



university of
 groningen

faculty of economics
 and business

2019007-OPERA

**Joint Condition-based Maintenance
 and Load-sharing Optimization for
 Multi-unit Systems with Economic
 Dependency**

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 Load-sharing Optimization for Multi-unit Systems with Economic Dependency

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 load-sharing optimization for multi-unit systems with economic dependency

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

Many production facilities consist of multiple units of equipment, such as pumps or turbines, that are jointly used to satisfy a given production target. Such systems often have overcapacity to ensure high levels of reliability and availability. The deterioration rates of the units typically depend on their production rates, implying that the operator can control deterioration by dynamically reallocating load among units. In this study, we examine the value of condition-based load-sharing decisions for multi-unit systems with economic dependency. We formulate the system as a Markov decision process and provide optimal joint condition-based maintenance and production policies. Our numerical results show that, dependent on the system characteristics, substantial cost savings of up to 40% can be realized compared to the optimal condition-based maintenance policy under equal load-sharing. The structure of the optimal policy particularly depends on the maintenance setup cost and the penalty that is incurred if the production target is not satisfied. For systems with high setup costs, the clustering of maintenance interventions is improved by synchronizing the deterioration of the units. On the contrary, for low setup costs, the deterioration levels are desynchronized and the maintenance interventions are alternated.

Keywords: condition-based maintenance, condition-based production, load-sharing, economic dependency, clustering, production-dependent deterioration

1. Introduction

Many production facilities consist of multiple units that are jointly used to satisfy a production target. These units deteriorate due to load and stress caused by production and eventually require maintenance in order to keep the system in, or bring it back to, an operating condition. The resulting maintenance expenses often constitute a substantial part of the total budget of production facilities, and can even form up to 70 percent of the total production costs (Bevilacqua and Braglia, 2000). Many studies aim to reduce these costs by developing condition-based maintenance policies and show that such policies reduce costs while improving availability and productivity.

Another option to improve the cost efficiency of production facilities is to control the deterioration of its units by adopting condition-based production policies (Uit het Broek et al., 2019a,b). Such policies exploit the relation between the production rate and the deterioration rate by dynamically adjusting the production rate based on condition information. Although others have shown the effectiveness of condition-based production policies for single-unit systems, there are, to the best of our knowledge, no studies devoted to condition-based production policies for multi-unit systems that consider dynamic reallocation of load among units. Optimal maintenance policies for multi-unit systems are often more advanced than for single-unit systems because of the various types of dependencies that exists between

*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)

units (Olde Keizer et al., 2017). It is therefore also expected that condition-based production policies will be different for multi-unit systems.

The most commonly studied dependency is positive economic dependency such as a fixed maintenance setup cost that is independent of the number of units that are maintained. In such cases, clustering maintenance interventions for various units is often more cost efficient than performing them separately. However, clustering maintenance for units with different degradation levels implies that maintenance is performed unnecessarily early for units with relatively low levels of deterioration. In such situations, an interesting question is whether it can be profitable to control the deterioration processes by reallocating load from highly deteriorated units among the other units. Hereby the operator can actively synchronize the deterioration levels of the units, which introduces opportunities to improve the clustering of maintenance interventions.

In this study, we present a first exploration of the benefits of condition-based load-sharing decisions for multi-unit systems with economic dependency. The deterioration rates of the units depend on their respective loads, implying that the operator can control their deterioration by dynamically reallocating load among units. We formulate the problem as a Markov decision process and use this to determine optimal maintenance and production policies. Our results show that condition-based load-sharing improves the effectiveness of condition-based maintenance policies, and that its effectiveness heavily depends on the degree of overcapacity. Throughout this study, we use the term *overcapacity* to refer to systems with some overcapacity, but for which all units are needed to satisfy the production target. Furthermore, by *redundancy* we refer to systems with sufficient overcapacity to still reach the production target if one or more machines are not functioning. Substantial cost savings up to 20% can be obtained for systems with overcapacity, and these savings increase up to 40% for systems with redundancy. The savings are the result of fewer failures, fewer maintenance actions per unit, improved maintenance clustering, and reduced risks of production shortages.

An insightful observation is that condition-based load-sharing policies are also effective for systems without economic dependency. For such systems, cost savings are possible by actively desynchronizing the deterioration levels of the units. Moreover, for many systems there are scenarios in which the most deteriorated unit takes over load from the least deteriorated unit. An interesting side effect of adopting condition-based load-sharing policies is that doing so not only reduces the expected cost but also its variance, implying higher financial robustness.

The remainder of this study is organized as follows. In Section 2, we discuss the literature on maintenance and production decisions and specifically address studies that consider multi-unit systems with dependency between the units. In Section 3, we formally describe the system that we consider. The Markov decision process formulation used to obtain optimal policies is given in Section 4. In Sections 5 – 7, we examine the structure of the optimal policies and the associated cost savings. We conclude and provide future research opportunities in Section 8.

2. Literature review

In this study, we introduce condition-based load-sharing decisions and combine this with condition-based maintenance, redundancy, and economic dependency. For extensive reviews on condition-based maintenance we refer to Alaswad and Xiang (2017) and De Jonge and Scarf (2019). For a review on condition-based maintenance for multi-unit systems with dependencies we refer to Olde Keizer et al. (2017). In the remainder of our literature review, we first discuss studies on condition-based maintenance that also include redundancy or economic dependency. Then we zoom in on studies with load sharing, which can be divided into failure-based and degradation-based load sharing. In both streams, the load sharing dynamics are exogeneously given and cannot be used as a feature to control the deterioration of

units. The last stream that we discuss are studies that examine condition-based production policies for single-unit systems.

The literature on condition-based maintenance for multi-unit systems is rich, and both redundancy (Lu and Jiang, 2007; Wang et al., 2009) and economic dependency (Castanier et al., 2005; Do et al., 2013; De Jonge et al., 2016) are addressed in various settings. Also the joint effect of redundancy and economic dependency is studied, including 1-out-of- N systems (Li et al., 2016), k -out-of- N systems (Olde Keizer et al., 2016), and series-parallel systems (Zhou et al., 2013). The aforementioned studies investigate condition-based maintenance policies for multi-unit systems with either redundancy, economic dependency, or both, but none of them include the effect of load sharing. The observation that research on the integration of condition-based maintenance with load sharing is lacking is also brought forward by Olde Keizer et al. (2017, 2018).

Others have addressed multi-unit systems with failure-based load sharing and degradation processes that can be monitored. Under failure-based load sharing, the total load is equally shared among all functioning units and thus the load faced by a unit can only change upon failure of another unit. Zhang et al. (2014, 2015) investigate maintenance policies with an opportunistic threshold for preventive maintenance. They consider a system whose units deteriorate with a nominal rate as long as all units are functioning and the deterioration rate of all units accelerate once at least one unit has failed. Marseguerra et al. (2002) analyze condition-based maintenance policies for series and parallel systems. They consider policies in which the maintenance decision for a unit only depends on its own health and not on the complete system state. Olde Keizer et al. (2018) examine optimal condition-based maintenance policies for 1-out-of- N systems with economic dependency and load sharing. They model the deterioration rate of units as a function of the number of functioning units. Their results show that it is important to base decisions on the complete system state and that load-sharing effects should not be ignored in making those decisions. They also find that postponing maintenance of failed units can be cost-effective in order to improve the clustering of maintenance tasks. Zhao et al. (2018) consider the reliability of a multi-unit system whose units deteriorate according to a Brownian motion. In these studies, the total load processed by the system is constant over time and is equally shared among the functioning units. Hence, reallocating load is triggered by failures only and is not used as an opportunity to dynamically control the deterioration processes of units.

Another research stream that includes load sharing is degradation-based load sharing. In contrast to failure-based load sharing, load is not reallocated upon failure, but the load of units gradually increases when the deterioration level of other units increases. Many settings are addressed in this research stream, including settings with condition monitoring and with economic dependency (see, e.g., Do et al., 2015, 2019; Rasmekomen and Parlikad, 2016; Zhou et al., 2016). Studies in this stream clearly differ from our research since, similar to the failure-based load sharing stream, the deterioration processes are not controlled by dynamically reallocating load among units.

All the above-mentioned studies consider condition-based maintenance policies for systems with load sharing. However, none of them utilize condition information to determine the load applied to a unit by controlling the production rate. We note that the static equal load-sharing rule as addressed by the above studies is realistic for many practical systems. For instance, if one cable of a cable-supported bridge fails, this increases the load faced by the other cables and an operator can not dynamically decide which cable should take over the load. In practice, however, there are also many examples where the operator can determine how the total load is allocated among units. This holds in particular for manufacturing systems, systems used in the process industry, and energy systems. For instance, for a wind farm only the total production at the system level is relevant, and not how the load is distributed among the individual turbines.

Uit het Broek et al. (2019b) study condition-based production rates for a single-unit system for which the next maintenance intervention is already scheduled. The production rate directly affects the deterio-

ration rate and can thus be used to control the deterioration process. Uit het Broek et al. (2019a) extend this to the joint optimization of condition-based maintenance and production. Their study shows that condition-based production and maintenance decisions can complement each other and that the effectiveness of both strongly depends on various characteristics of the system. There are some other studies on condition-based production policies, but these assume that the production rate does not affect the deterioration rate of the system (see, e.g., Iravani and Duenyas, 2002; Sloan, 2004). Although these studies consider condition-based production, none of them addresses the value of dynamically sharing load between multiple units.

We conclude that condition-based maintenance, redundancy, economic dependency, and load sharing are well studied in isolation, but are scarcely jointly addressed. Moreover, condition-based production rate decisions have received little attention, even in isolation of the other effects. To the best of our knowledge, there is no study that examines the value of dynamically redistributing load among units based on condition information. The aim of our study is to combine these effects and to explore the potentials of condition-based load-sharing control in such settings.

3. Problem description

We consider a multi-unit system consisting of n identical units. The production rate of each unit is adjustable over time and affects the deterioration rate of the unit. There is an economic dependency between the units as carrying out maintenance incurs a fixed setup cost, independent of the number of units that are maintained. Our aim is to determine the joint condition-based maintenance and production policy that minimizes the long-run cost rate while an overall production target is taken into account. In the remainder of this section, we first introduce the possible states of the system, followed by the admissible decisions and how these affect the state of the system. Thereafter the costs corresponding to these actions are given.

The deterioration process of each unit is continuously monitored and is described by a continuous-time continuous-state nondecreasing stochastic process $X_i = \{X_i(t) \geq 0 \mid t \geq 0\}$. In this process, deterioration level 0 represents that the unit is as-good-as new and the unit fails when its deterioration level exceeds a fixed failure level L . The condition of the complete system is denoted as $x = (x_1, \dots, x_n)$.

The operator can schedule maintenance interventions based on the condition of the system at any time. Maintenance actions are assumed to be perfect, that is, they restore the condition of a unit to the as-good-as-new condition. Furthermore, a planning time s is required to schedule a maintenance intervention. The maintenance actions themselves are assumed to require a negligible amount of time. This is often realistic as repair times are typically hours to days whereas expected lifetimes are often in the order of years. Planning maintenance, however, can take several months due to lengthy lead times for specialized tools and equipment, and therefore we do consider a planning time in our model. Moreover, the operator has the flexibility to decide which specific units are maintained after the planning time.

At any time, the operator can control the production rate u_i of each unit. The possible production rates $u_i \in [0, 1]$ range from the idle mode to producing at full speed. Naturally, units in the failed state cannot produce and their production rate is fixed to zero. The production rate of a unit directly affects its deterioration rate, that is, the average amount of additional deterioration per period. In accordance with Uit het Broek et al. (2019a,b), we use a function g that describes the production-deterioration relation (pd-relation in short). When a unit produces at rate u_i , its deterioration rate equals $g(u_i)$. Moreover, we assume that units deteriorate faster when producing at higher rates, and thus the pd-relation g is an increasing function. For clarity, we denote the minimum and maximum deterioration rate by $\mu_{\min} = g(0)$ and $\mu_{\max} = g(1)$, respectively.

The cost of maintaining a unit depends on its condition at the moment of maintenance. Maintaining a functioning unit is referred to as preventive maintenance and costs c_{pm} . Corrective maintenance is

required for a unit that has failed and costs $c_{\text{cm}} \geq c_{\text{pm}}$. The fixed setup cost for maintenance is denoted by c_{setup} and is incurred at the moment that a maintenance intervention is scheduled.

The costs corresponding to the production decision depend on the total production rate $\hat{u} = \sum_{i=1}^n u_i$ and a given production target κ . A constant penalty $\tilde{\pi}$ is incurred per time period that the target is not satisfied, i.e., if $\hat{u} < \kappa$. Moreover, there is a variable penalty per time period π_1 and a bonus per time period π_2 that are proportional to the shortage and overproduction, respectively. Note that this cost structure is flexible and allows us to study systems with both hard and soft production constraints, and systems where failures are severe or not. For instance, systems for which shortages are avoided at all costs and for which there is no benefit of overproduction (e.g., gas turbines that must provide a reliable gas flow with a steady pressure) can be analyzed by using an extremely high constant penalty $\tilde{\pi}$ and a bonus of $\pi_2 = 0$. Production facilities that purely maximize profit (i.e., production revenues minus maintenance costs) can be analyzed by setting the production target κ equal to the maximum production capacity, the constant penalty $\tilde{\pi}$ that is incurred when the target is not satisfied to zero, and the variable penalty for shortages π_1 equal to the production value. Systems that aim to minimize costs under a given production target while being able to sell overproduction for lower prices (e.g., offshore wind farms) can be analyzed by choosing positive values for $\tilde{\pi}$, π_1 , and π_2 .

4. Markov decision process formulation

We formulate the system as a Markov decision process (MDP) in order to determine optimal policies. An MDP is defined by a set of decision epochs, a finite set of system states, a finite set of admissible actions per state, and state- and action-dependent transition probabilities and immediate costs. In the remainder of this section, we first introduce the discretized states and the corresponding admissible actions. Thereafter, we give an overview of each decision epoch and the corresponding Bellman equations. We end the section with the algorithm that we use to determine optimal policies.

4.1. Discretization

Time is discretized into periods of length Δt , chosen in such a way that the planning time s is a multiple of Δt . The continuous set of production rates is discretized into m positive production rates plus the idle mode, that is, $U = \{i/m \mid i = 0, \dots, m\}$. The deterioration interval $[0, L]$ is discretized into $\bar{X} = \{(i + 0.5)\Delta X \mid i = 0, \dots, L/\Delta X\}$ where ΔX is selected such that L is a multiple of ΔX . The discrete deterioration levels are represented by an index ranging from 0 (index of the as-good-as-new condition) to ℓ (index of the failed condition).

Similar to Uit het Broek et al. (2019a,b), who considered single-unit systems, we let $F_{u,\Delta t}$ be the cumulative distribution function of the deterioration increment for a single unit in a period Δt while it produces at rate u . For ease of notation, we generally omit the subscripts u and Δt in the remainder of this study. We remark that we only consider nondecreasing deterioration processes and thus $F(0) = 0$. The probability of a zero deterioration increment is modeled as $F(0.5\Delta X)$. The probability to jump from deterioration level k to $k+i$, given that $i \geq 1$ and $k+i < \ell$, is modeled as $F((i+0.5)\Delta X) - F((i-0.5)\Delta X)$. The probability to jump from deterioration level k to the failed level ℓ equals the remaining probability $1 - F((i-0.5)\Delta X)$. Thus the probability $P(k, k+i)$ to jump from deterioration level k to deterioration level $k+i$ equals

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

4.2. The value functions

In order to ease the analysis, we model each decision epoch as a sequence of three consecutive stages, see Figure 1. In the first stage, the system state is observed and we determine whether a new maintenance intervention will be scheduled. In the second stage, we determine whether maintenance will be carried out. In the third stage, we choose the production rates of the units, and we model the deterioration of the system. We let v , w_1 , and w_2 denote the value functions at the start of the three stages, respectively. In what follows, we discuss the stages in more detail and we provide explicit formulations for the three value functions.




Figure 1: Order of events in each decision epoch.

Stage 1: Observe state and schedule maintenance

At the start of each epoch, the deterioration levels x and the remaining planning time until the next scheduled maintenance intervention, denoted as $\tau \in \{1, \dots, s/\Delta t, ns\}$, are observed. Notation ns is used to indicate that no maintenance intervention is currently scheduled. We remark that $\tau = 0$ is not possible at this stage as will be explained later.

When the next maintenance invention is already scheduled (i.e., $\tau \neq ns$), there is no decision to be made and the remaining planning time is reduced by one, thus $v(x, \tau | \tau \neq ns) = w_1(x, \tau - 1)$. When maintenance has not been scheduled yet (i.e., $\tau = ns$), the operator has to decide whether an intervention will be scheduled or not. In case no new intervention is scheduled, both the remaining planning time τ and the deterioration state x remain unaltered. In case maintenance will be scheduled, the maintenance setup cost c_{setup} is incurred and the remaining planning time is set to $\tau = s$. The value function thus equals $v(x, \tau | \tau = ns) = \min\{w_1(x, ns), c_{\text{setup}} + w_1(x, s)\}$.

Summarizing, the value function v equals

$$v(x, \tau) = \begin{cases} \min\{w_1(x, ns), c_{\text{setup}} + w_1(x, s)\} & \text{if } \tau = ns, \\ w_1(x, \tau - 1) & \text{otherwise.} \end{cases} \quad (1)$$

Stage 2: Carry out maintenance

The second step is to determine whether to carry out maintenance, which is only possible at the end of the planning time. Recall that $w_1(x, \tau)$ represents the value function after the decision whether to schedule a maintenance intervention has been made. Since maintenance can only be performed at the end of the planning time, there is no decision in this stage when $\tau \neq 0$. It follows that $w_1(x, \tau | \tau \neq 0) = w_2(x, \tau)$.

When $\tau = 0$, the operator decides which units are maintained. We denote the maintenance decision as $r = (r_1, \dots, r_n)$, where $r_i = 1$ if unit i is maintained and $r_i = 0$ if not. The set $\mathcal{R} = \{0, 1\}^n$ denotes the set of all possible maintenance decisions. Maintenance restores a unit to the as-good-as-new condition and thus the post-maintenance condition for unit i equals

$$x'_i(x_i, r_i) = \begin{cases} x_i & \text{if } r_i = 0, \\ 0 & \text{if } r_i = 1. \end{cases}$$

We denote the deterioration levels of the whole system after maintenance action r as $x'(x, r)$. Furthermore, regardless of decision r , the remaining planning time is reset to $\tau = \text{ns}$ to indicate that the next maintenance intervention is not scheduled yet. This also explains why an epoch can never start with $\tau = 0$. The direct costs incurred by performing maintenance action r depend on the system condition x and equals

$$\varphi_1(x, r) = \sum_{i=1}^n r_i \tilde{c}(x_i),$$

where $\tilde{c}(x_i) = c_{\text{pm}}$ if $x_i < \ell$ and $\tilde{c}(x_i) = c_{\text{cm}}$ if $x_i = \ell$. The value function w_1 given that $\tau = 0$ thus equals $w_1(x, \tau | \tau = 0) = \min_{r \in \mathcal{R}} \{\varphi_1(x, r) + w_2(x'(x, r), \text{ns})\}$.

Summarizing, the value function w_1 equals

$$w_1(x, \tau) = \begin{cases} w_2(x, \tau) & \text{if } \tau \neq 0, \\ \min_{r \in \mathcal{R}} \{\varphi_1(x, r) + w_2(x'(x, r), \text{ns})\} & \text{if } \tau = 0. \end{cases} \quad (2)$$

Stage 3: Production decision and deterioration

In the final stage, the operator selects the production rates of the functioning units while the production rates of failed units are fixed to zero. The function $w_2(x, \tau)$ represents the value function of the post-decision state after maintenance has been performed. Remark that the remaining planning time τ is only decremented in the first stage and that the second stage resets it to ns at the end of the planning time; hence, $\tau = 0$ is not possible at the start of this stage.

We let $\mathcal{U}(x)$ denote the set of all admissible production decisions given that the system is in deterioration condition x . The production decision $u \in \mathcal{U}(x)$ affects both the direct cost $\varphi_2(u)$ and the expected deterioration increments.

The direct costs consist of a possible fixed and variable penalty if the production target is not satisfied and a bonus in case of overproduction. To define the direct cost function, we let \mathbb{I}_A be an indicator function that equals one if condition A is true and zero otherwise. Recall that $\hat{u} = \sum_{i=1}^n u_i$ equals the total production rate of the system as defined in Section 3. Now we have

$$\varphi_2(\hat{u}) = \mathbb{I}_{\hat{u} < \kappa} (\tilde{\pi} + (\kappa - \hat{u})\pi_1) \Delta t + \mathbb{I}_{\hat{u} > \kappa} (\hat{u} - \kappa)\pi_2 \Delta t.$$

We let $\mathcal{X}'(x) = \{(x'_1, \dots, x'_n) \mid x_i \leq x'_i \leq \ell\}$ denote the set of all reachable deterioration states from state x . Note that, although the deterioration increment probabilities depend on the selected production rates u , the set of reachable states only depends on the current state x . The value function w_2 equals

$$w_2(x, \tau) = \min_{u \in \mathcal{U}(x)} \left\{ \varphi_2 \left(\sum_{i=1}^n u_i \right) + \sum_{x' \in \mathcal{X}'(x)} P_{u, \Delta t}(x, x') v(x', \tau) \right\}. \quad (3)$$

4.3. Modified policy iteration

We use modified policy iteration, an algorithm that combines value iteration with policy iteration, to find stationary ϵ -optimal policies for the value functions given in Section 4.2. In general, policy iteration spends most of the time in exactly solving the value functions for a given policy, whereas value iteration is computationally expensive because it considers all possible policies in each iteration and typically requires many iterations to converge. Puterman (1994, pp 386) describes a modified policy iteration algorithm to accelerate the convergence rate by combining both algorithms. The intuition behind this approach is to apply policy iteration but instead of solving the exact values for v , w_1 , and w_2 , the values are approximated by value iteration while the policy is kept fixed for a number of successive iterations.

The modified policy iteration that we use is provided in the Appendix in Algorithm 3. We let \bar{v} denote the value function after an iteration that starts with value function v . The algorithm starts with initializing $v(x, \tau) = 0$ for all x and τ , and iteratively updates the best actions and corresponding values for each state. The difference with the default value iteration algorithm is that not all admissible actions are considered in each iteration. Instead, the current best policy is fixed for a number of iterations, followed by a single iteration that considers all policies. The algorithm stops if the span, defined as $sp(w) = \max_{x,\tau} w(x, \tau) - \min_{x,\tau} w(x, \tau)$ where $w = \bar{v} - v$, is smaller than a given positive number $\epsilon > 0$. The optimal long-run cost rate g^* is then estimated as

$$g = (\min\{\bar{v} - v\} + \max\{\bar{v} - v\}) / 2, \quad (4)$$

for which holds that $|g - g^*| < \epsilon/2$.

5. Setup numerical experiments

We examine the value of dynamically reallocating load among units based on condition information by comparing the optimal joint condition-based load-sharing and maintenance policy to a policy that only uses condition information to schedule maintenance and that equally shares load among the functioning units. We refer to the former as the *condition-based load-sharing policy* and to the latter as the *equal load-sharing policy*. We note that this benchmark policy equals the optimal condition-based maintenance policy studied by Olde Keizer et al. (2018), which they showed to be much more effective for systems with load sharing than other commonly applied maintenance policies.

In Section 5.1, we introduce the gamma process that we use to model deterioration of the units. Thereafter, in Section 5.2, we introduce two base systems that are characterized by their production contracts that prescribe a production target and the associated penalties if the target is not satisfied. The first contract type models a system with some overcapacity and a small penalty if the fixed production target is not satisfied. The second contract type models a system that primarily focuses on reliability, which is done by including redundancy and incurring an extremely high penalty if the target is not satisfied.

The structure of the optimal policy under both contract types and the corresponding cost savings compared to the equal load-sharing policy will be discussed in Sections 6 and 7. In these sections, we will provide many illustrations of the optimal policies in order to give a clear insight into how the optimal policy is affected by the various system parameters.

5.1. Deterioration process

We use stationary gamma processes to model deterioration as these are suitable to model monotonically increasing deterioration processes such as wear, erosion, and fatigue (Van Noortwijk, 2009). Moreover, the gamma process is flexible and allows to examine deterioration processes with different characteristics as its rate and volatility can be controlled by two parameters.

A gamma process consist of independently gamma distributed increments. For the gamma density function, we use the same parametric form with shape $\alpha > 0$ and scale $\beta > 0$ as Uit het Broek et al. (2019a). For the deterioration increment per time unit Y we have $E[Y] = \alpha\beta$ and $\text{Var}(Y) = \alpha\beta^2$. The density function for the increments in a time interval with length Δt is obtained by rescaling the shape parameter to $\alpha\Delta t$.

We let the production rate affect the deterioration increments of a unit in such a way that the expected deterioration increment per time unit equals $E[Y | u_i] = g(u_i)$, the variance of the deterioration increments while producing at the maximum rate equals $\text{Var}(Y | u_i = 1) = \sigma_{\max}^2$, and the coefficient of variation is independent of the production rate. These three properties are obtained by setting $\alpha = \mu_{\max}^2 / \sigma_{\max}^2$ and $\beta(u_i) = g(u_i) \sigma_{\max}^2 / \mu_{\max}^2$. If verifying this, recall that $g(1) = \mu_{\max}$.

Table 1: Parameter values for the base case excluding those for the production contract

Parameter	Value	Interpretation
μ_{\min}	0.5	Deterioration rate when idle
μ_{\max}	5.0	Deterioration rate at maximum production rate
σ_{\max}	6.0	St.dev. deterioration increments at maximum rate
γ	2.0	Shape pd-relation
s	1.0	Planning time for maintenance
c_{setup}	3.0	Maintenance set-up costs
c_{pm}	5.0	Preventive maintenance costs
c_{cm}	20.0	Corrective maintenance costs
L	100.0	Failure level
η	20	Number of positive production rates
Δt	1.0	Partitioning size time horizon
Δx	1.0	Partitioning size deterioration levels
n	2.0	Number of units

Table 2: Parameter values for the two contract types

Parameter	Type I	Type II	Interpretation
κ	1.6	1.0	Production target
$\tilde{\pi}$	10.0	10^6	Fixed penalty for production shortages
π_1	1.0	1.0	Variable penalty for production shortages
π_2	0.0	0.0	Bonus for producing more than κ

5.2. Base systems

The parameter values for the base case considered in this study are listed in Tables 1 and 2. We model the pd-relation by $g(u_i) = \mu_{\min} + (\mu_{\max} - \mu_{\min})u_i^\gamma$, which allows to address concave ($0 < \gamma < 1$), linear ($\gamma = 1$), and convex ($\gamma > 1$) relations. The deterioration rate equals $\mu_{\min} = 0.5$ for idle units and $\mu_{\max} = 5.0$ for units that produce at the maximum rate. Thus, a unit also slowly deteriorates while being idle. In practice this happens, for instance due to corrosion, bearings that become slightly unbalanced as a result of one-sided pressure, or externally caused load due to weather conditions. Moreover, we focus on convex pd-relations as these are most conceivable for real-life systems. A convex pd-relation implies an incentive to share load equally among units because this results in the lowest average deterioration rate at the system level.

The two base systems share all parameter values except for the ones that describe the production contracts. Both systems consist of two units and thus their total capacity equals 2. Contract type I represents a production facility that has some overcapacity but no redundancy and that aims to meet the production target, although not at any cost. We model this system by setting a target below the maximum production capacity $\kappa = 1.6$, a fixed penalty $\tilde{\pi} = 10$, a variable penalty $\pi_1 = 1$, and no bonus for overproduction, i.e., $\pi_2 = 0$. Contract type II represents a system that primarily focuses on a reliable production output. This base system has a redundant unit and an extreme penalty if the target is not met. The redundant unit is modeled by setting the production target to $\kappa = 1$ and the fixed penalty is set to $\tilde{\pi} = 10^6$. There is no benefit of producing more than the target and thus $\pi_2 = 0$. One could argue that overproduction is even discouraged or impossible in such systems and thus that we should have $\pi_2 < 0$. Although this is true, by choosing $\pi_2 = 0$ there is no advantage of producing at higher rates while the system will deteriorate faster, and thus the optimal policies for $\pi_2 < 0$ and $\pi_2 = 0$ are the same. Moreover, the fixed penalty is substantial and thus the optimal policy always aims to avoid production shortages. For numerical reasons, however, we still set a small positive variable penalty $\pi_1 = 1$, which does not affect the observed optimal policy and its corresponding costs.

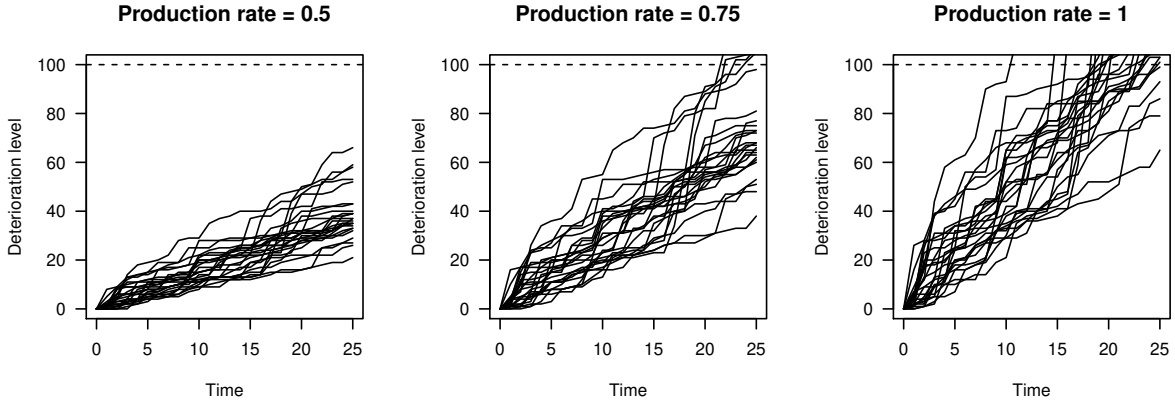


Figure 2: Effect of the production rate on the deterioration process in the base case

If a unit continuously produces at rate u_i , then its expected lifetime approximately equals $L/g(u_i)$ time units. Thus, if a unit would always produce at full speed, its expected lifetime approximately equals 20 time units. To provide some intuition on the deterioration process of the base system, Figure 2 depicts 25 sample paths of the deterioration process for different production rates. We clearly see that producing at lower rates increases the expected lifetime and results in more stable deterioration per time unit.

6. Results contract type I

In this section, we consider contract type I, which has some overcapacity, but no redundancy, and with a fixed penalty in case the production target is not satisfied. In Section 6.1, we zoom in on the optimal decisions for both the equal load-sharing and the condition-based load-sharing policies. Thereafter, in Section 6.2, we examine how the policies and their performances are affected by the maintenance setup cost, the volatility of the deterioration process, and the degree of overcapacity. In doing so, we define the gap = $|x_1 - x_2|$ as the absolute difference between the deterioration levels of the two units.

6.1. Optimal policy for the base system

Figure 3 shows the optimal decisions under the equal load-sharing (left) and the condition-based load-sharing (right) policy for the base system described in Section 5.2. The production rate of unit 1 is indicated by gray scale, ranging from idle (black) to producing at the maximum rate (white). The remaining areas at the top and right side indicate (in both text and color) when maintenance is scheduled, where the three subareas indicate which units are maintained at the end of the planning time. The optimal production rate of unit 2 immediately follows from that of unit 1 because the optimal policy exactly meets the production target whenever possible. This is intuitive since there is no incentive to produce more than the production target as $\pi_2 = 0$ while there is a penalty $\tilde{\pi} = 10$ if the target is not met.

In the considered system, there is a maintenance setup cost and thus there is an incentive to cluster the maintenance actions of both units. However, deterioration is stochastic and thus clustering maintenance implies that maintenance is either performed unnecessarily early for one unit or too late for the other. From a first inspection of the optimal policies provided in Figure 3, we immediately see that the maintenance decisions are fairly similar for both policies, whereas their production decisions differ a lot.

Figure 4 depicts the long-run stationary state distribution under the optimal policies. Such distributions show the probability to be in a certain state at an arbitrary moment in time, thereby providing insights on how the deterioration processes are expected to behave over time and on the expected gap.

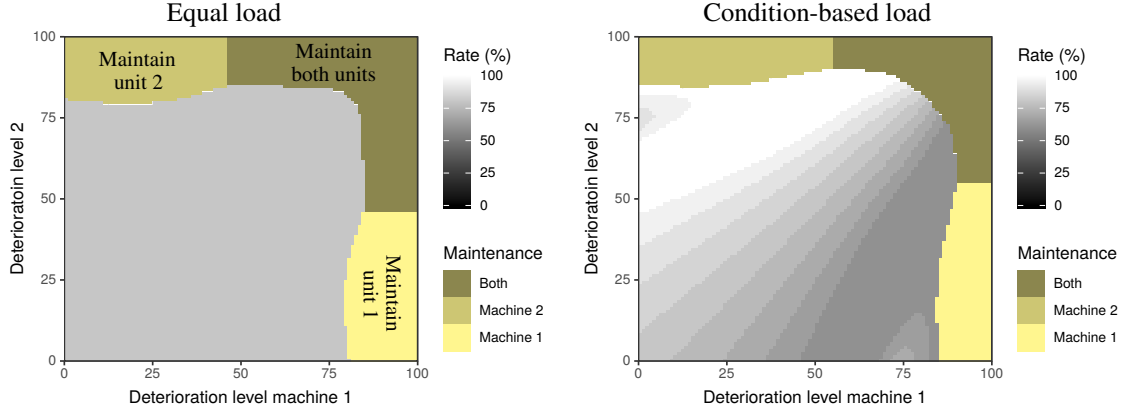


Figure 3: Optimal decisions for the base case under equal load-sharing (left) and condition-based load-sharing (right). Gray scale indicates the production rate of unit 1, ranging from idle (black) to the maximum rate (white). In the remaining areas, a maintenance intervention is scheduled. In these areas, the units continue producing until the end of the planning time, however, for clarity these production rates are not shown.

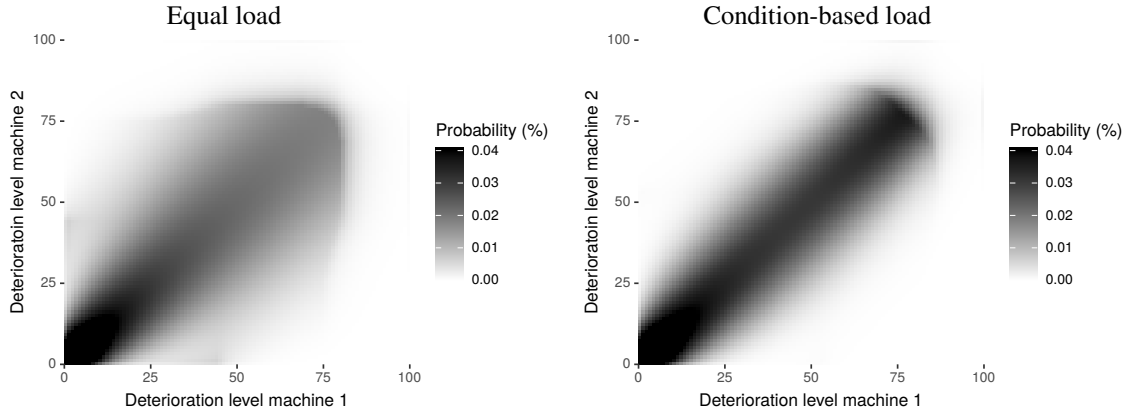


Figure 4: Heat map of the long-run stationary state distribution under both policies.

We see that under condition-based load-sharing, the deterioration processes are expected to move close along the diagonal, that is, the expected gap remains small when the units become further deteriorated. On the contrary, under equal load-sharing, it is likely that the gap becomes larger over time. It follows that the condition-based load-sharing policy uses the adjustable production rate to retain a small gap such that the maintenance interventions can be clustered without wasting remaining useful life.

Observations on the production decisions

The production decisions under the condition-based load-sharing policy can be characterized as follows. Load is only reallocated when the gap exceeds a certain threshold, and this threshold becomes smaller when the units get further deteriorated. Furthermore, the larger the current gap, the more skewed load will be shared among the units. This structure stems from the fact that sharing load unequally among units implies a higher average deterioration rate. For small gaps, it is quite likely that the deterioration processes will synchronize without intervening. Consequently, it is better to continue producing at the most efficient loads, that is, equally sharing load among units. If the processes do not synchronize, then the operator can still intervene at a later stage.

An exception to the above is a situation with a healthy and a highly deteriorated unit. In this case, the deteriorated unit takes over load from the healthy unit and synchronization is reached by only maintaining

the deteriorated unit. Performing maintenance immediately would waste remaining useful life of the deteriorated unit, whereas postponing it implies a larger gap after the maintenance action because the healthy unit also continues to deteriorate. By reallocating load, maintenance can be postponed until the deteriorated unit has depleted its remaining useful life while the other unit can retain its health. We note that this scenario is unlikely to occur because large gaps are generally corrected in an earlier stage.

Observations on the maintenance decisions

The maintenance decisions under both policies are largely similar. Maintenance is clustered if both units are highly deteriorated whereas only the most deteriorated unit is maintained if the deterioration levels differ too much. Furthermore, for a given deterioration level of the healthiest unit, the other unit is maintained according to a threshold policy.

A particular observation is that this threshold is first decreasing and then increasing in the deterioration level of the other unit. The threshold is non-constant because of two opposing incentives. On the one hand, maintenance for the deteriorated unit should be performed early because this synchronizes the deterioration levels. On the other hand, postponing the maintenance action better utilizes the useful life of the deteriorated unit. If the deterioration level of the healthy unit is very low, then maintenance of the deteriorated unit can be postponed without causing a too large gap after the maintenance action. The higher the deterioration level of the healthy unit, the earlier maintenance for the deteriorated unit should be performed in order to avoid a too large gap. This explains why the threshold first decreases as function of the deterioration level of the healthy unit. If the deterioration level of the healthy unit increases further, it becomes more likely that maintenance can be clustered in this cycle, which explains why the threshold eventually increases. We note that this effect, although not mentioned by others, is solely caused by the economic dependency and not by load sharing dynamics.

We also observe two structural differences between the two policies. Firstly, condition-based load-sharing allows to schedule maintenance interventions at higher deterioration levels. The reason is that lower production rates not only reduce the deterioration rates but also the volatility of the deterioration increment per period. With condition-based load-sharing, the most deteriorated unit typically produces at a lower speed, thereby reducing the risk of failure.

Secondly, because the equal load-sharing policy can only use maintenance to synchronize deterioration levels, it clusters maintenance for considerably more states than the condition-based load-sharing policy. For instance, if unit 1 is in the highly deteriorated state $x_1 = 90$, then the equal load-sharing and condition-based load-sharing policies opportunistically maintain the second unit for deterioration levels above 46 and 55, respectively. To understand this dynamic, let us consider the situation that the deterioration level of the second unit lies between these thresholds, e.g. $(x_1, x_2) = (90, 50)$. The second unit clearly has no need for maintenance whereas maintenance for the first unit cannot be postponed. By only maintaining the first unit, the system moves to state $(x_1, x_2) = (0, 50)$. Under equal load-sharing this implies that the next maintenance actions are again unlikely to be clustered, and thus it is better to synchronize their deterioration by maintaining both units, thereby wasting a substantial remaining useful life of the healthy unit. On the contrary, under condition-based load-sharing, the resulting gap can easily be synchronized before the next maintenance intervention, and thus it is not necessary to waste the remaining useful life of the healthy unit.

From the above effects, it follows that both policies use the maintenance decision to synchronize the deterioration levels of the units (e.g., by performing maintenance for a deteriorated unit earlier than actually necessary for this single unit). However, such interventions waste remaining useful life of units and is therefore significantly more expensive than using the more subtle option to synchronize the deterioration levels by reallocating load. We indeed observe that the condition-based load-sharing policy uses a maintenance intervention substantially less often to synchronize the deterioration levels.

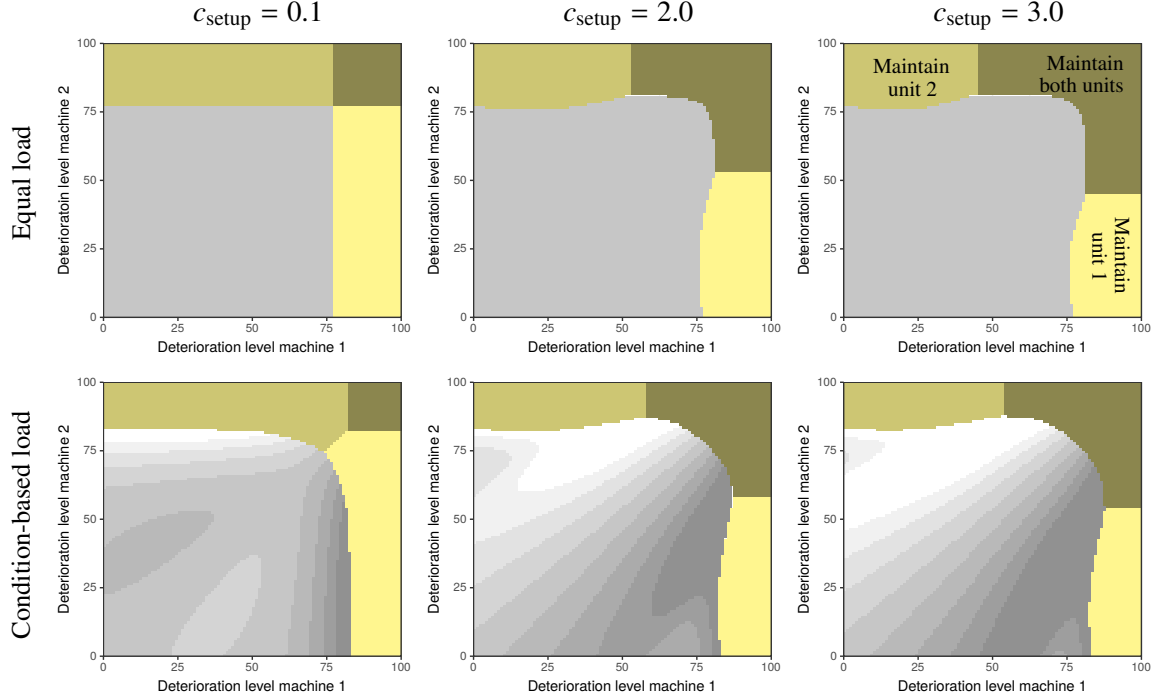


Figure 5: Effect of maintenance setup cost under the equal load-sharing (top) and condition-based load-sharing (bottom) policy. Gray scale indicates the production rate of unit 1, ranging from idle (black) to the maximum rate (white). In the remaining areas, a maintenance intervention is scheduled.

6.2. Parameter sensitivity

We continue by examining the effects of changing various parameter values on the structure of the optimal policy and on the corresponding cost savings of condition-based load-sharing compared to equal load-sharing. The results are obtained by taking the base system and adjusting the parameter values one by one.

Effect maintenance setup costs

Figure 5 shows the optimal policies for various maintenance setup costs for the equal load-sharing policy (top) and the condition-based production policy (bottom). Under equal load-sharing we observe that 1) the area in which the healthy unit is opportunistically maintained decreases in size if the setup cost decreases, and 2) for very low setup costs the maintenance decisions for the two units are independent of each other.

Now consider the condition-based load-sharing policy. As long as the setup costs are substantial (say $c_{\text{setup}} = 2$), the maintenance decisions are insensitive to an increase of the setup cost whereas the production decisions are affected. If the setup cost increases, clustering becomes more important and the optimal policy assigns more load to the healthy unit. For instance, suppose we have $x_1 = 10$ and $x_2 = 70$. Then, for $c_{\text{setup}} = 2$ the policy does not fully reallocate the load to the healthy unit ($u_1 = 90\%$ and $u_2 = 70\%$) in order to produce at a more efficient rate, whereas for $c_{\text{setup}} = 3$ the load is fully reallocated ($u_1 = 100\%$ and $u_2 = 60\%$). Further increasing the setup cost has almost no effect on the optimal policy because the maintenance actions are already virtually always clustered.

Figure 6 (left) shows how the cost saving of adopting condition-based load-sharing is affected by the maintenance setup cost. We indeed see that the cost saving first increases in the setup cost and then stabilizes. An interesting observation is that without a setup cost, the optimal production and maintenance

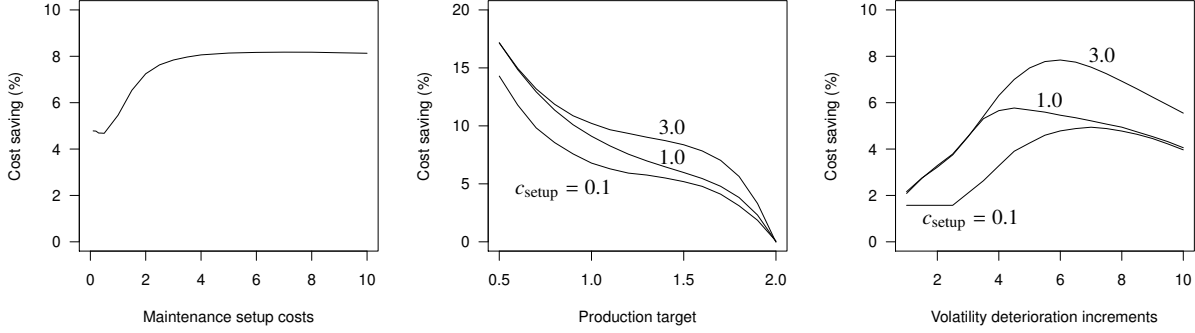


Figure 6: Relative cost savings of condition-based production decisions compared to equal load-sharing as function of the maintenance setup cost c_{setup} (left), the production target κ (middle), and the volatility of the deterioration process σ_{max} (right).

decisions of the units are still dependent, and cost savings around 5% are realized. In this case, the deterioration levels of the units are actively desynchronized and their maintenance interventions are alternated. Hereby, the useful lives of the units can be better utilized by slowing down the most deteriorated unit when it reaches the failure level.

Effect of the production target

Figure 7 shows optimal condition-based load-sharing decisions for different production targets (increasing from left to right) for both stable (bottom) and volatile deterioration (top). A lower production target implies more overcapacity, which gives the operator more flexibility to reallocate load among the units, resulting in two benefits. Firstly, because it is easier to synchronize large gaps, the optimal policy allows for larger gaps before load is reallocated. Secondly, the load of the most deteriorated unit can be reduced further, resulting in a considerably less conservative maintenance policy that utilizes the useful life of units more effectively.

In Figure 6 (middle), we see that the cost saving increases when the production target decreases, and that there is no cost saving if the target equals the maximum production capacity. Moreover, the cost savings are very sensitive to the production target if there is only little overcapacity, and it becomes less sensitive if the production target comes closer to 1. However, if the production target drops below 1, the cost saving becomes more sensitive again because the redundant unit provides new operational options (see also Section 7).

Effect of the volatility of the deterioration process

Now we consider the effect of the volatility of the deterioration process by comparing the policies presented in the top row of Figure 7 to those presented in the bottom row. We see that, regardless of the production target, the optimal production decisions close to the diagonal are not affected by the volatility of the deterioration process. The main difference is observed for large gaps that are not synchronized before the next maintenance intervention (i.e., top left and bottom right areas in the figures). For stable deterioration, the load is gradually shifted to the healthy unit if the gap increases. If the gap becomes too large, the most deteriorated unit suddenly takes over load from the healthy unit. For more volatile deterioration, this transition is less sudden and the area in which load is shifted to the deteriorated unit is smaller. This is the case because the likelihood of synchronization by chance is higher, and because accelerating a deteriorated unit too much results in unacceptable failure risks.

Figure 6 (right) depicts the effect of the volatility on the cost savings. For stable deterioration, the cost savings are small because the deterioration levels of the units are not expected to diverge. If the volatility increases, the expected gap at the end of the lifetime of the units increases too. Both policies still use a high degree of clustering, but the condition-based load-sharing policy better utilizes the useful lives of

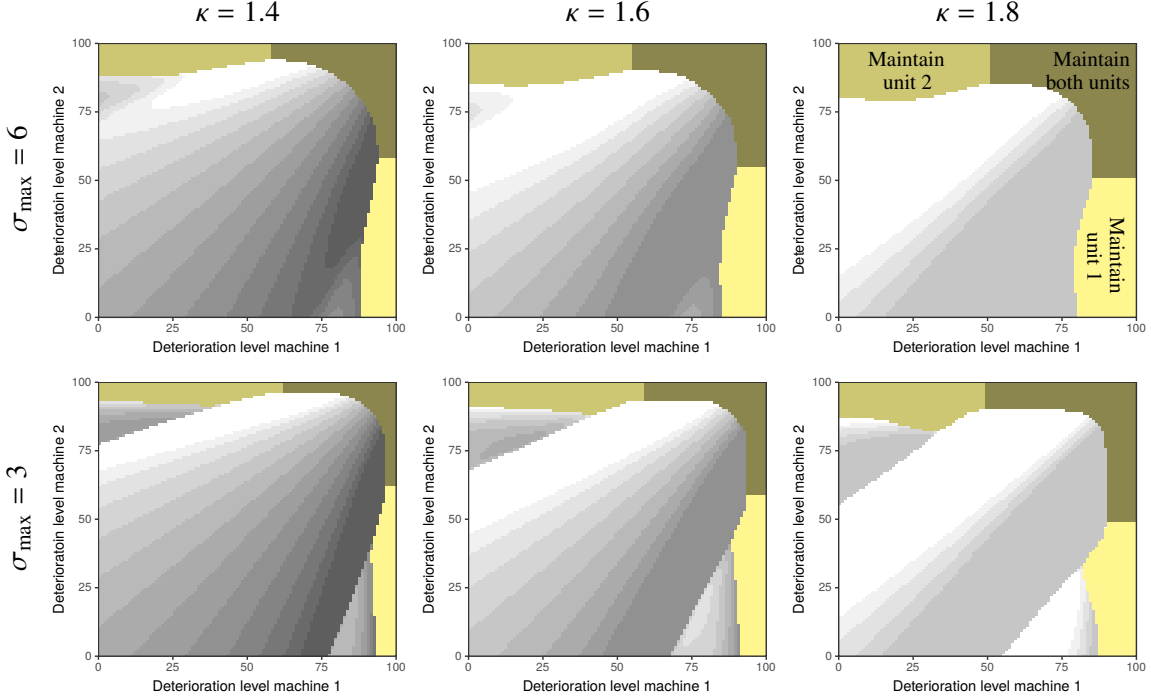


Figure 7: Effect of production target and of the volatility of the deterioration process. Gray scale indicates the production rate of unit 1, ranging from idle (black) to the maximum rate (white). In the remaining areas, a maintenance intervention is scheduled.

the units by synchronizing their deterioration levels. Consequently, the benefit of condition-based load-sharing increases if the volatility increases. Finally, if the deterioration process becomes highly volatile, then large gaps that cannot be corrected for by reallocating load become more likely, and we indeed see that the cost savings start to decline.

7. Results contract type II

We continue with the second contract type that applies to production facilities that must provide a constant and reliable production output. The key priority for such systems is to avoid production shortages, whereas minimizing operational costs is only a secondary objective. We model this by setting the fixed penalty for shortages to $\tilde{\pi} = 10^6$. Moreover, in practice, the reliability of such systems is often improved by including a redundant unit, which we model by setting the production target equal to the capacity of a single unit $\kappa = 1.0$.

Most interactions for this contract type are similar to those for contract type I as discussed in Section 6 and are therefore not repeated here. We do, however, observe different effects of the volatility of the deterioration process and of the maintenance setup cost, which we address in this section.

7.1. Effect volatility deterioration process

Compared to the base case, we lower the maintenance setup cost to $c_{\text{setup}} = 1$ in this section because this gives more clear-cut policies while it does not affect the structural insights that we obtain. In Section 7.2, we show that other maintenance setup costs result in similar insights.

Figure 8 shows the optimal condition-based load-sharing policy (top) and the corresponding long-run state distribution (bottom) for stable (left), medium volatile (middle), and highly volatile (right) deterioration. For stable deterioration, the risk that both units fail simultaneously is negligible. As a result, their maintenance can be clustered without risking excessive penalties for shortages. For medium




Figure 8: Effect of volatility of the deterioration process on the optimal production and maintenance decisions (top) and the corresponding long-run state distribution (bottom).

volatile deterioration, having two units with intermediate or high deterioration levels becomes too risky and the focus lies on minimizing the risk of production shortages. The deterioration levels of the units are actively desynchronized such that the gap is around 45. Note that, because of the redundancy, failure of one unit is allowed, and consequently the maintenance policy itself is actually less conservative than in all previously considered cases with the same volatility. For highly volatile deterioration, not only the maintenance interventions are alternated but also the usage of the units. Thereby, the system produces at a less efficient rate, but also always keeps one unit in a good condition. This reduces the risk of shortages if the other unit fails unexpectedly. Notice that maintenance is virtually never clustered, not even if both units are highly deteriorated. In such cases, only one unit is maintained to desynchronize the deterioration levels.

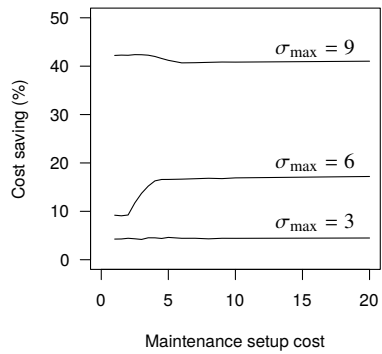


Figure 9: Effect of the maintenance setup costs on the cost savings compared to the equal load-sharing policy

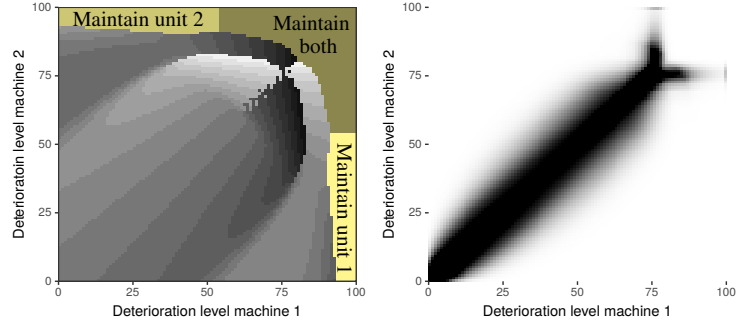


Figure 10: Optimal condition-based load-sharing policy and the corresponding long-run state distribution if $c_{\text{setup}} = 3$ and $\sigma_{\text{max}} = 6$.

7.2. Effect maintenance setup cost

Figure 9 shows the effect of the maintenance setup cost on the cost savings compared to the equal load-sharing policy. Higher maintenance setup costs have almost no effect on the optimal policy for stable deterioration ($\sigma_{\text{max}} = 3$) as for those maintenance actions are always clustered, and not on that for highly volatile deterioration ($\sigma_{\text{max}} = 9$) as for those maintenance actions are never clustered. Correspondingly, there is also no significant effect on the potential cost savings.

However, for medium volatile deterioration ($\sigma_{\text{max}} = 6$), the structure does change if we increase the maintenance setup cost, as can be seen in Figure 10. For low setup costs, the operator focuses on eliminating the risk of shortages by alternating the maintenance interventions. For higher setup costs, total costs can be reduced by synchronizing the deterioration levels of the units as long as these are in a good condition such that their maintenance can be clustered. When the units are highly deteriorated, their deterioration levels are desynchronized again to reduce the risk of simultaneous failure. The equal load-sharing policy can only reduce the risk of shortages by alternating the maintenance interventions and thus cannot share the maintenance setup cost among the units. This also explains the significant increase in cost savings that we observe in Figure 9.

8. Conclusion

We have investigated joint condition-based production and maintenance policies for multi-unit systems with economic dependency and whose units have adjustable production rates. The production rate of a unit affects its deterioration rate, implying that condition-based production policies can be used to control deterioration of the units. A production target at the system level is adopted and a penalty is incurred if this target is not satisfied. Condition-based production decisions enable the operator to (de)synchronize the deterioration levels of the units, thereby improving the clustering of maintenance interventions or reducing the risk that multiple units fail simultaneously.

We have formulated the system as a Markov decision process and used this to determine cost-minimizing joint production and maintenance policies. The benefits of dynamically reallocating load among units is examined by comparing the newly proposed policy to a policy that combines condition-based maintenance with a static production policy that shares load equally among all functioning units. Results show that cost savings up to 20% can be obtained for systems with overcapacity but no redundancy, and that these savings increase to 40% for systems with redundancy. The cost savings originate from fewer failures, reduced risks of production shortages, improved clustering opportunities, and fewer maintenance interventions per unit. Another promising observation is that adopting condition-based production policies not only reduces expected costs but also its variance.

For sufficiently high maintenance setup costs, the optimal policy aims to synchronize the deterioration levels of the units by assigning more load to the least deteriorated unit. The larger the difference in deterioration, the more load is assigned to the healthy unit. Moreover, for low deterioration levels, the optimal policy does not immediately adjust the production rates as the deterioration levels may synchronize themselves and otherwise there is still sufficient time left to correct the gap at a later stage. Interestingly, when the deterioration levels are far apart, the operator should not try to synchronize them before the next maintenance intervention and should even accelerate the most deteriorated unit. At the next maintenance intervention, maintenance will then only be carried out for the this unit, resulting in better synchronized deterioration levels after this maintenance intervention. Postponing the maintenance intervention implies a larger gap after maintenance because the healthy unit also continues to deteriorate, whereas performing it immediately results in wasting remaining useful life of the most deteriorated unit. These two aspects are better balanced by reallocating load to the most deteriorated one.

Another insightful result is that even without maintenance setup costs, the optimal production and maintenance decisions of the units are still dependent. The condition-based load-sharing policy actively alternates their maintenance interventions. Thereby, the most deteriorated unit can decelerate when its deterioration level approaches the failure level. This results in better utilization of the useful life of the units, which can result in cost savings up to 10%.

Condition-based load-sharing decisions seem to be particularly useful for systems with redundancy and severe consequences if the production target is not satisfied. Examples are facilities that must provide a reliable production flow such as gas turbines that maintain a constant gas pressure in the network. In such scenarios, we observe cost savings of up to 40% for numerous problem instances. The structure of the optimal production policy for stable deterioration processes is very different from that of volatile deterioration processes. For stable processes, the deterioration levels are synchronized as long as units are in a reasonable condition such that their maintenance interventions can be clustered. When the units are highly deteriorated, their deterioration levels are desynchronized to reduce the risk of simultaneous failure. On the contrary, for medium volatile processes, the focus lies solely on minimizing the risk of production shortages by alternating the maintenance inventions. For very volatile processes, also the usage of the units is alternated such that always at least one of the units is in a good condition. Thereby, the risk of a production shortage in case the most deteriorated unit fails is minimized.

Controlling deterioration of multi-unit systems by adopting condition-based production rates provide numerous opportunities for further research. We have considered situations with a fixed production target over time. However, for instance in the electricity industry, output can be traded on a day-ahead market, and thus the production target can be included as a short-term decision variable for such systems. We expect that the operational efficiency improves by placing lower output bids if many units are highly deteriorated. This provides more flexibility to reallocate load among units, and it allows to produce at lower rates which enables the operator to let the deterioration levels approach the failure level more closely.

It is also worthwhile to study systems with a maintenance capacity due to the need for scarce resources such as skilled technicians and specialized vessels. We expect that condition-based load-sharing policies can better cope with a restricted maintenance capacity than static production policies because the preferred maintenance moments of units can be desynchronized by dynamically adjusting their production rates.

Another direction is to examine the value of condition-based production planning for multi-unit systems in settings commonly faced by practitioners and often studied in the maintenance literature. Examples are uncertain failure levels, imperfect condition monitoring, fluctuating production prices, uncertain production capacities, (aperiodic) inspections, partial repairs, and non-homogeneous deterioration processes. We expect that condition-based load-sharing decisions in multi-unit systems can be particularly useful to cope with uncertain condition information because highly deteriorated units can be decelerated without losing production output on the system level.

Appendix A. The modified policy iteration algorithm

In this appendix, we provide an overview of the modified policy iteration (see Algorithms 1-3), and complement this with some additional remarks and tips to accelerate the convergence rate or to reduce the memory usage. In contrast to Puterman (1994), we do not store the values for each iteration in separate vectors but we reuse the same vector. This does not affect the flow of the algorithm, as only the values of the last iteration are used to compute the new values, while it results in substantial lower memory usage.

Initial experiments have shown that more than 95% of the total running time is caused by evaluating Equation (3) in line 3 of Algorithm 1. In particular, the second term within the minimization, i.e., the expectation that incorporates the deterioration increments, seems to be time-consuming. It clearly follows that all effort to improve the performance of the algorithm should focus on this for-loop.

The iterations of all for-loops in Algorithms 1 and 2 are independent of each other, and thus their iterations can easily run in parallel. For our initial experiments, we observed that running the first for-loop of Algorithm 1 results in an almost linear speed up in the number of CPU cores (regardless of the considered system) whereas the other for-loops are not worth running in parallel (and doing so can even slow down the algorithm due to increased processor overhead).

The considered system has identical units and thus the state space has some symmetry, e.g., for a two-unit system the value of state (x_1, x_2, τ) is the same as for (x_2, x_1, τ) . The easiest way to exploit this is by adjusting the given algorithms such that for states with $x_2 > x_1$ the value $\bar{w}_2(x_1, x_2, \tau)$ is copied from $\bar{w}_2(x_2, x_1, \tau)$ instead of being calculated. Initial experiments suggests that it beneficial to implement this in all for-loops in Algorithm 1 and 2, resulting in a speed-up of around 40% for two-unit systems.

Algorithm 1 updatePolicyAndValues

Input: Values for v

Output: Updated values for v and current best actions $d_0, d_1,$ and d_2 .

```
1: // For ease of notation,  $S = \{(x, \tau)\}$  denotes the set of all states
2: for all  $s \in S$  do
3:   Compute  $w_2(x, \tau)$  according to Equation (3)
4:   While computing  $\bar{w}_2(x, \tau)$ , store the minimizing argument as  $d_2(x, \tau)$ 
5: end for
6: for all  $s \in S$  do
7:   Compute  $w_1(x, \tau)$  according to Equation (2)
8:   While computing  $w_1(x, \tau)$ , store for  $\tau = 0$  the minimizing argument as  $d_1(x, \tau)$ 
9: end for
10: for all  $s \in S$  do
11:   Compute  $\bar{v}(x, \tau)$  according to Equation (1)
12:   While computing  $\bar{v}(x, \tau)$ , store for  $\tau = ns$  the minimizing argument as  $d_0(x, \tau)$ 
13: end for
14: return  $\bar{v}, d_0, d_1, d_2,$ 
```

Algorithm 2 evaluateFixedPolicy

Input: Convergence criterion ϵ , maximum number of iterations η , a policy (d_0, d_1, d_2) , and values for v

Output: Updated values for v under the given policy

```
1: // For ease of notation,  $S = \{(x, \tau)\}$  denotes the set of all states
2: while true do
3:   for all  $s \in S$  do
4:     Set  $u = d_2(x, \tau)$  and compute  $\bar{w}_2(x, \tau) = \varphi_2(\sum_{i=1}^n u_i) + \sum_{x' \in \mathcal{X}'(x)} P_{u, \Delta t}(x, x') v(x', \tau)$ 
5:   end for
6:   for all  $s \in S$  do
7:     Set  $r = d_1(x, \tau)$  and compute  $\bar{w}_1(x, \tau) = \begin{cases} \bar{w}_2(x, \tau) & \text{if } \tau \neq 0, \\ \varphi_1(x, r) + \bar{w}_2(x'(x, r), \text{ns}) & \text{if } \tau = 0. \end{cases}$ 
8:   end for
9:   for all  $s \in S$  do
10:    Compute  $\bar{v}(x, \tau) = \begin{cases} \bar{w}_1(x, \text{ns}) & \text{if } \tau = \text{ns} \text{ and } d_0(x, \tau) = 0, \\ \tilde{c} + \bar{w}_1(x, s) & \text{if } \tau = \text{ns} \text{ and } d_0(x, \tau) = 1, \\ \bar{w}_1(x, \tau - 1) & \text{otherwise.} \end{cases}$ 
11:   end for
12:
13:   if  $sp(\bar{v} - v) < \epsilon$  or  $\eta \leq 0$  then
14:     return  $\bar{v}$ 
15:   end if
16:
17:   Set  $v = \bar{v}$  and  $\eta = \eta - 1$ 
18: end while
```

Algorithm 3 Modified policy iteration

Input: Convergence criterion ϵ and maximum number of iterations to fix a policy η

Output: Optimal actions for each stage d_0, d_1, d_2 , and the corresponding average cost rate g

```
1: // For ease of notation,  $S = \{(x, \tau)\}$  denotes the set of all states
2: Set  $v(x, \tau) = 0$  for all  $s \in S$ 
3: while true do
4:   Set  $\bar{v}, d_0, d_1, d_2 = \text{updatePolicyAndValues}(v)$ 
5:
6:   if  $sp(\bar{v} - v) < \epsilon$  then
7:     Compute  $g$  according to Equation (4)
8:     return  $d_0, d_1, d_2, g$ 
9:   end if
10:
11:   Set  $v = \text{evaluateFixedPolicy}(\epsilon, \eta, d_0, d_1, d_2, \bar{v})$ 
12: end while
```

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.
- Bevilacqua, M., Braglia, M., 2000. The analytic hierarchy process applied to maintenance strategy selection. *Reliability Engineering & System Safety* 70 (1), 71–83.
- Castanier, B., Grall, A., Bérenguer, C., 2005. A condition-based maintenance policy with non-periodic inspections for a two-unit series system. *Reliability Engineering & System Safety* 87 (1), 109–120.
- De Jonge, B., Klingenberg, W., Teunter, R. H., Tinga, T., 2016. Reducing costs by clustering maintenance activities for multiple critical units. *Reliability Engineering & System Safety* 145, 93–103.
- De Jonge, B., Scarf, P. A., 2019. A review on maintenance optimization. *European Journal of Operational Research*, In press.

- Do, P., Assaf, R., Scarf, P., Iung, B., 2019. Modelling and application of condition-based maintenance for a two-component system with stochastic and economic dependencies. *Reliability Engineering & System Safety* 182, 86–97.
- Do, P., Barros, A., Bérenguer, C., Bouvard, K., Brissaud, F., 2013. Dynamic grouping maintenance with time limited opportunities. *Reliability Engineering & System Safety* 120, 51–59.
- Do, P., Scarf, P., Iung, B., 2015. Condition-based maintenance for a two-component system with dependencies. *IFAC-PapersOnLine* 48 (21), 946–951.
- Iravani, S., Duenyas, I., 2002. Integrated maintenance and production control of a deteriorating production system. *IIE Transactions* 34 (5), 423–435.
- Li, H., Deloux, E., Dieulle, L., 2016. A condition-based maintenance policy for multi-component systems with lévy copulas dependence. *Reliability Engineering & System Safety* 149, 44–55.
- Lu, L., Jiang, J., 2007. Analysis of on-line maintenance strategies for k-out-of-n standby safety systems. *Reliability Engineering & System Safety* 92 (2), 144–155.
- Marseguerra, M., Zio, E., Podofillini, L., 2002. Condition-based maintenance optimization by means of genetic algorithms and monte carlo simulation. *Reliability Engineering & System Safety* 77 (2), 151–165.
- Olde Keizer, M., Teunter, R., Veldman, J., 2016. Clustering condition-based maintenance for systems with redundancy and economic dependencies. *European Journal of Operational Research* 251 (2), 531–540.
- Olde Keizer, M., Teunter, R., Veldman, J., Babai, M., 2018. Condition-based maintenance for systems with economic dependence and load sharing. *International Journal of Production Economics* 195, 319–327.
- Olde Keizer, M. C. A., Flapper, S. D. P., Teunter, R. H., 2017. Condition-based maintenance policies for systems with multiple dependent components: A review. *European Journal of Operational Research* 261 (2), 405–420.
- Puterman, M., 1994. *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons, New York.
- Rasmekomen, N., Parlikad, A., 2016. Condition-based maintenance of multi-component systems with degradation state-rate interactions. *Reliability Engineering & System Safety* 148, 1–10.
- Sloan, T., 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.
- Uit het Broek, M. A. J., Teunter, R. H., De Jonge, B., Veldman, J., 2019a. Joint condition-based maintenance and condition-based production optimization, under review.
- Uit het Broek, M. A. J., Teunter, R. H., De Jonge, B., Veldman, J., Van Foreest, N. D., 2019b. 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 & System Safety* 94 (1), 2–21.
- Wang, L., Zheng, E., Li, Y., Wang, B., Wu, J., 2009. Maintenance optimization of generating equipment based on a condition-based maintenance policy for multi-unit systems. In: *2009 Chinese Control and Decision Conference*. IEEE, pp. 2440–2445.
- Zhang, Z., Wu, S., Lee, S., Ni, J., 2014. Modified iterative aggregation procedure for maintenance optimisation of multi-component systems with failure interaction. *International Journal of Systems Science* 45 (12), 2480–2489.
- Zhang, Z., Wu, S., Li, B., Lee, S., 2015. (n, N) type maintenance policy for multi-component systems with failure interactions. *International Journal of Systems Science* 46 (6), 1051–1064.
- Zhao, X., Liu, B., Liu, Y., 2018. Reliability modeling and analysis of load-sharing systems with continuously degrading components. *IEEE Transactions on Reliability* 67 (3), 1096–1110.
- Zhou, Y., Lin, T., Sun, Y., Ma, L., 2016. Maintenance optimisation of a parallel-series system with stochastic and economic dependence under limited maintenance capacity. *Reliability Engineering & System Safety* 155, 137–146.
- Zhou, Y., Zhang, Z., Lin, T., Ma, L., 2013. Maintenance optimisation of a multi-state series-parallel system considering economic dependence and state-dependent inspection intervals. *Reliability Engineering & System Safety* 111, 248–259.



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. Trimborn, 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 the 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. Gijzenberg, 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.

2019007-OPERA: Broek, M.A.J. uit het, R.H. Teunter, B. de Jonge, J. Veldman, Joint Condition-based Maintenance and Load-sharing Optimization for Multi-unit Systems with Economic Dependency



www.rug.nl/feb