University of Groningen

Condition-Based Production Planning

uit het Broek, Michiel A. J.; Teunter, Ruud H.; de Jonge, Bram; Veldman, Jasper; Van Foreest, Nicky D.

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Final author's version (accepted by publisher, after peer review)

[Link to publication in University of Groningen/UMCG research database](#)

Condition-Based Production Planning: Adjusting Production Rates to Balance Output and Failure Risk

Michiel A. J. uit het Broek, Ruud H. Teunter,
Bram de Jonge, Jasper Veldman, Nicky D. van Foreest
University of Groningen, Faculty of Economics and Business, Operations,
Nettelbosje 2, 9747 AE Groningen, a.j.uit.het.broek@rug.nl

Problem definition: Many production systems deteriorate over time as a result of load and stress caused by production. The deterioration rate of these systems typically depends on the production rate, implying that the equipment's deterioration rate can be controlled by adjusting the production rate. We introduce the use of condition monitoring to dynamically adjust the production rate in order to minimize maintenance costs and maximize production revenues. We study a single-unit system for which the next maintenance action is scheduled upfront. **Academic / Practical Relevance:** Condition-based maintenance decisions are frequently seen in the literature. However, in many real-life systems, maintenance planning has limited flexibility and cannot be done last minute. As an alternative, we are the first to propose using condition information to optimize the production rate, which is a more flexible short-term decision. **Methodology:** We derive structural optimality results from the analysis of deterministic deterioration processes. A Markov decision process formulation of the problem is used to obtain numerical results for stochastic deterioration processes. **Results:** The structure of the optimal policy strongly depends on the (convex or concave) relation between the production rate and the corresponding deterioration rate. Condition-based production rate decisions result in significant cost savings (by up to 50%), achieved by better balancing the failure risk and production output. For several systems a win-win scenario is observed, with both reduced failure risk and increased expected total production. Furthermore, condition-based production rates increase robustness and lead to more stable profits and production output. **Managerial Implications:** Using condition information to dynamically adjust production rates provides opportunities to improve the operational performance of systems with production-dependent deterioration.

Key words: Optimal Production Control, Condition Monitoring, Adjustable Production Rate, Reliability, Maintenance Cost, Operational Decision Making, Productivity
History: November 5, 2018

1. Introduction

Many production systems deteriorate over time as a result of load and stress caused by production. Recent advances in modern sensor techniques have created opportunities for monitoring such systems to improve their operational decision making. Researchers have designed sophisticated condition-based maintenance policies that rely on various types of condition information, providing insights on how to reduce maintenance cost and increase equipment reliability. However, in many real-life systems, maintenance planning has limited flexibility and cannot be done last minute because arranging tools, parts, and technicians takes time.

A more short-term operational option is to control the equipment's deterioration process by adjusting the production rate. This applies, for instance, to wind turbine gearboxes and generators that deteriorate faster at higher speeds (Feng et al. 2013, Zhang et al. 2015), conveyor belts that fail more often when used at higher rotational speeds (Nourelfath and Yalaoui 2012), trucks that fail earlier when heavier loaded (Filus 1987), large computer clusters that fail more often under higher workloads (Ang and Tham 2007, Iyer and Rossetti 1986), and cutting tools that wear faster at higher speeds (Dolinšek et al. 2001). In addition, the recent advent of the Internet of Things (IoT) allows to control production rates remotely and in real-time, thereby making it practically viable to exploit this relation between production and deterioration. Surprisingly, to the best of

our knowledge, no studies exist that focus on controlling the equipment's deterioration process by adapting the production rate based on condition information.

As mentioned, maintenance operations often require many resources to be mobilized, making it difficult (if at all possible) to reschedule maintenance activities. For this reason, plants like paper mills, power plants, and refineries typically perform periodic maintenance activities at so-called turnarounds. At these turnarounds, the entire system is shut down, which also allows for maintenance activities to be clustered. However, equipment may deteriorate more slowly or faster than expected and, therefore, it can be profitable to adjust the production rate based on condition information in order to increase production or to avoid a possible costly failure before the next turnaround.

This study explores the benefits of condition-based production rate optimization between consecutive maintenance operations that are planned well in advance. We analyze these benefits for a single-unit system with a single measurable condition. The system is overhauled at prespecified times and we study optimal production decisions between two consecutive maintenance activities. The decision maker can, based on the current deterioration level and the remaining time until the next maintenance moment, adjust the production rate and thereby the deterioration rate. Our main contribution is to introduce and explore the concept of condition-based production rate decisions for systems with an adjustable production rate that directly affects the deterioration rate.

Structural insights and exact analytical solutions are derived for deterministic deterioration processes, and a numerical analysis shows that stochastic systems behave similarly. Our results reveal that the system's profitability, which is driven by a subtle trade-off between production revenues and maintenance costs, can be significantly improved by dynamically adjusting the production rate. An encouraging result is that we observe win-win scenarios for several systems, with both reduced failure risk and increased expected production. Another insightful result is that the structure of the optimal policy is very different for systems with decreasing or increasing marginal deterioration rates as a function of the production rate. For decreasing marginal deterioration rates, the optimal policy is to produce at the maximum rate and to switch off the system when the deterioration level reaches a time-dependent threshold. For increasing marginal deterioration rates, it is optimal to aim at a constant production rate.

The remainder of this paper is organized as follows. In Section 2 we discuss the relevant literature on production planning and on the use of condition monitoring for operational decision making. In Section 3 we provide a formal problem description. In Section 4 we analytically study the system with a deterministic deterioration process. In Section 5 we use a Markov decision process formulation to validate the insights from deterministic systems for systems with a stochastic deterioration process. We conclude and provide suggestions for future research in Section 6.

2. Literature
In the current literature, production and maintenance decisions are often optimized separately (e.g., Shen et al. 2014, Iravani and Krishnamurthy 2007). The literature on production planning under uncertainty, including equipment failure uncertainty, is reviewed by Mula et al. (2006) and an extensive review on the use of condition monitoring for maintenance decisions is conducted by Alaswad and Xiang (2017). A review that addresses the joint optimization of production and maintenance is carried out by Sethi et al. (2002). In what follows, we distinguish three streams of literature on the interaction between production and system failures. The first introduces adjustable production rates and assumes that higher production rates result in increased failure risk. The second considers production-dependent deterioration without condition monitoring. Third, we discuss the literature on the general use of condition monitoring for operational decision making, and also zoom in on condition monitoring for production decisions.

There are many studies on systems with an adjustable production rate and production-dependent failure rates. Liberopoulos and Caramanis (1994) consider a single-unit system with constant demand, where corrective maintenance is performed upon failure. Their objective is to find a policy that minimizes backorder and inventory costs. Boukas et al. (1995) include preventive maintenance decisions into this system, and Hu et al. (1994) show that the reliability of the system can be improved by reducing the production rate. Martinelli (2005, 2007) studies the structure of optimal

production policies under production-dependent failures with two failure rates, and later generalizes this to more general failure rate functions (Martinelli 2010). Recent extensions include a system with two machines (Francie et al. 2014) and a run-based maintenance policy for the production scheduling problem (Lu et al. 2015). These studies assume that failure rates only depend on the age and the current production rate. Thus the production rate only affects the current failure risk and has no effect on the future failure behavior of the system.

In many practical situations, the production rate does not only affect the current failure probability but also results in permanent deterioration to the system, referred to as production-dependent deterioration. Zied et al. (2011) analyze production-dependent deterioration by accelerating the system's aging proportional to the production rate. They consider a single-unit system with stochastic demand and optimize a block-based maintenance policy. Between maintenance actions, the adjustable production rate is used to balance inventory cost and failure risk. Ayed et al. (2012) extend the system to two units. De Jonge and Jakobsons (2018) consider block-based maintenance optimization for a machine for which the usage is random and that only deteriorates when it is turned on. These studies include production-dependent deterioration, but do not consider the potential of monitoring the actual deterioration level of the system, as we do.

It is well known that the use of condition monitoring can significantly improve operational decision making. For example, condition-based maintenance results in improved system reliability and lower maintenance costs (Makis and Jiang 2003, Kim and Makis 2013, Liu et al. 2017). The literature on condition-based maintenance, not restricted to production systems, is rich and deterioration processes depending on time, the current deterioration level, and exogenously given operational modes are covered (Liu et al. 2013, Khaleghei and Makis 2016, Samuelson et al. 2017). A current trend in the literature is to study the use of condition monitoring for other operational decisions such as improved stock keeping of spare parts (Olde Keizer et al. 2017, Zhang and Zeng 2017), managing rentals like cars (Slaugh et al. 2016), and determining optimal production lot-sizes (Peng and Van Houtum 2016). The latter study uses condition monitoring to determine whether a new lot is started or preventive maintenance is performed. These studies, in contrast to ours, exclude the possibility to actively influence the deterioration process by adjusting the production rate.

Others have considered the use of condition information to schedule the production of multiple product types. Sloan and Shanthikumar (2000) study a single-unit system in which the yield differs between products and is affected by the deterioration level of the system. Condition information is used to decide which product to produce and when to perform maintenance. Batun and Maillart (2012) reconsider this study and point out an error in the objective function. Kazaz and Sloan (2008, 2013) extend the system by incorporating products with different production times, and, as a consequence, different expected deterioration increments. These studies assume that maintenance can be performed at any time, that is, after a negligible planning time, whereas in our system maintenance is only performed at prespecified maintenance times. Furthermore, their focus is on production scheduling and the condition information is not used to adjust the production rate.

There are very few studies that do consider systems with adjustable production rates and condition monitoring. However, in these studies the production rate has no influence on the deterioration rate of the system. Iravani and Duenyas (2002) minimize inventory holding costs by producing at a slower rate if the deterioration level is low. Sloan (2004) extends the setting by introducing stochastic demand.

We conclude that the interaction between production decisions and failure behavior of systems is well studied, but that the potential value of using condition monitoring to adjust the production rate, and thereby the deterioration rate, has been ignored.

3. Problem Description
We consider a single-unit system with a single condition parameter. The production rate of the system is adjustable over time, and the deterioration rate (i.e., the average amount of additional deterioration per time unit) depends on the current production rate. Maintenance is only performed at prespecified maintenance moments and the next maintenance action is scheduled at time T. The deterioration process of the system is continuously monitored and described by a continuous-time

stochastic process $X = \{X(t) \,|\, t \geq 0\}$. Deterioration level 0 indicates that the system is as-good-as-new and failure occurs when the deterioration level exceeds a fixed failure level L.

At any time t and deterioration level $X(t)$, the decision maker can control the production rate $u(t, X(t))$ of the system, which ranges between 0 (no production) and 1 (maximum production). If the system has failed, it cannot produce and the production rate is fixed at 0. For a given state (t, x), the set of feasible production rates $\mathcal{A}(x)$ is thus equal to

$$\mathcal{A}(x) = \begin{cases} [0,1] & \text{if } x < L, \\ \{0\} & \text{otherwise.} \end{cases}$$

In words, the decision maker can only control the production rate as long as the system is functioning, i.e., if it is in a state in the set $\mathcal{S} = \{(t, x) \,|\, 0 \leq t \leq T, \ 0 \leq x \leq L\}$. We consider a policy to be admissible if it specifies a feasible production rate for each state in the set $\mathcal{S}$ and, to facilitate the proofs in the next section, if it has a finite, but arbitrary number of jumps over time. Clearly, imposing the latter constraint does not have any practical implications. We let $\mathcal{A}$ denote the set of all admissible policies.

The deterioration rate of the system depends on the production rate and is denoted by $g(u)$. We refer to this function g as the production-deterioration relation (pd-relation in short). It is natural to assume that there are no production rates for which the condition of the system improves, hence g is assumed to be nonnegative. We let $g_{\min} = \min_{u \in [0,1]} g(u)$ and $g_{\max} = \max_{u \in [0,1]} g(u)$ refer to the minimum and maximum deterioration rate, respectively. Notice that we distinguish systems that do deteriorate for all production rates ($g_{\min} > 0$) and systems that do not ($g_{\min} = 0$). We remark that systems may deteriorate (although slowly) even if the system is idle, for instance due to bearings that may become slightly unbalanced due to one-sided pressure or as a result of corrosion. Furthermore, we note that pd-relations are most likely to be increasing in practice, but our analysis does not require this assumption.

The production revenue generated by the system is proportional to the production rate and equals $u\pi$ per time unit when producing at rate u. The cost of performing maintenance depends on the deterioration level at the moment of maintenance. If the system is still functioning, preventive maintenance at a cost c_{pm} is carried out, whereas more expensive corrective maintenance at a cost c_{cm} has to be performed if the system has failed. Thus the maintenance cost as a function of the deterioration level $X(T)$ at the moment of maintenance equals

$$c(X(T)) = \begin{cases} c_{\text{pm}} & \text{if } X(T) < L, \\ c_{\text{cm}} & \text{otherwise.} \end{cases}$$

This cost structure is commonly used in the maintenance literature (see, e.g., Liu et al. 2017, De Jonge et al. 2017, Zhang and Zeng 2017). It is often realistic, for instance when maintenance means the replacement of a unit, implying that its cost is fixed as long as the unit has not failed. A corrective replacement is often more expensive than a preventive replacement, for instance if unit failure results in damage to other units as well. Furthermore, we assume that maintenance will always be carried out at the scheduled moment, regardless of the deterioration level. This is justified when maintenance has to be planned well in advance. Moreover, under the optimal policy, production rates will be high as long as deterioration is low, implying that it is very unlikely that the deterioration level is low at the maintenance moment.

The expected profit until the next maintenance moment for a given state (t, x) and a given production policy $\boldsymbol{u} = \{u(\tau, X(\tau)) \,|\, 0 \leq \tau \leq T, 0 \leq X(\tau) \leq L\}$ equals

$$J(\boldsymbol{u}; t, x) = E\left[\pi \int_t^T u(\tau, X(\tau)) \, \mathrm{d}\tau - c(X(T))\right]. \tag{1}$$

Our aim is to maximize this expected total profit. However, for some scenarios, the supremum $J^*(t, x) = \sup_{\boldsymbol{u} \in \mathcal{A}} J(\boldsymbol{u}; t, x)$ cannot be attained because of the discontinuity in the maintenance cost function $c(X(T))$. We therefore determine an admissible policy whose objective value is arbitrarily

close to the supremum $J^*(t,x)$. Thus, for any $\epsilon > 0$, we determine an admissible policy $\boldsymbol{u}^* \in \mathcal{A}$ such that $J(\boldsymbol{u}^*; t, x) > J^*(t, x) - \epsilon$.

We note that the above described setting with a single condition parameter is not only applicable to single-unit systems, but also to multi-unit systems in which one of the units requires a considerably higher maintenance frequency than the other units. For such systems, the length of the maintenance interval will be based on the unit that deteriorates fastest and that requires the highest maintenance frequency. This critical unit will then be maintained after each maintenance interval, and these maintenance moments will be used as opportunities to sometimes maintain the other, more slowly deteriorating units as well (depending on their respective deterioration levels). In such settings, the critical unit is also the main driver for the dynamic production rate.

4. Deterministic Deterioration

In this section we consider a deterministic deterioration process with deterioration increments that are fixed for a given production rate. Although most deterioration processes behave stochastically in practice, studying a deterministic deterioration process allows us to derive analytical insights into the structure of the optimal production control. In Section 5 we will show that the same structure is observed for stochastic processes. Given the deterministic deterioration process and an initial state $(t, x) \in \mathcal{S}$, the problem reduces to

$$J^*(t,x) = \sup_{\boldsymbol{u} \in \mathcal{A}} \left\{ \pi \int_t^T u(\tau, X(\tau))\ \mathrm{d}\tau - c(X(T)) \right\}, \tag{2}$$

subject to

$$X(t) = x,$$
$$\dot{X}(\tau) = g(u(\tau, X(\tau))),$$

where $\dot{X}(\tau)$ denotes the right derivative of $X(\tau)$. In the remainder of this section, we assume that the pd-relation g is continuously differentiable. Moreover, for a given state $(t, x) \in \mathcal{S}$, a policy $\boldsymbol{u}$ fixes the trajectory of the deterioration process $X(\tau)$ for $\tau \geq t$. Thus, for a given state (t, x), we can describe the optimal production rate as function of time, $u(\tau)$ for $\tau \geq t$, instead of a function of both time and the deterioration level.

The first step of our analysis is to partition the set $\mathcal{S}$ into three subsets as illustrated in Figure 1. When the deterioration level is relatively low compared to the remaining time until the maintenance moment, the system remains functioning regardless of the production rate. We let the set $\mathcal{S}_1$ contain all states in which the system will be functioning at time T even if the production rate with the highest deterioration rate is chosen, i.e.,

$$\mathcal{S}_1 = \{(t,x) \in \mathcal{S} \mid x + (T-t)\, g_{\max} < L\}.$$

Recall that the highest deterioration rate $g_{\max} = \max_{u \in [0,1]} g(u)$ does not necessarily correspond to the maximum production rate $u = 1$. On the other hand, for systems that deteriorate for all production rates, we know that the system will fail before maintenance takes place if the deterioration level is close to the failure level given the remaining time until maintenance. We let the set $\mathcal{S}_3$ contain all states in which the system will fail with certainty before time T, i.e.,

$$\mathcal{S}_3 = \{(t,x) \in \mathcal{S} \mid x + (T-t)\, g_{\min} \geq L\}.$$

The set $\mathcal{S}_2$ contains the remaining states, in which the system can either be functioning or failed upon maintenance, depending on the selected production rates, thus $\mathcal{S}_2 = \mathcal{S} \setminus (\mathcal{S}_1 \cup \mathcal{S}_3)$.

For all practical cases we have $g_{\max} < \infty$ implying that the set $\mathcal{S}_1$ cannot be empty. The set $\mathcal{S}_2$ is empty if and only if $g_{\min} = g_{\max}$, which is the case if the production rate has no influence on the deterioration rate. The set $\mathcal{S}_3$ is empty if and only if $g_{\min} = 0$. A practical scenario with $g_{\min} = 0$ is a system that does not deteriorate when idle, that is, $g(0) = 0$.
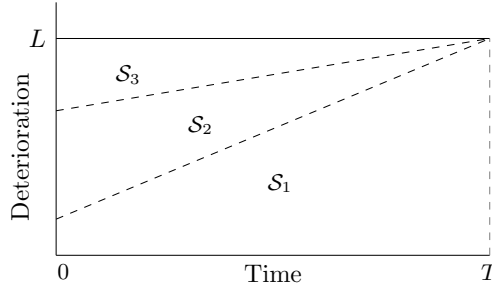
Figure 1   **Schematic overview of the subsets of** $\mathcal{S} = \{(t,x) \mid 0 \le t \le T, \ 0 \le x \le L\}$.

The remainder of this section is organized as follows. In Section 4.1, we show that there is a policy whose objective value is arbitrarily close to the supremum (2), even if we restrict the decision maker to control the production rate at prespecified times only. In Section 4.2, we derive the optimal policy for states in $\mathcal{S}_3$. In Section 4.3, we find the optimal policy under the restriction that the system does not fail. This is obviously the optimal policy starting in states in $\mathcal{S}_1$, as the system cannot fail from those states. It also provides the best policy that avoids failure for states in $\mathcal{S}_2$. However, for those states we further have to consider policies that deliberately let the system fail, which is done in Section 4.4. We end with an illustrative example in Section 4.5.

4.1. Prespecified Decision Moments

We show that the supremum (2) can be approached arbitrarily close even if the decision maker is only allowed to control the production rate at prespecified moments, as long as the maximum duration between two consecutive decisions is sufficiently short. This allows us to use prespecified partitionings in the proofs of subsequent sections.

Let P be a partitioning of the time interval $[0,T]$ into n subintervals $[t_{i-1}, t_i)$ where $i \in \mathcal{I} = \{1, \ldots, n\}$ and $0 = t_0 < t_1 < \ldots < t_n = T$. The length of interval i equals $\delta_i = t_i - t_{i-1}$ and the longest interval has length $\delta_{\max} = \max_{i \in \mathcal{I}} \delta_i$. In each interval, the decision maker can only set a single production rate $\hat{u}_{P,i}$. We denote the restricted policy as $\hat{\boldsymbol{u}}_P = (\hat{u}_{P,1}, \ldots, \hat{u}_{P,n})$ and the set of all admissible policies on a given partitioning as $\hat{\mathcal{A}}_P$. For notational ease, we drop the subscript P for $\hat{u}_{P,i}$, $\hat{\boldsymbol{u}}_P$, and $\hat{\mathcal{A}}_P$ in the remainder of the study and distinguish the restricted policy by the hat symbol.

Recall that the optimal value of the problem equals $J^*(t,x) = \sup_{\boldsymbol{u} \in \mathcal{A}} J(\boldsymbol{u}; t, x)$. Furthermore, the set of feasible policies $\mathcal{A}$ is nonempty since the policy $u(\tau) = 0$ for $\tau \ge t$ is always feasible. Then, by definition of the supremum, there is a policy in $\mathcal{A}$ whose corresponding objective value is arbitrarily close to the supremum. That is, for every $\epsilon > 0$ there is a policy $\boldsymbol{u}^* \in \mathcal{A}$ such that

$$J(\boldsymbol{u}^*; t, x) > J^*(t, x) - \epsilon. \tag{3}$$

For a given partitioning, we construct a policy $\hat{\boldsymbol{u}}^*$ corresponding to $\boldsymbol{u}^*$ such that, in each interval $i \in \mathcal{I}$, $\hat{\boldsymbol{u}}^*$ takes the production rate of $\boldsymbol{u}^*$ with the lowest corresponding deterioration rate, i.e.,

$$\hat{u}_i^* = \underset{u^*(\tau)}{\arg \min} \left\{ g(u^*(\tau)) : \tau \in [t_{i-1}, t_i] \right\}.$$

Because $\boldsymbol{u}^*$ is feasible, it immediately follows that $\hat{\boldsymbol{u}}^*$ is feasible. Note that $J(\hat{\boldsymbol{u}}^*; t, x) \le J(\boldsymbol{u}^*; t, x)$ since $\hat{\boldsymbol{u}}^* \in \hat{\mathcal{A}}$, $\boldsymbol{u}^* \in \mathcal{A}$, and $\hat{\mathcal{A}} \subset \mathcal{A}$.

Policy $\boldsymbol{u}^*$ is Riemann-integrable over time since the production rate has a bounded range and a finite number of jumps (see Rudin 1976, Theorem 11.33). By definition of Riemann-integrability (see Abbott 2015, Theorem 8.1.2), it follows that for every $\hat{\epsilon} > 0$ there is a $\delta > 0$ such that for any partitioning with $\delta_{\max} < \delta$ we have

$$J(\hat{\boldsymbol{u}}^*; t, x) > J(\boldsymbol{u}^*; t, x) - \hat{\epsilon}. \tag{4}$$

It follows from (3) and (4) that there are partitionings for which the objective value of the restricted policy $\hat{\boldsymbol{u}}^*$ corresponding to $\boldsymbol{u}^*$ is arbitrarily close to the supremum (2). Only the length of the longest interval is relevant and thus we can use partitions both with equally and with unequally sized intervals. Hence, for fine enough partitionings, maximizing the objective value with prespecified decision moments is equivalent to maximizing the objective value of the unrestricted policy.

4.2. Optimal Policy with Unavoidable Failure

We first derive the optimal policy for states $(t, x) \in \mathcal{S}_3$, i.e., states in which the system will fail with certainty before the maintenance action at time T. Recall that the set $\mathcal{S}_3$ is empty for $g_{\min} = 0$ and thus in this section we have $g_{\min} > 0$. We first show the optimal policy for general pd-relations, and then simplify this for linear, strictly concave, and strictly convex pd-relations.

PROPOSITION 1. *For each state* $(t, x) \in \mathcal{S}_3$, *a sufficient condition for optimality is that* $u(\tau) \in \arg\sup_{u \in [0,1]}\{u/g(u)\}$ *for all* $\tau \geq t$.

Proof. Suppose the system is in a state $(t, x) \in \mathcal{S}_3$. Maximizing profit until maintenance is equivalent to maximizing production revenues until maintenance since the system will fail with certainty and thus the maintenance cost $c(X(T))$ equals c_{cm}.

We partition the deterioration interval $[x, L]$ into n equally large subintervals of size $\delta = (L - x)/n$ and restrict the decision maker to set a single production rate in each interval (see Figure 2). We partition the deterioration levels instead of the time horizon since each possible deterioration trajectory will reach the failure level L before the end of the time horizon. The rate in interval i is denoted by $\tilde{u}_i^{(n)}$ where $i \in \mathcal{I} = \{1, \ldots, n\}$. The restricted policy is denoted as $\tilde{\boldsymbol{u}}^{(n)} = (\tilde{u}_1^{(n)}, \ldots, \tilde{u}_n^{(n)})$ and the set of all admissible policies as $\tilde{\mathcal{A}}^{(n)}$. For notational purposes, we drop the superscript for $\tilde{u}_i^{(n)}$, $\tilde{\boldsymbol{u}}^{(n)}$, and $\tilde{\mathcal{A}}^{(n)}$ in the remainder of the proof.

The total revenue for a given policy $\tilde{\boldsymbol{u}}$ is the sum of the revenues in the individual intervals, i.e.,

$$J(\tilde{\boldsymbol{u}}; t, x) = \pi \sum_{i \in \mathcal{I}} \tilde{u}_i \tau(\tilde{u}_i), \tag{5}$$

where $\tau(\tilde{u}_i)$ is the time spent in interval i when producing at rate $\tilde{u}_i$. Since the system deteriorates with rate $g(\tilde{u}_i)$ we have $\tau(\tilde{u}_i) = \delta/g(\tilde{u}_i)$. Substituting this into objective function (5) gives

$$J^*(t, x) = \sup_{\tilde{\boldsymbol{u}} \in \tilde{\mathcal{A}}} J(\tilde{\boldsymbol{u}}; t, x) = \delta\pi \sup_{\tilde{\boldsymbol{u}} \in \tilde{\mathcal{A}}} \left\{ \sum_{i \in \mathcal{I}} \frac{\tilde{u}_i}{g(\tilde{u}_i)} \right\}.$$

The terms within the summation depend on interval i only, thus are independent of the decisions made in the other intervals. It follows that we can interchange the summation and the supremum, hence

$$J^*(t, x) = \delta\pi \sum_{i \in \mathcal{I}} \left(\sup_{\tilde{u}_i \in [0,1]} \left\{ \frac{\tilde{u}_i}{g(\tilde{u}_i)} \right\} \right) = (L - x)\pi \sup_{u \in [0,1]} \left\{ \frac{u}{g(u)} \right\},$$

where the last equality follows by substituting $\delta = (L - x)/n$. Observe that the objective value that we attain is independent of the partitioning that is used. We conclude that the optimal value is attained if $u(\tau) \in \arg\sup\{u/g(u)\}$ for all $\tau \geq t$. $\square$

COROLLARY 1. *For any state for which a failure cannot be avoided, there is an optimal policy that produces at a constant rate until failure.*

Proposition 1 states that if the system is in state $(t, x) \in \mathcal{S}_3$, then any policy that maximizes $u/g(u)$ until the unavoidable failure is optimal. This result is intuitive since the decision maker already knows that the system will fail and therefore aims to maximize the production gained for each additional unit of deterioration. From now on, we refer to the set of production rates that maximize $u/g(u)$ as the set of efficient rates U_{eff}. The set U_{eff} is independent of time and thus there exist an optimal policy that produces at a constant rate until failure (Corollary 1). However, this
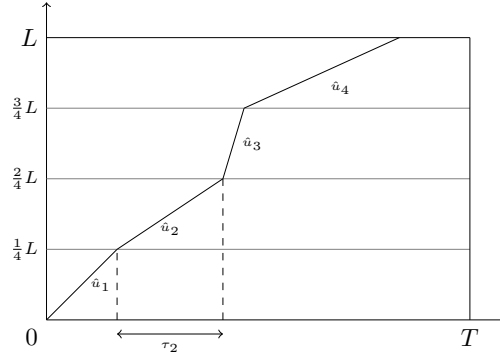
Figure 2 **Overview of the restricted decisions and a possible trajectory for** $n = 4$ **and** $(t, x) = (0, 0)$.

policy is not necessarily unique since U_{eff} can contain multiple elements. For example, consider a pd-relation with $g(u) > u$ for $u \in [0, 0.5)$ and $g(u) = u$ for $u \in [0.5, 1]$. The set of efficient rates for this pd-relation equals $U_{\text{eff}} = [0.5, 1]$ and thus any admissible policy, also the non-constant ones, with $u(\tau) \in [0.5, 1]$ for $\tau \geq t$ is optimal.

The following lemma simplifies the result of Proposition 1 if we make specific assumptions on the form of the pd-relation g. The proofs of this and all subsequent lemmas can be found in the appendix.

LEMMA 1. *Consider a pd-relation g with $g_{\min} > 0$.*
a) *If the pd-relation is strictly concave or linear, then the set of efficient rates is $U_{eff} = \{1\}$.*
b) *If the pd-relation is strictly convex, then there is only one efficient rate. The efficient rate can be found by first solving $z = \arg\{g(u) = u\, g'(u)\}$ and consequently setting $U_{eff} = \{\min(1, z)\}$.*

From the above lemma, we know that if the system is in state $(t, x) \in \mathcal{S}_3$ and the pd-relation g is linear, strictly concave, or strictly convex, then the optimal policy is unique and constant over time. Furthermore, for the former two, the optimal policy is to produce at the maximum rate until failure.

4.3. Optimal Policy with Maximum Deterioration Constraint

For states $(t_1, x_1) \in \mathcal{S}_2$ we have the option to prevent failure or to deliberately let the system fail. In this section we consider the case in which failure is prevented by introducing a constraint that describes a maximum allowed deterioration level x_2 at time t_2, that is, $X(t_2) \leq x_2$ where $t_1 < t_2 \leq T$ and $x_1 < x_2 < L$. This scenario directly solves both the optimal policy for any state in $\mathcal{S}_1$ (as the system is guaranteed not to fail from any such state) and the case in which the decision maker decides to avoid a failure while being in a state in $\mathcal{S}_2$, namely by using $t_2 = T$ and $x_2 = L - \epsilon$ where ϵ is an arbitrarily small positive number. In addition, the insights hold for any functioning state and thus also for states in $\mathcal{S}_3$. The results are also used in the next section to derive the optimal policy for the case in which the decision maker decides to let the system fail while being in a state in $\mathcal{S}_2$.

We partition the time interval $[t_1, t_2]$ into n equally large subintervals with length $\delta = (t_2 - t_1)/n$. Furthermore, the decision maker is restricted to set a single production rate in each interval, denoted by $\hat{u}_i$ for $i \in \mathcal{I} = \{1, \dots, n\}$. The policy is denoted as $\hat{\boldsymbol{u}} = (\hat{u}_1, \dots, \hat{u}_n)$. We can use this partitioning since the objective value of the restricted policy can approach the supremum arbitrarily closely if n is large enough (see Section 4.1).

Under the maximum allowed deterioration constraint, it is clearly optimal to maximize production revenues. The total revenue equals the sum of the revenues in the individual intervals, i.e.,

$$J(\hat{\boldsymbol{u}}; t_1, x_1) = \delta \pi \sum_{i \in \mathcal{I}} \hat{u}_i. \tag{6}$$

The deterioration increase in interval i equals $\delta\, g(\hat{u}_i)$, and the total deterioration increase over the time interval $[t_1, t_2]$ is allowed to be at most $x_2 - x_1$. It follows that policy $\hat{\boldsymbol{u}}$ must satisfy

$$\sum_{i \in \mathcal{I}} g(\hat{u}_i) \leq \frac{x_2 - x_1}{\delta}, \tag{7}$$

which we refer to as the maximum deterioration constraint. With (6) and (7) we can formulate our optimization problem as $\max \sum_{i \in \mathcal{I}} \hat{u}_i$ subject to $\hat{u}_i \in [0,1]$ and $\sum_{i \in \mathcal{I}} g(\hat{u}_i) \leq c$, where c is some constant. It trivially follows that whenever the policy $\hat{\boldsymbol{u}} = (1, \ldots, 1)$ is feasible, it is the unique optimal policy; thus the optimal policy for any state $(t,x) \in \mathcal{S}_1$ is to produce at the maximum rate until maintenance.

To obtain structural insights into the optimal policy, we use the necessary conditions for optimality described by the Karush-Kuhn-Tucker (KKT) conditions. These conditions imply a set of necessary constraints on the Lagrange multipliers of the dual problem which must be satisfied by the optimal policy. We note that the KKT conditions rely on the assumption that the pd-relation g is continuously differentiable. Let ν be the multiplier corresponding to the maximum deterioration constraint (7), let $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n)$ be a vector with the multipliers corresponding to the constraints $\hat{u}_i \geq 0$ for $i \in \mathcal{I}$, and let $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_n)$ be a vector with the multipliers corresponding to the constraints $\hat{u}_i \leq 1$ for $i \in \mathcal{I}$. The KKT conditions state that for the optimal policy $\hat{\boldsymbol{u}}^*$ there must exist values for ν, $\boldsymbol{\mu}$, and $\boldsymbol{\lambda}$ such that

$$1 - \nu g'(\hat{u}_i) - \lambda_i + \mu_i = 0, \quad \text{for } i \in \mathcal{I}, \tag{8a}$$

$$\nu \left(\sum_{i \in \mathcal{I}} g(\hat{u}_i) - c \right) = 0, \tag{8b}$$

$$\lambda_i(\hat{u}_i - 1) = 0, \quad \text{for } i \in \mathcal{I}, \tag{8c}$$

$$\mu_i \hat{u}_i = 0, \quad \text{for } i \in \mathcal{I}, \tag{8d}$$

$$\nu, \, \boldsymbol{\lambda}, \, \boldsymbol{\mu} \geq 0, \tag{8e}$$

where g' denotes the derivative of g, constraint (8a) is a necessary condition for being in an extreme point of the feasible set, constraints (8b - 8d) represent the complementary slackness conditions, and constraint (8e) implies dual feasibility. Considering the KKT conditions for objective (6) with constraint (7) results in the following properties (see proofs in the appendix).

LEMMA 2. *Suppose we have a pd-relation g, the system is in a functioning state $(t_1, x_1) \in \mathcal{S}$, and there is a maximum deterioration constraint $X(t_2) \leq x_2$ where $x_1 < x_2 < L$ and $t_1 < t_2 \leq T$. Let $\hat{\boldsymbol{u}}^* = (\hat{u}_1, \ldots, \hat{u}_n)$ be an optimal policy, where $\hat{u}_i$ denotes the production rate in time interval $i \in \mathcal{I} = \{1, \ldots, n\}$.*
  a) *If there is an $i \in \mathcal{I}$ such that $\hat{u}_i < 1$, then the maximum deterioration constraint is binding.*
  b) *If the policy $\hat{\boldsymbol{u}} = (1, \ldots, 1)$ is feasible, then it is the unique optimal policy.*
  c) *For all $i, j \in \mathcal{I}$ for which $\hat{u}_i, \hat{u}_j \in (0, 1)$, we have $g'(\hat{u}_i) = g'(\hat{u}_j)$.*
  d) *For all $i \in \mathcal{I}$ for which $\hat{u}_i < 1$, we have $g'(\hat{u}_i) > 0$.*
  e) *If g' is a one-to-one function, then for all $i, j \in \mathcal{I}$ for which $\hat{u}_i, \hat{u}_j \in (0, 1)$, we have $\hat{u}_i = \hat{u}_j$.*

LEMMA 3. *For all pd-relations g, there is an optimal policy with at most two production rates.*

Lemma 2a states that production rates below the maximum rate are only used if this is enforced by the maximum deterioration constraint. It immediately follows that the policy $\hat{\boldsymbol{u}} = (1, \ldots, 1)$ is the unique optimal policy whenever it is feasible (Lemma 2b). This is in fact trivial, since no policy can produce more than producing at the maximum rate over the whole time interval. It follows that the optimal policy for any state $(t, x) \in \mathcal{S}_1$ is to produce at the maximum rate until maintenance.

Lemma 2c follows from the condition that, as stated by Lemma 2a, the maximum deterioration constraint must be binding for any policy that uses intermediate production rates. If a policy uses two intermediate production rates $\hat{u}_i$ and $\hat{u}_j$ for which $g'(\hat{u}_i) > g'(\hat{u}_j)$, then the decision maker can improve the generated revenue by marginally decreasing the production rate in interval i while marginally increasing the production rate in interval j. Lemma 2d implies that rates $u < 1$ for which $g'(u) \leq 0$ cannot be part of an optimal policy. This is intuitive, since for these rates one can increase the production rate while the system would deteriorate slower.

Lemma 2e states that for pd-relations with a one-to-one derivative (e.g., strictly convex functions), the optimal policy can only contain a single intermediate rate. Thus, for this class of pd-relations, we know that the complexity of the optimal policy reduces to the class of policies that take at

most three values over time, namely the minimum and maximum rate and one intermediate rate. Lemma 3 states that for any pd-relation there is an optimal policy that uses at most two different production rates over time.

So far we did not make assumptions on the structure of the pd-relation. In the following sections, we use the previous lemmas to derive exact closed-form optimal policies for strictly convex, strictly concave, and linear pd-relations for any functioning state $(t_1, x_1) \in \mathcal{S}$ with a given maximum deterioration constraint $X(t_2) \le x_2$ where $t_1 < t_2 \le T$ and $x_1 < x_2 < L$.

4.3.1.   Strictly Convex pd-relations First we show that the optimal production rate for strictly convex pd-relations is constant over time, independent of the partitioning that is used.

LEMMA 4. *The optimal rate is constant over time for strictly convex pd-relations.*

We know that either producing at the maximum rate is the unique optimal policy or the maximum deterioration constraint is binding (see Lemma 2). Substituting $\delta = (t_2 - t_1)/n$ and a constant for the rate (see Lemma 4) into the maximum deterioration constraint (7) gives

$$\hat{\boldsymbol{u}}_i^* = g^{-1}\left(\frac{x_2 - x_1}{t_2 - t_1}\right) \quad \text{for all } i \in \mathcal{I}. \tag{9}$$

Notice that the full inverse of g can have two solutions when g is first decreasing and then increasing. By Lemma 2d, it directly follows that the solution in the increasing part is the optimal one. We conclude that, under a maximum deterioration constraint, the optimal policy equals $\hat{\boldsymbol{u}}^* = (1, \dots, 1)$ if this is feasible and otherwise the optimal rate is as given in (9). We summarize this finding in the following proposition.

PROPOSITION 2. *Suppose the pd-relation g is strictly convex, the system is in a functioning state $(t_1, x_1) \in \mathcal{S}$, and there is a given maximum deterioration constraint $X(t_2) \le x_2$ where $x_1 < x_2 < L$ and $t_1 < t_2 \le T$. Then the unique optimal policy is constant over time and equals $u(\tau) = 1$ for $\tau \ge t_1$ if feasible or, otherwise,*

$$u(\tau) = \max\left\{g^{-1}\left(\frac{x_2 - x_1}{t_2 - t_1}\right)\right\} \text{ for } \tau \ge t_1.$$

When the system is in a functioning state $(t_1, x_1) \in \mathcal{S}_2$ and the decision maker wants to avoid the failure, i.e., we have the constraint $X(T) \le L - \epsilon$ where ϵ is an arbitrarily small positive number, then by Proposition 2 we know that $u(\tau) = \max\{g^{-1}((L - \epsilon - x_1)/(T - t_1))\}$ for $\tau \ge t_1$ is the unique optimal policy. Thus the optimal policy is to produce at the highest rate such that the system just not fails upon the moment of maintenance.

4.3.2.   Strictly Concave pd-relations The derivative of strictly concave functions is one-to-one and thus the optimal policy for concave functions can use at most one intermediate rate (see Lemma 2e). Lemma 5 strengthens this result and states that for strictly concave pd-relation the optimal policy uses the intermediate rate in at most a single time interval.

LEMMA 5. *For strictly concave pd-relations, the optimal policy has at most one time interval in which an intermediate production rate is used.*

So the optimal policy has at most one time interval in which the system produces at an intermediate rate. In all other intervals, the system is either idle or producing at the maximum rate. Producing at the maximum rate generates more revenue than being idle and thus the optimal policy produces at the maximum rate as much as possible. Remark that the possible single time interval in which the system produces at an intermediate rate becomes negligible as the partitioning becomes finer. Moreover, the specific time intervals in which the system is producing is irrelevant and thus the optimal policy is not unique. We summarize this finding in the following proposition.

PROPOSITION 3. *Suppose the pd-relation g is strictly concave, the system is in a functioning state $(t_1, x_1) \in \mathcal{S}$, and there is a maximum deterioration constraint $X(t_2) \le x_2$ where $x_1 < x_2 < L$ and $t_1 < t_2 \le T$. Then an optimal policy is to produce at the maximum rate and then switch off the system at the latest moment in time such that $X(t_2) \le x_2$.*

4.3.3.   Linear pd-relations We now consider a linear pd-relation $g(u) = a + bu$, where $a \geq 0$ and $a + b \geq 0$ since the pd-relation is nonnegative. Note that by definition of $\mathcal{S}_2$, there is a production policy that satisfies the maximum production constraint (7). The special case $b = 0$ implies that all production rates result in the same deterioration rate, which in turn implies that producing at the maximum rate is feasible in this case. If producing at the maximum rate is not feasible, then the maximum deterioration constraint must be binding (see Lemma 2a). Substituting the linear pd-relation g into the maximum deterioration constraint (7) gives

$$\sum_{i \in \mathcal{I}} \hat{u}_i = \frac{n}{b} \left(\frac{x_2 - x_1}{t_2 - t_1} - a \right). \tag{10}$$

Equation (10) provides a necessary condition for optimality, which does not need to be sufficient. However, all policies that satisfy this necessary condition clearly result in the same profit, and therefore all are optimal. Hence, given that the system is in a functioning state $(t, x) \in \mathcal{S}_2$ and the decision maker wants to avoid failure, producing at the maximum rate is optimal if feasible, and otherwise any policy that keeps the deterioration level just below the failure level at the moment of maintenance is optimal.

4.4.   Optimal Policy with Deliberate Failure

As mentioned at the start of Section 4.3, the optimal policy that avoids failure is clearly also the overall optimal policy starting from states in $\mathcal{S}_1$. For the state $(t_1, x_1) \in \mathcal{S}_2$, a comparison is needed to the best policy for which the system fails, which is done in this section.

First notice that, for systems with $g_{\min} = 0$, a failure can be avoided by switching to the production rate corresponding to $g_{\min}$ just before the system fails. The resulting production loss is negligible while the savings in maintenance costs are not. It immediately follows that for such systems, the optimal policy always prevents a failure and thus we can apply the propositions from Section 4.3. In the remainder of the subsection, we therefore consider a system with $g_{\min} > 0$.

For these systems there are states for which failure cannot be avoided and thus the set $\mathcal{S}_3$ is nonempty. Observe that the system's state always enters $\mathcal{S}_3$ before it fails. We let $(t_2, x_2) \in \mathcal{S}_3$ refer to the state at which the system's state transits from $\mathcal{S}_2$ to $\mathcal{S}_3$.

The problem can be seen as an optimal stopping problem in which the stopping time equals the failure time. Let the decision variable t_{fail} denote the time of failure. We have $t_{\text{fail}} \geq t_1 + (L - x_1)g_{\max}$ since the system cannot fail earlier. Furthermore, for a given t_{fail}, the properties derived in Section 4.3 can be used by using the maximum deterioration constraint with $X(t_{\text{fail}}) \leq L$.

Suppose the pd-relation is strictly convex. Then Proposition 1 implies that the optimal rate after the system transits to $\mathcal{S}_3$ equals $u^*(\tau) \in U_{\text{eff}}$ for $\tau \geq t_2$, which is unique by Lemma 1. Lemma 4 states that, for strictly convex pd-relations, the optimal rate between two functioning states in $\mathcal{S}$ is constant over time and thus the optimal rate cannot change at time t_2. Hence, for strictly convex pd-relations, the unique optimal policy in case of a deliberate failure is to produce at the most efficient rate, that is $u^*(\tau) \in U_{\text{eff}}$ for $\tau \geq t_1$.

Suppose the pd-relation is strictly concave. Then by Lemma 1, the optimal policy produces at the maximum rate as soon as the state enters $\mathcal{S}_3$, that is $u^*(\tau) = 1$ for $\tau \geq t_2$. Proposition 3 implies that the optimal policy to move from (t_1, x_1) to (t_2, x_2) is to produce at the maximum rate for as long as possible and then switch off the system. The revenue produced after time t_2 decreases as we postpone the time the system enters $\mathcal{S}_3$. Furthermore, the revenue produced between (t_1, x_1) and (t_2, x_2) is constant since the optimal policy produces at the maximum rate until the state hits the boundary between $\mathcal{S}_2$ and $\mathcal{S}_3$. Hence, for strictly concave pd-relations, the unique optimal policy in case of a deliberate failure is $u^*(\tau) = 1$ for $\tau \geq t_1$.

Suppose the pd-relation is linear. Then the same argument as for strictly concave pd-relations holds and it follows that the unique optimal policy in case of a deliberate failure is $u^*(\tau) = 1$ for $\tau \geq t_1$.

4.5. Illustrative Example

We end the section with an illustration and discussion of the optimal policy for both strictly concave and strictly convex pd-relations. We focus on the structure of the optimal policy and postpone the discussion of the cost savings compared to producing at the maximum rate to Section 5.

We study a system with failure level $L = 10$, where maintenance takes place after $T = 10$ time units. The revenue per time unit is $\pi = 2.5$ when producing at the maximum rate. The preventive and corrective maintenance costs are $c_{\mathrm{pm}} = 10$ and $c_{\mathrm{cm}} = 15$, respectively. We consider the pd-relation $g(u) = \mu_{\min} + (\mu_{\max} - \mu_{\min}) u^{\alpha}$, which is concave for $0 < \alpha < 1$ and convex for $\alpha > 1$. The parameters $\mu_{\min}$ and $\mu_{\max}$ describe the deterioration rate when the system is idle or producing at the maximum rate, respectively. In this example we use $\alpha = 0.5$ (concave) and $\alpha = 1.6$ (convex), $\mu_{\min} = 0.4$, and $\mu_{\max} = 1.5$. The optimal policy for both the convex and the concave pd-relation are depicted in Figure 3. The optimal production rate is indicated by grey scale, which ranges from white (maximum rate) to black (idle). In the figures, we partition the set $\mathcal{S}_2$ (in which failure is possible but avoidable) into $\mathcal{S}_{2A}$ in which failure is avoided and $\mathcal{S}_{2B}$ in which the failure is not prevented.
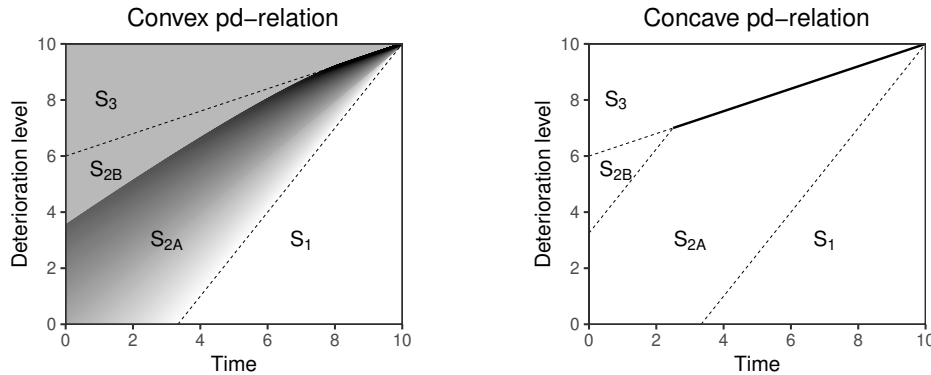

Figure 3 Optimal policy for a convex and concave pd-relation. The optimal production rate is indicated by grey scale, which ranges from white (maximum rate) to black (idle). The dashed lines represents the boundaries between the different subsets of $\mathcal{S}$.

First, consider the system with a strictly convex pd-relation ($\alpha = 1.6$). For any given state, the optimal policy is to produce at a constant rate until maintenance or a failure occurs. However, depending on the time and deterioration level, any rate can be optimal. Within $\mathcal{S}_1$ the system is always functioning at the moment of maintenance and the optimal policy is to produce at the maximum rate. In $\mathcal{S}_3$ and $\mathcal{S}_{2B}$, failure is not prevented and the most efficient rate $u_{\mathrm{eff}} \approx 0.7$ is used in order to maximize the production. For any state in $\mathcal{S}_{2A}$, failure is prevented by producing at the highest constant rate for which the system does not fail. Moreover, within $\mathcal{S}_{2A}$ there are states for which a lower production rate both prevents failure and increases the production compared to producing at the maximum rate.

Now consider the optimal policy for the strictly concave pd-relation ($\alpha = 0.5$) reflected on the right hand side of Figure 3. The optimal policy is either to produce at the maximum rate until maintenance or failure occurs, or to switch off the system in order to avoid failure. Within $\mathcal{S}_1$ the system cannot fail and production is maximized by producing at the maximum rate until maintenance is performed. Within $\mathcal{S}_3$ and $\mathcal{S}_{2B}$, failure is not prevented and the system produces at the most efficient rate, which is the maximum rate for any concave pd-relation. Within $\mathcal{S}_{2A}$, failure is prevented by producing at the maximum rate as long as possible and then switch off the system.

5. Stochastic Deterioration

In this section we show that key insights obtained from the deterministic system carry over to more realistic systems with stochastic deterioration. We first discuss the similarities and differences

between the optimal policies for stochastic and deterministic systems. Thereafter, we illustrate the benefits of using condition-based production decisions by aid of a numerical example. The results are obtained by formulating the system as a Markov decision process (MDP). Based on the differences between the optimal policies for deterministic and stochastic systems, we then propose two heuristic adaptations of the optimal deterministic policy for use in a stochastic setting.

5.1. Markov Decision Process

An MDP is defined by a set of states, a set of possible actions for each state, and state and action dependent transition probabilities and rewards. We have a finite horizon problem and use backward induction to determine optimal policies for the MDP (see Puterman 1994).

We first discretize the state space, the time horizon, and the range of production rates. The state space $[0, L]$ is partitioned into n equally sized intervals of length $\Delta X = L/n$, and is then discretized to the ordered set of midpoints $\{(i - 0.5)\Delta X \mid i = 1, \ldots, n\}$ of these intervals. All deterioration levels above L are combined into a single state with index $n + 1$ that represents the failed state. Time is discretized to $\{i\Delta t \mid i = 0, \ldots, m\}$ with $\Delta t = T/m$. The discrete production rates are constrained to the set $\{i/\eta \mid i = 0, \ldots, \eta\}$ so that the rates are equally distributed over the interval $[0, 1]$, including the boundaries.

We let $F_{u,\Delta t}$ denote the distribution function of the additional amount of deterioration during a single time period with length Δt when producing at rate u. To obtain the transition probabilities in the discretized process, we model the probability of staying in the same deterioration state as $F_{u,\Delta t}(0.5\Delta X)$. The probability of moving from state k to state $k + i$, where $i \geq 1$ and $k + i \leq n$, equals $F_{u,\Delta t}((i + 0.5)\Delta X) - F_{u,\Delta t}((i - 0.5)\Delta X)$. The probability of moving from the kth state to the failed state $n + 1$ equals $1 - F_{u,\Delta t}((i - 0.5)\Delta X)$ where $i = n - k + 1$. Summarizing, the transition probability matrix of the discrete deterioration process can be written as

$$P_u(k, k+i) = \begin{cases} 0 & \text{if } i < 0, \\ F_{u,\Delta t}(0.5\Delta X) & \text{if } i = 0, \\ F_{u,\Delta t}((i + 0.5)\Delta X) - F_{u,\Delta t}((i - 0.5)\Delta X) & \text{if } 0 < i < n - k + 1, \\ 1 - F_{u,\Delta t}((i - 0.5)\Delta X) & \text{if } i = n - k + 1. \end{cases}$$

5.2. Base System

As a base case for our numerical analysis, we consider a system with the following parameter values (see also Table 1). Maintenance is scheduled at time $T = 100$ and the system has failure level $L = 100$. At the maximum rate, the generated revenue per time unit is $\pi = 0.1$. The preventive maintenance cost is $c_{\mathrm{pm}} = 2$ and the corrective maintenance cost is $c_{\mathrm{cm}} = 6$. To approximate the continuous deterioration process, we partition the state space and the time horizon into small subintervals with respective lengths $\Delta X = 0.1$ and $\Delta t = 1$. At each decision epoch, the decision maker can choose from $\eta = 100$ different production rates. We consider the same parametric form $g(u) = \mu_{\min} + (\mu_{\max} - \mu_{\min})\,u^{\alpha}$ for the pd-relation as in Section 4.5, and we consider $\alpha = 0.5$ (concave) and $\alpha = 3$ (convex), $\mu_{\min} = 0.15$, and $\mu_{\max} = 0.8$.

We model the underlying continuous deterioration process by a gamma process. This process is appropriate for modeling monotonically increasing deterioration processes such as wear, erosion, and fatigue (Van Noortwijk 2009, Alaswad and Xiang 2017). We use the same parametric form with a shape and a scale parameter for the gamma process as De Jonge et al. (2017). To relate the scale and shape parameter of the gamma process to the pd-relation g, we impose the following three properties. First, when producing at rate u, the additional amount of deterioration per time unit has mean $g(u)$. Second, when producing at the maximum rate, the standard deviation of the deterioration increments equals $\sigma_{\max}$. Third, the coefficient of variation of the deterioration increments is the same for all production rates. For the given parametric form of the pd-relation, these properties are obtained by setting the shape parameter equal to $k = \mu_{\max}^2/\sigma_{\max}^2$ and the scale parameter as a function of the production rate equal to $\theta(u) = g(u) \cdot \sigma_{\max}^2/\mu_{\max}^2$. We note in passing that we also modeled the deterioration process as a compound Poisson and as a Brownian motion with positive drift. As the results were comparable, we did not include these.

Table 1   **Base system used in the numerical analysis**

Parameter	Value	Interpretation
T	100	Length time horizon
L	100	Failure level
η	100	Number of production rates
π	0.1	Revenue per time unit at maximum rate
c_{pm}	2	Preventive maintenance cost
c_{cm}	6	Corrective maintenance cost
α	0.5, 3	Shape pd-relation
μ_{min}	0.15	Deterioration rate when idle
μ_{max}	0.80	Deterioration rate at maximum production rate
σ_{max}	$\sqrt{5}$	St.dev. deterioration increment at maximum rate
Δt	1.0	Partitioning size time horizon
ΔX	0.1	Partitioning size deterioration level

5.3. Structure of Optimal Policy

We compare the structure of the optimal policy for stochastic systems with the optimal policy for deterministic systems as derived in Section 4. The optimal policies for both the deterministic and the stochastic system with both a concave and a convex pd-relation are shown in Figure 4. The production rate is indicated by grey scale, which ranges from white (maximum rate) to black (idle). We note that we have carried out a more extensive comparison in an online appendix to this paper.

The optimal policies for the concave pd-relation are shown in Figures 4a and 4b. For both systems the optimal policy is to either produce at full speed or not to produce at all. We also see that under both policies the system is switched off at the latest possible moment in time. Although the policies are similar, there is one structural difference. The deterministic system is only switched off on a line segment whereas the stochastic system is switched off within a larger bandwidth that mainly lies below the line segment of the deterministic system. This larger bandwidth is due to two reasons. First, also when the system is idle, the deterioration process is stochastic, and thus the deterioration process would immediately move off a line segment. Given the bandwidth, it is very likely that when the system is turned off, it will remain idle until the moment of maintenance. Second, a jump process such as the gamma process would jump over a line segment. Notice that since the uncertainty of the future deterioration trajectory decreases as we approach the moment of maintenance, the bandwidth decreases at the end of the time horizon.

The optimal policies for the convex pd-relation are shown in Figures 4c and 4d. We directly observe the similar structure for the two systems, which is also confirmed by the additional comparisons in the online appendix. For high deterioration levels (the areas in the top-left part of the figures above the darkest areas), both the deterministic and the stochastic system maximize the expected production by producing at the most efficient rate, i.e., the rate that maximizes $u/g(u)$. For low deterioration levels, both systems maximize production by producing at the maximum rate. For intermediate deterioration levels, both systems reduce the production rate. Hereby, the stochastic system reduces the risk of a failure and the deterministic system avoids the failure with certainty. However, we also observe two structural differences between the deterministic and the stochastic system. First, for the intermediate deterioration levels the stochastic system produces at a slightly lower rate than the deterministic system as it needs some safety margin in order to deal with the uncertainty of the deterioration process. This also explains the reduced production rate for the stochastic system at the end of the time horizon for highly deteriorated states. The second difference is that, for the deterministic system, there is a sudden transition from producing slowly to producing at the most efficient rate when we go from the region where failure is avoided to the region where failure is not prevented (because this either would cause too much production losses or the failure is unavoidable). These separate regions cannot be distinguished for the stochastic system, which results in a gradual change in the production rate for this system. To lower the risk of failure, the production rates for the stochastic system are also somewhat lower around this transition.

5.4. Cost Savings by Condition-Based Production

In order to assess the benefits of an adjustable production rate based on available condition information, we will compare the optimal policy to three benchmark policies that differ in their degree
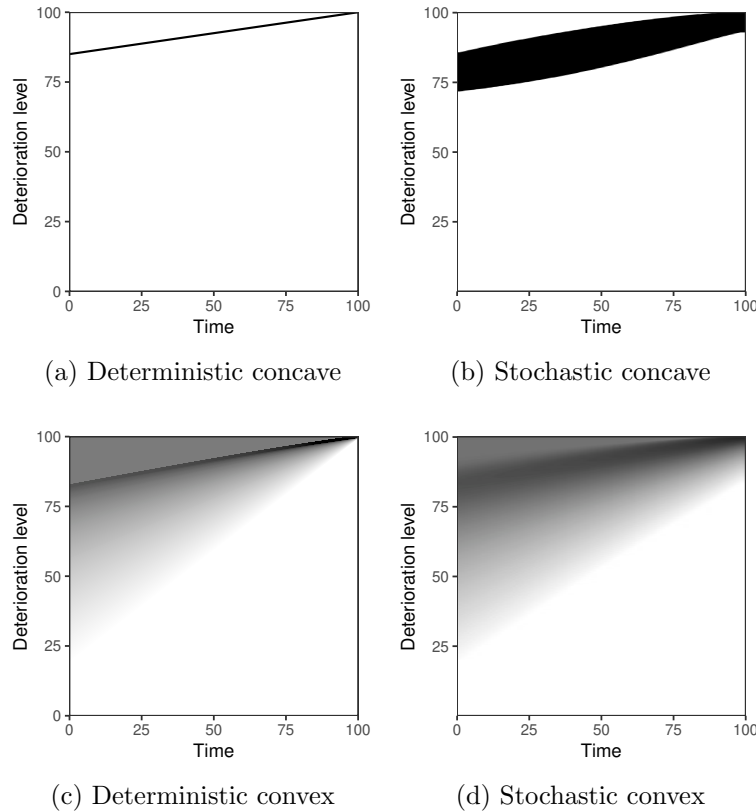
(a) Deterministic concave

(b) Stochastic concave

(c) Deterministic convex

(d) Stochastic convex

Figure 4   **Optimal policies for both a concave and a convex pd-relation. White indicates that the system produces at its maximum rate and black indicates that the system does not produce.**

of flexibility. The *max-rate policy* has no adjustable production rate and the system produces at the maximum rate until failure occurs or maintenance is performed. The *fixed policy* allows the decision maker to set a single production rate at the start of the time horizon, i.e., $t = 0$, which cannot be changed afterwards. This policy does not use condition monitoring but can use the system characteristics such as the pd-relation. The *on-off policy* allows the decision maker to choose a single production rate at the start of the time horizon. During the time horizon, the decision maker can switch between this fixed rate and turning the system off. This policy uses both condition monitoring and the knowledge on the pd-relation. The *optimal policy* allows the decision maker to set any production rate at any time.

To numerically compare the performance of the policies we mainly focus on $\bar{J}(p) = \pi T - J(p)$, where πT is the revenue that can be attained when it would be possible to always produce at the maximum rate, and $J(p)$ is the expected profit obtained under policy p. The function $\bar{J}(p)$ can be interpreted as the total cost, consisting of the maintenance cost and the loss of revenue compared to always producing at the maximum rate. This allows a more clearcut evaluation of comparative policy performance than focusing on the total profit, since all considered policies produce at (almost) the maximum rate and thus all policies generate (almost) maximum revenue for most of the planning horizon. For completeness, however, we also report the expected profit, the expected production, the failure probability, and the standard deviations of the expected profit and production.

Table 2 shows the performance of the four policies when these are applied to the base system described in Section 5.2. We first consider the system with the convex pd-relation ($\alpha = 3.0$). Compared to the max-rate policy, the optimal policy significantly reduces the probability of a failure, namely from 16.81% to 1.66%, while the drop in expected production is only 0.6% (from 96.63 to 96.02). By better balancing the production output and the failure risk, the optimal policy is able to reduce total cost by 18.9% (from 2.91 to 2.36). We also observe lower standard deviations of both the profit (0.85 instead of 2.08) and the realized production (5.61 instead of 7.17), implying a more reliable production system.

Table 2 **Performance measures for the base system for various policies.**

		Convex			Concave		
	Max-rate	Fixed	On-off	Optimal	Fixed	On-off	Optimal
Expected profit	6.99	7.05	7.26	7.54	6.99	7.26	7.26
St. dev. profit	2.08	1.65	1.53	0.85	2.08	1.53	1.53
Expected production	96.63	94.69	94.46	96.02	96.63	94.46	94.46
St. dev. production	7.17	5.12	10.58	5.61	7.17	10.58	10.58
Failure probability (%)	16.81	10.55	4.61	1.66	16.81	4.61	4.61
Total cost	2.91	2.85	2.64	2.36	2.91	2.64	2.64
- Maintenance cost	2.67	2.42	2.18	2.07	2.67	2.18	2.18
- Revenue loss	0.24	0.43	0.46	0.30	0.24	0.46	0.46

The more restrictive on-off policy is also able to significantly reduce the failure probability to 4.61%. However, this policy has a lower expected production and is much less able to reduce the uncertainty of the expected profit and actually increases the standard deviation of the expected production. The higher uncertainty is due to the all or nothing nature of this policy. In most scenarios the policy does produce at the fixed rate. However, in case the system deteriorates faster than expected, the policy can only react by switching off the system completely while the optimal policy can respond subtler by slightly reducing the production rate.

Now consider the concave pd-relation ($\alpha = 0.5$). The performance of the optimal policy and of the on-off policy are the same because the former only uses the minimum and maximum rate as seen in the previous section. Furthermore, the total cost is reduced by 9.3% (from 2.91 to 2.64) and thus an adjustable production rate is less effective for concave pd-relations than for convex pd-relations. This is because convex pd-relations have production rates that are more efficient than the maximum rate while for a concave pd-relation the most efficient rate always is the maximum rate (see Lemma 1). For the same reason, the fixed policy does not outperform the max-rate policy for concave pd-relations.

5.5. Parameter Sensitivity

We continue by analyzing the effect of different system parameters on the performance of the production policies. The results in this section are obtained by deviating one parameter value at a time compared to the base system with the convex pd-relation ($\alpha = 3.0$). For the first parameter, we discuss the cost savings as well as the underlying trade-off between the failure risk and the expected production. Because this trade-off is similar for all parameters, we only show the effect on the resulting cost savings for the other parameters.

Let us first consider the influence of the revenue per time unit π when producing at the maximum rate. When this revenue is small compared to the maintenance cost, the main priority of the decision maker is to avoid a possible failure. Instead, when the production revenue is large, the main goal is to maximize production. The effect of the revenue parameter on the failure probability and the expected production for the four different policies is depicted in Figure 5. The relative cost savings of the fixed, on-off, and optimal policy compared to the max-rate policy are shown in Figure 6a.

When the revenue parameter is low, all policies are able to realize a significant cost saving by reducing the production and thereby the failure probability. However, when the revenue parameter increases, the cost savings of the fixed and on-off policy rapidly diminish whilst the optimal policy results in a cost saving for all values of the production revenue. The optimal policy reduces cost even for high revenues since, in case the system deteriorates faster than expected, this policy can increase production by postponing the failure.

For almost all values of the revenue parameter, the optimal policy has both the lowest failure probability and the highest expected production. At first this seems counterintuitive, since the failure probability is reduced by lowering the production rate. However, the optimal policy can postpone the decision to focus either on avoiding a possible failure or on maximizing production until condition information becomes available.
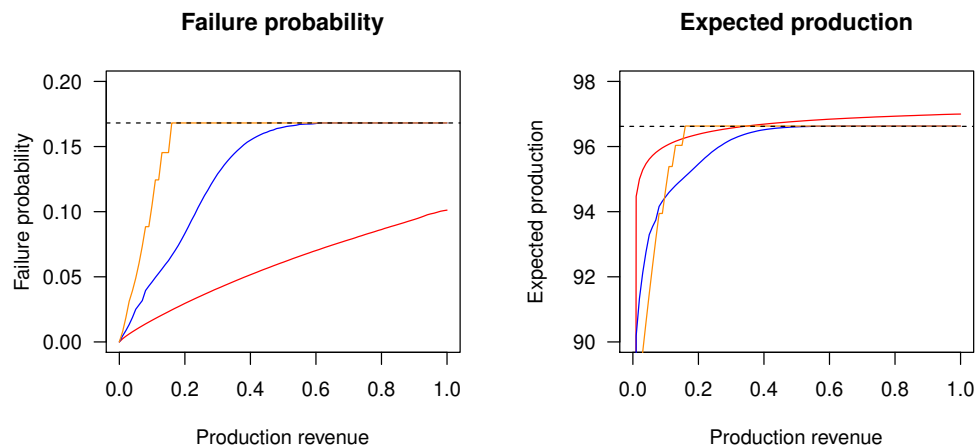
Failure probability **Expected production**

Figure 5 **Effect of the production revenue π on the failure probability and the expected production of the max-rate (- - -), fixed (—△—), on-off (—◇—), and optimal (—○—) policy.**
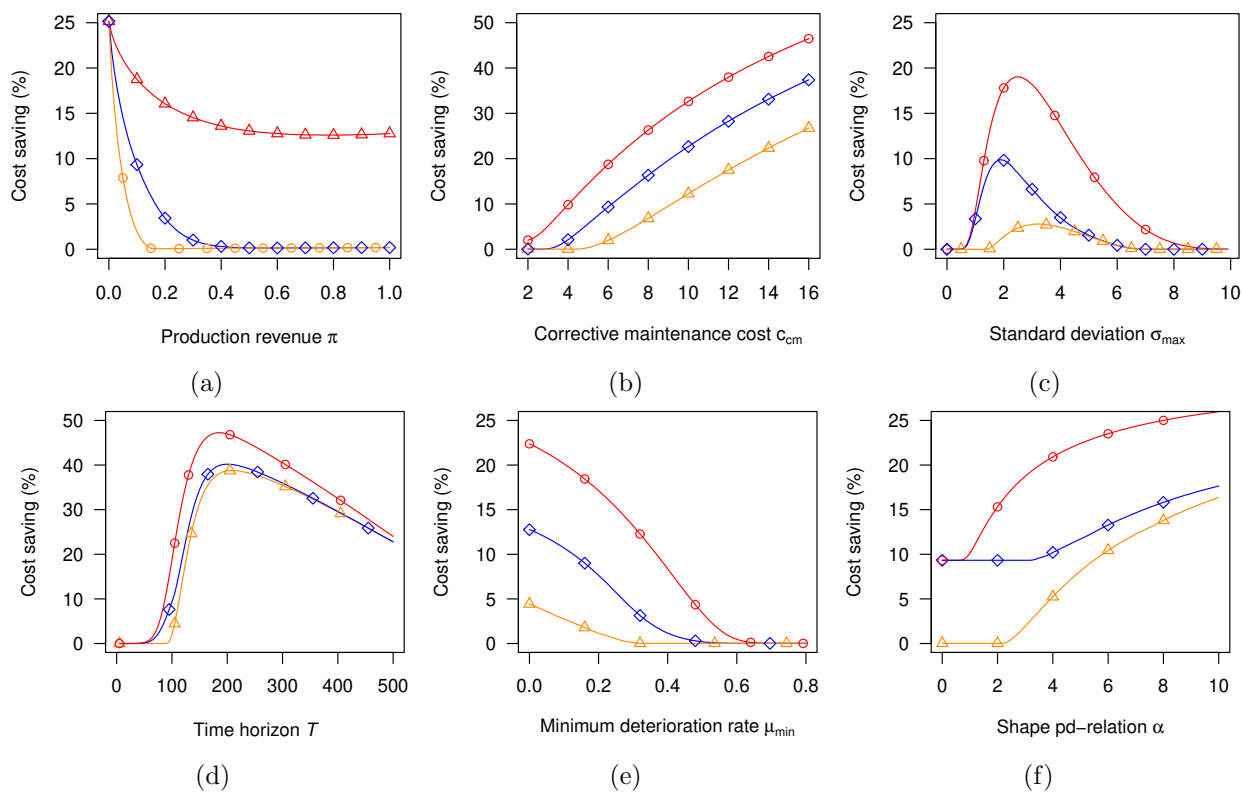

(a) (b) (c)

(d) (e) (f)

Figure 6 **Effect of different parameters on the relative cost saving compared to the max-rate policy for the fixed (—△—), on-off (—◇—), and optimal (—○—) policy.**

Figure 6b depicts the effect of the cost of corrective maintenance on the relative cost savings. We see that the benefit of the adjustable production rate increases as corrective maintenance becomes more expensive compared to preventive maintenance. Furthermore, the optimal policy realizes a cost saving even when preventive maintenance and corrective maintenance have the same cost (recall $c_{pm} = 2$), since the optimal policy can also increase the expected production.

Next, we consider the volatility of the deterioration process and its effect on the cost savings compared to the max-rate policy, as shown in Figure 6c. For $\sigma_{max} = 0$, the deterioration process is deterministic and the system will be functioning at the turnaround when producing at the maximum

rate over the whole time horizon. It is optimal to produce at the maximum rate and therefore condition-based production decisions do not result in cost savings. For relatively small values of $\sigma_{\max}$, the deterioration process has many small jumps of different sizes. The condition-based production policies can effectively react to these jumps by adjusting the production rate after the occurrence of a jump. For large values of $\sigma_{\max}$, the size of the deterioration increments increases and failure is most likely caused by a single jump. Condition monitoring cannot provide information on the timing of this jump and therefore the value of condition-based production diminishes as the volatility of the deterioration process increases.

Figure 6d shows the effect of the time horizon on the cost saving compared to the max-rate policy. The largest cost savings are realized for time horizons of moderate lengths. For short time horizons, any policy produces at the maximum rate during the entire horizon, implying that they are all equivalent. For extremely long time horizons, failure cannot be avoided, and all policies will improve the production revenues by producing at the most efficient rate.

The effect of the pd-relation parameters $\mu_{\min}$ and α are given in Figures 6e and 6f. The deterioration rate in the idle mode ranges from $\mu_{\min} = 0$ (no deterioration when idle) to $\mu_{\min} = \mu_{\max} = 0.8$ (same deterioration rate for all production rates). We see that the optimal policy results in a significant cost saving, even if the deterioration rate in idle mode is half of that in full mode (e.g., due to exogenous conditions like weather). The fixed policy only results in a small cost saving if $\mu_{\min}$ is low. For $\alpha < 1$, the pd-relation is concave, implying that the optimal policy is an on-off policy (see Section 5.3). The fixed policy only results in a cost saving for larger values of α. Furthermore, for all convex pd-relations ($\alpha > 1$), the optimal policy performs significantly better than the fixed policy and the on-off policy.

5.6. Heuristics based on Deterministic Deterioration

The results until now have shown that, in settings with stochastic deterioration, the operational performance can be considerably improved by using the optimal policy compared to simpler, less flexible policies that we have considered so far. However, determining this policy can be computationally expensive and its complexity might hinder practical implementation. In the previous sections, we observed that the deterministic and stochastic policies share many similarities. Building on this observation, we construct two simple heuristics for the stochastic system. We focus on convex pd-relations, as those are arguably the most realistic, but a similar approach can be used to develop heuristics for concave pd-relations.

The main difference between the optimal policy for the deterministic and for the stochastic case is that the latter produces at a slightly lower rate for intermediate deterioration levels (see Section 5.3). The main premise of the heuristics is therefore to produce at a slightly lower rate than the optimal rate for the deterministic system, which is denoted by $u(t, x)$. The first heuristic achieves this in the most straightforward way, by subtracting some constant α_1 from the optimal production rate, i.e., $u_1(t, x) = u(t, x) - \alpha_1$ (negative rates are set to 0). A downside of doing this is that the system never operates at the maximum speed, even if it is in a very good state (given the time left until maintenance). To avoid this, but still build in safety, the second heuristic adds some constant α_2 to the current deterioration level and then computes the production rate, i.e., $u_2(t, x) = u(t, x + \alpha_2)$, where $x + \alpha_2$ is capped by failure level L. The heuristics select a production rate for each state and thus the corresponding expected cost can be evaluated by solving a finite time Markov chain.

The effectiveness of the heuristics is analyzed by using the base case and deviating various system parameters one by one. The cost savings of both heuristics, with the parameters α_1 or α_2 optimized per instance, compared to the max-rate policy are shown in Figure 7. We clearly see that the second heuristic outperforms the first heuristic. Apparently, it is indeed important to be able to produce at maximum speed if the system is in a relatively good state. The second heuristic has a near-optimal performance and strongly outperforms the fixed and the on-off introduced in the previous sections. Since the heuristic is only a simple modification of the deterministic policy, its performance underlines our intuition that the structure of the optimal policy for the deterministic systems (partially) carries over to stochastic systems.

We remark that applying these heuristics still requires optimization of the policy parameters, which is not straightforward. Figure 8 shows the cost savings realized by the second heuristic as a
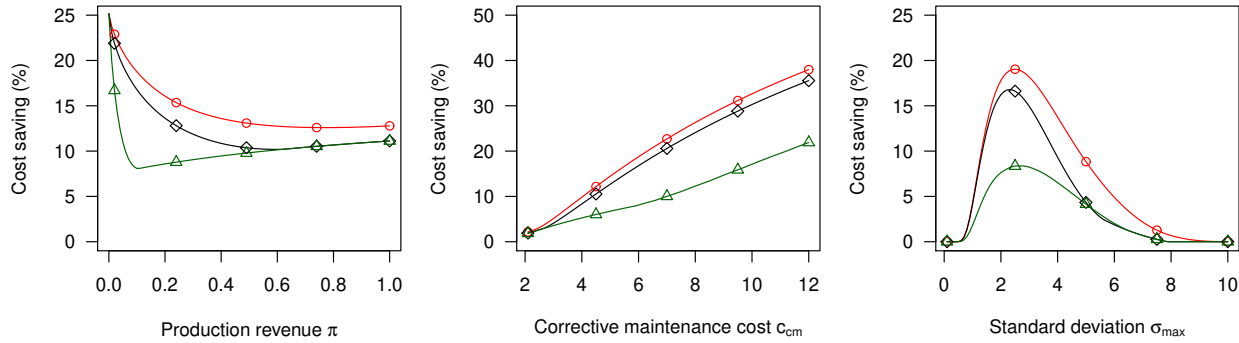
Figure 7 **Effect of various system parameters on the relative cost saving compared to the max-rate policy for the optimal policy (—○—), and the first (—△—) and second (—◇—) heuristic.**
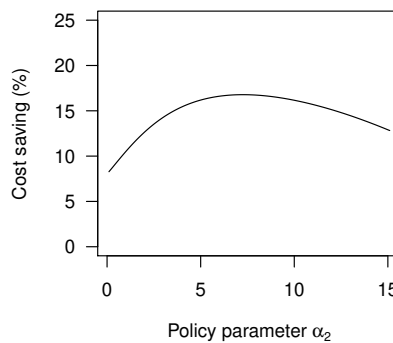

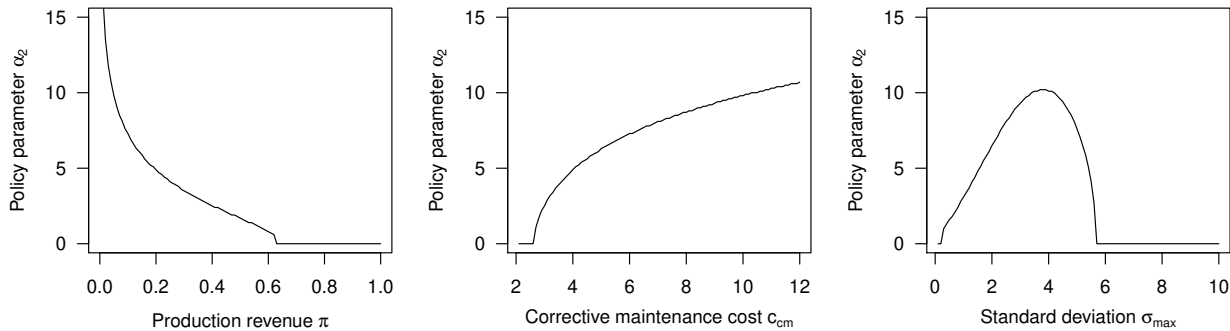**Figure 8** **Relative cost savings of the second heuristics as function of the policy parameter α_2 compared to the max-rate policy.**

Figure 9 **Optimized policy parameter α_2 for the second heuristic.**

function of the policy parameter α_2 when we apply this heuristic to the base case. We observe that the relative cost saving is a convex function of the policy parameter. Furthermore, the heuristic results in a considerable cost saving for all reasonable values of α_2. For large values, the heuristic is too conservative, resulting in an unnecessary low failure risk by sacrificing too much production. For small values, the heuristic is too optimistic, resulting in high production revenues but high maintenance costs as well. Moreover, for $\alpha_2 = 0$ the heuristic exactly equals the deterministic policy and thus even applying the optimal deterministic policy results in a cost saving of 8.3%. The graph only shows the result for the base case, however, similar results are observed for other problem instances.

Figure 9 shows how the optimal value of α_2 is affected by the system parameters π, c_{cm}, and σ_{max} (only α_2 is shown in order to keep the exposition concise). We observe that the heuristic is more conservative when the production revenue is low or when the corrective maintenance cost is high.

The effect of the volatility of the deterioration process is more complex. At first, with increasing volatility, the heuristic becomes more conservative. However, if the deterioration process becomes very volatile, it is better to always produce at a high rate as a failure is most likely caused by a large deterioration jump that cannot be prevented by producing at a slightly lower rate. We conclude that the relation between the system parameters and the optimal heuristic parameter is complex. Approximate closed-form relations can be developed and tested to ease the implementation of the heuristic, but we consider this to be outside the scope of this study.

6. Conclusion

This study is the first to introduce condition-based production rate decisions that affect the deterioration rate of a system. We recognize that the advent of inexpensive sensor technology and the recent advances in the Internet of Things offer opportunities to remotely monitor the equipment's deterioration level and to control its usage in real-time. Based on the available condition information, the adjustable production rate can be used to control the deterioration process and thereby we can balance the risk of a failure with the production revenues.

Exact analytical solutions are derived for deterministic deterioration processes, which reveal several structural insights. Firstly, for all systems it can be beneficial to avoid a failure by reducing the production rate. The optimal production policy depends on the specific system, and in particular on the relation between the production rate and the deterioration rate. Secondly, if a failure cannot be avoided, the production can be increased by producing at a more efficient rate. Likewise, even if a failure can be prevented, it is sometimes better to maximize production and thereby let the system fail. Thirdly, there exist win-win scenarios in which production rate adjustments both prevent a failure and increase the production.

The numerical analysis, based on a Markov decision process formulation of the problem, shows that the structural insights largely carry over to stochastic systems. Optimizing production rates based on condition information reduces the total cost by up to 50% for the considered cases. Simpler policies, such as an on-off policy that only switches between a single fixed rate and the idle mode, are much less effective in reducing cost. We conclude that using condition monitoring to dynamically adjust production rates over time provides significant opportunities to improve the operational efficiency of production systems. The profitability of the system increases by reducing the expected maintenance cost while increasing the expected production. Furthermore, although the optimal dynamic policy may be complex, we also developed a simple heuristic that performs well and is much easier to implement.

Based on the promising results of our exploratory study, we conclude that there is ample scope for further research with three main avenues. The first avenue is to jointly and dynamically optimize maintenance timing and production rates based on condition information. Our results show that using condition-based production rates reduces the uncertainty of the deterioration process since one can respond to the volatility of the process by dynamically adjusting the production rate. This reduced uncertainty provides opportunities for maintenance policies to be less conservative by scheduling fewer maintenance actions. For example, a lower maintenance frequency may be used for block-based maintenance policies, and the threshold maintenance age could be increased for age-based maintenance policies.

Secondly, one could consider multi-unit systems where the production per time unit must be in a specified range (e.g., due to given supply contracts). Examples of such systems include offshore wind farms (commonly consisting of many turbines) that have to produce a minimum amount of electricity, and compressors in a gas network that together have to maintain a reliable gas pressure. For such systems, condition-based production rates create opportunities to improve the clustering of maintenance actions for several units. For example, one can decelerate the production speed of highly deteriorated turbines and accelerate it for turbines in a good condition. Thereby, the deterioration processes of these turbines are better synchronized and their maintenance can be clustered to reduce cost.

The third avenue is to elaborate on the single-unit system either by incorporating settings commonly seen both in practice and in the maintenance literature, such as fluctuating production

revenues, imperfect condition monitoring, uncertain failure levels, aperiodic inspections and preventive repair costs that depend on the deterioration level, or by studying analytical properties for the system with stochastic deterioration. For future research on energy production, considering price fluctuations seems particularly promising as price changes occur very frequently, and in fact even negative prices exist in times with overproduction. The operational efficiency may improve by producing at a high rate when prices are high and switch off the system when prices are low. Lastly, an interesting research direction is to consider deterioration processes that not only depend on the current production rate but also on the age of the system, the current deterioration level, and environmental conditions such as weather.

Appendix A: Lemmas

LEMMA 1. *Consider a pd-relation g with $g_{\min} > 0$.*

a) *If the pd-relation is strictly concave or linear, then the set of efficient rates is $U_{\text{eff}} = \{1\}$.*

b) *If the pd-relation is strictly convex, then there is only one efficient rate. The efficient rate can be found by first solving $z = \arg\{g(u) = u\,g'(u)\}$ and consequently setting $U_{\text{eff}} = \{\min(1, z)\}$.*

Proof.   The set of efficient production rates is defined as $U_{\text{eff}} = \arg\max\{u/g(u)\}$. Using the quotient rule we get

$$\frac{\mathrm{d}}{\mathrm{d}u} \frac{u}{g(u)} = \frac{g(u) - g'(u)u}{g(u)^2}.$$

The denominator is clearly always positive and thus the sign of the derivative is determined by the sign of the numerator. We define $k(u) = g(u) - u\,g'(u)$ and thereby $k'(u) = -u\,g''(u)$.

For strictly concave pd-relations we have $g''(u) < 0$ and thus $k'(u) \geq 0$ for $0 \leq u \leq 1$. Furthermore, $g_{\min} > 0$ implies $g(0) > 0$ and thereby $k(0) > 0$. Combining these two observations implies $k(u) > 0$ for $0 \leq u \leq 1$. It follows that the derivative of $u/g(u)$ is always positive. We conclude that $u/g(u)$ is maximized by the maximum production rate and thus for strictly concave pd-relations we have $U_{\text{eff}} = \{1\}$.

For linear pd-relations $g(u) = a + bu$, where $a > 0$ and $a + b \geq 0$ since $g_{\min} > 0$, we have $k(u) = a$ and $k'(u) = 0$. It immediately follows that $u/g(u)$ is maximized by the maximum production rate and thus for linear pd-relations we have $U_{\text{eff}} = \{1\}$.

For strictly convex pd-relations we have $g''(u) > 0$ and thus $k'(u) < 0$ for $u > 0$. Next, $g_{\min} > 0$ implies $g(0) > 0$ and thereby $k(0) > 0$. Combining $k(0) > 0$ and $k'(u) < 0$ implies that the derivative of $u/g(u)$ is first positive and at some point becomes negative. Hence, there is only one rate that maximizes $u/g(u)$. $\square$

LEMMA 2. *Suppose we have a pd-relation g, the system is in a functioning state $(t_1, x_1) \in \mathcal{S}$, and there is a maximum deterioration constraint $X(t_2) \leq x_2$ where $x_1 < x_2 < L$ and $t_1 < t_2 \leq T$. Let $\hat{\boldsymbol{u}}^* = (\hat{u}_1, \ldots, \hat{u}_n)$ be an optimal policy, where $\hat{u}_i$ denotes the production rate in time interval $i \in \mathcal{I} = \{1, \ldots, n\}$.*

a) *If there is an $i \in \mathcal{I}$ such that $\hat{u}_i < 1$, then the maximum deterioration constraint is binding.*

b) *If the policy $\hat{\boldsymbol{u}} = (1, \ldots, 1)$ is feasible, then it is the unique optimal policy.*

c) *For all $i, j \in \mathcal{I}$ for which $\hat{u}_i, \hat{u}_j \in (0, 1)$, we have $g'(\hat{u}_i) = g'(\hat{u}_j)$.*

d) *For all $i \in \mathcal{I}$ for which $\hat{u}_i < 1$, we have $g'(\hat{u}_i) > 0$.*

e) *If g' is a one-to-one function, then for all $i, j \in \mathcal{I}$ for which $\hat{u}_i, \hat{u}_j \in (0, 1)$, we have $\hat{u}_i = \hat{u}_j$.*

Proof.   (a) Consider a policy $\hat{\boldsymbol{u}}$ for which $\hat{u}_i < 1$ and suppose the maximum deterioration constraint is not binding. Then the complementary slackness (8b) implies $\nu = 0$. Substituting this into (8a) gives $\lambda_i = \mu_i + 1$. From the complementary slackness given by (8c) and (8d) it follows that at least one of λ_i or μ_i equals zero. When $\lambda_i = 0$ we have $\mu_i = -1$, which violates (8e). It follows $\mu_i = 0 \;\Rightarrow\; \lambda_i = 1 \;\Rightarrow\; \hat{u}_i = 1$. However, $\hat{u}_i = 1$ contradicts with the given policy where $\hat{u}_i < 1$. Concluding, for any policy $\hat{\boldsymbol{u}}$ with elements $\hat{u}_i < 1$, the maximum deterioration constraint as given in (7) is binding.

(b) Suppose $\hat{\boldsymbol{u}} = (1, \ldots, 1)$ is feasible, then this trivially is the unique optimal policy as no policy can produce more than always producing at the maximum rate.

(c) Suppose we have $0 < \hat{u}_i < 1$ for all $i \in \mathcal{I}_1$ where $\mathcal{I}_1 \subseteq \mathcal{I}$. The complementary slackness implies $\lambda_i = \mu_i = 0$ and substituting these values into (8a) gives $\nu = 1/g'(\hat{u}_i)$. Since ν is a constant, we have $g'(\hat{u}_i) = g'(\hat{u}_j)$ for all $i, j \in \mathcal{I}_1$.

(d) Consider a policy $\hat{\boldsymbol{u}}$ for which $g'(\hat{u}_i) \leq 0$. Then the total production can be increased by marginally increasing $\hat{u}_i$ while the corresponding deterioration $g(\hat{u}_i)$ is non-increasing.

(e) This directly follows from (c). $\square$

LEMMA 3. *For all pd-relations g, there is an optimal policy with at most two production rates.*

Proof. Suppose we have a policy that uses three or more different production rates. Then we can select three arbitrary time intervals from this policy with different production rates, and improve the policy by replacing one of the rates by a combination of the other two rates. As we select the three rates arbitrarily, we can repeat the same procedure until only two rates are left. Notice that for any optimal policy which uses intermediate production rates, the maximum deterioration constraint must be binding (see Lemma 2a). Thus a policy can only improve by increasing its total production without increasing the corresponding total deterioration.

Select three arbitrary intervals with different production rates from the given policy. Assume, without loss of generality, that $u_1 < u_2 < u_3$ and $g(u_1) < g(u_2) < g(u_3)$. The duration of the intervals is denoted by τ_1, τ_2, and τ_3. In the remainder of this proof, we refer to $g(u_i)$ as g_i.

We first partition the second time interval into two subintervals with lengths $(1-\alpha)\tau_2$ and $\alpha\tau_2$. In the first subinterval, the new policy produces with rate u_1. In the second subinterval, the new policy produces with rate u_3. The deterioration of the given policy and the new policy equal $\tau_1 g_1 + \tau_2 g_2 + \tau_3 g_3$ and $(\tau_1 + (1-\alpha)\tau_2)g_1 + (\tau_3 + \alpha\tau_2)g_3$, respectively. Equating the two deterioration levels gives $\alpha = (g_2 - g_1)/(g_3 - g_1)$. Notice that $0 < \alpha < 1$ since $g_1 < g_2 < g_3$. The total production of the given and new policy equal $\tau_1 u_1 + \tau_2 u_2 + \tau_3 u_3$ and $(\tau_1 + (1-\alpha)\tau_2)u_1 + (\tau_3 + \alpha\tau_2)u_3$, respectively. Substituting α implies that the new policy produces at least as much as the given policy if

$$g_2 \geq g_1 + (u_2 - u_1)\frac{g_3 - g_1}{u_3 - u_1}.$$

Secondly, we shorten the duration of interval 1 and 3 by α and β, respectively. The duration of interval 2 is increased with $\alpha + \beta$. Equating the deterioration level of the given and the new policy gives $\alpha = \beta(g_3 - g_1)/(g_2 - g_1)$. We select β such that the length of exactly one of the intervals becomes zero and the other length remains nonnegative. Substituting α into the new total production function implies that the new policy produces at least as much as the given policy if

$$g_2 \leq g_1 + (u_2 - u_1)\frac{g_3 - g_1}{u_3 - u_1}.$$

Observe that always one of the two conditions is satisfied. It follows that for any given policy which uses more than two production rates, we can construct a policy with two production rates that produces at least as much as the given policy. We conclude that, for any pd-relation g, there is an optimal policy with at most two production rates. □

LEMMA 4. *The optimal rate is constant over time for strictly convex pd-relations.*

Proof. The proof is structured as follows. First, we show that an optimal policy with rates smaller than the maximum rate cannot contain the maximum rate. Second, we show that an optimal policy with rates larger than the minimum rate cannot contain the minimum rate. Combining this with other lemmas implies that the production rate is constant over time.

We divide $\mathcal{I}$ into two proper subsets $\mathcal{I}_1$ and $\mathcal{I}_2$ such that $\hat{u}_i < 1$ for $i \in \mathcal{I}_1$ and $\hat{u}_j = 1$ for $j \in \mathcal{I}_2$. The complementary slackness give $\lambda_i = 0$ and $\mu_j = 0$. Substituting these values into (8a) implies $\nu = (1 + \mu_i)/g'(\hat{u}_i)$ and $\nu = (1 - \lambda_j)/g'(1)$. Equating the two expressions for ν gives $\lambda_j = 1 - (1 + \mu_i)\,g'(1)/g'(\hat{u}_i)$. Since $\lambda_j \geq 0$ we must have $g'(\hat{u}_i)/g'(1) \geq 1$, which is not possible for strictly convex functions g. Hence, for strictly convex functions g, any policy that contains elements $\hat{u}_i < 1$ and $\hat{u}_j = 1$ cannot be optimal.

We divide $\mathcal{I}$ into two proper subsets $\mathcal{I}_1$ and $\mathcal{I}_2$ such that $\hat{u}_i > 0$ for $i \in \mathcal{I}_1$ and $\hat{u}_j = 0$ for $j \in \mathcal{I}_2$. The complementary slackness give $\mu_i = 0$ and $\lambda_j = 0$. Substituting these values into (8a) implies $\nu = (1 - \lambda_i)/g'(\hat{u}_i)$ and $\nu = (1 + \mu_j)/g'(0)$. Equating the two expressions for ν gives $\lambda_i = 1 - (1 + \mu_j)\,g'(\hat{u}_i)/g'(0)$. Since $\lambda_i \geq 0$ we must have $g'(0)/g'(\hat{u}_i) \geq 1$, which is not possible for strictly convex functions g. Hence, for strictly convex functions g, any policy that contains elements $\hat{u}_i > 0$ and $\hat{u}_j = 0$ cannot be optimal.

The derivative of a strictly convex pd-relation is an one-to-one function. By Lemma 2e it follows that the optimal policy can only contain a single intermediate rate. Combining this with the previous paragraphs implies that the optimal rate is constant over time, independent of the partitioning that is used. □

LEMMA 5. *For strictly concave pd-relations, the optimal policy $\hat{\mathbf{u}}^*$ has at most one time interval in which an intermediate production rate is used.*

Proof. Divide the set $\mathcal{I}$ into three subsets $\mathcal{I}_1$, $\mathcal{I}_2$, and $\mathcal{I}_3$ such that $\hat{u}_i = 0$ for $i \in \mathcal{I}_1$, $0 < \hat{u}_j < 1$ for $j \in \mathcal{I}_2$, and $\hat{u}_k = 1$ for $k \in \mathcal{I}_3$. We denote the overall policy by $\hat{\mathbf{u}}_A$ and the cardinality of the subsets by η_1, η_2, and η_3. We assume that $\mathcal{I}_2$ contains at least two elements while the other two may be empty sets. Suppose, without loss of generality, that 1 and 2 are elements of $\mathcal{I}_2$.

We have $\hat{u}_j = \hat{u}_l$ for all $j, l \in \mathcal{I}_2$ (see Lemma 2e) and $\sum_{i \in \mathcal{I}} g(\hat{u}_i) = c$ (see Lemma 2a). The objective value for policy $\hat{\boldsymbol{u}}_A$ equals $J(\hat{\boldsymbol{u}}_A) = \hat{u}_1 + \hat{u}_2 + (\eta_2 - 2) \cdot \hat{u}_1 + \eta_3$. The maximum deterioration constraint can be written as $\eta_1 \, g(0) + g(\hat{u}_1) + g(\hat{u}_2) + (\eta_2 - 2) \, g(\hat{u}_1) + \eta_3 \, g(1) = c$, which is rewritten to

$$g(\hat{u}_1) + g(\hat{u}_2) = c - \eta_1 \, g(0) - (\eta_2 - 2) \, g(\hat{u}_1) - \eta_3 \, g(1).$$

Now notice there exist two numbers $0 < \epsilon < \delta$ such that

$$g(\hat{u}_1) + g(\hat{u}_2) = g(\hat{u}_1 - \epsilon) + g(\hat{u}_2 + \delta),$$

which holds since g is strictly concave and $g'(u_j) > 0$ for $j \in \mathcal{I}_2$ (see Lemma 2d). The objective value for the new policy equals $J(\hat{\boldsymbol{u}}_B) = (\hat{u}_1 - \epsilon) + (\hat{u}_2 + \delta) + (\eta_2 - 2) \cdot \hat{u}_1 + \eta_3$. We have $J(\hat{\boldsymbol{u}}_A) < J(\hat{\boldsymbol{u}}_B)$ since $\epsilon < \delta$ and thus $\hat{\boldsymbol{u}}_A$ cannot be optimal. We only assumed that $\hat{\boldsymbol{u}}_A$ uses an intermediate rate in at least two time intervals. Hence, for strictly concave pd-relations g, the optimal policy uses an intermediate rates in at most one time interval. $\square$

References

Abbott, S. 2015. *Understanding Analysis: Additional Topics*. Springer, New York, 249–303.

Alaswad, S., Y. Xiang. 2017. A review on condition-based maintenance optimization models for stochastically deteriorating system. *Reliability Engineering & System Safety* **157** 54–63.

Ang, C. W., C. K. Tham. 2007. Analysis and optimization of service availability in a ha cluster with load-dependent machine availability. *IEEE Transactions on Parallel and Distributed Systems* **18**(9) 1307–1319.

Ayed, S., D. Sofiene, R. Nidhal. 2012. Joint optimisation of maintenance and production policies considering random demand and variable production rate. *International Journal of Production Research* **50**(23) 6870–6885.

Batun, S., L. M. Maillart. 2012. Reassessing tradeoffs inherent to simultaneous maintenance and production planning. *Production and Operations Management* **21**(2) 396–403.

Boukas, E. K., Q. Zhang, G. Yin. 1995. Robust production and maintenance planning in stochastic manufacturing systems. *IEEE Transactions on Automatic Control* **40**(6) 1098–1102.

De Jonge, B., E. Jakobsons. 2018. Optimizing block-based maintenance under random machine usage. *European Journal of Operational Research* **265**(2) 703–809.

De Jonge, B., R. H. Teunter, T. Tinga. 2017. The influence of practical factors on the benefits of condition-based maintenance over time-based maintenance. *Reliability Engineering & System Safety* **158** 21–30.

Dolinšek, S., B. Šuštaršič, J. Kopač. 2001. Wear mechanisms of cutting tools in high-speed cutting processes. *Wear* **250**(1) 349–356.

Feng, Y., Y. Qiu, C. J. Crabtree, H. Long, P. J. Tavner. 2013. Monitoring wind turbine gearboxes. *Wind Energy* **16**(5) 728–740.

Filus, J. 1987. The load optimization of a repairable system with gamma-distributed time-to-failure. *Reliability Engineering* **18**(4) 275–284.

Francie, K. A., K. Jean-Pierre, D. Pierre, S. Victor, P. Vladimir. 2014. Stochastic optimal control of manufacturing systems under production-dependent failure rates. *International Journal of Production Economics* **150** 174–187.

Hu, J. Q., P. Vakili, G. X. Yu. 1994. Optimality of hedging point policies in the production control of failure prone manufacturing systems. *IEEE Transactions on Automatic Control* **39**(9) 1875–1880.

Iravani, S. M. R., I. Duenyas. 2002. Integrated maintenance and production control of a deteriorating production system. *IIE Transactions* **34**(5) 423–435.

Iravani, S. M. R., V. Krishnamurthy. 2007. Workforce agility in repair and maintenance environments. *Manufacturing & Service Operations Management* **9**(2) 168–184.

Iyer, R. K., D. J. Rossetti. 1986. A measurement-based model for workload dependence of cpu errors. *IEEE Transactions on Computers* **100**(6) 511–519.

Kazaz, B., T. W. Sloan. 2008. Production policies under deteriorating process conditions. *IIE Transactions* **40**(3) 187–205.

Kazaz, B., T. W. Sloan. 2013. The impact of process deterioration on production and maintenance policies. *European Journal of Operational Research* **227**(1) 88–100.

Khaleghei, A., V. Makis. 2016. Reliability estimation of a system subject to condition monitoring with two dependent failure modes. *IIE Transactions* **48**(11) 1058–1071.

Kim, M. J., V. Makis. 2013. Joint optimization of sampling and control of partially observable failing systems. *Operations Research* **61**(3) 777–790.

Liberopoulos, G., M. Caramanis. 1994. Production control of manufacturing systems with production rate-dependent failure rates. *IEEE Transactions on Automatic Control* **39**(4) 889–895.

Liu, B., S. Wu, M. Xie, W. Kuo. 2017. A condition-based maintenance policy for degrading systems with age-and state-dependent operating cost. *European Journal of Operational Research* **263**(3) 879–887.

Liu, X., J. Li, K. N. Al-Khalifa, A. S. Hamouda, D. W. Coit, E. A. Elsayed. 2013. Condition-based maintenance for continuously monitored degrading systems with multiple failure modes. *IIE Transactions* **45**(4) 422–435.

Lu, Z., W. Cui, X. Han. 2015. Integrated production and preventive maintenance scheduling for a single machine with failure uncertainty. *Computers & Industrial Engineering* **80** 236–244.

Makis, V., X. Jiang. 2003. Optimal replacement under partial observations. *Mathematics of Operations Research* **28**(2) 382–394.

Martinelli, F. 2005. Control of manufacturing systems with a two-value, production-dependent failure rate. *Automatica* **41**(11) 1943–1948.

Martinelli, F. 2007. Optimality of a two-threshold feedback control for a manufacturing system with a production dependent failure rate. *IEEE Transactions on Automatic Control* **52**(10) 1937–1942.

Martinelli, F. 2010. Manufacturing systems with a production dependent failure rate: structure of optimality. *IEEE Transactions on Automatic Control* **55**(10) 2401–2406.

Mula, J., R. Poler, J. P. Garcia-Sabater, F. C. Lario. 2006. Models for production planning under uncertainty: A review. *International Journal of Production Economics* **103**(1) 271–285.

Nourelfath, M., F. Yalaoui. 2012. Integrated load distribution and production planning in series-parallel multi-state systems with failure rate depending on load. *Reliability Engineering & System Safety* **106** 138–145.

Olde Keizer, M. C. A., R. H. Teunter, J. Veldman. 2017. Joint condition-based maintenance and inventory optimization for systems with multiple components. *European Journal of Operational Research* **257**(1) 209–222.

Peng, H., G. J. Van Houtum. 2016. Joint optimization of condition-based maintenance and production lot-sizing. *European Journal of Operational Research* **253**(1) 94–107.

Puterman, M. 1994. *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons, New York.

Rudin, W. 1976. *Principles of mathematical analysis*. McGraw-Hill, New York.

Samuelson, A., A. Haigh, M. M. O'Reilly, N. G. Bean. 2017. Stochastic model for maintenance in continuously deteriorating systems. *European Journal of Operational Research* **259**(3) 1169–1179.

Sethi, S. P., H. Yan, H. Zhang, Q. Zhang. 2002. Optimal and hierarchical controls in dynamic stochastic manufacturing systems: A survey. *Manufacturing & Service Operations Management* **4**(2) 133–170.

Shen, W., I. Duenyas, Kapuscinski R. 2014. Optimal pricing, production, and inventory for new product diffusion under supply constraints. *Manufacturing & Service Operations Management* **16**(1) 28–45.

Slaugh, V. W., B. Biller, S. R. Tayur. 2016. Managing rentals with usage-based loss. *Manufacturing & Service Operations Management* **18**(3) 429–444.

Sloan, T. W. 2004. A periodic review production and maintenance model with random demand, deteriorating equipment, and binomial yield. *Journal of the Operational Research Society* **55**(6) 647–656.

Sloan, T. W., J. G. Shanthikumar. 2000. Combined production and maintenance scheduling for a multiple-product, single-machine production system. *Production and Operations Management* **9**(4) 379–399.

Van Noortwijk, J. M. 2009. A survey of the application of gamma processes in maintenance. *Reliability Engineering & System Safety* **94**(1) 2–21.

Zhang, T., R. Dwight, K. El-Akruti. 2015. *Condition Based Maintenance and Operation of Wind Turbines*. Springer International Publishing, Cham, 1013–1025.

Zhang, X., J. Zeng. 2017. Joint optimization of condition-based opportunistic maintenance and spare parts provisioning policy in multiunit systems. *European Journal of Operational Research* **262**(2) 479–498.

Zied, H., D. Sofiene, R. Nidhal. 2011. Optimal integrated maintenance/production policy for randomly failing systems with variable failure rate. *International Journal of Production Research* **49**(19) 5695–5712.