Production and Repair Decisions with Time-Consuming Repair and a Deadline

Medini R. Singh
Columbia Business School
Columbia University, New York, NY 10027
e-mail: ms2149@columbia.edu

Jing-Sheng Song
The Fuqua School of Business
Duke University, Durham, NC 27708
e-mail: jssong@duke.edu

Candace A. Yano
IEOR Department and Haas School of Business
University of California, Berkeley, CA 94720
e-mail: yano@ieor.berkeley.edu

Augusto Moreno-Beltran
IEOR Department
University of California, Berkeley, CA 94720
e-mail: chuktah@hotmail.com

November 26, 2004

Abstract

We investigate optimal production and repair decisions when a customer order is due at the end of a finite horizon. The machine stochastically deteriorates over time when operating, and the yield distribution (output of good product) depends on the machine state. Repair takes time but brings the machine to a stochastically better state (not necessarily to the perfect state). Both production and repair incur costs. Revenue is earned for units produced by the end of the horizon, including possible salvage value for units exceeding the demand. There is also a state-dependent reward for the machine condition at the end of the horizon. In each period, we must decide whether to produce, repair, or do nothing, with the objective of maximizing expected profit during the horizon.

We show that the optimal policy for the single-period problem can be characterized by a single critical inventory level for each machine state. However, the simple structure does not necessarily extend to the multi-period problem. We develop a heuristic policy that combines a threshold repairing policy based on the machine state with a myopic production control based on the inventory level. The heuristic performs very well and is much easier to implement than the optimal policy.

1 Introduction

In the day-to-day operation of manufacturing plants, production managers often face the dilemma of whether to invest time to repair or improve the condition of a machine so as to increase its effective output rate, or whether to continue operating the machine in a degraded condition. Making such decisions is especially difficult when time is of the essence, either because of relatively short windows of opportunity for meeting market demand, or because customer orders have imminent due dates. Problems of this type are especially difficult when repair takes time (in addition to costing money) and may be imperfect, leading to a tradeoff between time and money. Such a tradeoff is difficult to evaluate in a finite horizon setting because the tradeoff function changes with the proximity of the horizon.

We address a finite-horizon production and repair problem in which both the production and repair processes take time and may be imperfect. The vast majority of the literature on repair and/or replacement policies are based on an infinite horizon, instantaneous repair, or perfect repair, or a combination of these assumptions. Relaxing all three of these assumptions makes the problem more realistic, but also much more difficult to analyze.

Our work on this problem was motivated by our observations of equipment operators and production managers in a variety of discrete parts manufacturing industries in such time-stress situations. Typically, a rough production schedule is established and customer orders are processed one at a time on a given machine, possibly due to setups or simply for administrative convenience, with more urgent orders being processed before those that are due later. The machine operator is primarily concerned about completing the most urgent customer order by its due date. What we have observed, time and again, is the strong tendency to continue producing, even if the effective output rate of the machine is low. The main deterrent (for the machine operator) to stopping the machine for recalibration or maintenance is not the cost but the time required for repair, during which no production can take place. We would like to determine how best to operate systems under these realistic conditions.

There is a single product which is produced on a single production facility ("machine") that has a fixed capacity in each period. To model scenarios with imminent deadlines, in our finite-horizon framework we assume each time period is short enough so that it is practical to devote the entire period to a single activity. For ease of exposition, we do not discount costs, but it is straightforward to modify our formulation to admit discounted costs if desired.

There are a finite number of machine states and we assume the machine state is known at the beginning of each period. The state of the machine at the beginning of next period depends only on its current state and our current decision (produce, repair, or do nothing). We recognize that machine states may be difficult to define precisely in practice, and that they may not be fully or accurately observable. In our discussions with experienced machine operators, however, we learned that equipment often have certain characteristic modes of operating, and these modes could be construed as machine states. The modes may be defined by such factors as the level or nature of the noise being emitted by the equipment, or levels of various factors (e.g., pressure, temperature, chemical concentration) that are detectable from gauges. Alternatively, modes (states) may be defined by characteristics of the output. Recognizing that the definition and observation of machine states are, in practice, more of an art than a science, the management of these manufacturing firms nevertheless were interested in understanding the tradeoff between getting some good parts now versus (possibly) more later, which is a tradeoff that they considered to be of critical importance.

The yield distribution (distribution of good output) depends on the machine state and is stochastically decreasing in the state index. The machine can be repaired, consuming both time and money. Repair may be imperfect, but it will bring the machine to some (stochastically) better state in the next time period. A state-independent fixed cost is incurred for each period of repair. This reflects situations where a repairer or repair crew is needed, and most of the cost of repair is the true or opportunity cost of the associated labor.

There is a state-independent fixed production cost in each production period. In most instances, this consists of the cost of materials and processing for the fixed lot, which typically does not depend on the quality of the output. Recall, however, that the yield is state-dependent, so the expected cost per unit of good output increases as the machine state becomes worse. There is also a state-dependent reward for the machine condition at the end of the horizon, reflecting the impact of the machine state on future profits.

An order of D units is due at the end of the horizon. A per-unit revenue is earned for each good unit up to D, and a salvage value is earned for each good unit in excess of D. For simplicity, we assume that all revenue is earned at the end of the horizon. The objective is to maximize the expected total profit over the horizon.

The remainder of the paper is organized as follows. In Section 2, we provide a review of the related literature. In Section 3, we present our assumptions and a formulation. Then, in Section 4, we derive structural results that characterize the optimal policy for the single-period problem. Section 5 contains some structural results for the dynamic problem. In both Sections 4 and 5, we further explore characteristics of the optimal solution through numerical examples and present related insights. In Section 6, we utilize these insights and other observations from numerical results to devise a heuristic policy. The heuristic combines a myopic rule for stopping production as accumulated inventory approaches the order quantity with a simple threshold (machine state) policy that specifies whether to produce or repair in the earlier part of the horizon. We report the results of extensive computational results which show that the heuristic performs quite well. The paper ends with some final thoughts in Section 7. All proofs appear in Appendix B.

2 Literature Review

There is a large literature on problems involving maintenance decisions. We review models in which one may perform repair prior to failure of the machine, and repair is imperfect and/or non-instantaneous. See McCall (1963), Pierskalla and Voelker (1976), Sherif and Smith (1981), Valdez-Flores and Feldman (1989), and Wang (2002) for reviews of the broader literature on repair/replacement policies. A few papers permit varying degrees of repair (e.g., Chikte and Deshmukh 1981, Kijima 1989, Zuckerman 1986, and Statje and Zuckerman 1991), but many of these assume that maintenance is performed only upon failure.

Early research (Ross 1971, Smallwood and Sondik 1972, Rosenfield 1976 and White 1978) considers perfect, non-instantaneous repair and costly inspection. Most are infinite horizon models, and most assume imperfect inspection and thus emphasize the inspection decision, with the production versus replacement (i.e., perfect repair) decision being secondary.

More closely related to our work is a line of research that presumes inspection is instantaneous and free, and concentrates on whether to produce or repair. Venkatesan (1984) examines a finite horizon problem with random demands and perfect repair. The objective function includes fixed and variable production, inventory, shortage and repair costs, where the cost of replacement is higher if the equipment has failed. Venkatesan derives the structure of the optimal policy for the single-period problem, which has four regions. Because the dynamic programming value functions are non-convex, he derives additional conditions on these value functions in which the four- region policy carries over to multiple periods. Unfortunately, these conditions cannot be tested in advance. Bobos and Protonotarious (1978) consider an infinite horizon problem with different degrees of repair, and each repair action has a probabilistic outcome. They show that the optimal policy has a multiple-control-limit form, under the assumption that the expected deterioration of the machine decreases as one applies more effective maintenance. The form of the control-limit policy is as follows: the "no action" decision is optimal for a set of "best" states, and successively more extensive repair actions are optimal for a set of "next best" states.

Douer and Yechiali (1994) consider a similar problem with several degrees of instantaneous repair, including replacement. Both deterministic repair and uncertain repair are analyzed, and the cost of repair depends on the machine state prior to repair. They show that under certain conditions on the costs and the transition probabilities (for deterioration), the optimal policy has a generalized control limit rule: repair or replace only when the state of the system is worse than a threshold. The conditions for optimality of the generalized control limit policy bear resemblance to some of the conditions that we require in our problem; we elaborate on this point later. Although the Douer and Yechiali model is more general in some respects than many earlier papers, the authors assume that the probability that a planned repair to a state j ends up in state v does *not* depend on the pre-repair state i. This imposes considerable structure on the repair transition matrix, as we explain in more detail later. Indeed, the assumptions made by Bobos and Protonarios in this regard are less restrictive, and parallel our assumptions. Sloan and Shanthikumar (2000) generalize the model of the Douer and Yechiali model to consider repair transitions that may depend on the pre-repair state. As in Douer and Yechiali, repair is instantaneous. They show that the optimal policy has a control-limit form.

A recent article by Iravani and Duenyas (2002) is also related. Similar to our study, they too consider integrated decisions of maintenance and production in a production system with a deteriorating machine. However, the details of their model are different from ours. First, they consider infinite horizon. Second, they model the production and repair times as exponential random variables. Instead of assuming random yield and imperfect repair, they assume the machine state affects the production and repair processes through state-dependent production and repair rates. They show that the optimal control policy has a complex form, and they propose a double-threshold policy and derive exact and approximate methods for evaluating and optimizing this type of policy. Numerical results show that the proposed policy is quite effective. We refer the reader to the references in this paper for other related literature.

As we have seen, the vast majority of research addresses steady-state, infinite horizon problems with perfect and/or instantaneous repair. In an infinite horizon framework, *either* perfect or instantaneous repair allows strong characterizations of the optimal solution, making the problem relatively easy to solve. Our paper differs from the aforementioned papers in several ways. The key complication is the combination of non-instantaneous repair, a finite horizon and finite demand, but the presence of random yields adds another complication. Together, these features create time tradeoffs that do not occur in infinite horizon problems, because repair efforts compete with production for the available time. In addition, these factors make the implied value of repair time-dependent. Moreover, the imperfect nature of the repair process, when combined with the finite horizon, causes time dependencies in the "riskiness" of the repair decision, with early repair decisions being less risky than those performed later in the horizon.

3 Assumptions and Problem Formulation

In this section, we state more detailed assumptions about problem data and parameters,and then go on to formulate the problem. A list of basic notation follows:

- k stage index, $k = 1, ..., K$; stage 1 is the last period in the horizon and k represents number of periods remaining.
- i machine state index, $i = 1, ..., I$; state 1 is the best and state I is the worst.
- x inventory on hand.
- r_{ij} probability that machine improves from state i to state j after one period of repair.
- p_{ij} probability that machine degrades from state i to state j after one period of production.
- U_i production yield when the machine is in state i, a random variable with mean μ_i, pmf $h_i(\cdot)$ and cdf $H_i(\cdot)$; $0 \leq U_i \leq q$.
- D order quantity due at end of the horizon.
- q quantity processed in one period if production takes place.
- f cost per period of repair.
- w cost per period of production.
- π revenue per good unit produced.
- δ salvage value per good unit in excess of D at the end of the horizon.
- η_i the value of being in state i at the end of the horizon.

We make the following assumptions, which are qualitatively reasonable:

A.1 The revenue per unit is greater than the salvage value per unit, i.e., $\pi \geq \delta$. This ensures that we satisfy demand before salvaging units.

A.2 The machine provides a stochastically larger yield when it is in a better state. That is, $U_i \geq_{st} U_j$ for $i < j$, or

$$E[g(U_i)] \geq E[g(U_j)] \qquad \text{for all increasing function } g \text{ and } i < j.$$

A corollary to this assumption is that μ_i is non-increasing in i.

A.3 A worse machine state yields a smaller terminal value. That is, η_i decreases in i.

A.4 There is an incentive to produce. That is, the expected reward from production, when entire output is used for satisfying demand, exceeds the cost of production and associated machine deterioration,

$$\pi\mu_i > w + \left(\eta_i - \sum_{j=i}^{I} p_{ij}\eta_j\right) \qquad \text{for all } i. \tag{1}$$

Otherwise, one would never produce in state i. Note that condition (1) does not imply that it is optimal to produce in every state. For example, one may prefer inaction over production if there is enough inventory on hand. Similarly, one may prefer repair over production if repair is economically more attractive. Condition (1) simply means that production should be considered among the set of possible actions in each state.

A.5 It is not profitable to produce simply to salvage the output. That is, the expected reward from salvage does not exceed the cost of production plus machine deterioration,

$$\delta\mu_i < w + \left(\eta_i - \sum_{j=i}^{I} p_{ij}\eta_j\right) \qquad \text{for all } i.$$

A.6 The machine state after producing for one period is stochastically worse if the machine was in a worse state when production began. That is,

$$\sum_{j=k}^{I} p_{ij} \qquad \text{is nondecreasing in } i \text{ for all } k.$$

In other words, the machine deterioration matrix, $\mathbf{P}$, is increasing failure rate (IFR).

A.7 The machine state after repair is stochastically better if the machine was in a better state before the repair. That is,

$$\sum_{j=1}^{k} r_{ij} \qquad \text{is nonincreasing in } i \text{ for all } k.$$

Since $\sum_{j=1}^{I} r_{ij} = 1$, this condition also implies that $\sum_{j=k}^{I} r_{ij}$ is nondecreasing in i for all k, i.e., the repair matrix, $\mathbf{R}$, is also IFR.

A.8 The difference between the expected machine state reward resulting from repairing and from producing is larger if the machine condition is worse. That is,

$$\sum_{j=1}^{I}(p_{ij} - r_{ij})\eta_j \quad \text{decreases in } i.$$

Because assumption A.8 involves matrices **P** and **R**, we further explore its relationship to assumptions A.6 and A.7 in Appendix A.

In practical settings, the machine condition rarely improves by itself during production. Similarly, a repair seldom makes the machine state worse. For these reasons, **P** is usually upper-triangular and **R** lower-triangular. A good example of machine deterioration and repair processes that satisfy assumptions A.6 and A.7 results from a machine with N identical Bernoulli components. The machine state is related to the number of failed components. Matrix **P** depends upon the probability of component failure and matrix **R** depends upon the probability of successful repair. We use such a deterioration and repair process in the numerical examples in Sections 4 and 5.

We formulate the problem as a dynamic program. It is clear that we need two state variables to keep track of the system dynamics. One is the current machine state i and the other is the current inventory level x. We assume that revenue is collected only at the end of the horizon. In practice, this is reasonable, because the end of the planning horizon is defined by the due date of the customer order. Because we are treating a short-horizon problem, we do not discount costs, but it is straightforward to modify our formulation to admit discounted costs if desired.

Define $V_k(i, x)$ be the *optimal* total expected profit for stages $k, \ldots, 1$, given that the system state is (i, x) at stage k. Also, let $P_k(i, x)$, $R_k(i, x)$ and $N_k(i, x)$ be the total expected profit when at stage k the system state is (i, x) and a decision is made to produce, to repair, or do nothing, respectively, and an optimal policy is used for stages $k, \ldots, 1$. Then, we can write the dynamic programming recursion equations as follows:

$$\begin{aligned} V_0(i, x) &= \eta_i + \pi \min(D, x) + \delta(x - D)^+ \\ V_k(i, x) &= \max\{P_k(i, x), R_k(i, x), N_k(i, x)\}, \end{aligned} \tag{2}$$

where

$$\begin{aligned} P_k(i, x) &= \sum_{j=i}^{I} p_{ij} E\left[V_{k-1}(j, x + U_i)\right] - w, \\ R_k(i, x) &= \sum_{j=1}^{i} r_{ij} V_{k-1}(j, x) - f, \\ N_k(i, x) &= V_{k-1}(i, x), \qquad k \geq 1. \end{aligned}$$

4 Analysis of Single-Period Problem

In this section we focus on the last period, i.e., $k = 1$. Although an important motivation for solving the static problem is to provide a foundation for the dynamic problem, the static problem is important in its own right. Many plant managers have found themselves in a situation where an order must be shipped today to meet a customer deadline, but the order is not complete and the associated equipment is operable but not performing well. The decision to repair, rather than attempting to complete the order, would necessitate making a "short" shipment, or using premium transportation or overtime, or calling the customer to advise of a delay and incurring a loss of goodwill (or in some cases, monetary penalties). The optimal decision should, of course, depend on the extent of the shortage and the state of the machine. It turns out that the decisions are most complex when only a few units remain to be produced and the machine is performing neither well nor poorly. When only a few additional units are needed, the common response is to continue producing, but this is often suboptimal. We explore these issues in this section.

Observe that

$$\begin{aligned} V_0(i,x) &= \eta_i + \pi[D-(D-x)^+] + \delta(x-D)^+ \\ &= \eta_i + \pi D + \delta(x-D)^+ - \pi(D-x)^+ \\ &= \eta_i + \pi x + (\delta-\pi)(x-D)^+. \end{aligned}$$

So

$$\begin{aligned} P_1(i,x) &= \sum_{j=i}^{I} p_{ij}\eta_j + \pi(x+\mu_i) + (\delta-\pi)E\left\{(x+U_i-D)^+\right\} - w, \\ R_1(i,x) &= \sum_{j=1}^{i} r_{ij}\eta_j + \pi x + (\delta-\pi)(x-D)^+ - f, \\ N_1(i,x) &= \eta_i + \pi x + (\delta-\pi)(x-D)^+ . \end{aligned}$$

Define function $F(i,x) = P_1(i,x) - R_1(i,x)$, which represents the relative value of performing production instead of repair in state (i,x). Similarly, define function $G(i,x) = N_1(i,x) - R_1(i,x)$ which represents the relative value of no action compared to repair in state (i,x). The optimal decision can be determined by the signs and relative values of these functions. Whenever F is positive and greater than G, it is optimal to produce. Similarly, whenever G is positive and greater than F, doing nothing is optimal. Finally, whenever both F and G are negative, repair is optimal. Clearly, the nature of functions F and G is critical in determining the form of the optimal policy. These functions can be written as:

$$\begin{aligned} F(i,x) &= (\delta-\pi)\ E\left\{(x+U_i-D)^+ - (x-D)^+\right\} \\ &\qquad + \pi\mu_i + \sum_{j=i}^{I} p_{ij}\eta_j - w - \sum_{j=1}^{i} r_{ij}\eta_j + f, \end{aligned} \tag{3}$$

$$G(i,x) \quad = \quad \eta_i - \sum_{j=1}^{i} r_{ij}\eta_j + f. \tag{4}$$

Note from (4) that $G(i,x) = G(i)$ is independent of x. Also, it may not be decreasing in i even though η_i and $\sum\limits_{j=1}^{i} r_{ij}\eta_j$ are both decreasing i.

Classification of States

Definition: A machine state i is called

(i) a *good* state if

$$\sum_{j=1}^{i} r_{ij}\eta_j - \sum_{j=i}^{I} p_{ij}\eta_j + w - f \leq \delta\mu_i, \tag{5}$$

(ii) a *bad* state if

$$\sum_{j=1}^{i} r_{ij}\eta_j - \sum_{j=i}^{I} p_{ij}\eta_j + w - f \geq \pi\mu_i, \text{ or} \tag{6}$$

(iii) an *intermediate* state if

$$\delta\mu_i < \sum_{j=1}^{i} r_{ij}\eta_j - \sum_{j=i}^{I} p_{ij}\eta_j + w - f < \pi\mu_i. \tag{7}$$

An intuitive explanation for the classification of machines states can be given as follows. If the machine is in state i, a decision to repair it will cost f but will improve its expected value by $\sum_{j=1}^{i} r_{ij}\eta_j - \eta_i$. Similarly, a decision to produce will incur a production cost w and a cost due to expected deterioration in machine value, $\eta_i - \sum_{j=i}^{I} p_{ij}\eta_j$, but will bring a reward based on the output. A good state is such that the net profit resulting from repair is no larger than the net profit resulting from production even if the output is completely in excess of the demand in the current planning horizon. This can be seen by rearranging equation (5) as

$$\left(\sum_{j=1}^{i} r_{ij}\eta_j - \eta_i\right) - f \leq \delta\mu_i - \left(w + \left(\eta_i - \sum_{j=i}^{I} p_{ij}\eta_j\right)\right).$$

Similarly, a bad state is such that the net profit resulting from production when all the output is useful for satisfying demand (not for salvage) is no larger than the net profit resulting from repair. This can be seen by rearranging equation (6) as

$$\pi\mu_i - \left(w + \left(\eta_i - \sum_{j=i}^{I} p_{ij}\eta_j\right)\right) \leq \left(\sum_{j=1}^{i} r_{ij}\eta_j - \eta_i\right) - f.$$

The machine is in an intermediate state if it is neither in a good nor in a bad state.

Proposition 1 *The state classification is such that*

(i) If i is a good state, then $i-1$ is also a good state.

(ii) If i is a bad state, then $i+1$ is also a bad state.

(iii) The sets of good, bad and intermediate states are mutually exclusive and compact.

Critical Numbers

Let $x^*(i)$ be the smallest value of x such that $F(i,x) < 0$, or equivalently $P(i,x) < R(i,x)$. If $x < x^*(i)$, one would prefer production over repair; if $x \geq x^*(i)$, one would prefer repair over production. We define $x^*(i) = \infty$ if $F(i,x) > 0 \quad \forall x$, and similarly define $x^*(i) = 0$ if $F(i,x) \leq 0 \ \forall x$.

Let $\tilde{x}(i)$ be the smallest value of x such that $F(i,x) < G(i)$, or equivalently $P(i,x) < N(i,x)$. If $x < \tilde{x}(i)$, one would prefer production over inaction; if $x \geq \tilde{x}(i)$, one would prefer inaction over production.

In the following two propositions, we show that these critical numbers exist and are unique. We then characterize the optimal policy for the single-period problem.

Proposition 2 *For each intermediate state i, there exists a unique critical number $x^*(i)$ such that*

(i) $D - q < x^*(i) < D, \ \forall i$

(ii) $x^*(i) > x^*(i+1), \ \forall i$.

Proposition 3 *For each state i, there exists a unique critical number $\tilde{x}(i)$ such that $D - q < \tilde{x}(i) < D, \ \forall i$.*

Note that $\tilde{x}(i)$ is finite for all machine states and unlike $x^*(i)$, $\tilde{x}(i)$ is not necessarily decreasing in i.

Theorem 1 *The optimal policy for the single period problem has the following structure:*

(i) If i is a **good** *machine state satisfying (5) then*

  (a) regardless of inventory level x, it is never optimal to repair in machine state i, and

  (b) there exists a unique critical number, $\tilde{x}(i)$, such that it is optimal to produce if $x < \tilde{x}(i)$ and do nothing if $x \geq \tilde{x}(i)$.

(ii) If i is a **bad** *machine state satisfying (6) then it is optimal to repair in state i regardless of the inventory level x.*

(iii) If i is an **intermediate** *machine state satisfying (7) then there exist two critical numbers, $\tilde{x}(i)$ and $x^*(i)$, the smaller of which determines the optimal policy such that*

(a) if $\tilde{x}(i) < x^*(i)$ then it is optimal to produce if $x < \tilde{x}(i)$ and do nothing if $x \geq \tilde{x}(i)$,

(b) if $\tilde{x}(i) > x^*(i)$ then it is optimal to produce if $x < x^*(i)$ and repair if $x \geq x^*(i)$,

(c) if $\tilde{x}(i) = x^*(i)$ then it is optimal to produce if $x < \tilde{x}(i) = x^*(i)$). The optimal policy for $x \geq \tilde{x}(i) = x^*(i)$ is to do nothing if $G(i) > 0$, repair if $G(i) < 0$ and either do nothing or repair (between which one is indifferent) if $G(i) = 0$.

Theorem 1 implies a few clean and intuitive features of the optimal policy as a function of inventory level x for a given machine state:

(i) For any machine state i, the optimal policy is given by a single critical value of x. That is, as x increases, the decision changes at most once. As we shall see, this is not necessarily the case for the multi-period problem.

(ii) Suppose it is optimal to repair (do nothing) in state (i, x). Then it is also optimal to repair (do nothing) in state $(i, x + 1)$.

(iii) If it is optimal to produce in state (i, x) then it is also optimal to produce in state $(i, x-1)$.

However, counter-intuitively, for each inventory level x, the optimal policy cannot be characterized by a single critical number on machine state. Implications include:

(i) Since the critical numbers $x^*(i)$ and $\tilde{x}(i)$ may cross each other several times, the set of intermediate states may have several disjointed repair and do-nothing regions. As the state gets worse, the optimal decision may change several times, as shown in Figure 1.

(ii) One may conjecture that if it is optimal to repair the machine in state (i, x) then it is also optimal to repair in state $(i + 1, x)$, i.e., in a worse machine state at the same inventory level. This is not necessarily the case.

(iii) One may conjecture that if it is optimal to produce in state (i, x), it is also optimal to produce in state $(i - 1, x)$, i.e., in a better machine state at the same inventory level. This is not necessarily the case since $\tilde{x}(i)$ is not monotonic.

We now consider conditions that simplify the optimal policy. These conditions are more restrictive but they may hold in some realistic situations. They also shed further light on the nature of the optimal policy.

Suppose repair becomes more rewarding as the machine state gets worse. That is, the expected gain in terminal value of the machine is higher if repair is carried out in a worse state. Then, for the intermediate states, the optimal policy has a simpler structure, as characterized by the following proposition.

Proposition 4 *If* $\sum_{j=1}^{i} r_{ij}\eta_j - \eta_i$ *is increasing in* i *then*

(i) the sequence of critical numbers, $\tilde{x}(i)$*, crosses the sequence of critical numbers* $x^*(i)$ *at most once, and it does so from below,*

(ii) for any inventory level, if repair is optimal in a machine state, then it is also optimal in all worse machine states.

Disjoint repair regions, as shown in Figure 1, are ruled out by Proposition 4. The optimal policy for intermediate machine states has a simple structure: For the "better" intermediate states, produce if $x < \tilde{x}(i)$ and do nothing if $x \geq \tilde{x}(i)$; for the "worse" intermediate states, produce if $x < x^*(i)$ and do nothing if $x \geq x^*(i)$. The transition from the "better" to "worse" intermediate states occurs when $\tilde{x}(i)$ crosses $x^*(i)$. Of course, it is possible that all intermediate states fall into the same classification, in which case the policy is even simpler.

Now suppose production is more detrimental to the machine if it is carried out in a worse machine state. That is, the expected loss in terminal value of the machine is higher if production is carried out in a worse state. The following proposition establishes the monotonicity of critical numbers, $\tilde{x}(i)$, under this condition.

Proposition 5 *If* $\eta_i - \sum_{j=i}^{I} p_{ij}\eta_j$ *is increasing in* i *then*

(i) $\tilde{x}(i)$ *is decreasing in* i*, and*

(ii) for any inventory level, if production is optimal in a machine state, then it is also optimal in all better machine states.

It follows from Proposition 5, that for any inventory level, if doing nothing is optimal in a machine state, then for a worse machine state, production cannot be optimal.

Finally, the simplest policy results when repair is more rewarding *and* production is more detrimental, as machine state gets worse. Under these circumstances, both Propositions 4 and 5 apply. Note that both sequences of critical numbers, $\tilde{x}(i)$ and $x^*(i)$, are decreasing, but they

i	1	2	3	4	5	6	7	8	9	10
ρ_i	0.95	0.90	0.85	0.80	0.75	0.70	0.60	0.50	0.40	0.30
η_i	54	52	49	45	40	34	27	19	10	0

Table 1: Parameters for Example 1

may still cross each other once. As the state gets worse, the policy changes from "Produce" $\rightarrow$ "Do Nothing" $\rightarrow$ "Repair". Of course, one or more of these regions may be empty, so we may have "Produce" $\rightarrow$ "Repair", "Do Nothing" $\rightarrow$ "Repair", etc., depending upon the problem data and the inventory level.

We conclude this section by presenting two numerical examples.

Example 1. Suppose the machine has $(I-1)$ components, each of which can fail during production with probability p. Similarly, each failed components can be restored to "as good as new" with probability r during a repair. Machine state 1 corresponds to all components in the best condition while machine state I corresponds to all components in a failed state. In general, machine state i represents $(I-i)$ functional components and $(i-1)$ failed components. Since both the failure and repair of each component is an independent Bernoulli trial, the machine transition probabilities are given by a Binomial distribution. Let $B(n,p,x) = \binom{n}{x} p^x(1-p)^{n-x}$. Then the machine deterioration matrix, $\mathbf{P}$, is given by: $p_{II}=1$, $p_{ij}=0$ for $i>j$, and $p_{ij}=B(I-i,p,j-i)$ otherwise. Similarly, the machine repair matrix, $\mathbf{R}$, is given by: $r_{11}=1$, $r_{ij}=0$ for $i<j$, and $r_{ij}=B(i-1,r,j-1)$ otherwise. Here we have assumed that production only leads to deterioration of the machine state while repair only leads to an improvement of the machine state. That is, failed components do not return to good condition by themselves during production and good components do not fail during repair. We have assumed, for this example, that $I=10$, $p=0.6$ and $r=0.4$.

Assume that each unit of product is good with probability ρ_i if production is carried out in state i. Thus, the production yield is a Binomial random variable; that is, $\Pr[U_i=u] = \binom{q}{u} \rho_i^u(1-\rho_i)^{q-u}$. Parameter ρ_i decreases with i and so does η_i. The associated parameter values used in our example are shown in Table 1.

We assume that demand is 100 units and in each period a batch of 25 units can be processed. The costs of production and repair are \$12 and \$30, respectively. It is assumed that each unit of good product yields \$2 in revenue while excess good products are salvaged for \$0.50 per unit.

i	1	2	3	4	5	6	7	8	9	10
$\overline{F}(i)$	41.74	37.32	32.90	28.48	24.06	19.64	12.72	5.80	-1.12	-8.04
$\underline{F}(i)$	6.12	3.57	1.03	-1.52	-4.07	-6.61	-9.78	-12.95	-16.12	-19.29
$G(i)$	30	28.8	26.76	23.88	20.16	15.6	10.2	3.96	-3.12	-11.04
$x^*(i)$	∞	∞	∞	99	98	96	94	92	**0**	**0**
$\tilde{x}(i)$	**85**	**84**	**83**	**83**	**84**	**86**	**87**	**88**	91	95

Table 2: Optimal Policy for Example 1

We now compute the optimal policy for the last period. The values of $\overline{F}(i)$ and $\underline{F}(i)$ are shown in Table 2, where $\overline{F}(i)$ and $\underline{F}(i)$ are upper and lower bounds, respectively, on $F(i)$ for any value of x. (See Lemma 1 in Appendix B for more formal definitions.) From the signs of $\overline{F}(i)$ and $\underline{F}(i)$ one can conclude that machine states 1–3 are good, 4–8 are intermediate and 9–10 are bad. Since $G(i)$ is decreasing in i, the optimal policy will have the property that if it is optimal to repair in state i, then it is also optimal to repair in all machine states worse than i. The values of $x^*(i)$ and $\tilde{x}(i)$ are shown in the last two rows of the table. Note that all critical numbers lie between 75 and 100 (i.e., between $D - q$ and D).

The optimal policy for the good states is governed by critical numbers $\tilde{x}(i)$. For example, the optimal policy for machine state 1 is to do nothing if inventory is at or above 85 and to produce otherwise. The optimal policy for the bad states is to repair. For this example, it turns out that $\tilde{x}(i) < x^*(i)$ for all intermediate states and hence $\tilde{x}(i)$ are the deciding critical numbers for intermediate states. For example, the optimal policy for machine state 4 is to do nothing if inventory is at or above 83 and to produce otherwise.

The optimal policy structure for the last period is shown in Figure 2 where the black area represents the repair region, the dark grey represents the production region and the light grey area is the no-action region. Note that $\tilde{x}(i)$ decreases and then increases as the state index increases. The optimal policy, when on-hand inventory is 83 units, is to produce if machine is in states 1, 2, 5, 6, 7, 8, to do nothing if machine is in states 3 or 4 and to repair if machine is in states 9 or 10. It is quite counterintuitive that one would do nothing in one state, but produce in a worse machine state at the same level of on-hand inventory.

Example 2. The production and repair costs are modified to $w = 6, f = 10$; all other data remain the same as in Example 1. The results for the last period are shown in Table 3.

It is evident from the signs of $\overline{F}(i)$ and $\underline{F}(i)$ that machine states 1–6 are intermediate and 7–10 are bad. Because of the relatively low cost of repair, there is no good state in this case. As

i	1	2	3	4	5	6	7	8	9	10
$\overline{F}(i)$	27.74	23.32	18.9	14.48	10.06	5.64	-1.28	-8.2	-15.12	-22.04
$\underline{F}(i)$	-7.89	-10.43	-12.98	-15.52	-18.07	-20.61	-23.78	-26.95	-30.12	-33.29
$G(i)$	10	8.8	6.76	3.88	0.16	-4.4	-9.8	-16.04	-23.12	-31.04
$x^*(i)$	95	94	92	90	88	**87**	**0**	**0**	**0**	**0**
$\tilde{x}(i)$	**89**	**88**	**87**	**88**	**88**	90	91	93	96	99

Table 3: Optimal Policy for Example 2

always, the optimal policy for the bad states is to repair. For intermediate states, the optimal policy is governed by $\min\{x^*(i), \tilde{x}(i)\}$ as highlighted in the table. Machine state 5 is interesting since $x^*(i) = \tilde{x}(i) = 88$. According to part (iii.c) of Theorem 1 , the deciding critical number for machine state 5 is $\tilde{x}(5)$ since $G(5) = 0.16 > 0$. Consequently, the optimal policy for state 5 is to do nothing if inventory is at or above 88 and to produce otherwise. The optimal policy for machine state 6, on the other hand, is to repair if inventory is at or above 87 and to produce otherwise. The optimal policy structure for the last period is shown Figure 3.

5 Analysis of Multi-Period Problem

It turns out, unfortunately, that the characteristics of the optimal policies for the single-period problem discussed above do not extend to multi-period problems. In particular, the optimal policy is not necessarily characterized by a single critical number inventory level x for each i because the functions $P_k(i, x)$ and $R_k(i, x)$ may cross several times as x increases.

We characterize properties of the value function and show that set of the states (if any) for which the "do nothing" alternative is optimal expands as one moves farther from the end of the horizon.

Proposition 6 *$V_k(i, x)$ is non-increasing in i and non-decreasing in k and x.*

Theorem 2 below states that for any period, if it is optimal to do nothing in a state (i, x), then it is also optimal to do nothing in that state for all future periods. That is, the do-nothing policy region can only expand as we approach the due date (or equivalently, shrink as we move earlier in the horizon). Therefore, as we solve the dynamic program, for stage k, the do-nothing alternative needs to be considered only for those states (i, x) for which it is optimal to do nothing at stage $k - 1$. Thus, as one moves farther from the end of the horizon, the policy regions become simpler because the do-nothing alternative is eliminated for an increasing percentage of the states.

Theorem 2 *If* $N_k(i,x) > P_k(i,x)$ *and* $N_k(i,x) > R_k(i,x)$, *then* $N_{k-1}(i,x) > P_{k-1}(i,x)$ *and* $N_{k-1}(i,x) > R_{k-1}(i,x)$.

Because it is difficult to provide a further characterization of the optimal solution for the multi-period problem, the solution must be obtained by dynamic programming, which would be intractable for problems with many machine states or many time periods. To illustrate the structure of optimal solutions, we solved the examples in the previous section for $k = 2$, 3, 4, and 20.

The optimal policy structures in Example 2 for $k = 2, 3$ and 4 are shown in Figures 4-6, respectively. It is clear that a single critical number policy is no longer optimal. For example, for $k = 2$ and $i = 10$, the optimal decision changes from repair to production then to repair and again to production and finally to repair as inventory level increases. Note also that, as claimed in Theorem 2, the do-nothing region "shrinks" as one moves farther away from the due date. No such structure exists for the production or repair regions. The change in the optimal policy in Example 1 (not shown here) as one moves away from the due date demonstrates similar behavior.

6 Heuristic Policy

From the analysis in the previous sections, it is clear that a simple characterization of the optimal policy is difficult to obtain. For this reason, we decided to explore implementable heuristic policies that take advantage of our understanding of and observations about the structure of optimal solutions.

We implemented the dynamic programming procedure described earlier and solved problem instances for a variety of problem parameters and time horizons. We observed the following patterns in the optimal solutions:

a) For periods early in the horizon when demand is far from being satisfied, in all of our test problems, the optimal policy has the form of a simple threshold policy in which one produces if the machine state is i^* or better, and repairs if the machine state is worse than i^*, where value of i^* does not depend upon the level of accumulated production, and is typically identical to the threshold level for the infinite-horizon policy.

b) When only a few units of demand remain to be satisfied (e.g., less than the single-period production capacity), the option of doing nothing (in the current and in all future periods) may become a viable alternative. Thus, for any state of the machine, the key decision is whether to stop or continue. The optimal solutions for our test problems are similar to those that would be made using a myopic policy.

c) When the remaining demand is one up to a few periods' worth of production, the optimal policy usually remains a threshold policy with a structure as described above, but the threshold may oscillate between two adjacent machine states (depending upon the number of periods remaining in the horizon) because of end-of-horizon effects. However, numerical results suggest that the penalty for choosing one of these two machine states and implementing a constant threshold is very small. Moreover, in all of our test problems, one of these two machine states is the same as the early-in-the-horizon optimal machine state threshold.

These observations led us to the following heuristic procedure:

1. While there are at least two periods remaining in the horizon,

   1a. **Threshold Machine State** Implement a simple, machine-state threshold policy with threshold state i^* (derived from the optimal infinite horizon policy or an approximation thereof): repair the machine if the machine state is i^* or higher.

   1b. **Myopic Production Control** For each machine state i less than i^*, produce until the accumulated inventory reaches the threshold $\tilde{x}_1(i)$, where $\tilde{x}_1(i)$ is the smallest value of x for which $P_1(i,x) - N_1(i,x)$ is negative.

2. In the last period in the horizon, use the optimal single-period policy.

We now describe how to obtain i^*. Recall that we are interested in periods early in the horizon when demand is far from being satisfied. Consequently, we assume that all good units can be used to satisfy demand, and we focus on the expected long-run profit. Suppose we follow a threshold repair policy with a candidate threshold machine state i^0. Then for any machine state i less than i^0, production occurs, and the resulting change in the long-run profit function is

$$\sum_{j=i}^{I} p_{ij}\eta_j - \eta_i + \pi\mu_i - w. \tag{8}$$

For any machine state i greater than i^0, repair occurs, and the resulting change in the long-run profit function is

$$\sum_{j=1}^{i} r_{ij}\eta_j - \eta_i - f. \tag{9}$$

Letting the machine state transition matrix consists of rows 1 through $i^0 - 1$ of P and rows i^0 through i of R, we can then compute the expected long-run profit for the policy with the threshold i^0. The threshold i^* is one that achieves the highest expected long-run profit among all i^0. (Observe that if the true values of η_i were known in the context of producing the product at hand, then the value function changes in (8) and (9) would be exact. Recognizing that the functions are approximations, what we presented here is an approximate steady-state analysis.)

Parameter	Low Value	Base Value	High Value
Production cost per period	10	12	15
Repair cost per period	20	30	60
P(component failure)	0.2	0.4	0.6
P(successful component repair)	0.4	0.6	0.8
P(good unit) by machine state	{.90 .89 .82 .75 .68 .61 .54 .47 .41 .32}	{.95 .92 .86 .80 .74 .68 .62 .56 .51 .44}	{1.0 .95 .90 .85 .80 .75 .70 .65 .60 .55}
Machine state value		{54 52 49 45 40 34 27 19 10 0}	{62 58 51 44 37 30 23 16 9 0}

Table 4: Parameters for Test Problems

6.1 Tests of the Heuristic Procedure

Some problem parameters were held fixed for all problem instances, whereas others were varied. The production capacity in each period is 25 and the total demand is 100. The revenue per unit is 2.0 and the salvage value per unit is 0.5. The machine is modeled as a collection of 10 components that are either in an operational or failed state. The components fail independently, and the state of the machine is defined by the number of failed components. If repair takes place, repair is attempted on all failed components with independent success rates. We solved test problems comprising all 486 combinations of problem parameters shown in Table 4. For each combination of parameters, we solved the dynamic program and implemented the heuristic policy for 5- and 10-period horizons *starting in each possible machine state*. In general, during a 5-period horizon, it is difficult to satisfy the full customer order, whereas during a 10-period horizon, there is a high probability of satisfying the order.

In practice, the manufacturing facility would be "solving" a series of finite horizon problems, so the distribution of machine states at the beginning of each customer order would be difficult to ascertain. As an approximation, we assume that the distribution of initial machine states is the same as the distribution of machine states that arises from the best threshold policy, and compare the optimal and heuristic solutions using the weights implied by this distribution.

The heuristic produced solutions within 2% of optimality for 98% of the 5-period problems and 90% of the 10-period problems. For a 5% optimality tolerance, the corresponding percentages are nearly 100% and 96%, respectively. The performance of the heuristic was adversely affected by factors that make it easier to satisfy the full order, because in these cases, the heuristic policy, with only a limited number of control parameters, cannot be finely tuned enough to mimic the optimal policy, particularly near the completion of the order. The performance of the heuristic also deteriorates slightly as more variability is introduced into the system, for example, due to lower component repair probabilities or a wider range of yields for different machine states.

Nevertheless, the heuristic performed well across a wide range of conditions.

7 Conclusions

We have studied the tradeoffs that arise in a finite horizon model when repair takes time but allows the machine to provide a higher output rate, on average. To our knowledge, this is the first study to examine this combination of realistic factors. We show that the typical response of continuing to produce as the due date nears, despite the deteriorating condition of the machine, is often suboptimal. We also show that in the single-period problem, there is categorization of machine states. In certain machine states, it is optimal to take time to repair. For some combinations of machine states and inventory levels, it may be optimal to do nothing. We show that for each machine state, the optimal decision is characterized by a single critical inventory level.

The general findings remain qualitatively the same in the multi-period problem. That is, it is not always optimal to produce even if the deadline is approaching. One should take into account other important factors such as the machine condition and the inventory level. Unfortunately, unlike in the single-period problem, a simple characterization of the optimal policy structure does not exist for the multi-period problem. Using numerical examples, we show that for a given machine state, it may be optimal to produce at inventory level x, repair at inventory level $x + 1$, and produce again at inventory level $x + 2$. These results arise because of very complex interactions among the terminal values of the various machine states, the machine deterioration transition matrix, and the machine improvement transition matrix for repair. Thus, finding the optimal solution may be computationally challenging. However, these insights and observations of patterns from numerical results were instrumental in allowing us to devise an implementable heuristic.

The heuristic performs very well and is much easier to implement than the optimal policy, because it requires only $I+1$ parameters: an inventory threshold at which production should stop for each machine state, and a machine state threshold that specifies whether to produce or repair until the stopping rule comes into effect. We observed in our computational study that ignoring the end-of-horizon effects may have a large adverse effect on profits. The proposed heuristic policy is able to regain nearly all of the potential loss.

Further research is needed to consider such factors as partial observability of the machine state (perhaps through inspection of the output), multiple customer orders with different due dates, and different modes of repair.

Acknowledgment: This research has been supported in part by National Science Foundation grant GER/HRD 93-96288 to the University of California, Berkeley.

A Relationship of Assumption A.8 to Assumptions A.6 and A.7

If matrix $(\mathbf{P}-\mathbf{R})$ is IFR, then assumption A.8 is satisfied for all decreasing η_i. But assumption A.8 may hold even if $(\mathbf{P}-\mathbf{R})$ is not IFR, particularly if η_i decreases rapidly with i. Note that an IFR $\mathbf{P}$ and an IFR $\mathbf{R}$ (as proposed in A.6 and A.7, respectively) are neither necessary nor sufficient for $(\mathbf{P}-\mathbf{R})$ to be IFR. To see this, consider

$$\mathbf{P} = \begin{pmatrix} 0.4 & 0.1 & 0.2 & 0.3 \\ & 0.3 & 0.3 & 0.4 \\ & & 0.5 & 0.5 \\ & & & 1 \end{pmatrix}, \qquad \mathbf{R1} = \begin{pmatrix} 1 & & & \\ 0.5 & 0.5 & & \\ 0.4 & 0.3 & 0.3 & \\ 0.3 & 0.2 & 0.1 & 0.4 \end{pmatrix},$$

$$\mathbf{R2} = \begin{pmatrix} 1 & & & \\ 0.6 & 0.4 & & \\ 0.65 & 0.1 & 0.25 & \\ 0.7 & 0.1 & 0.15 & 0.05 \end{pmatrix}, \quad \mathbf{R3} = \begin{pmatrix} 1 & & & \\ 0.6 & 0.4 & & \\ 0.6 & 0.1 & 0.3 & \\ 0.6 & 0.1 & 0.2 & 0.1 \end{pmatrix}.$$

It is easy to verify that $\mathbf{P}$ and $\mathbf{R1}$ are IFR but $(\mathbf{P}-\mathbf{R1})$ is not IFR. On the other hand, $(\mathbf{P}-\mathbf{R2})$ is IFR even though $\mathbf{R2}$ is not IFR. Finally, $\mathbf{P}$, $\mathbf{R3}$, and $(\mathbf{P}-\mathbf{R3})$ are all IFR.

When $\mathbf{P}$ is IFR, the probability that the machine will end up in a set of worst states increases if the production is carried out in a worse state. Similarly, when $\mathbf{R}$ is IFR, the probability that the machine will end up in a set of best states increases if the repair is carried out in a better state. In case of $\mathbf{R2}$, as the machine deteriorates from state 2 to state 4, its chances of recovery to state 1 after repair increases from 60% to 70%. As a result, the matrix $\mathbf{R2}$ cannot be IFR.

When $(\mathbf{P}-\mathbf{R})$ is IFR, the *difference between* the probability of deterioration due to production and the probability of improvement due to repair increases as the state gets worse. Note that if $(\mathbf{P}-\mathbf{R})$ is IFR, then

$$\begin{aligned} & \sum_{j=k}^{I} (p_{ij} - r_{ij}) && \text{is nondecreasing in } i \text{ for all } k \\ \textit{or,} \quad & \sum_{j=k}^{I} p_{ij} - \left(1 - \sum_{j=1}^{k-1} r_{ij}\right) && \text{is nondecreasing in } i \text{ for all } k \\ \textit{or,} \quad & \sum_{j=k}^{I} p_{ij} + \sum_{j=1}^{k-1} r_{ij} && \text{is nondecreasing in } i \text{ for all } k. \end{aligned}$$

That is, the probability that the machine deteriorates to a set of machine states (say, k to I) in case of production, plus the probability that the machine improves to the complimentary set of machine states (1 to $k-1$) in case of repair, increases if one starts in a worse state. In this sense, the matrix $(\mathbf{P}-\mathbf{R})$ is IFR if the *composite rate* of deterioration and improvement increases as the state gets worse.

B Proofs

In this appendix we first establish a lemma, which contains some properties useful in several proofs.

Lemma 1 *The following properties are true for functions F and G:*

(P.1) For $x \leq (D-q)$, $F(i,x)$ is invariant in x, i.e., $F(i,x) = \overline{F}(i)$, where

$$\overline{F}(i) = \sum_{j=i}^{I} p_{ij}\eta_j - \sum_{j=1}^{i} r_{ij}\eta_j + \pi\mu_i + f - w. \tag{10}$$

(P.2) For $x \geq D$, $F(i,x)$ is invariant in x, i.e., $F(i,x) = \underline{F}(i)$, where

$$\underline{F}(i) = \sum_{j=i}^{I} p_{ij}\eta_j - \sum_{j=1}^{i} r_{ij}\eta_j + \delta\mu_i + f - w. \tag{11}$$

(P.3) $F(i,x)$ is strictly decreasing in x for $(D-q) \leq x \leq D$.

(P.4) $\overline{F}(i) \geq F(i,x) \geq \underline{F}(i)$ $\forall x$.

(P.5) $F(i,x)$ is decreasing in i for each x.

(P.6) $\overline{F}(i) - \underline{F}(i)$, is positive and decreasing in i, $\forall i$.

(P.7) For each i, $G(i,x)$ is independent of x, i.e., $G(i,x) = G(i)$.

(P.8) $\underline{F}(i) < G(i) < \overline{F}(i)$, $\forall i$

Proof: (P.1) and (P.2) follows from (3) and the fact that

$$E\left\{(x+U_i-D)^+ - (x-D)^+\right\} = \begin{cases} 0 & \text{if } x \leq D-q, \\ \mu_i & \text{if } x \geq D. \end{cases}$$

To prove (P.3) note that, for $x \leq D$, $E\left\{(x+U_i-D)^+ - (x-D)^+\right\} = E\left\{(x+U_i-D)^+\right\}$, which is strictly increasing in x for $x \geq (D-q)$. Since $\pi > \delta$, the first term in (3) is strictly decreasing in x and hence $F(i,x)$ is strictly decreasing in x in the interval $(D-q) \leq x \leq D$. (P.4) follows from (P.3) and the definition of $\overline{F}(i)$ and $\underline{F}(i)$.

To prove (P.5), note that

$$E\left\{(x+U_i+D)^+\right\} = E\left\{(x+U_i) - \min(x+U_i, D)\right\} = x + \mu_i - E\left\{\min(x+U_i, D)\right\},$$

and $(x - D)^+ = x - \min(x, D)$. Hence (3) can be rewritten as

$$\begin{aligned} F(i,x) &= (\pi - \delta)\, E\left\{\min(x + U_i, D)\right\} + \delta\mu_i \\ &\quad + \sum_{j=i}^{I} p_{ij}\eta_j - \sum_{j=1}^{i} r_{ij}\eta_j - (\pi - \delta)\min(x, D) - w + f. \end{aligned} \tag{12}$$

Note that the first two terms on the r.h.s. are positive and decreasing in i due to assumptions A.1 and A.2. The next two terms are decreasing in i due to assumption A.8. Hence $F(i,x)$ is decreasing in i.

To prove (P.6), we can subtract (11) from (10) to obtain

$$\overline{F}(i) - \underline{F}(i) = (\pi - \delta)\mu_i\,,$$

which is positive and decreasing in i.

(P.7) follows directly from (4).

To prove (P.8), note from assumption A.5 that

$$\begin{aligned} \delta\mu_i - w + \sum_{j=i}^{I} p_{ij}\eta_j &< \eta_i \\ \text{or, } \delta\mu_i - w + \sum_{j=i}^{I} p_{ij}\eta_j - \sum_{j=1}^{i} r_{ij}\eta_j + f &< \eta_i - \sum_{j=1}^{i} r_{ij}\eta_j + f \\ \text{or, } \underline{F}(i) &< G(i). \end{aligned}$$

Likewise, from Assumption A.4

$$\begin{aligned} \pi\mu_i + \sum_{j=i}^{I} p_{ij}\eta_j - w &> \eta_i \\ \text{or, } \pi\mu_i + \sum_{j=i}^{I} p_{ij}\eta_j - w - \sum_{j=1}^{i} r_{ij}\eta_j + f &> \eta_i - \sum_{j=1}^{i} r_{ij}\eta_j + f \\ \text{or, } \overline{F}(i) &> G(i). \ \square \end{aligned}$$

Proof of Proposition 1: To show (i), note that if i is a good state then $\underline{F}(i) \geq 0$. But, according to (P.5), $F(i,x)$ is decreasing in i. That is, $F(i-1,x) > F(i,x)$, $\forall x$. Applying this fact for $x \geq D$ means that $\underline{F}(i-1) > \underline{F}(i)$. Hence $\underline{F}(i) \geq 0$ implies that $\underline{F}(i-1) \geq 0$, i.e., $i-1$ is also a good state. Similarly, to prove (ii), one can use the fact that $F(i,x)$ is decreasing in i to show that $\overline{F}(i+1) < \overline{F}(i)$. Then $\overline{F}(i) \leq 0$ implies that $\overline{F}(i+1) \leq 0$. Finally, the mutually exclusive nature of the sets follows from the definition. To prove the compactness property in (iii), we use (i) and (ii). Suppose the set of good states are not compact. Then between two good states, say i, and k, $i < k$, lies a state j, $i < j < k$, that is not good. But if k is a good state, then according to (ii), so are $k-1, k-2, ..., 1$. This rules out the possibility of a non-good state $j < k$. Hence, the set of good states is compact. Similarly, (ii) can be used to argue that the set of bad states is also compact. Together, these results imply that the set of intermediate states must also be compact. □

Proof of Proposition 2: Note from (P.3) that $F(i, x)$ is strictly decreasing in x and from the definition of an intermediate state that $F(i, D-q) = \overline{F}(i) > 0 > \underline{F}(i) = F(i, D)$. Hence $F(i, x)$ crosses zero exactly once in the range $D - q < x < D$. This proves (i) and establishes the uniqueness of $x^*(i)$. To prove (ii), note from (P.5) that $F(i, x)$ is decreasing in i. Hence, $F(i+1, x^*(i)) \leq F(i, x^*(i)) < 0$. But $x^*(i+1)$ is the smallest x such that $F(i+1, x) < 0$, so $x^*(i+1) \leq x^*(i)$. □

Proof of Proposition 3: Note from (P.8) that $F(i, D-q) = \overline{F}(i) > G(i) > \underline{F}(i) = F(i, D)$. Hence from (P.3) and (P.7), $F(i, x)$ must cross $G(i)$ exactly once in the range $D - q < x < D$. □

Proof of Theorem 1: Note that the optimal decision in state (i, x) is to

Produce	if $F(i, x) > 0$ and $F(i, x) > G(i)$,	(the sign of G does not matter),
Repair	if $F(i, x) < 0$ and $G(i) < 0$,	(the relative magnitudes of F and G are immaterial), and
Do nothing	if $G(i) > 0$ and $F(i, x) < G(i)$,	(the sign of F does not matter).

If i is a good state then $\underline{F}(i) \geq 0$ and (P.4) implies that $F(i, x) > 0$, $\forall x$. That is, $P_1(i, x) > R_1(i, x)$ $\forall x$, and it is never optimal to repair in machine state i. Moreover, if $x \geq \tilde{x}(i)$, then $F(i, x) < G(i)$ and doing nothing is optimal; otherwise production is optimal. The existence of a unique $\tilde{x}(i)$ is guaranteed for each machine state by Proposition 3.

If i is a bad state then $\overline{F}(i) \leq 0$ and (P.4) implies that $F(i, x) < 0$, $\forall x$. Moreover, according to (P.8), $G(i) < \overline{F}(i)$ $\forall i$, hence $G(i) < 0$ for a bad state. Together these imply that repair is optimal in a bad state.

For any intermediate state i, the existence of $x^*(i)$ and $\tilde{x}(i)$ is guaranteed by propositions 2 and 3, respectively. Consider first the case (iii.a), i.e., when $\tilde{x}(i) < x^*(i)$. As x increases, $F(i, x)$ crosses $G(i)$ before it crosses zero. This implies that $G(i) > 0$ and repair is ruled out as an optimal decision (doing nothing is better than repair in this case). For $x < \tilde{x}(i)$, it must be that $F(i, x) > G(i) > 0$ and it is optimal to produce. Similarly, for $x \geq \tilde{x}(i)$, it must be that $F(i, x) < G(i)$ and it is optimal to do nothing.

Next consider the case (iii.b), i.e., when $\tilde{x}(i) > x^*(i)$. As x increases, $F(i, x)$ crosses zero before it crosses $G(i)$. This implies that $G(i) < 0$ and the optimal decision cannot be to do nothing (repair is better than doing nothing). For $x < x^*(i)$, it must be that $F(i, x) > 0 > G(i)$

and it is optimal to produce. Similarly, for $x \geq x^*(i)$, it must be that $F(i,x) < 0$ and it is optimal to repair.

Finally consider the case of $\tilde{x}(i) = x^*(i)$ (case iii.c). As x increases, $F(i,x)$ crosses $G(i)$ and zero at the same (discrete) value of x. No information about the sign of $G(i)$ can be inferred from the two critical numbers. We must use the numerical value of $G(i)$ to distinguish between the two possible cases. For $x < \tilde{x}(i) = x^*(i)$, it must be that $F(i,x) > G(i)$ and $F(i,x) > 0$ and is optimal to produce. For $x \geq \tilde{x}(i) = x^*(i)$, the optimal policy depends upon the sign of $G(i)$. If $G(i) > 0$, then following the argument for case (iii.a), it is optimal to do nothing. Similarly, if $G(i) < 0$, then following the argument for case (iii.b), it is optimal to repair. Finally, if $G(i) = 0$ one is indifferent between repair and doing nothing. □

Proof of Proposition 4: First note, from (4), that if $\sum_{j=1}^{i} r_{ij}\eta_j - \eta_i$ is increasing in i, then $G(i)$ is decreasing in i. Also, from (P.5), $F(i,x)$ is decreasing in i.

To prove (i), recall that $F(i,x)$ is decreasing in x. This implies that (a) if $G(i) > 0$ then $\tilde{x}(i) \leq x^*(i)$, (b) if $G(i) = 0$ then $\tilde{x}(i) = x^*(i)$, and (c) if $G(i) < 0$ then $\tilde{x}(i) \geq x^*(i)$. Since $G(i)$ is decreasing in i, it can cross zero at most once. Hence $\tilde{x}(i)$ crosses $x^*(i)$ at most once and it does so from below.

To prove (ii), suppose that for an inventory level x, repair is optimal in machine state i. Then $R_1(i,x) > P_1(i,x)$ and $R_1(i,x) > N_1(i,x)$, i.e. both $F(i,x)$ and $G(i)$ negative. Since $F(i,x)$ and $G(i)$ are decreasing in i, it follows that both $F(i+1,x)$ and $G(i+1)$ are also negative. Hence repair is optimal in state $(i+1,x)$. □

Proof of Proposition 5: Let $H(i,x) = F(i,x) - G(i,x) = P_1(i,x) - N_1(i,x)$, which represents the relative value of production compared to inaction. Using (12), one can rewrite $H(i,x)$ as

$$\begin{aligned} H(i,x) &= (\pi - \delta) E\left\{\min\left(x + U_i, D\right)\right\} + \delta\mu_i \\ &\quad - \left(\eta_i - \sum_{j=i}^{I} p_{ij}\eta_j\right) - (\pi - \delta)\min\left(x, D\right) - w. \end{aligned}$$

The first two terms on the r.h.s. are positive and decreasing in i due to assumptions A.1 and A.2. The third term is decreasing in i due to the condition postulated in proposition 5. Hence $H(i,x)$ is decreasing in i. Also, from (P.3) and (P.7), it follows that $H(i,x)$ is strictly decreasing in x for $(D-q) \leq x \leq D$.

To prove (i), note that $\tilde{x}(i)$ is the smallest value of x such that $H(i,x) < 0$. Since $H(i,x)$ is decreasing in i, $H(i+1, \tilde{x}(i)) \leq H(i, \tilde{x}(i)) < 0$. But $\tilde{x}(i+1)$ is the smallest x such that $H(i+1,x) < 0$, so $\tilde{x}(i+1) \leq \tilde{x}(i)$.

To prove (ii), suppose that for inventory level x, production is optimal in machine state i. Then $P_1(i,x) > R_1(i,x)$ and $P_1(i,x) > N_1(i,x)$, i.e. both $F(i,x)$ and $H(i,x)$ are positive. Since $F(i,x)$ and $H(i,x)$ are decreasing in i, it follows that both $F(i-1,x)$ and $H(i-1,x)$ are also positive. Hence production is optimal in state $(i-1,x)$. $\square$

Proof of Proposition 6: Since

$$\begin{aligned} V_k(i,x) &= \max\{P_k(i,x), R_k(i,x), N_k(i,x)\} \\ &= \max\{P_k(i,x), R_k(i,x), V_{k-1}(i,x)\} \\ &\geq V_{k-1}(i,x), \end{aligned}$$

$V_k(i,x)$ is non-decreasing in k.

To prove that $V_k(i,x)$ is non-decreasing in x, first note that $V_0(i,x)$ is increasing in x. Suppose $V_{k-1}(i,x)$ is non-decreasing in x. Then, by definition, $N_k(i,x)$ is non-decreasing in x. Since $p_{ij} \geq 0$ and $r_{ij} \geq 0$, $P_k(i,x)$ and $R_k(i,x)$ are also non-decreasing in x. The maximum of three non-decreasing functions, $\max\{P_k(i,x), R_k(i,x), N_k(i,x)\}$ is also non-decreasing in x.

To show that $V_k(i,x)$ is non-increasing in i, first note that $V_0(i,x)$ is decreasing in i. Now suppose $V_{k-1}(i,x)$ is non-increasing in i. Then $N_k(i,x)$ is also non-increasing in i by definition. Since matrix $\mathbf{R}$ is IFR and $V_{k-1}(i,x)$ is non-increasing in i, $\sum_{j=1}^{i} r_{ij} V_{k-1}(j,x)$ is also non-increasing in i and hence $R_k(i,x)$ is also non-increasing in i. Finally, to show that $P_k(i,x)$ is also non-increasing in i, rewrite

$$\begin{aligned} P_k(i,x) &= \sum_{j=i}^{I} p_{ij} \sum_{u=0}^{q} h_i(u) V_{k-1}(j,x+u) - w \\ &= \sum_{u=0}^{q} h_i(u) J_{k-1}(i,x+u) - w \end{aligned}$$

where $J_{k-1}(i,x) = \sum_{j=i}^{I} p_{ij} V_{k-1}(j,x)$. Note that $J_{k-1}(i,x)$ is non-decreasing in x since $V_{k-1}(i,x)$ is non-decreasing in x as shown above. This, together with the fact that $U_i \geq_{st} U_{i+1}$, $\forall i$, implies that

$$\begin{aligned} J_{k-1}(i,x+U_i) &\geq_{st} J_{k-1}(i,x+U_{i+1}), \\ \text{or, } E[J_{k-1}(i,x+U_i)] &\geq E[J_{k-1}(i,x+U_{i+1})] \\ \text{or, } \sum_{u=0}^{q} h_i(u) J_{k-1}(i,x+u) &\geq \sum_{u=0}^{q} h_{i+1}(u) J_{k-1}(i,x+u). \end{aligned}$$

Hence

$$P_k(i,x) = \sum_{u=0}^{q} h_i(u) J_{k-1}(i,x+u) - w$$

$$
\begin{aligned}
&\geq \sum_{u=0}^{q} h_{i+1}(u) J_{k-1}(i, x+u) - w \\
&\geq \sum_{u=0}^{q} h_{i+1}(u) J_{k-1}(i+1, x+u) - w \\
&= P_k(i+1, x)
\end{aligned}
$$

The last inequality follows from the fact that $J_{k-1}(i,x)$ is non-increasing in i because matrix $\mathbf{P}$ is IFR and $V_{k-1}(i,x)$ is non-increasing in i. Now, since $P_k(i,x)$, $R_k(i,x)$ and $N_k(i,x)$ are all non-increasing in i, their maximum, $V_k(i,x)$, is also non-increasing in i. $\square$

Proof of Theorem 2: If $N_k(i,x) > P_k(i,x)$ then, substituting their respective definitions, we have that

$$
\begin{aligned}
N_k(i,x) &= V_{k-1}(i,x) > \sum_{j=i}^{I} p_{ij} E\left[V_{k-1}(j, x+U_i)\right] - w \\
&= \sum_{j=i}^{I} p_{ij} \sum_{u=0}^{q} h_i(u) V_{k-1}(j, x+u) - w \\
&= \sum_{j=i}^{I} p_{ij} \sum_{u=0}^{q} h_i(u) \max\{P_{k-1}(j,x+u), R_{k-1}(j,x+u), N_{k-1}(j,x+u)\} - w
\end{aligned}
$$

Likewise, if $N_k(i,x) > R_k(i,x)$, then

$$
\begin{aligned}
N_k(i,x) &= V_{k-1}(i,x) > \sum_{j=1}^{i} r_{ij} V_{k-1}(j,x) - f \\
&= \sum_{j=1}^{i} r_{ij} \max\{P_{k-1}(j,x), R_{k-1}(j,x), N_{k-1}(j,x)\} - f
\end{aligned}
$$

But $N_{k-1}(j,x) = V_{k-2}(j,x)$ and $V_{k-1}(i,x) = \max\{P_{k-1}(i,x), R_{k-1}(i,x), V_{k-2}(i,x)\}$. So, by making appropriate substitutions, we have that

$$
\begin{aligned}
\max &\quad \{P_{k-1}(i,x), R_{k-1}(i,x), V_{k-2}(i,x)\} \\
> &\quad \sum_{j=i}^{I} p_{ij} \sum_{u=0}^{q} h_i(u) \max\{P_{k-1}(j,x+u), R_{k-1}(j,x+u), V_{k-2}(j,x+u)\} - w \\
\geq &\quad \sum_{j=i}^{I} p_{ij} \sum_{u=0}^{q} h_i(u) V_{k-2}(j,x+u) - w = \sum_{j=i}^{I} p_{ij} E\left[V_{k-2}(j,x+U_i)\right] - w = P_{k-1}(i,x)
\end{aligned}
$$

and

$$
\begin{aligned}
\max\{P_{k-1}(i,x), R_{k-1}(i,x), V_{k-2}(i,x)\} &> \sum_{j=1}^{i} r_{ij} \max\{P_{k-1}(j,x), R_{k-1}(j,x), V_{k-2}(j,x)\} - f \\
&\geq \sum_{j=1}^{i} r_{ij} V_{k-2}(j,x) - f = R_{k-1}(i,x)
\end{aligned}
$$

These inequalities together imply that $N_{k-1}(i,x) = V_{k-2}(i,x) > P_{k-1}$ and $N_{k-1}(i,x) = V_{k-2}(i,x) > R_{k-1}$. Thus, the result holds. $\square$

References

[1] Bobos, A.G. and E.N. Protonotarios. Optimal Systems for Equipment Maintenance and Replacement under Markovian Deterioration. *European Journal of Operational Research* **2** (1978), 257-264.

[2] Chitke, S.D. and S.D. Deshmukh. Preventive Maintenance and Replacement under Additive Damage. *Naval Research Logistics Quarterly* **28** (1981), 33-46.

[3] Douer, N. and U. Yechiali. Optimal Repair and Replacement in Markovian Systems. *Communications in Statistics–Stochastic Models* **10** (1994), 253-270.

[4] Iravani, S. and I. Duenyas. Integrated Maintenance and Production Control of a Deteriorating Production System. *IIE Transactions* **34** (2002), 423-435.

[5] Kijima, M. Some Results for Repairable Systems with General Repair. *Journal of Applied Probability* **26** (1989), 89-102.

[6] McCall, J.J. Maintenance Policies for Stochastically Failing Equipment: A Survey. *Management Science* **10** (1963), 85-97.

[7] Pierskalla, W.P and J.A. Voelker. A Survey of Maintenance Models: The Control and Surveillance of Deteriorating Systems. *Naval Research Logistics Quarterly* **23** (1976), 353-388.

[8] Rosenfield, D. Markovian Deterioration with Uncertain Information. *Operations Research* **24** (1976), 141-155.

[9] Ross, S.M. Quality Control under Markovian Deterioration. *Management Science* **17** (1971), 587-596.

[10] Sherif, Y.S. and M.L. Smith. Optimal Maintenance Models for Systems Subject to Failure–A Review. *Naval Research Logistics Quarterly* **28** (1981), 47-74.

[11] Sloan, T. and J.G. Shanthikumar. Combined Production and Maintenance Scheduling for a Multiple-Product, Single-Machine Production System. *Production and Operations Management* **9** (2000), 379-399.

[12] Smallwood, R.D. and E.J. Sondik. The Optimal Control of Partially Observable Markov Processes over a Finite Horizon. *Operations Research* **21** (1973), 1071-1088.

[13] Stadje, W. and D. Zuckerman. Optimal Maintenance Strategies for Repairable Systems with General Degree of Repair. *Journal of Applied Probability* **28** (1991), 384-396.

[14] Valdez-Flores, C. and R.M. Feldman. A Survey of Preventive Maintenance Models for Stochastically Deteriorating Single-Unit Systems. *Naval Research Logistics* **36** (1989), 419-446.

[15] Venkatesan, M. Production-Inventory with Equipment Replacement–PIER. *Operations Research* **32** (1984), 1286-1295.

[16] Wang, H. A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research* **139** (2002), 469-489.

[17] White, C.C., III. Optimal Inspection and Repair of a Production Process Subject to Deterioration. *Journal of the Operational Research Society* **29** (1978), 235-243.

[18] Zuckerman, D. Optimal Maintenance and Replacement of Equipment: A Diffusion Approximation. *Naval Research Logistics Quarterly* **33** (1986), 469-477.

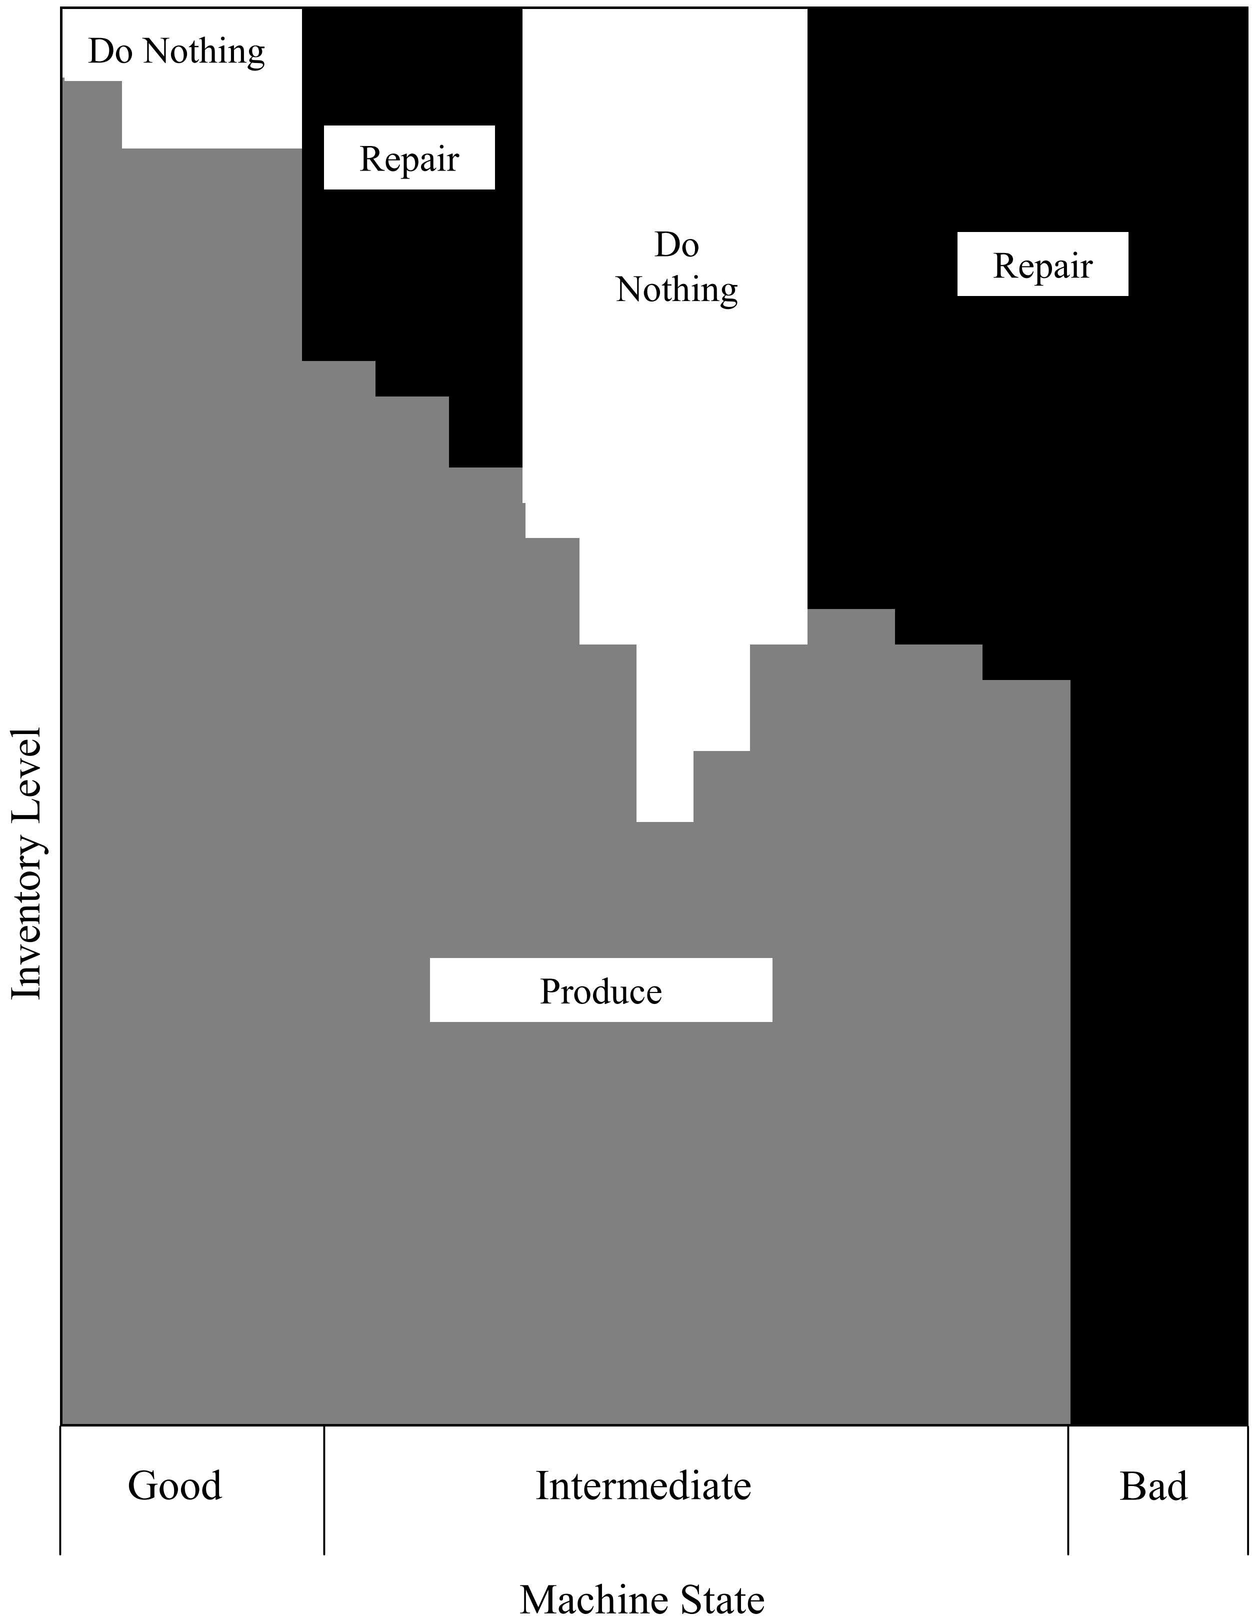

Figure 1: A hypothetical policy for the single period problem

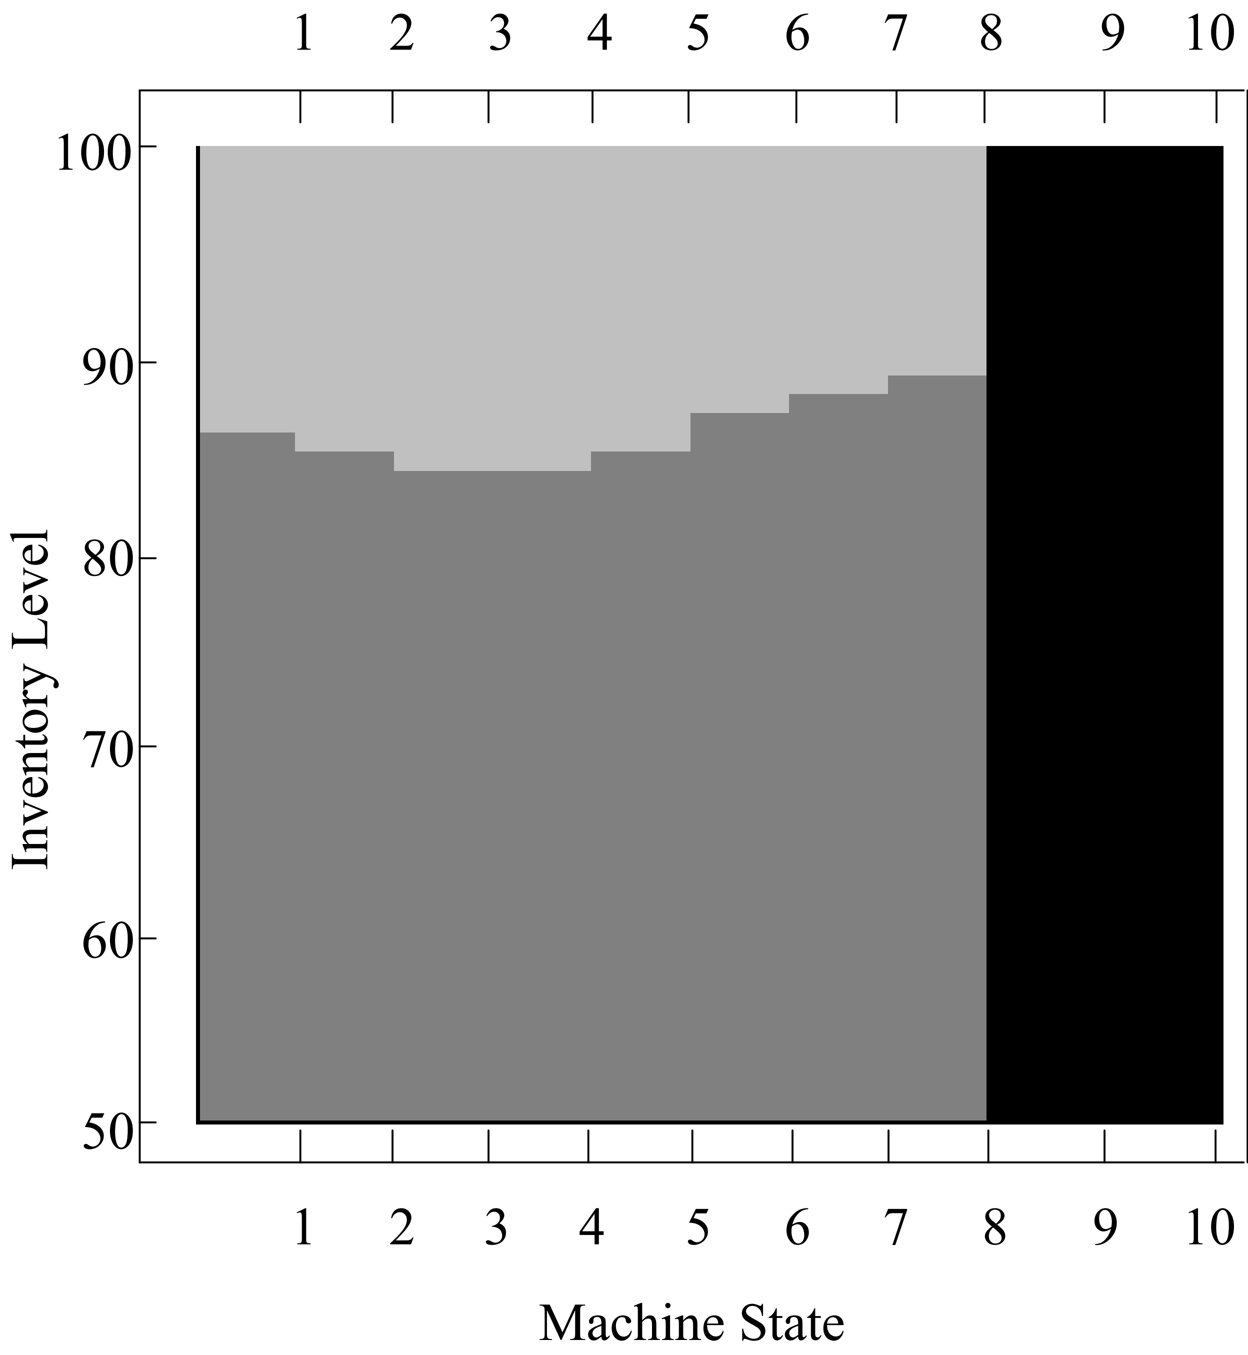

Figure 2: Optimal policy for the last period

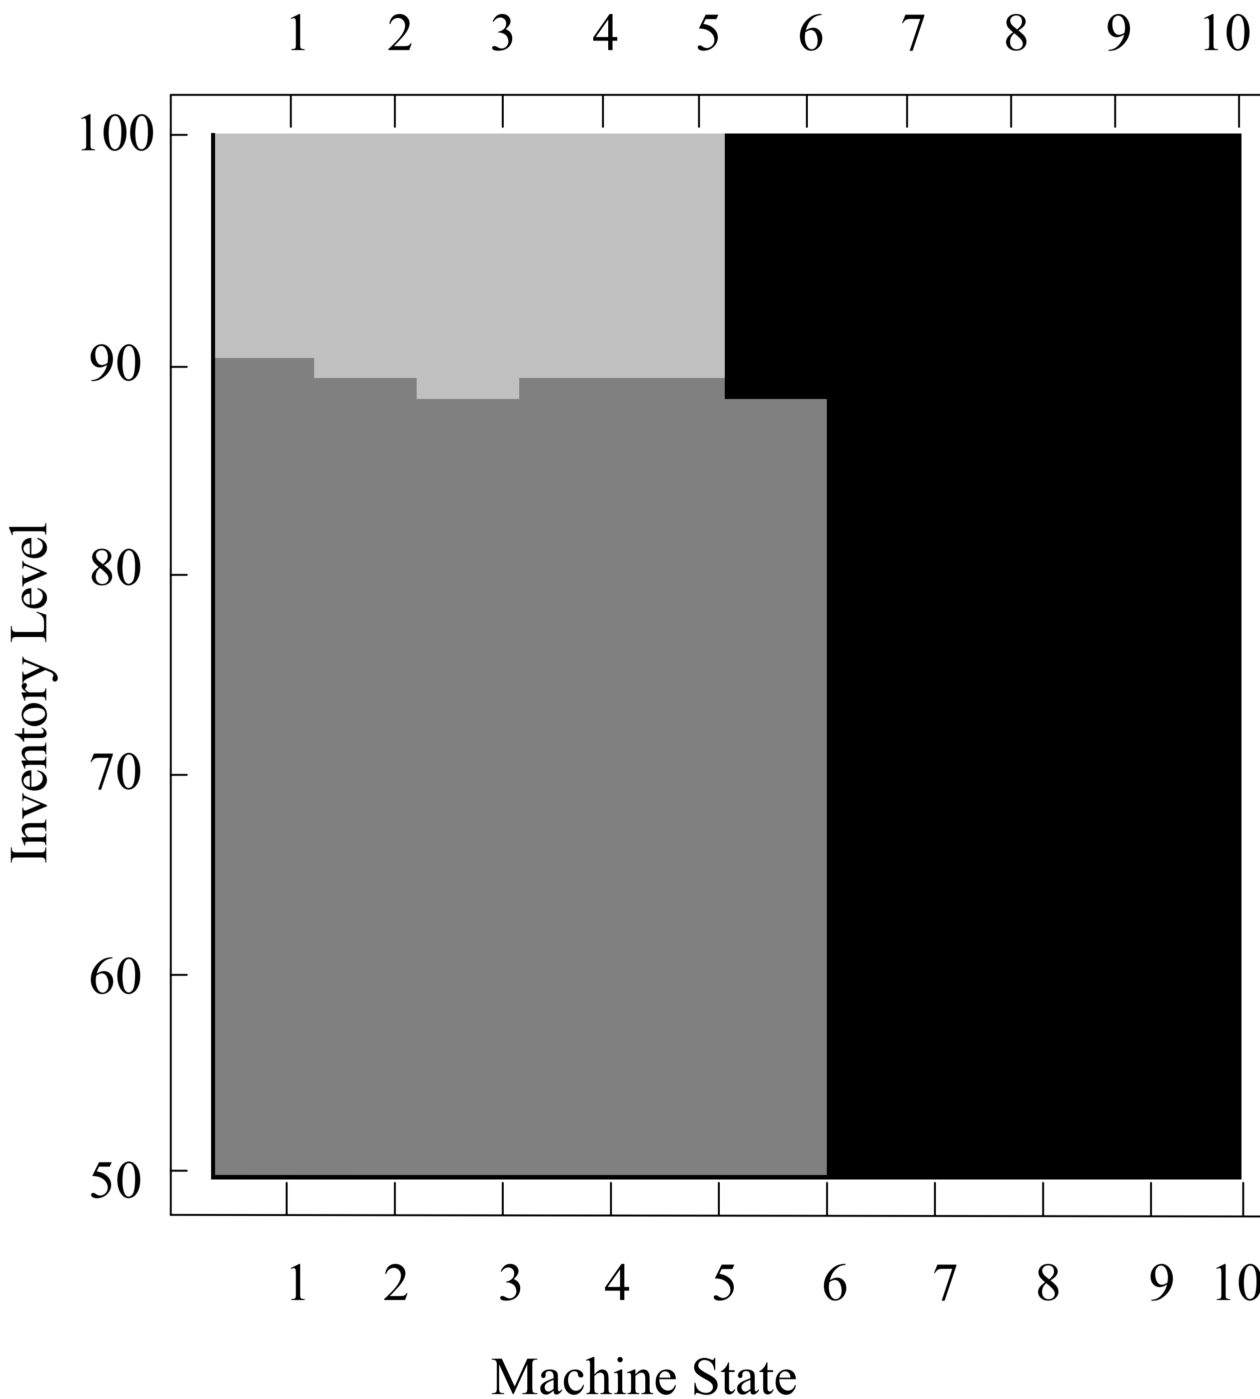

Figure 3: Optimal policy structure for the last period

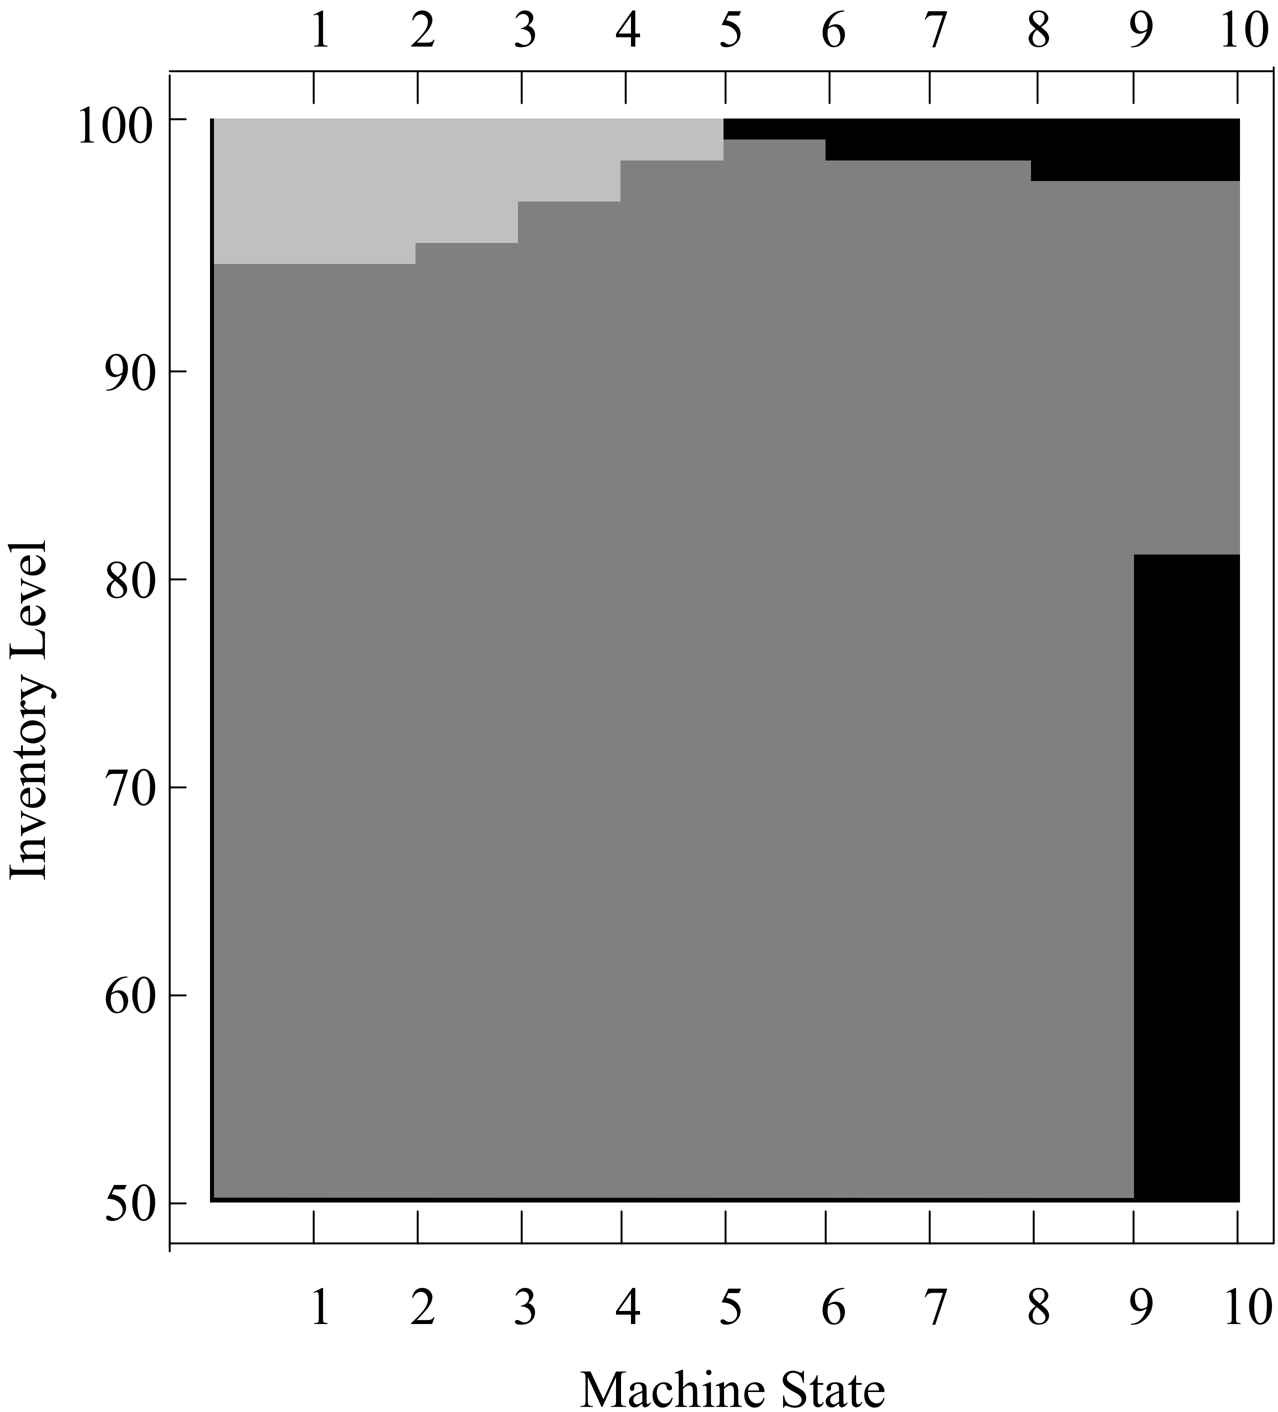

Figure 4: Optimal policy structure for two periods from the due date

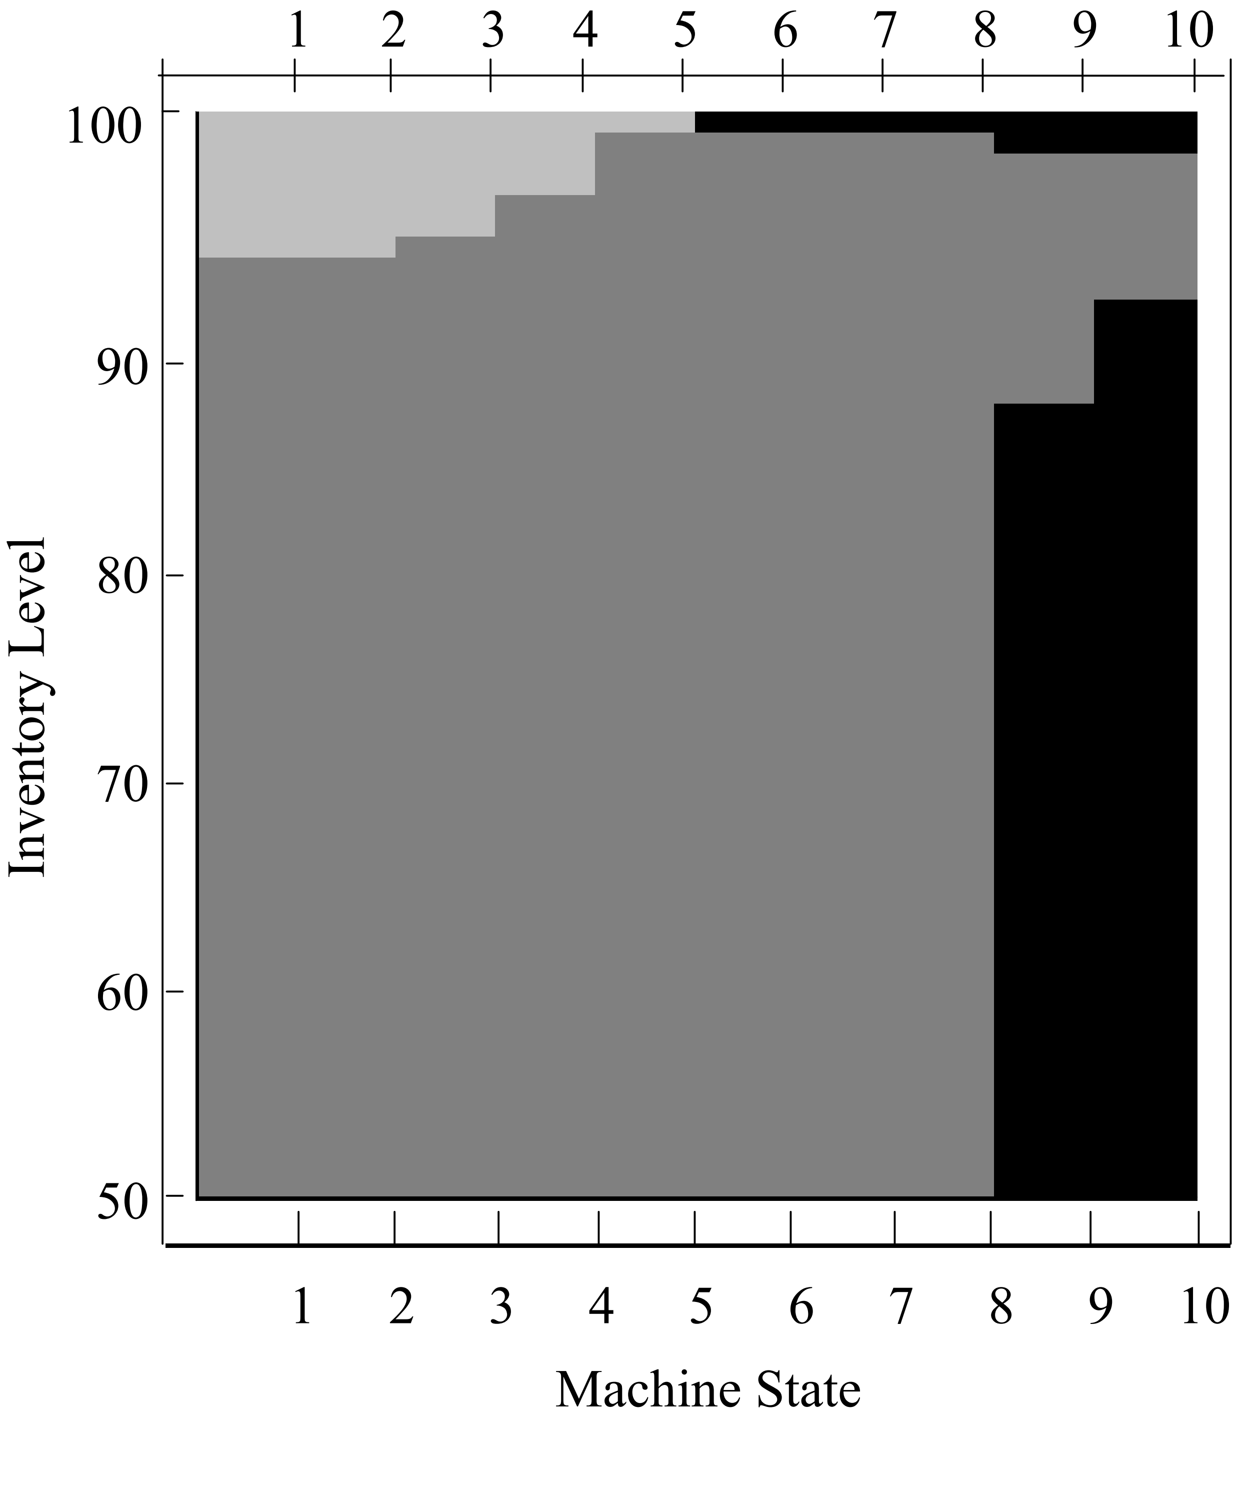

Figure 5: Optimal policy structures for three periods from the due date

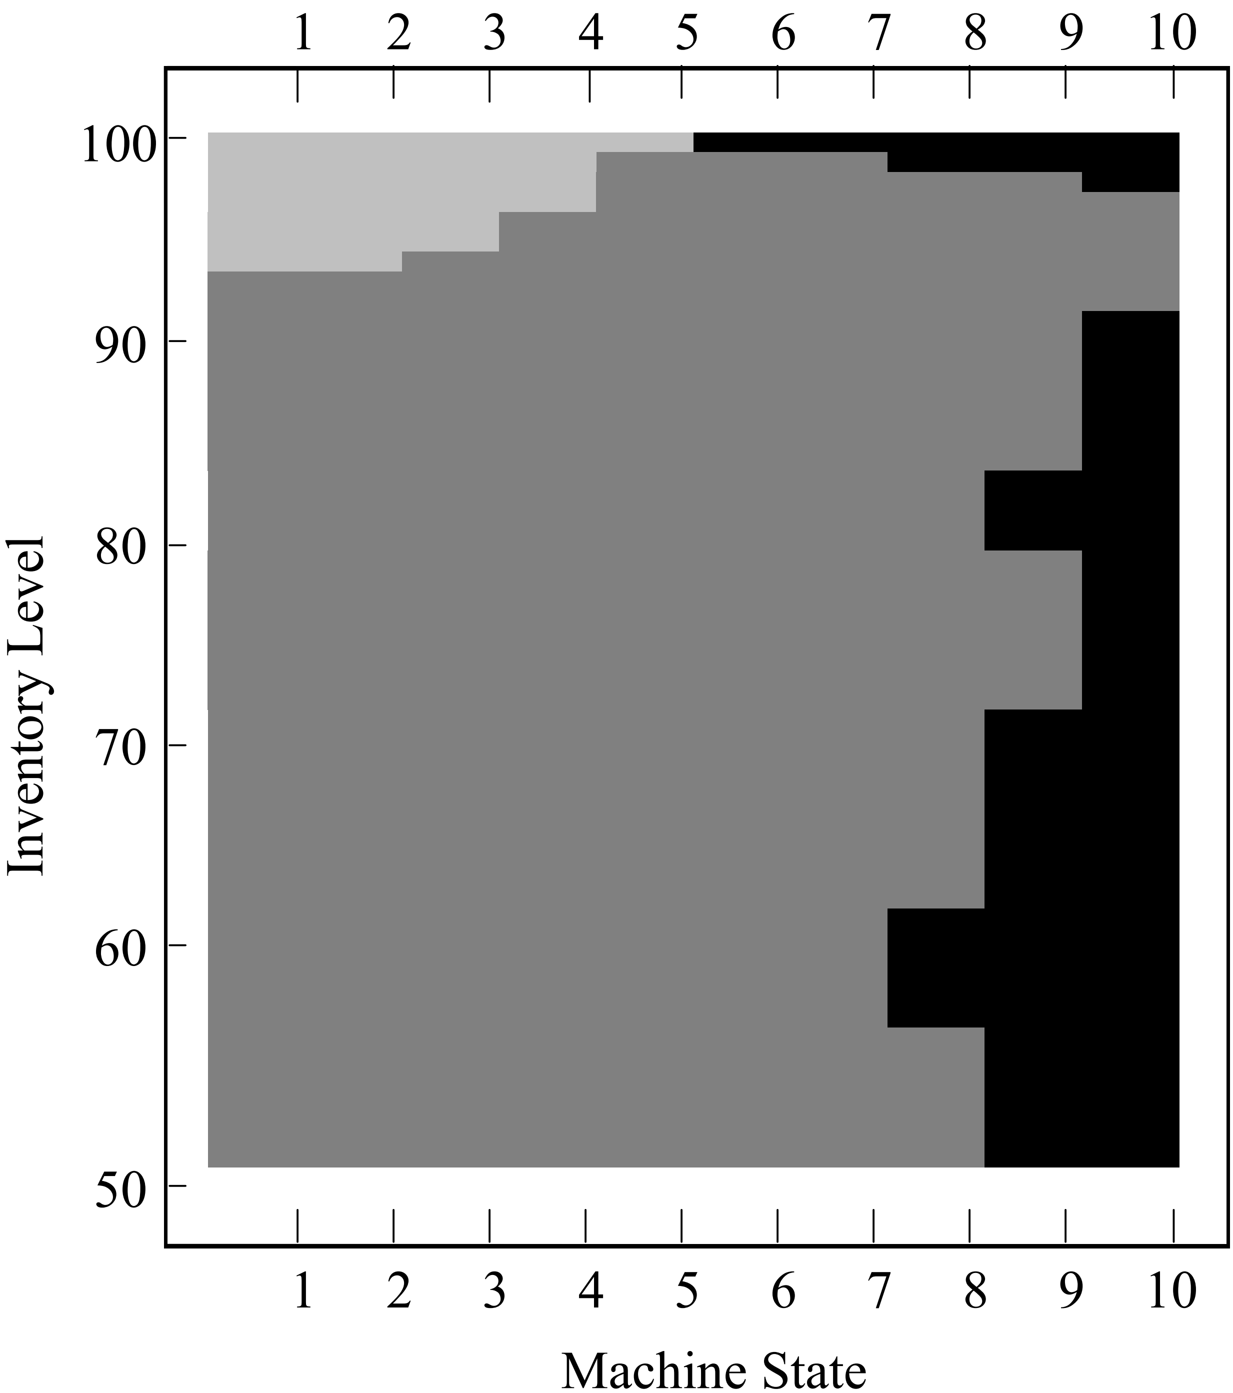

Figure 6: Optimal policy structures for four periods from the due date