the state (n, m) in system III. Hence

$$\tau \geq_{st} \eta. \tag{27}$$

By Equations (24)–(27), and the fact that the coupling time is the same on D_1, D_2 and G_2 we get:

$$\begin{aligned} C(\eta, D_2) + C(\eta, G_2) - C(\eta, D_1) &\leq C(\eta, G_1) \\ &\leq C(\tau, G_1), \end{aligned} \tag{28}$$

and

$$C(\eta, D_2) + C(\eta, G_2) \leq C(\eta, D_1) + C(\tau, G_1). \tag{29}$$

Combining (23) and (29) we get:

$$\begin{aligned} &E(\rho^{\eta} I(G_2))[V(n+m, 0) - V(n+m-1, 0)] \\ &\geq E(\rho^{\tau} I(G_1))[V(n+m+1, 0) - V(n+m, 0)]. \end{aligned} \tag{30}$$

$V(n + m + 1, 0) - V(n + m, 0) < 0$ from Equations (20) and (21). Using this fact in Equation (30) and using (26) and (27) it follows that (30) is a contradiction of Lemma 1, that $V(n, 0)$ is convex in n. □

Proof of Lemma 3. The proof is by contradiction and is similar to Lemma 2. We sketch this proof. For more details refer Nurani et al. (1995). Assume that the optimal policy is not to produce in the state (n, m) and to produce in the state $(n + 1, m)$. Consider one pair of systems I and II, having initial inventory of $(n + 1, m)$ and $(n + 1, m + 1)$, respectively. System I follows the control policy of system II. System II is assumed to be operated using the optimal policy. Assume that the distribution of m grey items is the same in the two systems, and if there is a bad item in system I then the number of bad items in system II is one more than that in system I. Define η to be the number of transitions needed for these two systems to hit the same state through detection of a defective item and $\{D_1\}$ to be the set of events on which it occurs. Let τ_n be the time for systems I and II to hit the states $(n + m, 0)$ and $(n + m + 1, 0)$, respectively, and τ be the number of transitions. Let $\{G_1\}$ be the set of events on which this happens. We have

$$\begin{aligned} &V(n+1, m+1) - V(n+1, m) \geq P(II, [0, \infty]) \\ &- P(I, [0, \infty]) = C(\eta, D_1) + C(\tau, G_1) + E(\rho^{\tau} I(G_1)) \\ &\times [V(n+m+1, 0) - V(n+m, 0)]. \end{aligned} \tag{31}$$

Consider two more systems, called system III and IV, starting with the initial inventory (n, m) and $(n, m + 1)$, respectively. η is the number of transitions needed for the systems to hit the same state through detection of a defective item, and let $\{D_2\}$ be the set of events in which this happens. Next, let ν_n be the time for system III and IV to hit the states $(n + m - 1, 0)$ and $(n + m, 0)$, respectively. Let $\{G_2\}$ be the set of events on which this happens. Then we have

$$\begin{aligned} V(n, m+1) - V(n, m) &\leq P(IV, [0, \infty]) - P(III, [0, \infty]) \\ &= C(\eta, D_2) + C(\eta, G_2) + E(\rho^{\eta} I(G_2)) \\ &\times [V(n+m, 0) - V(n+m-1, 0)]. \end{aligned} \tag{32}$$

As in Lemma (2), we can show:

$$C(\eta, D_2) + C(\eta, G_2) \leq C(\eta, D_1) + C(\tau, G_1), \tag{33}$$

and also that

$$\Pr\{G_1\} \leq \Pr\{G_2\}, \tag{34}$$

because the probability of having all the items good is the same in systems II and IV. Also we have

$$\tau \geq_{st} \eta, \tag{35}$$

because we do not produce in systems III and IV, whereas we may produce in system II in the starting state $(n+1, m+1)$. Using (31), (32), (33), (34), (35), and the assumptions we get:

$$E(\rho^{\eta} I(G_2))[V(n+m, 0) - V(n+m-1, 0)]$$

$$\geqslant E(\rho^{\tau} I(G_1))[V(n+m+1, 0) - V(n+m, 0)]. \quad (36)$$

$V(n+m+1, 0) - V(n+m, 0) < 0$ from the assumption that we produce in the state $(n+1, m)$ and (31). Using this fact in (36), and using (34) and (35) it follows that (36) is a contradiction of lemma 1 that $V(n, 0)$ is convex in n. □

ACKNOWLEDGMENT

The authors thank the Area Editor, the Associate Editor and an anonymous referee for their insightful comments and suggestions. Their comments helped improve the exposition of the paper. The second part of Theorem 2 was motivated by one of the anonymous referees' comments. The first author would like to thank the department of Industrial Engineering and Operations Research, University of California at Berkeley, where the paper was written while he was visiting. His research has been supported in part by KLA Instruments Corporation. The research of the third author has been supported in part by the NSF grant DDM-9113008 and by the Sloan grant for "Competitive Semiconductor Manufacturing."

REFERENCES

Bertsekas, D. P. 1983. *Dynamic Programming: Deterministic and Stochastic Models.* Prentice-Hall, Englewood Cliffs, NJ.

Bielecki, T. and P. R. Kumar. 1988. Optimality of Zero-Inventory Policies for Unreliable Manufacturing Systems. *Opns. Res.* **36,** 532–541.

Buzacott, J. A. and J. G. Shanthikumar. 1993. *Stochastic Models of Manufacturing Systems.* Prentice-Hall, Englewood Cliffs, NJ.

Cohen, M. A. and D. Pekelman. 1979. Optimal Inventory Ordering Policy with Tax Payments under FIFO and LIFO Accounting Systems. *Mgmt. Sci.* **25,** 729–743.

Fogarty, D. W., J. H. Blackstone, Jr., and T. R. Hoffmann. 1991. *Production and Inventory Management.* Second edition, South-Western, Cincinnati, OH.

Hsu, L.-F. and C. S. Tapiero. 1992. Integration of Process Monitoring, Quality Control and Maintenance in an M/G/1 Queue-like Production System. *Int. J. Prod. Res.* **30,** 2363–2379.

Karlin, S. and H. M. Taylor. 1975. *A First Course in Stochastic Processes.* Second Edition. Academic Press, Inc., CA.

Lippman, S. 1975. Applying a New Device in the Optimization of Exponential Queueing Systems. *Opns. Res.* **23,** 687–710.

Monden, Y. 1981. Adaptable Kanban System Helps Toyota Maintain Just-in-Time Production. *Industrial Engineering,* **13,** 28–46.

Nahmias, S. 1982. Perishable Inventory Theory: A Review. *Opns. Res.* **30,** 680–708.

Nurani, R. K., S. Seshadri, and J. G. Shanthikumar. 1995. Optimal Control of a Single Stage Production System Subject to Uncertain Process Quality. Working Paper.

Porteus, E. L. 1986. Optimal Lot Sizing, Process Quality Improvement and Setup Cost Reduction. *Opns. Res.* **34,** 137–144.

Silver, E. A. and R. Peterson. 1985. *Decision Systems for Inventory Management and Production Planning.* John Wiley & Sons, NY.

Taguchi, G., E. A. Elsayed, and T. Hsiang. 1989. *Quality Engineering in Production Systems.* McGraw Hill, NY.

Tapiero, C. S. and L.-F. Hsu. 1987. Quality Control of the M/M/1 Queue. *Int. J. Prod. Res.* **25,** 447–455.

Yano, C. A. and H. L. Lee. 1995. Lot-Sizing With Random Yields: A Review. *Opns. Res.* **43,** 311–334.