



university of
 groningen

faculty of economics
 and business

2019006-OPERA

**Joint Condition-based Maintenance
 and Condition-based Production
 Optimization**

October 2019

Michiel A.J. uit het Broek
 Ruud H. Teunter
 Bram de Jonge
 Jasper Veldman



SOM is the research institute of the Faculty of Economics & Business at the University of Groningen. SOM has six programmes:

- Economics, Econometrics and Finance
- Global Economics & Management
- Innovation & Organization
- Marketing
- Operations Management & Operations Research
- Organizational Behaviour

Research Institute SOM
Faculty of Economics & Business
University of Groningen

Visiting address:
Nettelbosje 2
9747 AE Groningen
The Netherlands

Postal address:
P.O. Box 800
9700 AV Groningen
The Netherlands

T +31 50 363 9090/7068/3815

www.rug.nl/feb/research



Joint Condition-based Maintenance and Condition-based Production Optimization

Michiel A.J. uit het Broek

University of Groningen, Faculty of Economics and Business, Department of Operations
a.j.uit.het.broek@rug.nl

Ruud H. Teunter

University of Groningen, Faculty of Economics and Business, Department of Operations

Bram de Jonge

University of Groningen, Faculty of Economics and Business, Department of Operations

Jasper Veldman

University of Groningen, Faculty of Economics and Business, Department of Operations

Joint condition-based maintenance and condition-based production optimization

Michiel A. J. uit het Broek*, Ruud H. Teunter, Bram de Jonge, Jasper Veldman

Department of Operations, Faculty of Economics and Business, University of Groningen, the Netherlands

Abstract

Developments in sensor equipment and the Internet of Things increasingly allow production facilities to be monitored and controlled remotely and in real time. Organizations can exploit these opportunities to reduce costs by employing condition-based maintenance (CBM) policies. Another recently proposed option is to adopt condition-based production (CBP) policies which control the deterioration of equipment remotely and in real time by dynamically adapting the production rate. This study compares their relative performance and introduces a fully dynamic condition-based maintenance and production (CBMP) policy that integrates both policies. Numerical results show that the cost-effectiveness of the policies strongly depends on system characteristics such as the planning time for maintenance, the cost of corrective maintenance, and the rate and volatility of the deterioration process. Integrating condition-based production decisions into a condition-based maintenance policy substantially reduces the failure risk, while fewer maintenance actions are performed. Interestingly, in some situations, the combination of condition-dependent production and maintenance even yields higher cost savings than the sum of their separate cost savings. Moreover, particularly condition-based production is able to cope with incorrect specifications of the deterioration process. Overall, there is much to be gained by making the production rate condition dependent, also, and sometimes even more so, if maintenance is already condition-based. These insights provide managerial guidance in selecting CBM, CBP, or the fully flexible CBMP policy.

Keywords: Production control, condition monitoring, adjustable production rate, production-dependent deterioration, condition-based operations

1. Introduction

Maintenance activities are a major cost driver for modern production facilities. For instance, manufacturing firms typically face maintenance costs ranging between 15-40% of their total expenses (Wireman, 2014), and for power plants and offshore wind farms maintenance costs constitute up to 30% of the total costs (Blanco, 2009; Gräber, 2004). Consequently, efficient and effective operations and maintenance strategies are of crucial importance for the profitability and competitiveness of firms. Various developments that provide opportunities to improve operational decision making are decreasing prices of monitoring equipment, advances in the Internet of Things (IoT), and improved machine learning techniques to process large amounts of condition information (Feng and Shanthikumar, 2018; Choi et al., 2018). These developments enable operators to monitor and control production facilities remotely and in real time.

In light of these developments, many studies aim to reduce maintenance costs by implementing flexible maintenance policies that schedule maintenance based on condition information. Such policies try to schedule maintenance just before imminent failure, thereby avoiding wastage of remaining useful life of equipment and lowering the number of unexpected failures. The effectiveness of such condition-based maintenance (CBM) policies, however, heavily depends on the characteristics of the deterioration process and on logistical planning times (De Jonge et al., 2017).

*Corresponding author: Nettelbosje 2, 9747 AE Groningen, the Netherlands.

Email addresses: a.j.uit.het.broek@rug.nl (Michiel A. J. uit het Broek), r.h.teunter@rug.nl (Ruud H. Teunter), b.de.jonge@rug.nl (Bram de Jonge), j.veldman@rug.nl (Jasper Veldman)

An alternative policy, that reduces operational costs but does not require a flexible maintenance schedule, is to control the deterioration of equipment by dynamically adapting the production rate based on condition information (Uit het Broek et al., 2019). This approach exploits the fact that machines typically deteriorate faster at higher production rates. Such condition-based production (CBP) policies can be used in combination with static maintenance schedules (e.g., perform maintenance every year) and use the adjustable production rate to better balance production revenues and maintenance costs.

Instead of implementing either a dynamic maintenance or a dynamic production policy, one can also integrate both into a fully dynamic policy with condition-based maintenance and condition-based production (CBMP). Surprisingly, despite the abundance of condition monitoring in practice and the fact that both CBM and CBP are studied, CBMP policies have not been considered in the literature to the best of our knowledge. An example of a real-life system where CBMP is expected to be valuable is an offshore wind farm. Turbine components such as gearboxes and generators deteriorate over time and their condition is closely monitored (e.g., by measuring noise, vibration, and temperature). As long as a turbine is in good condition, it can produce at a high rate while the condition information is used to schedule maintenance. However, initiating maintenance requires a considerable planning time to arrange spare parts, skilled technicians, and specialized vessels. During this planning time, the condition information cannot be used to improve maintenance decisions anymore, but it can still be used in real time to dynamically adjust the production rate of the turbine.

Our first contribution is to analyze the structure of the optimal CBMP policy, leading to insights on how maintenance and production flexibility can best be jointly exploited. Our second contribution is to compare the performance of the three policies (i.e., CBM, CBP, and CBMP), providing guidance on determining when the combination is particularly beneficial. We also address whether a flexible production planning (which is required for CBP) replaces the need for a flexible maintenance planning (which is required for CBM).

As our research is exploratory, we consider a single-unit system with a perfectly observable single deterioration parameter. The deterioration rate of the system is assumed to depend on the adjustable production rate and thus the deterioration process can be controlled by adapting the production speed. We obtain numerical results by formulating the problem as a Markov decision process. The results show that the two dynamic decisions (partially) complement each other for a wide range of systems. Condition-based maintenance aims to schedule maintenance just-in-time whereas condition-based production extends the equipment lifetime by reducing the short-term failure risk if needed. CBMP combines both benefits by reducing the failure risk for high deterioration levels, thereby creating the opportunity to apply a less conservative maintenance policy. All policies that use condition information result in reduced failure risks, fewer maintenance interventions, and lower costs. Moreover, integrating the two dynamic decisions improves the concept of just-in-time maintenance, since CBMP performs maintenance at higher deterioration levels than CBM while realizing fewer failures. Furthermore, the two dynamic decisions can enhance each other's cost-effectiveness as for various systems CBMP realizes cost savings larger than the sum of the cost savings obtained by implementing CBM and CBP in isolation. However, there are also systems for which both CBM and CBP reduce cost while implementing them together does not result in further cost benefits.

The remainder of this study is organized as follows. In Section 2, we discuss the relevant literature on maintenance and production decisions. In Section 3, we formally introduce the problem that we consider. The Markov decision processes that we use to determine optimal policies are described in Section 4. In Section 5, we discuss the optimal policies and compare their effectiveness. We conclude and provide future research suggestions in Section 6.

2. Literature review

Despite the growing interest in the interaction between maintenance and production decisions, both fields are typically considered in isolation. Ding and Kamaruddin (2015) provide a general review on maintenance, whereas Alaswad and Xiang (2017) focus on condition-based maintenance. The literature on production decisions under uncertainty, such as uncertain machine failures, is reviewed by Mula et al. (2006). The most recent review that focuses on the interaction between production and maintenance is conducted by Sethi et al. (2002). In the remainder of this section, we first discuss studies on maintenance, mainly focussing on those that compare dynamic maintenance policies (i.e., condition-based) with static policies (i.e., block-based or age-based). Thereafter we discuss studies on production decisions that either affect the failure behavior of equipment or that include condition monitoring.

The maintenance literature is extensive and various types of maintenance interventions are considered. A first distinction that can be made is between preventive and corrective maintenance. Preventive maintenance actions can be time-based or condition-based, and policies of both types have been studied and compared in many studies (Paté-Cornell et al., 1987; McKone and Weiss, 2002; Makis and Jiang, 2003; Crowder and Lawless, 2007; Panagiotidou and Tagaras, 2010; Bouvard et al., 2011; Huynh et al., 2011; Zio and Compare, 2013; Kim and Makis, 2013; Liu et al., 2013; De Jonge et al., 2017). These studies show that condition-based maintenance policies reduce operational cost and improve system reliability. The above studies address the value of dynamic maintenance policies, but, in contrast to our study, the value of actively controlling the deterioration process is not studied.

The production literature shows a growing interest in the interaction between production decisions and the failure behavior of systems. Within this area, we can distinguish three main research streams. The first and largest stream assumes that failure risks only depend on the age of the system and the current production rate (see, e.g., Hu et al., 1994; Boukas et al., 1995; Martinelli, 2005, 2007, 2010; Nourelfath and Yalaoui, 2012; Francie et al., 2014; Lu et al., 2015). These studies assume that producing at higher rates only increases the current failure risk, instead of resulting in permanent deterioration of the system.

The second stream takes into account production-dependent deterioration. Zied et al. (2011) model the relation by accelerating aging as a linear function of the production rate. They consider inventory, backlog, and maintenance cost, and sequentially optimize the production and the maintenance policy. Ayed et al. (2012) extend this setting to multi-unit systems. Rivera-Gómez et al. (2018) study a similar setting and incorporate the possibility that defective products are produced. The above studies consider production-dependent deterioration, but estimate this by an operational age and do not monitor the actual deterioration level of the system. Uit het Broek et al. (2019) do address the value of monitoring the deterioration process for a single-unit system with production-dependent deterioration, but restrict their attention to condition-based production decisions for a predetermined block-based maintenance policy.

The third stream incorporates condition information into policies that optimize both maintenance and production decisions. However, in this stream the deterioration rate cannot be controlled by adjusting the production rate. Boukas and Liu (2001) study a single-unit system with three functioning states and include inventory, backlog, and maintenance costs. They use a continuous flow formulation to optimize the production and maintenance rate. Iravani and Duenyas (2002) study a similar system, but use a semi-Markov decision process formulation instead. The setting is extended with stochastic demand by Sloan (2004). Although the above studies optimize the production rate under the presence of condition information, the adjustable production rate has no effect on the deterioration rate. Sloan and Shanthikumar (2000) study the use of condition monitoring for a single-unit system that can produce various items, where yields are affected by the deterioration level. Their focus is to determine which item to produce next and when to initiate maintenance. Batun and Maillart (2012) point out that the previous study overestimates the cost savings for some settings and revise the model given by Sloan and Shanthikumar (2000). Extensions including different production times per item are covered as well (Sloan and Shanthikumar, 2002; Kazaz

and Sloan, 2008, 2013). Peng and Van Houtum (2016) determine optimal production lot-sizes and use condition information to determine whether maintenance is initiated or a new lot is produced.

Summarizing, although there is a growing interest in the interaction between production and maintenance, production rate decisions based on condition information are rarely considered. Existing studies have mainly focused on the effect of production decisions on the failure rate, and not on how they affect deterioration. The few studies that do take the permanent effect of the production rate into account, generally model this by an operational age and do not consider monitoring of the actual deterioration level. The limited number of studies that use condition information for production decisions either ignore the relation between the production rate and the deterioration rate, or focus on systems that can produce various products that are differently affected by the deterioration level. To the best of our knowledge, no existing study compares condition-based maintenance to condition-based production for systems with production-dependent deterioration. Furthermore, there is no study that jointly optimizes condition-based maintenance and condition-based production rate decisions for such systems.

3. Problem description

We study a single-unit system whose condition can be described by a single deterioration 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 period) depends on the production rate. The deterioration process is described by a nondecreasing continuous-time continuous-state stochastic process $X = \{X(t) \mid t \geq 0\}$. Deterioration level 0 indicates that the unit is as-good-as-new, whereas deterioration levels exceeding L indicate that the unit has failed.

There are two maintenance actions that restore the system to the as-good-as-new state. Preventive maintenance at a cost c_{pm} can be performed as long as the unit is functioning, and more expensive corrective maintenance at a cost c_{cm} is required once the unit has failed. Both maintenance actions require a negligible amount of time but do need a fixed planning time s . We let $\theta(t, x) \in \{0, 1\}$ be the decision variable that denotes whether maintenance is initiated at time t (since the last system renewal) and condition x . The complete maintenance policy is denoted as $\theta = \{\theta(t, x) \mid t \geq 0, 0 \leq x \leq L\}$. We let $c(X(t))$ be the maintenance cost as a function of deterioration level $X(t)$, that is,

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

The system can produce at different production rates that range from 0 (idle) to 1 (maximum rate). The possible production rates are given by the set $U = [0, 1]$. When the system has failed, it cannot produce and the production rate is fixed at 0. When the system does not produce at its maximum rate, there is a revenue loss that is proportional to the production rate $u \in U$, which equals $(1 - u)\pi$ per time period. We let $\mathbf{u} = \{u(t, x) \mid t \geq 0, 0 \leq x \leq L\}$ denote the production policy.

The deterioration rate depends on the production rate and is denoted by $g(u)$. We refer to this function as the production-deterioration relation (pd-relation for short). The pd-relation g is assumed to be increasing because the system is assumed to deteriorate faster for higher production rates. We let $\mu_{\text{min}} = g(0)$ and $\mu_{\text{max}} = g(1)$ denote the minimum and maximum deterioration rate, respectively. Moreover, for a given production rate u , we consider a stationary deterioration process, i.e., the deterioration increments do not depend on the current deterioration level.

It is left to express the costs as a function of a given maintenance policy θ and a given production policy \mathbf{u} . For a given maintenance policy θ and a given realization of the deterioration process X , we can derive all moments τ_i , $i \in \mathbb{N}$, at which maintenance is performed. Recall that maintenance requires a fixed planning time s , and thus the first maintenance action is performed s time units after maintenance is

initiated for the first time, that is, $\tau_0 = \inf\{t \geq 0 \mid \theta(t, X(t)) = 1\} + s$. We find the subsequent maintenance moments by $\tau_i = \inf\{t > \tau_{i-1} + s \mid \theta(t, X(t)) = 1\}$, $i \in \mathbb{N} \setminus \{0\}$.

For a given θ and \mathbf{u} , the total expected cost up to time t equals

$$J(\theta, \mathbf{u}, t) = E \left[\sum_{i=0}^{\infty} I_{\{\tau_i \leq t\}} \cdot c(X(\tau_i)) \right] + \pi \cdot E \left[\int_0^t 1 - u(s, X(s)) \, ds \right],$$

where $I_{\{\tau_i \leq t\}}$ is the indicator function that equals one if $\tau_i \leq t$ and zero otherwise. The first term represents the expected maintenance costs and the second term the expected revenue losses. We define the long-run average cost as

$$J(\theta, \mathbf{u}) = \limsup_{t \rightarrow \infty} \frac{J(\theta, \mathbf{u}, t)}{t}.$$

The minimal long-run average cost equals $J^* = \inf_{\theta, \mathbf{u}} J(\theta, \mathbf{u})$. Our aim is to determine a joint maintenance and production policy (θ^*, \mathbf{u}^*) that minimizes the long-run average costs, that is, to determine θ^* and \mathbf{u}^* such that $J(\theta^*, \mathbf{u}^*) = J^*$.

3.1. Control strategies

We define a strategy as a combination of a production policy and a maintenance policy. The various strategies that we consider differ in their flexibility regarding the maintenance and production decisions. We consider two maintenance policies (referred to as *block-based* and *condition-based*) and two production policies (referred to as *max-rate* and *condition-based rate*).

We remark that in our initial exploration we also considered a production policy with a fixed, but not necessarily maximum production rate. However, in most considered situations the best fixed rate turned out to be the maximum rate and so discussing this policy would have limited value.

In the remainder of this study, we let the *fixed maintenance and production* (FMP) strategy refer to a block-based maintenance policy combined with the max-rate production policy. This strategy has no flexibility and does not use condition information. The *condition-based production* (CBP) strategy employs a block-based maintenance policy combined with an adjustable production rate. The *condition-based maintenance* (CBM) strategy refers to a condition-based maintenance policy combined with the max-rate production policy. The *condition-based maintenance and production* (CBMP) strategy combines condition-based maintenance with condition-based production.

Maintenance policies

The block-based maintenance policy is a static maintenance policy that fixes all maintenance actions in advance. Under this policy, the decision maker selects a block length T and maintenance is performed every T time units. We refer to such a time interval of T time units as a block. If the unit is functioning at the end of a block, preventive maintenance is performed; otherwise more expensive corrective maintenance is required. Additional maintenance actions during a block are not possible. Thus, a system failure results in production losses until the end of the current block. Notice that this maintenance policy is not affected by the planning time since it fixes all maintenance actions in advance.

The condition-based maintenance policy is flexible and allows the decision maker to plan maintenance interventions based on the condition information. Recall that there is a maintenance planning time s between initiating and performing maintenance. The case $s = 0$ implies that maintenance can be carried out instantaneously. At the end of the planning time, preventive maintenance is carried out if the system is still functioning, whereas corrective maintenance is performed if the system has failed. System failure before or during the planning time results in production losses until the end of the planning time.

Production policies

The production policies define a set of admissible production rates, denoted by $\mathcal{A}(x)$, as a function of the current deterioration level $x \in X$. Under the max-rate policy, the system produces at its maximum rate as long as it is functioning. Thus the set of admissible production rates equals

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

The condition-based production policy is fully flexible and allows the decision maker to control the production rate at any time, and so

$$\mathcal{A}(x) = \begin{cases} U & \text{if } x < L, \\ \{0\} & \text{if } x \geq L. \end{cases}$$

4. Markov decision process formulation

In this section, we provide the Markov decision process (MDP) formulations that we use to determine optimal policies. Before introducing the MDP formulations, we discretize the state space, the time horizon, and the set of admissible production rates (see Section 4.1). The structure of the resulting MDP formulations depends on the maintenance policy that is used. For the block-based maintenance policy, the optimal control policy can be determined by formulating an MDP with a finite time horizon (see Section 4.2). For the condition-based maintenance policy, the optimal control policy can be determined by formulating an MDP with an infinite time horizon (see Section 4.3).

4.1. Discretization

We partition the continuous deterioration interval $[0, L]$ into m equally sized intervals of length $\Delta X = L/m$, and then discretize this to the ordered set of midpoints $\bar{X} = \{(i + 0.5)\Delta X \mid i = 0, \dots, m - 1\}$. All deterioration levels above L are merged into a single state with index m that indicates system failure. The time horizon is discretized into periods with length Δt , and there is a decision epoch at the start of each period. The continuous set of production rates is discretized into $n + 1$ production rates that are uniformly distributed between the minimum and maximum production rate, that is, $\bar{U} = \{i/n \mid i = 0, \dots, n\}$. Note that we can approximate a continuous system arbitrarily close by setting the step sizes sufficiently small, although smaller step sizes also result in increased computation times.

We let $F_{\Delta t, u}$ denote the distribution function of the additional amount of deterioration during a time period Δt when producing at rate u . For ease of notation, we drop the subscript Δt in the remainder of this study. To obtain the transition probabilities in the discretized process, we set the probability of staying in the same deterioration state to $F_u(0.5\Delta X)$; the probability of moving from state k to state $k + i$, where $i \geq 1$ and $k + i < m$, to $F_u((i + 0.5)\Delta X) - F_u((i - 0.5)\Delta X)$; and the probability of moving from state k to the failed state m to $1 - F_u((i - 0.5)\Delta X)$ where $i = m - k$. Summarizing, the transition probabilities of the discrete deterioration process are

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

4.2. MDP for block-based maintenance

The block-based maintenance policy results in a renewal process in which the maintenance actions are the renewal points. Each block starts at the as-good-as-new deterioration level, regardless of the actions and deterioration realizations in previous blocks. As a result, for a given block length T , it is sufficient to minimize the total expected costs during a single block. Thus, systems with a block-based maintenance policy can be optimized based on an MDP with a finite horizon by using backward induction (see Puterman, 1994, ch.4).

We let $\tau \in \{0, \Delta t, 2\Delta t, \dots, T\}$ denote the remaining time until the end of the current block, i.e., maintenance is performed when $\tau = 0$. Let $V(x, \tau)$ denote the total expected cost in the remainder of the current block given that the current deterioration level of the system is $x \in \bar{X}$. At the end of the block, maintenance is performed and thus

$$V(x, 0) = \begin{cases} c_{\text{pm}} & \text{if } x < m, \\ c_{\text{cm}} & \text{if } x = m. \end{cases}$$

In all other periods, a production rate u can be selected. This affects the production loss $(1 - u)\pi\Delta t$ and the expected future costs. Thus, for $\tau > 0$ we get

$$V(x, \tau) = \min_{u \in \bar{U}(x)} \left\{ (1 - u)\pi\Delta t + \sum_{i=0}^{m-x} P_u(x, x+i) V(x+i, \tau - \Delta t) \right\},$$

where $\bar{U}(x)$ is the discretized set of admissible production rates for a given production policy, as described in Section 3.1 and 4.1.

The expected total costs during a single block of length T equals $V(0, T)$, and the long-run average cost per time unit equals $V(0, T)/T$. We find the optimal block length T^* by solving the system up to a sufficiently large value τ_{max} and selecting $T^* \in \arg \min_{0 < T \leq \tau_{\text{max}}} \{V(0, T)/T\}$.

4.3. MDP for condition-based maintenance

For systems with a condition-based maintenance policy, we consider an MDP with an infinite horizon and we use value iteration to find ϵ -optimal policies. We first introduce some notation and show that the assumptions required for value iteration are satisfied. Thereafter, we introduce the value functions used by the value iteration algorithm (see Puterman, 1994, ch.8).

We let $\tau \in S := \{0, \Delta t, 2\Delta t, \dots, s, \text{ns}\}$ denote the remaining planning time until the next scheduled maintenance intervention. For modeling purposes, we have added the additional element 'ns' to S that indicates that maintenance is *not scheduled* yet.

The state of the system is described by the current deterioration level $x \in \bar{X}$ and the remaining planning time $\tau \in S$. Both the state space and the action space are finite. Furthermore, the direct cost function per time period is the sum of production loss $(1 - u)\pi\Delta t$ and maintenance cost. Both parts are bounded and thus the direct cost function itself is bounded. Furthermore, it can easily be seen that we have a unichain model, that is, the transition matrix corresponding to every deterministic stationary policy consists of a single recurrent class plus a possibly empty set of transient states (see Puterman, 1994, ch.8). It follows that a stationary average optimal policy exists, which can be obtained by value iteration.

Let $V_n(x, \tau)$ denote the value function after n iterations of the value iteration algorithm. At the end of the planning time (i.e., $\tau = 0$), maintenance is performed and the deterioration level immediately jumps to the as-good-as-new level. After maintenance is performed, the production decision is the same as if we started the time period at deterioration level $x = 0$, and thus

$$V_n(x, 0) = \begin{cases} c_{\text{pm}} + V_n(0, s) & \text{if } x < m, \\ c_{\text{cm}} + V_n(0, s) & \text{if } x = m. \end{cases}$$

During the planning time (i.e., $0 < \tau \leq s$), the decision maker can only decide on the production rate u . The remaining planning time is reduced by Δt , regardless of the production decision. The value function during the planning time equals

$$V_n(x, \tau) = \min_{u \in \bar{U}(x)} \left\{ (1-u)\pi\Delta t + \sum_{i=0}^{m-x} P_u(x, x+i) V_{n-1}(x+i, \tau - \Delta t) \right\}.$$

Before maintenance is scheduled (i.e., $\tau = ns$), the decision maker can decide on the production rate and whether or not to schedule maintenance. If maintenance is scheduled, the remaining planning time is set to $\tau = s$. Otherwise, the remaining planning time remains $\tau = ns$. The value function before maintenance is scheduled equals

$$V_n(x, ns) = \min_{\tau \in \{ns, s\}, u \in \bar{U}(x)} \left\{ (1-u)\pi\Delta t + \sum_{i=0}^{m-x} P_u(x, x+i) V_{n-1}(x+i, \tau) \right\}.$$

We initialize the value iteration algorithm by setting $V_0(x, \tau) = 0$ for all $x \in \bar{X}$ and $\tau \in S$. Let the span of the value function be defined as $sp(V_n) = \max_{x, \tau} V_n(x, \tau) - \min_{x, \tau} V_n(x, \tau)$, and let $W_n = V_{n+1} - V_n$. The value iteration algorithm stops when $sp(W_n) < \epsilon$. The ϵ -optimal policy is retrieved by performing one more iteration while storing the arguments that minimize the value functions. The corresponding long-run average cost g^* is estimated as $g' = 0.5(\min W_n + \max W_n)$, for which we know that $|g' - g^*| < \epsilon/2$.

5. Numerical analysis

We continue by analyzing the effectiveness of the various strategies. Our approach is to first consider a base case and then present a sensitivity analysis where we vary the different system parameters. The base case is chosen somewhat arbitrarily, but in such a way that the effects are representative for a large set of instances that we initially considered. Using the base case, we describe through what mechanisms CBP reduces cost, how the performance of CBM compares to that of CBP, and how the two interact with each other. A sensitivity analysis then continues to explore when CBP and CBM are particularly effective in isolation or in combination (i.e., CBMP). Before discussing the results, we first introduce the deterioration process and the base case parameter values that we consider.

5.1. Deterioration process

Various stochastic processes have been suggested in the literature to model deterioration, including compound Poisson processes, Brownian motions with drift, and gamma processes. In this study, we use stationary gamma processes since these are the most appropriate to model monotonically increasing deterioration such as erosion, wear, and fatigue (Van Noortwijk, 2009; Alaswad and Xiang, 2017). The stationary gamma process is a flexible process for which the deterioration rate and volatility can be controlled by two parameters. This enables us to study a wide scope of systems with different deterioration characteristics. The gamma process is a continuous-time continuous-state process, and we will discretize it as described in Section 4.1.

The increments of a gamma process are independently gamma distributed. Denoting the shape parameter by $\alpha > 0$ and the scale parameter by $\beta > 0$, the gamma density function for the increment per time unit is given by

$$f_{\alpha\beta}(x) = \frac{x^{\alpha-1} \exp(-x/\beta)}{\Gamma(\alpha)\beta^\alpha},$$

where $\Gamma(\alpha) = \int_0^\infty z^{\alpha-1} \exp(-z) dz$ is the gamma function. The deterioration increment Y per time unit has mean $E[Y] = \alpha\beta$ and variance $\text{Var}(Y) = \alpha\beta^2$. The density function corresponding to increments per period with length Δt is obtained by scaling the shape parameter to $\alpha\Delta t$.

We relate the shape and scale parameters to the pd-relation g such that the deterioration increment per time unit has the following three properties. First, the deterioration increments have mean $E[Y | u] = g(u)$. Second, if the system produces at the maximum rate, the variance of the deterioration increments equals $\text{Var}(Y | u = 1) = \sigma_{\max}^2$. Third, the coefficient of variation (i.e., the standard deviation divided by the mean) of the deterioration increments is not affected by the production rate. It can easily be verified that this is accomplished by setting the shape parameter equal to $\alpha = \mu_{\max}^2 / \sigma_{\max}^2$ and the scale parameter, as a function of the production rate, equal to $\beta(u) = g(u) \cdot \sigma_{\max}^2 / \mu_{\max}^2$ (recall that $g(1) = \mu_{\max}$).

5.2. Base case system

The base case parameter values are listed in Table 1. The preventive maintenance cost is $c_{\text{pm}} = 20$ and the corrective maintenance cost is $c_{\text{cm}} = 100$. For real-life systems, for instance wind turbines, corrective maintenance is often much more costly than preventive maintenance, because of collateral damage and the lower salvage value of failed components. The revenue when producing at the maximum rate is normalized to $\pi = 1$ per time unit. The system fails at deterioration level $L = 100$, and maintenance requires a planning time of $s = 5$ time units.

For the pd-relation we consider the same parametric form as Uit het Broek et al. (2019), that is, $g(u) = \mu_{\min} + (\mu_{\max} - \mu_{\min})u^\gamma$, which is concave for $0 < \gamma < 1$, linear for $\gamma = 1$, and convex for $\gamma > 1$. The parameter $\mu_{\min} = 0.1$ describes the deterioration rate when the system is idle, and $\mu_{\max} = 1.5$ is the deterioration rate when the system produces at the maximum rate. For the base case we set $\gamma = 1.5$, and also in the sensitivity study we mainly focus on convex pd-relations as those are conceivably most likely to be encountered in real-life systems.

The expected time until failure when producing at the maximum rate equals approximately $L/\mu_{\max} \approx 67$ time units. The system only has to produce $c_{\text{pm}}/\pi = 20$ time units at the maximum rate to compensate for the preventive maintenance cost and thus the system is expected to be profitable. However, an expensive failure is expected to result in a loss-making cycle because producing at the maximum rate for $c_{\text{cm}}/\pi = 100$ time units is required to compensate for this.

To closely approximate the continuous deterioration process, we partition the time horizon and the deterioration levels into small intervals with respective lengths $\Delta t = 1.0$ and $\Delta X = 0.05$. Note that ΔX should be small compared to the expected deterioration increment per time period, that is, $\Delta X \ll \Delta t \cdot \mu_{\max}$. The continuous action space is discretized into $n = 50$ non-idle production rates.

5.3. Cost savings for the base case system

We examine the effectiveness of flexible production and flexible maintenance decisions by comparing the four control strategies introduced in Section 3.1. Table 2 compares the performance of the four considered strategies when applied to the base case. Overall, we observe that the fixed strategy is by far the worst, that CBM is slightly more effective than CBP considering expected costs, and that introducing a flexible production policy has the positive side-effect of lowering the cost variance. We also see that having both a flexible maintenance and a flexible production policy reduces the expected cost even further. Next, we will explain these results by studying the optimal policies more closely.

Condition-based production

Because CBM has been studied much more extensively than CBP, we start by discussing the optimal structure of the CBP strategy. With an adjustable production rate, the optimal block length turns out to be $T = 60$. The results for this policy are largely in line with the findings of Uit het Broek et al. (2019), although they predetermined rather than optimized the block length. We therefore keep this discussion

Table 1: Base case system used in the numerical analysis

Parameter	Value	Interpretation
c_{pm}	20.00	Preventive maintenance cost
c_{cm}	100.00	Corrective maintenance cost
π	1.00	Production revenue at maximum rate
L	100	Failure level
s	5	Planning time for maintenance
γ	1.50	Shape pd-relation
μ_{min}	0.10	Mean of increments when idle
μ_{max}	1.50	Mean of increments at maximum rate
σ_{max}	3.00	Standard deviation increments
n	50	Number of non-idle production rates
Δt	1.00	Time discretization
ΔX	0.05	Discretization size deterioration

Table 2: Performance statistics for the four strategies when applied to the base case

	Fixed	CBP	CBM	CBMP
PM threshold T or M	$T = 42$	$T = 60$	$M = 70.20$	$M = 78.80$
Mean cost per time unit	0.562	0.424	0.409	0.379
St.dev. cost per time unit	4.343	2.834	3.353	2.957
Mean production per time unit	0.995	0.922	0.999	0.977
Mean time to maintenance	42.00	60.00	53.31	59.19
Mean time between failures	995.12	6365.37	2456.39	4525.96
Mean deterioration at maintenance	62.46	81.43	79.75	86.11

fairly short and refer to their study for a further discussion of the CBP policy structure for a predetermined block length. Figure 1 shows the production rate in grey scale, ranging from black (no production) to white (maximum rate). The solid line indicates the expected deterioration trajectory, whereas the dashed lines indicate a region that contains the deterioration level with 95% certainty for a given point in time.

Three areas can be distinguished in Figure 1. Firstly, for low deterioration levels compared to the remaining time until maintenance (white lower triangular area), a failure is unlikely and the system produces at the maximum rate. Secondly, for intermediate to high deterioration levels (grey area), the production rate is gradually reduced in order to reduce the deterioration rate and thereby the likelihood of failure. Additionally, producing at a lower rate reduces the volatility of the deterioration process, thereby also reducing the risk of a sudden failure caused by an extreme deterioration increment. Thirdly, for extremely high deterioration levels compared to the remaining time to maintenance (small white upper triangular area), failure is almost certain, and production is maximized by producing at the maximum rate. It has to be noted, however, that the effect of this third area on the long-run average cost is negligible as its states are almost never reached.

We also note from the upper dashed line in Figure 1 that CBP does not immediately slow down production when the system deteriorates faster than expected. As long as there is sufficient time to prevent failure at a later stage, it is better to continue producing at the maximum rate, since it is possible that deterioration in the remainder of the block is lower than expected. If this does not happen, then production can still be slowed down, whereas lost production cannot be made up for.

Another observation from Figure 1 is that towards the end of the block, the expected deterioration trajectory (solid line) enters the area where the production rate is reduced. Thus, even under the expected




Figure 1: The optimal production rate for the CBP strategy in grey scale, ranging from black (no production) to white (maximum rate). The solid line indicates the expected deterioration trajectory, and the dashed lines indicate a region that contains the deterioration level with 95% certainty at a given point in time.




Figure 2: Effect of block length T on the costs per period, the failure risk per block, and the mean production per period for the fixed strategy (dashed blue) and the CBP strategy (solid red). The dots indicate the optimal block length for both strategies.

deterioration path, production is slowed down towards the end of the block. This reduces the volatility of the deterioration process, and thereby reducing the risk of an instant failure by a large deterioration increment.

Figure 2 shows the effect of the block length on the average cost per time period, the probability that the system fails during a block, and the mean production per period for both the fixed strategy (dashed) and the CBP strategy (solid). For block lengths exceeding 40, the failure risk under the fixed strategy rapidly increases and therefore this strategy cannot permit to schedule longer blocks. The CBP strategy controls the failure risk by reducing the production rate if and when needed. This allows CBP to be less conservative and schedule considerably longer blocks ($T^* = 60$ instead of $T^* = 42$), while simultaneously increasing the expected time between two consecutive failures from 995 to 6365 periods, reducing the average cost per time period by 25% compared to the fixed strategy (0.424 instead of 0.562).

Besides the cost reduction, CBP also achieves a significantly lower cost standard deviation (2.83 instead of 4.34) by reducing the number of expensive failures. Note that in terms of cost, the CBP strategy is also less sensitive to small changes in the block length, which offers a practical advantage as in real-life systems the exact maintenance moment may be uncertain.

Condition-based maintenance and production

We continue with the two strategies with condition-based maintenance (i.e., CBM and CBMP) and compare their performance to the strategies with a static maintenance policy (i.e., FMF and CBP). Introducing only a flexible maintenance policy (CBM) reduces the cost by 27% compared to the fixed strategy (from 0.562 to 0.409), and is thus slightly more cost effective than CBP for the base case. The main advantage of CBM over CBP is its higher production output, as it does not slow down production to control deterioration. On the other hand, CBM performs more maintenance actions than CBP (mean time to maintenance of 53 instead of 60), and still fails around 2.6 times as often. This is partially caused by the inability of CBM to respond to large deterioration increments during the planning time, whereas CBP can lower the production speed at any time. Failures imply high maintenance cost and thus the higher number of failures also explains the higher standard deviation of the cost for CBM compared to CBP.

The cost reduction by introducing both an adjustable production rate and a flexible maintenance policy (i.e., CBMP) is 33% (from 0.562 to 0.378), implying that condition-based production and condition-based maintenance partially complement each other. The adjustable production rate improves the condition-based maintenance policy in three ways. First, the failure risk during the planning time is lower because the operator can respond to the actual deterioration. This clearly reduces the expected maintenance cost per cycle. Second, exploiting this benefit, a less conservative preventive maintenance threshold is used (78.8 instead of 70.2), resulting in fewer maintenance actions. Third, if during the planning time the system deteriorates so fast that a failure becomes (almost) unavoidable, then the expected production losses are minimized by producing at a more efficient rate until failure occurs.

Recall that CBP lowers the production rate towards the end of a block in order to reduce the risk of a failure even if deterioration is as expected. A similar observation applies to the CBMP strategy, but now production is slowed down towards the end of the planning period, unless deterioration is much lower than expected. Moreover, for various systems, production is even already slowed down before maintenance is scheduled. We conclude that the ability to vary the production rate is exploited in a wide range of scenarios.

5.4. Parameter sensitivity

In the previous section, we have seen that both CBM and CBP realize considerable cost savings in the base case, and applying them together is even more effective. However, the joint cost savings are less than the sum of the separate savings, showing that CBM and CBP only partly complement each other for the base case. In this section, we perform a sensitivity study to obtain further insights into the comparative performance of CBM and CBP and into whether the two can enhance each other's performance. Besides cost savings, we also consider other performance measures such as expected production, mean time between maintenance, and probability of a failure. The results are obtained by studying the base case system while deviating various parameters one by one.

Planning time

We first consider the planning time s required to carry out maintenance, see Figure 3. The fixed strategy and the CBP strategy use block-based maintenance policies that are not affected by the planning time, and we indeed see that the costs for these strategies are independent of s . CBP realizes a cost saving of 25% compared to the fixed strategy, regardless of the planning time. CBM realizes a cost saving of 34% when there is no planning time, and its effectiveness obviously decreases in the length of the planning time. If the planning time equals the optimal block length under the fixed strategy (i.e., $T = 42$), then CBM immediately schedules maintenance upon each maintenance action and is not able to realize a cost saving compared to the fixed strategy.

CBMP utilizes the condition information to both schedule maintenance and adapt the production rate. However, similar to the CBM policy, the value of condition information for the maintenance planning




Figure 3: Effect of planning time s on the total average cost and on the relative cost saving compared to the fixed strategy. The results are given for the fixed strategy (dotted blue), CBM (dot dashed orange), CBP (dashed red), and CBMP (solid black).

decreases if the planning time increases. As a result, CBMP converges to CBP as s increases. For short planning times, the major part of the cost saving is due to the condition-based maintenance decisions and we indeed see that CBM and CBMP result in almost the same cost if there is no planning time (0.374 versus 0.371). A marginal cost difference exists because CBMP can use a slightly higher maintenance threshold (83.80 versus 81.65) by reducing the production rate already before maintenance is scheduled, thereby reducing the volatility of the deterioration increments and thus reducing the risk of an instant failure.

We observe that for planning times of at least 37 periods, the cost saving of CBMP is larger than the sum of cost savings of CBM and CBP. Thus, the two dynamic decisions can enhance each other's performance.

We conclude that the planning time strongly affects the effectiveness of the various policies. For short planning times, CBM outperforms CBP, whereas for longer planning times CBP is preferred. Although CBM is preferred when the planning time is short, CBP realizes a considerable cost saving as well and can thus be a viable alternative for CBM based on factors not considered in this study. On the other hand, for long planning times, CBM does not reduce costs and is thus not a viable alternative for CBP in that case. Moreover, CBMP is by far the most beneficial for intermediate planning times.

Maintenance cost and revenue

Figure 4 shows how the cost savings of the condition-based strategies are affected by the other system parameters. We first assess the effect of the corrective maintenance cost. If c_{cm} equals c_{pm} , CBP only marginally reduces cost compared to the fixed strategy by 1% while CBM realizes a considerable cost saving of 18%. The small cost saving realized by CBP is because the total production can be increased in the rare case a failure is virtually inevitable. Thus, if failures do not induce additional costs on top of the preventive maintenance cost, then CBM is clearly more effective than CBP. This is intuitive as it is not beneficial to avoid a failure by slowing down production, which is the main benefit of an adjustable production rate.

For increasing corrective maintenance costs, avoiding failures becomes more important. CBM can only achieve this by being conservative and initiating maintenance at a lower deterioration threshold, i.e., the maintenance policy has to be robust for above average deterioration. CBP, on the other hand, can use a more optimistic maintenance policy and predominantly uses the adjustable production rate to avoid failure which only causes additional costs when above average deterioration is observed. For instance, if c_{cm} increases from 25 to 250, CBM decreases the maintenance threshold from 86.2 to 63.9 while CBP only reduces the block length from 62 to 60. Moreover, for extremely high corrective maintenance costs (say $c_{cm} \geq 400$), CBP becomes insensitive to changes in the corrective maintenance costs as the failure risk




Figure 4: Effect of various parameters on the relative cost savings of CBP (dashed red), CBM (dot dashed orange), and CBMP (solid black) compared to the fixed strategy.

becomes negligible while costs under CBM continue to increase. Finally, as long as corrective maintenance is more expensive than preventive maintenance, the adjustable production rate complements the condition-based maintenance policy by reducing the failure risk during the planning time.

Considering the effect of the production revenue π , we observe that CBM becomes slightly more cost-effective when the production revenue increases, while CBP becomes significantly less effective. CBP reduces the maintenance cost by decelerating the production rate, which is more expensive for high values of π . For extremely high production revenues, it is optimal to only focus on maximizing production and avoid failures by scheduling very short blocks, thus CBP converges to the fixed strategy.

Volatility deterioration increments

The third graph shows the effect of the volatility of the deterioration increments. For standard deviations close to zero, the deterioration process is stable, implying that condition information has limited value. When the volatility starts to increase, all condition-based strategies effectively use the condition information and considerably outperform the fixed strategy. CBP reduces the failure risk during a block by decreasing the production rate when the deterioration level is high relative to the remaining time to maintenance. Doing so allows this strategy to be less conservative and schedule longer blocks than the fixed strategy (e.g., $T = 62$ versus $T = 47$ for $\sigma_{\max} = 2$). However, the flexible production rate comes with revenue losses, while the flexibility of CBM does not induce additional costs. As a result, CBP is slightly less effective than CBM for small volatilities.

For intermediate volatilities, the uncertainty during the planning time becomes considerable, and CBM has to significantly lower the threshold at which maintenance is initiated to cope with this uncertainty. However, the volatility is still low enough to be handled by adjusting the production rate. As a result, combining both dynamic decisions is most advantageous for an intermediate volatility.

For high volatilities, the system does not gradually deteriorate but failures are likely to be caused by

a single extreme deterioration increment. Such shocks arrive suddenly and condition monitoring does not provide any information on the arrival of these shocks. CBP reduces their magnitude by producing at a lower rate, such that the system is not expected to fail from a single shock, and immediately switches to the idle mode after a shock occurred. For extremely high volatilities, reducing the size of the shocks does not outweigh the induced production losses and we indeed observe that the cost savings of CBP diminish if the volatility continues to increase. CBM and CBMP, on the other hand, produce at the maximum rate until the system fails and then schedule maintenance, thereby even reducing costs for extremely volatile processes.

Parameters pd-relation

We continue with the effect of the pd-relation parameters γ , μ_{\min} , and μ_{\max} , see the bottom three graphs of Figure 4. Recall that γ describes whether the pd-relation is concave ($0 < \gamma < 1$) or convex ($\gamma > 1$). The minimum deterioration rate ranges from no deterioration when idle ($\mu_{\min} = 0$) to maximum deterioration when idle ($\mu_{\min} = \mu_{\max} = 1.5$), in which case the production rate does not affect the deterioration rate. The major insights are that (1) all strategies realize a significant cost saving for all values of γ , (2) for concave pd-relations CBM is preferred while CBP is better for ‘more convex’ pd-relations, (3) CBM and CBP complement each other for all values for γ , and (4) condition-based production decisions improve the CBM policy even if μ_{\min} is high compared to μ_{\max} . Finally, the cost savings of CBM decrease as μ_{\max} increases, whereas the cost savings of CBP actually increase. Higher values for μ_{\max} imply that the system deteriorates faster and, as a result, both the fixed strategy and CBM have to schedule more maintenance interventions, thereby considerably increasing costs. CBP overcomes this by slightly reducing the production rate. For very high values of μ_{\max} , CBP already reduces the production rate for low deterioration levels in order to produce at a more efficient rate. Note that, similar to the planning time parameter s , CBM converges to the fixed strategy if μ_{\max} continues to increase.

5.5. Parameter estimation errors

So far we have assumed that all system parameters are known with certainty, whereas in practice various parameters must be estimated. In this section, we compare the robustness of the considered strategies with respect to incorrect estimations of the parameters of the deterioration process. Although all system parameters are uncertain to some extent, we focus on the mean and the standard deviation of the deterioration increments per time unit as they are crucial for maintenance and production decisions. We vary the estimated mean $\hat{\mu}_{\max}$ from 0.2 (almost no deterioration) to 3.0 (twice as fast as the base system), while the coefficient of variation is kept constant. The estimated standard deviation $\hat{\sigma}_{\max}$ ranges from 0.2 (very stable deterioration) to 6.0 (twice as volatile as the base system).

Figure 5 shows the average costs of all four strategies when they are optimized based on the estimated values $\hat{\mu}_{\max}$ (top) and $\hat{\sigma}_{\max}$ (bottom) but are applied to our base system with $\mu_{\max} = 1.5$ and $\sigma_{\max} = 3.0$. For the base system (left panels), the three main observations are that the fixed strategy is significantly more sensitive to estimation errors in the mean than the condition-based strategies, that the three condition-based strategies are similarly affected by estimation errors for both the mean and the standard deviation, and that underestimating parameters is often worse than overestimating them. Note that for large underestimations of the deterioration rate, the cost of the fixed strategy actually decreases if the estimation error increases further. This is because for large underestimations, the fixed strategy schedules too long blocks and the system is likely to fail every block, resulting in very high maintenance costs. The production revenues do not outweigh the corrective maintenance costs and it is better to schedule even fewer maintenance interventions, which happens if the estimated value decreases further.

We also consider different values of the system parameters compared to the base case to obtain further insights. This confirms that the fixed strategy is most sensitive to estimation errors, but also that CBM is often considerably more sensitive for incorrect estimating of the deterioration rate than CBP. This is




Figure 5: Effect of mis-estimating the deterioration rate $\mu_{\max} = 1.5$ (top) and the standard deviation $\sigma_{\max} = 3.0$ (bottom) on the costs realized by FMP (dotted blue), CBM (dot dashed orange), CBP (dashed red), and CBMP (solid black). The left panels show the results for the base system, the middle panels increase the corrective maintenance cost from 100 to 300, and the right panels also increase the planning time from 5 to 20.

especially the case for more expensive corrective maintenance and for longer maintenance planning times. The middle panels (with the corrective maintenance cost increased from 100 to 300) and the right panels (with also the planning time increased from 5 to 20) illustrate this. Observe that higher corrective maintenance cost makes CBM more sensitive to underestimating the rate, while longer planning times makes CBM sensitive to overestimations.

The agility of CBP compared to that of CBM explains its more robust performance with regard to estimations errors. CBP can quickly react to the actual deterioration level if this differs from what was expected. This particularly allows, in case of underestimation of the deterioration rate, to better prevent expensive corrective maintenance actions by continuously adjusting the production rate based on the actual deterioration level. Furthermore, overestimating the mean deterioration rate leads to unnecessarily early maintenance interventions under CBM, which are relatively costly compared to the slightly reduced production rates under CBP.

Summarizing, the fixed strategy is very sensitive to estimation errors, as it plans maintenance without condition information and is not able to adapt. CBM does adapt by planning maintenance based on the observed level of deterioration, but can no longer react after maintenance has been planned. As a result, CBM is more robust than the fixed strategy but still quite sensitive to estimation errors if the planning time is long and/or corrective maintenance is expensive. Under such conditions, CBP and CBMP achieve a much lower cost when the deterioration process is mis-estimated, by adjusting the production rate according to the actual deterioration level.

6. Conclusion

Ongoing developments in the fields of online condition monitoring and real-time decision making create opportunities to operate industrial systems more efficiently by implementing condition-based maintenance (CBM) or condition-based production (CBP). The latter controls the deterioration of equipment by dynamically adjusting the production rate, thereby better balancing production revenues and maintenance costs. We have compared the effectiveness of employing the two condition-based policies in isolation and assessed the benefits of integrating them into a fully flexible policy (CBMP).

We considered a single-unit system that gradually deteriorates and for which the condition is continuously monitored. Exact numerical results were obtained by formulating the system as a Markov decision process. Varying the system parameters revealed a number of valuable insights. First, for almost all considered settings, all three dynamic strategies clearly outperform the simpler strategy that combines block-based maintenance with a fixed production rate. Second, the cost reduction mechanisms of CBM and CBP are quite different. CBP typically reduces the failure risk significantly at the expense of lower expected production, whereas CBM policies are characterized by higher expected production but substantially more failures too. This also explains why CBM is more effective for very high production revenues, and CBP is more effective for large corrective maintenance costs. The comparative performance of CBM and CBP depends on what effect dominates. Moreover, CBMP improves the trade-off between production and maintenance costs further, resulting in more production than CBP and fewer failures than CBM. In many settings, the two cost reducing effects are best combined by employing the CBMP strategy. In fact, sometimes the savings of the combined strategy are more than the sum of the savings of CBM and CBP.

We also examined how the various policies perform under wrongly estimated parameter values for the deterioration process, as these are often uncertain for real-life systems. This showed once more that condition information should be taken into account when available, since all dynamic strategies were considerably more robust to estimation errors of the deterioration rate than the fixed strategy. We also found that CBM is much less robust for large estimation errors than CBP if the planning time is long or corrective maintenance is expensive, as the latter can dynamically correct for wrong decisions. It follows that it is also beneficial to adopt the dynamic strategies if there is uncertainty of the deterioration parameters.

There are numerous research opportunities within this direction. One direction is to include unreliable condition information or uncertain parameter values for the deterioration process directly into the optimization. It is expected that condition-based production decisions can better cope with such uncertainties compared to condition-based maintenance as the effect of a wrong production decision is less severe than a wrong maintenance decision. After a production rate is set, this can be revised when more condition information becomes available whereas an expensive maintenance action cannot be made undone.

Moreover, incorporating condition-based production rates into multi-unit systems seems a promising direction. Production facilities often face contracts with considerable penalties if a minimum total production target is not met. For systems with many units (e.g., offshore wind farms), some units may be idle, but if too many fail simultaneously then penalties are incurred. For such systems, condition-based production decisions can be used to desynchronize the deterioration levels of various units in order to minimize the risk that multiple units fail shortly after each other. Similarly, adjustable production rates can also be applied to do the opposite, namely synchronizing the deterioration level of (a subset of) the units such that their maintenance can be clustered, thereby reducing setup costs. Clearly, the optimal policy is situation dependent and deserves attention.

Another direction is to study multi-unit systems with limited maintenance capacity, for instance, due to the need for specialized equipment or the limited availability of skilled technicians. When multiple units require maintenance around the same time this can result in long downtimes due to the limited maintenance capacity. With a dynamic production planning, units can produce at different rates such that their preferred maintenance moments can be spread, thereby reducing peak demand for scarce maintenance equipment.

References

- Alaswad, S., Xiang, Y., 2017. A review on condition-based maintenance optimization models for stochastically deteriorating system. *Reliability Engineering & System Safety* 157, 54–63.
- Ayed, S., Sofiene, D., Nidhal, R., 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., Maillart, L. M., 2012. Reassessing tradeoffs inherent to simultaneous maintenance and production planning. *Production and Operations Management* 21 (2), 396–403.
- Blanco, M. I., 2009. The economics of wind energy. *Renewable and sustainable energy reviews* 13 (6-7), 1372–1382.
- Boukas, E. K., Liu, Z. K., 2001. Production and maintenance control for manufacturing systems. *IEEE Transactions on Automatic Control* 46 (9), 1455–1460.
- Boukas, E. K., Zhang, Q., Yin, G., 1995. Robust production and maintenance planning in stochastic manufacturing systems. *IEEE Transactions on Automatic Control* 40 (6), 1098–1102.
- Bouvard, K., Artus, S., Bérenguer, C., Cocquempot, V., 2011. Condition-based dynamic maintenance operations planning & grouping. application to commercial heavy vehicles. *Reliability Engineering and System Safety* 96 (6), 601–610.
- Choi, T.-M., Wallace, S. W., Wang, Y., 2018. Big data analytics in operations management. *Production and Operations Management* 27 (10), 1868–1883.
- Crowder, M., Lawless, J., 2007. On a scheme for predictive maintenance. *European Journal of Operational Research* 176 (3), 1713–1722.
- De Jonge, B., Teunter, R. H., Tinga, T., 2017. The influence of practical factors on the benefits of condition-based maintenance over time-based maintenance. *Reliability Engineering and System Safety* 158, 21–30.
- Ding, S.-H., Kamaruddin, S., 2015. Maintenance policy optimization—literature review and directions. *The International Journal of Advanced Manufacturing Technology* 76 (5), 1263–1283.
- Feng, Q., Shanthikumar, J. G., 2018. How research in production and operations management may evolve in the era of big data. *Production and Operations Management* 27 (9), 1670–1684.
- Francie, K. A., Jean-Pierre, K., Pierre, D., Victor, S., Vladimir, P., 2014. Stochastic optimal control of manufacturing systems under production-dependent failure rates. *International Journal of Production Economics* 150, 174–187.
- Gräber, U., 2004. Advanced maintenance strategies for power plant operators—introducing inter-plant life cycle management. *International Journal of Pressure Vessels and Piping* 81 (10-11), 861–865.
- Hu, J. Q., Vakili, P., Yu, G. X., 1994. Optimality of hedging point policies in the production control of failure prone manufacturing systems. *IEEE Transactions on Automatic Control* 39 (9), 1875–1880.
- Huynh, K., Barros, A., Berenguer, C., Castro, I., 2011. A periodic inspection and replacement policy for systems subject to competing failure modes due to degradation and traumatic events. *Reliability Engineering & System Safety* 96 (4), 497–508.
- Iravani, S. M. R., Duenyas, I., 2002. Integrated maintenance and production control of a deteriorating production system. *IIE Transactions* 34 (5), 423–435.
- Kazaz, B., Sloan, T. W., 2008. Production policies under deteriorating process conditions. *IIE Transactions* 40 (3), 187–205.
- Kazaz, B., Sloan, T. W., 2013. The impact of process deterioration on production and maintenance policies. *European Journal of Operational Research* 227 (1), 88–100.
- Kim, M. J., Makis, V., 2013. Joint optimization of sampling and control of partially observable failing systems. *Operations Research* 61 (3), 777–790.
- Liu, X., Li, J., Al-Khalifa, K. N., Hamouda, A. S., Coit, D. W., Elsayed, E. A., 2013. Condition-based maintenance for continuously monitored degrading systems with multiple failure modes. *IIE Transactions* 45 (4), 422–435.

- Lu, Z., Cui, W., Han, X., 2015. Integrated production and preventive maintenance scheduling for a single machine with failure uncertainty. *Computers & Industrial Engineering* 80, 236–244.
- Makis, V., Jiang, X., 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.
- McKone, K., Weiss, E., 2002. Guidelines for implementing predictive maintenance. *Production and Operations Management* 11 (2), 109–124.
- Mula, J., Poler, R., Garcia-Sabater, J. P., Lario, F. C., 2006. Models for production planning under uncertainty: A review. *International Journal of Production Economics* 103 (1), 271–285.
- Nourelfath, M., Yalaoui, F., 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.
- Panagiotidou, S., Tagaras, G., 2010. Statistical process control and condition-based maintenance: a meaningful relationship through data sharing. *Production and operations management* 19 (2), 156–171.
- Paté-Cornell, M., Lee, H., Tagaras, G., 1987. Warnings of malfunction: the decision to inspect and maintain production processes on schedule or on demand. *Management Science* 33 (10), 1277–1290.
- Peng, H., Van Houtum, G. J., 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.
- Rivera-Gómez, H., Montaña-Arango, O., Corona-Armenta, J. R., Garnica-González, J., Hernández-Gress, E. S., Barragán-Vite, I., 2018. Production and maintenance planning for a deteriorating system with operation-dependent defectives. *Applied Sciences* 8 (2), 165.
- Sethi, S. P., Yan, H., Zhang, H., Zhang, Q., 2002. Optimal and hierarchical controls in dynamic stochastic manufacturing systems: A survey. *Manufacturing & Service Operations Management* 4 (2), 133–170.
- 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., Shanthikumar, J. G., 2000. Combined production and maintenance scheduling for a multiple-product, single-machine production system. *Production and Operations Management* 9 (4), 379–399.
- Sloan, T. W., Shanthikumar, J. G., 2002. Using in-line equipment condition and yield information for maintenance scheduling and dispatching in semiconductor wafer fabs. *IIE Transactions* 34 (2), 191–209.
- Uit het Broek, M. A. J., Teunter, R. H., De Jonge, B., Veldman, J., Van Foreest, N. D., 2019. Condition-based production planning: Adjusting production rates to balance output and failure risk. *Manufacturing & Service Operations Management*, In press.
- Van Noortwijk, J. M., 2009. A survey of the application of gamma processes in maintenance. *Reliability Engineering and System Safety* 94 (1), 2–21.
- Wireman, T., 2014. *Benchmarking Best practices for maintenance, reliability and asset management*, 3rd Edition. Industrial Press Inc, South Norwalk, CT.
- Zied, H., Sofiene, D., Nidhal, R., 2011. Optimal integrated maintenance/production policy for randomly failing systems with variable failure rate. *International Journal of Production Research* 49 (19), 5695–5712.
- Zio, E., Compare, M., 2013. Evaluating maintenance policies by quantitative modeling and analysis. *Reliability Engineering & System Safety* 109, 53–65.



List of research reports

14001-OPERA: Germs, R. and N.D. van Foreest, Optimal control of production-inventory systems with constant and compound poisson demand

14002-EEF: Bao, T. and J. Duffy, Adaptive vs. educative learning: Theory and evidence

14003-OPERA: Syntetos, A.A. and R.H. Teunter, On the calculation of safety stocks

14004-EEF: Bouwmeester, M.C., J. Oosterhaven and J.M. Rueda-Cantuche, Measuring the EU value added embodied in EU foreign exports by consolidating 27 national supply and use tables for 2000-2007

14005-OPERA: Prak, D.R.J., R.H. Teunter and J. Riezebos, Periodic review and continuous ordering

14006-EEF: Reijnders, L.S.M., The college gender gap reversal: Insights from a life-cycle perspective

14007-EEF: Reijnders, L.S.M., Child care subsidies with endogenous education and fertility

14008-EEF: Otter, P.W., J.P.A.M. Jacobs and A.H.J. den Reijer, A criterion for the number of factors in a data-rich environment

14009-EEF: Mierau, J.O. and E. Suari Andreu, Fiscal rules and government size in the European Union

14010-EEF: Dijkstra, P.T., M.A. Haan and M. Mulder, Industry structure and collusion with uniform yardstick competition: theory and experiments

14011-EEF: Huizingh, E. and M. Mulder, Effectiveness of regulatory interventions on firm behavior: a randomized field experiment with e-commerce firms

14012-GEM: Bressand, A., Proving the old spell wrong: New African hydrocarbon producers and the 'resource curse'

14013-EEF: Dijkstra P.T., Price leadership and unequal market sharing: Collusion in experimental markets

14014-EEF: Angelini, V., M. Bertoni, and L. Corazzini, Unpacking the determinants of life satisfaction: A survey experiment

14015-EEF: Heijdra, B.J., J.O. Mierau, and T. Trimbom, Stimulating annuity markets

14016-GEM: Bezemer, D., M. Grydaki, and L. Zhang, Is financial development bad for growth?

14017-EEF: De Cao, E. and C. Lutz, Sensitive survey questions: measuring attitudes regarding female circumcision through a list experiment

14018-EEF: De Cao, E., The height production function from birth to maturity



14019-EEF: Allers, M.A. and J.B. Geertsema, The effects of local government amalgamation on public spending and service levels. Evidence from 15 years of municipal boundary reform

14020-EEF: Kuper, G.H. and J.H. Veurink, Central bank independence and political pressure in the Greenspan era

14021-GEM: Samarina, A. and D. Bezemer, Do Capital Flows Change Domestic Credit Allocation?

14022-EEF: Soetevent, A.R. and L. Zhou, Loss Modification Incentives for Insurers Under Expected Utility and Loss Aversion

14023-EEF: Allers, M.A. and W. Vermeulen, Fiscal Equalization, Capitalization and the Flypaper Effect.

14024-GEM: Hoorn, A.A.J. van, Trust, Workplace Organization, and Comparative Economic Development.

14025-GEM: Bezemer, D., and L. Zhang, From Boom to Bust in de Credit Cycle: The Role of Mortgage Credit.

14026-GEM: Zhang, L., and D. Bezemer, How the Credit Cycle Affects Growth: The Role of Bank Balance Sheets.

14027-EEF: Bružikas, T., and A.R. Soetevent, Detailed Data and Changes in Market Structure: The Move to Unmanned Gasoline Service Stations.

14028-EEF: Bouwmeester, M.C., and B. Scholtens, Cross-border Spillovers from European Gas Infrastructure Investments.

14029-EEF: Lestano, and G.H. Kuper, Correlation Dynamics in East Asian Financial Markets.

14030-GEM: Bezemer, D.J., and M. Grydaki, Nonfinancial Sectors Debt and the U.S. Great Moderation.

14031-EEF: Hermes, N., and R. Lensink, Financial Liberalization and Capital Flight: Evidence from the African Continent.

14032-OPERA: Blok, C. de, A. Seepma, I. Roukema, D.P. van Donk, B. Keulen, and R. Otte, Digitalisering in Strafrechtketens: Ervaringen in Denemarken, Engeland, Oostenrijk en Estland vanuit een Supply Chain Perspectief.

14033-OPERA: Olde Keizer, M.C.A., and R.H. Teunter, Opportunistic condition-based maintenance and aperiodic inspections for a two-unit series system.

14034-EEF: Kuper, G.H., G. Sierksma, and F.C.R. Spieksma, Using Tennis Rankings to Predict Performance in Upcoming Tournaments

15001-EEF: Bao, T., X. Tian, X. Yu, Dictator Game with Indivisibility of Money



- 15002-GEM: Chen, Q., E. Dietzenbacher, and B. Los, The Effects of Ageing and Urbanization on China's Future Population and Labor Force
- 15003-EEF: Allers, M., B. van Ommeren, and B. Geertsema, Does intermunicipal cooperation create inefficiency? A comparison of interest rates paid by intermunicipal organizations, amalgamated municipalities and not recently amalgamated municipalities
- 15004-EEF: Dijkstra, P.T., M.A. Haan, and M. Mulder, Design of Yardstick Competition and Consumer Prices: Experimental Evidence
- 15005-EEF: Dijkstra, P.T., Price Leadership and Unequal Market Sharing: Collusion in Experimental Markets
- 15006-EEF: Anufriev, M., T. Bao, A. Sutin, and J. Tuinstra, Fee Structure, Return Chasing and Mutual Fund Choice: An Experiment
- 15007-EEF: Lamers, M., Depositor Discipline and Bank Failures in Local Markets During the Financial Crisis
- 15008-EEF: Oosterhaven, J., On de Doubtful Usability of the Inoperability IO Model
- 15009-GEM: Zhang, L. and D. Bezemer, A Global House of Debt Effect? Mortgages and Post-Crisis Recessions in Fifty Economies
- 15010-I&O: Hooghiemstra, R., N. Hermes, L. Oxelheim, and T. Randøy, The Impact of Board Internationalization on Earnings Management
- 15011-EEF: Haan, M.A., and W.H. Siekman, Winning Back the Unfaithful while Exploiting the Loyal: Retention Offers and Heterogeneous Switching Costs
- 15012-EEF: Haan, M.A., J.L. Moraga-González, and V. Petrikaite, Price and Match-Value Advertising with Directed Consumer Search
- 15013-EEF: Wiese, R., and S. Eriksen, Do Healthcare Financing Privatisations Curb Total Healthcare Expenditures? Evidence from OECD Countries
- 15014-EEF: Siekman, W.H., Directed Consumer Search
- 15015-GEM: Hoorn, A.A.J. van, Organizational Culture in the Financial Sector: Evidence from a Cross-Industry Analysis of Employee Personal Values and Career Success
- 15016-EEF: Te Bao, and C. Hommes, When Speculators Meet Constructors: Positive and Negative Feedback in Experimental Housing Markets
- 15017-EEF: Te Bao, and Xiaohua Yu, Memory and Discounting: Theory and Evidence
- 15018-EEF: Suari-Andreu, E., The Effect of House Price Changes on Household Saving Behaviour: A Theoretical and Empirical Study of the Dutch Case
- 15019-EEF: Bijlsma, M., J. Boone, and G. Zwart, Community Rating in Health Insurance: Trade-off between Coverage and Selection



15020-EEF: Mulder, M., and B. Scholtens, A Plant-level Analysis of the Spill-over Effects of the German *Energiewende*

15021-GEM: Samarina, A., L. Zhang, and D. Bezemer, Mortgages and Credit Cycle Divergence in Eurozone Economies

16001-GEM: Hoorn, A. van, How Are Migrant Employees Manages? An Integrated Analysis

16002-EEF: Soetevent, A.R., Te Bao, A.L. Schippers, A Commercial Gift for Charity

16003-GEM: Bouwmeester, M.C., and J. Oosterhaven, Economic Impacts of Natural Gas Flow Disruptions

16004-MARK: Holtrop, N., J.E. Wieringa, M.J. Gijsenberg, and P. Stern, Competitive Reactions to Personal Selling: The Difference between Strategic and Tactical Actions

16005-EEF: Plantinga, A. and B. Scholtens, The Financial Impact of Divestment from Fossil Fuels

16006-GEM: Hoorn, A. van, Trust and Signals in Workplace Organization: Evidence from Job Autonomy Differentials between Immigrant Groups

16007-EEF: Willems, B. and G. Zwart, Regulatory Holidays and Optimal Network Expansion

16008-GEF: Hoorn, A. van, Reliability and Validity of the Happiness Approach to Measuring Preferences

16009-EEF: Hinloopen, J., and A.R. Soetevent, (Non-)Insurance Markets, Loss Size Manipulation and Competition: Experimental Evidence

16010-EEF: Bekker, P.A., A Generalized Dynamic Arbitrage Free Yield Model

16011-EEF: Mierau, J.A., and M. Mink, A Descriptive Model of Banking and Aggregate Demand

16012-EEF: Mulder, M. and B. Willems, Competition in Retail Electricity Markets: An Assessment of Ten Year Dutch Experience

16013-GEM: Rozite, K., D.J. Bezemer, and J.P.A.M. Jacobs, Towards a Financial Cycle for the US, 1873-2014

16014-EEF: Neuteleers, S., M. Mulder, and F. Hindriks, Assessing Fairness of Dynamic Grid Tariffs

16015-EEF: Soetevent, A.R., and T. Bružikas, Risk and Loss Aversion, Price Uncertainty and the Implications for Consumer Search

16016-HRM&OB: Meer, P.H. van der, and R. Wielers, Happiness, Unemployment and Self-esteem



16017-EEF: Mulder, M., and M. Pangan, Influence of Environmental Policy and Market Forces on Coal-fired Power Plants: Evidence on the Dutch Market over 2006-2014

16018-EEF: Zeng, Y., and M. Mulder, Exploring Interaction Effects of Climate Policies: A Model Analysis of the Power Market

16019-EEF: Ma, Yiqun, Demand Response Potential of Electricity End-users Facing Real Time Pricing

16020-GEM: Bezemer, D., and A. Samarina, Debt Shift, Financial Development and Income Inequality in Europe

16021-EEF: Elkhuizen, L, N. Hermes, and J. Jacobs, Financial Development, Financial Liberalization and Social Capital

16022-GEM: Gerritse, M., Does Trade Cause Institutional Change? Evidence from Countries South of the Suez Canal

16023-EEF: Rook, M., and M. Mulder, Implicit Premiums in Renewable-Energy Support Schemes

17001-EEF: Trinks, A., B. Scholtens, M. Mulder, and L. Dam, Divesting Fossil Fuels: The Implications for Investment Portfolios

17002-EEF: Angelini, V., and J.O. Mierau, Late-life Health Effects of Teenage Motherhood

17003-EEF: Jong-A-Pin, R., M. Laméris, and H. Garretsen, Political Preferences of (Un)happy Voters: Evidence Based on New Ideological Measures

17004-EEF: Jiang, X., N. Hermes, and A. Meesters, Financial Liberalization, the Institutional Environment and Bank Efficiency

17005-EEF: Kwaak, C. van der, Financial Fragility and Unconventional Central Bank Lending Operations

17006-EEF: Postelnicu, L. and N. Hermes, The Economic Value of Social Capital

17007-EEF: Ommeren, B.J.F. van, M.A. Allers, and M.H. Vellekoop, Choosing the Optimal Moment to Arrange a Loan

17008-EEF: Bekker, P.A., and K.E. Bouwman, A Unified Approach to Dynamic Mean-Variance Analysis in Discrete and Continuous Time

17009-EEF: Bekker, P.A., Interpretable Parsimonious Arbitrage-free Modeling of the Yield Curve

17010-GEM: Schasfoort, J., A. Godin, D. Bezemer, A. Caiani, and S. Kinsella, Monetary Policy Transmission in a Macroeconomic Agent-Based Model

17011-I&O: Bogt, H. ter, Accountability, Transparency and Control of Outsourced Public Sector Activities



- 17012-GEM: Bezemer, D., A. Samarina, and L. Zhang, The Shift in Bank Credit Allocation: New Data and New Findings
- 17013-EEF: Boer, W.I.J. de, R.H. Koning, and J.O. Mierau, Ex-ante and Ex-post Willingness-to-pay for Hosting a Major Cycling Event
- 17014-OPERA: Laan, N. van der, W. Romeijnders, and M.H. van der Vlerk, Higher-order Total Variation Bounds for Expectations of Periodic Functions and Simple Integer Recourse Approximations
- 17015-GEM: Oosterhaven, J., Key Sector Analysis: A Note on the Other Side of the Coin
- 17016-EEF: Romensen, G.J., A.R. Soetevent: Tailored Feedback and Worker Green Behavior: Field Evidence from Bus Drivers
- 17017-EEF: Trinks, A., G. Ibikunle, M. Mulder, and B. Scholtens, Greenhouse Gas Emissions Intensity and the Cost of Capital
- 17018-GEM: Qian, X. and A. Steiner, The Reinforcement Effect of International Reserves for Financial Stability
- 17019-GEM/EEF: Klasing, M.J. and P. Milionis, The International Epidemiological Transition and the Education Gender Gap
- 2018001-EEF: Keller, J.T., G.H. Kuper, and M. Mulder, Mergers of Gas Markets Areas and Competition amongst Transmission System Operators: Evidence on Booking Behaviour in the German Markets
- 2018002-EEF: Soetevent, A.R. and S. Adikyan, The Impact of Short-Term Goals on Long-Term Objectives: Evidence from Running Data
- 2018003-MARK: Gijsenberg, M.J. and P.C. Verhoef, Moving Forward: The Role of Marketing in Fostering Public Transport Usage
- 2018004-MARK: Gijsenberg, M.J. and V.R. Nijs, Advertising Timing: In-Phase or Out-of-Phase with Competitors?
- 2018005-EEF: Hulshof, D., C. Jepma, and M. Mulder, Performance of Markets for European Renewable Energy Certificates
- 2018006-EEF: Fosgaard, T.R., and A.R. Soetevent, Promises Undone: How Committed Pledges Impact Donations to Charity
- 2018007-EEF: Durán, N. and J.P. Elhorst, A Spatio-temporal-similarity and Common Factor Approach of Individual Housing Prices: The Impact of Many Small Earthquakes in the North of Netherlands
- 2018008-EEF: Hermes, N., and M. Hudon, Determinants of the Performance of Microfinance Institutions: A Systematic Review
- 2018009-EEF: Katz, M., and C. van der Kwaak, The Macroeconomic Effectiveness of Bank Bail-ins



2018010-OPERA: Prak, D., R.H. Teunter, M.Z. Babai, A.A. Syntetos, and J.E. Boylan, Forecasting and Inventory Control with Compound Poisson Demand Using Periodic Demand Data

2018011-EEF: Brock, B. de, Converting a Non-trivial Use Case into an SSD: An Exercise

2018012-EEF: Harvey, L.A., J.O. Mierau, and J. Rockey, Inequality in an Equal Society

2018013-OPERA: Romeijnders, W., and N. van der Laan, Inexact cutting planes for two-stage mixed-integer stochastic programs

2018014-EEF: Green, C.P., and S. Homroy, Bringing Connections Onboard: The Value of Political Influence

2018015-OPERA: Laan, N. van der, and W. Romeijnders, Generalized alpha-approximations for two-stage mixed-integer recourse models

2018016-GEM: Rozite, K., Financial and Real Integration between Mexico and the United States

2019001-EEF: Lugalla, I.M., J. Jacobs, and W. Westerman, Drivers of Women Entrepreneurs in Tourism in Tanzania: Capital, Goal Setting and Business Growth

2019002-EEF: Brock, E.O. de, On Incremental and Agile Development of (Information) Systems

2019003-OPERA: Laan, N. van der, R.H. Teunter, W. Romeijnders, and O.A. Kilic, The Data-driven Newsvendor Problem: Achieving On-target Service Levels.

2019004-EEF: Dijk, H., and J. Mierau, Mental Health over the Life Course: Evidence for a U-Shape?

2019005-EEF: Freriks, R.D., and J.O. Mierau, Heterogeneous Effects of Conditional School Resources on Health and Test Scores: Evidence from the Netherlands.

2019006-OPERA: Broek, M.A.J. uit het, R.H. Teunter, B. de Jonge, J. Veldman, Joint Condition-based Maintenance and Condition-based Production Optimization.



www.rug.nl/feb