How good must failure predictions be to make local spare parts stock superfluous?

Dursun, I.; Akcay, Alp; van Houtum, Geert-Jan

DOI
10.1016/j.ijpe.2023.109060

Publication date
2023

Document Version
Final published version

Published in
International Journal of Production Economics

Int. J. Production Economics 267 (2024) 109060

Contents lists available at ScienceDirect

International Journal of Production Economics

journal homepage: www.elsevier.com/locate/ijpe

How good must failure predictions be to make local spare parts stock superfluous?

İpek Dursun [a,b,1], Alp Akcay [a], Geert-Jan van Houtum [a,*]

[a] Eindhoven University of Technology, Department of Industrial Engineering and Innovation Sciences, P.O. Box 513, 5600 MB, Eindhoven, The Netherlands
[b] Delft University of Technology, Faculty of Technology, Policy and Management, Jaffalaan 5, 2628 BX, Delft, The Netherlands

ARTICLE INFO

Keywords:
Spare parts
Failure predictions
Prediction accuracy
Demand lead time

ABSTRACT

Thanks to Industry 4.0 technologies, predictive algorithms can provide advance demand information on spare parts demand. Understanding how the goodness of predictions affects on-hand inventory and costs is important for decision makers before integrating these models into existing systems. We consider a spare parts inventory problem for multiple technical systems that are supported by one local stockpoint. Each system has a single critical component that is subject to random failures. Signals are generated to predict component failures. The signal that corresponds to a failure is generated a certain amount of time before the failure, referred to as the demand lead time. However, not every signal results in a failure and some failures are undetected. A component is replaced from the stock when a failure occurs. In case of stock-outs, an emergency shipment takes place. We formulate a discrete-time Markov decision process model to optimize the replenishment decisions with the objective of minimizing the long-run average cost per period. We investigate the effect of precision (i.e., the fraction of true signals among all signals) and sensitivity (i.e., the fraction of detected failures among all failures) of the predictions and the demand lead time on the costs, order-up-to levels, average on-hand inventory and emergency shipments under the optimal policy. In the worst case, the precision, sensitivity or demand lead time is zero. We show analytically that the optimal policy and optimal costs only depend on the sensitivity and the demand lead time through their product. In numerical experiments, we observe a Pareto principle for the reduction of costs in precision (e.g., a 30% perfectness in precision brings a 70% reduction in optimal cost compared to the worst case) and an inverse Pareto principle in the product of sensitivity and demand lead time (e.g., 70% perfectness in the sensitivity or demand lead time only brings 30% reduction in optimal cost compared to the worst case). Finally, we observe that the local spare parts stock only becomes superfluous when the signals are really close to perfect.

1. Introduction

Spare parts management is important for the timely maintenance and repair of technical systems. When a maintenance activity requires a spare part and it is not available, a part can be delivered via an emergency shipment. However, the technical system is then down for a longer time, which is costly, and the emergency shipment itself is generally also expensive. However, keeping stock for spare parts results in inventory holding costs, which involves opportunity costs, warehousing costs and/or costs of spare parts becoming obsolete. Prediction of demand for spare parts becomes crucial to balance the trade-off between inventory holding costs and costs of emergency shipments. Industry 4.0 provides many opportunities in terms of information technologies to the decision makers (Tortorella et al., 2021). It is possible to generate signals (by predictive algorithms) that say that certain components are having problems and may fail soon. These signals constitute so-called advance demand information (ADI) for the spare parts stock.

The ideal situation for maintenance of technical systems is that all failures are predicted, no false predictions are generated, and the predictions are made sufficiently far in advance. In that case, for all upcoming failures, a spare part can be sent to the system from a central location and the failing component can be replaced from the stock immediately. There would be no expensive local stocks of spare parts and no emergency shipments. However, ADI is not always perfect in practice. We refer to ADI as *perfect* if the actual demand is equal to the number of signals that indicate upcoming failures and the ADI is available far enough in advance. In all other cases, the ADI is *imperfect*.

* Corresponding author.
E-mail addresses: i.dursun@tudelft.nl (İ. Dursun), a.e.akcay@tue.nl (A. Akcay), g.j.v.houtum@tue.nl (G.J. van Houtum).

[1] This work has started during the employment at Eindhoven University of Technology and has been completed during the employment at Delft University of Technology.

https://doi.org/10.1016/j.ijpe.2023.109060
Received 27 October 2022; Received in revised form 31 July 2023; Accepted 25 September 2023
Available online 11 October 2023

The Internet of Things and Artificial Intelligence (AI) can bring us closer to the ideal situation of having perfect ADI. But how close do we need to be to that ideal situation in order to have sufficiently low spare parts stocks and total costs? It is important to understand how the closeness to the ideal situation affects the inventory levels and costs before integrating ADI into decision-making.

We investigate the reduction in spare parts stock and total costs for a setting with multiple technical systems that are supported by a local stockpoint with spare parts. We consider a setting with a single critical component in an infinite time horizon with periodic reviews. A signal can be generated in advance of a failure, and we refer to the time between the generation of the signal and the actual failure as the *demand lead time*. Based on the total number of signals, a replenishment takes place at the beginning of a period. When a component fails, it must be replaced. Signals are imperfect. That means false positive signals (i.e., a signal not leading to a failure) and false negative signals (i.e., unpredicted failures) are possible. The fraction of signals that result in an actual failure is called *precision*, and the fraction of failures for which a signal is generated is called *sensitivity*. In the worst case, all generated signals are false(zero precision), none of the failures is predicted by signals (zero sensitivity), or each failure happens at the same moment as when the signal is generated (zero demand lead time).

We formulate a Markov decision process model with precision, sensitivity, and demand lead time as input parameters. We derive an optimal policy for the spare parts inventory that minimizes the long-run average cost per period. Next, we compare the optimal costs for a given precision, sensitivity, delay time against the optimal costs in the worst case situation. Subsequently, we analyze how the optimal costs and the optimal spare parts stock reduce as precision, sensitivity and demand lead time go from the worst case to the ideal case.

We summarize the main contributions of this study as follows: (1) For a given precision level, the optimal costs depend on the sensitivity and the demand lead time only through the product of these two terms (see Theorem 1). To the best of our knowledge, such a theoretical result has not been found before in the existing literature. (2) The Pareto principle holds for precision, e.g., under a perfect sensitivity and a perfect demand lead time, 30% perfectness in precision (i.e., precision is equal to 30% of the perfect precision) brings 70% reduction in optimal costs compared to the worst case optimal costs. (3) The opposite of the Pareto principle holds for the product of sensitivity and demand lead time, e.g., under perfect precision, 70% perfectness in the product of sensitivity and demand lead time (i.e., the product is equal to 70% of the product of perfect sensitivity and perfect demand lead time) brings only 30% reduction in optimal costs compared to the worst case optimal costs. (4) We analyze the combined effect of precision and the product of sensitivity and demand lead time on optimal costs and inventory levels. To obtain a significant cost reduction, you need to have a high sensitivity and a high demand lead time while the precision can be moderate. (5) Local spare parts stocks only become superfluous when the signals are really close to perfect.

Adoption of new technologies comes with large investment costs (Tortorella et al., 2021; Barraza-Barraza et al., 2014). Examining the effect of Industry 4.0 implementations is crucial when adopting new technologies into business operations (Tortorella et al., 2023). Our findings provide managerial insights to decision-makers about the cost reduction for integrating predictive algorithms for different levels of precision, sensitivity and demand lead time. They also show that it is more important to have a high sensitivity and demand lead time than a high precision when designing predictive algorithms.

We organize the rest of the paper as follows. In Section 2, we provide the literature review on the related work. In Section 3, we present the model description. In Section 4, we show our main theoretical result (see Theorem 1), which implies a reduction of the main problem. In Section 5, we formulate an MDP model for this reduced problem. In Section 6, we analyze two special cases of the reduced problem. This section is followed by computational experiments and a sensitivity analysis in Section 7. Finally, in Section 8, we conclude the paper.

2. Literature review

There are two streams of literature on inventory control problems related to our work with imperfect ADI. The first stream concerns single-item, infinite-horizon inventory control problems with imperfect ADI. In this stream, the imperfect ADI usually is in the form of imperfect information provided directly by customers on their orders or estimations/signals generated for the future demands. The second stream focuses on spare parts inventory control problems with ADI that is obtained from condition-based monitoring of the machines installed in the field.

Different inventory control problems are studied under the first literature stream. There are also different types of information provided on future orders/sales. In this stream, van Donselaar et al. (2001) study the effect of imperfect ADI for inventory systems in a project-based supply chain. Thonemann (2002) investigates the effect of sharing imperfect ADI within a multi-echelon supply chain on average costs, mean basestock levels and variations of the production quantities. Results are shown for the value of ADI as a function of the order probability and information quality. Tan et al. (2007) consider an inventory control problem under imperfect ADI. They show that the optimal policy is of the order-up-to type and the order level is a function of the number of imperfect ADI signals. Benjaafar et al. (2011) consider a production–inventory system under imperfect ADI. In this problem, customers provide updates on their orders but the times between the consecutive updates are random. Song and Zipkin (2012) analyze a capacity/inventory planning problem under imperfect ADI for a single product with seasonal demand. Bernstein and DeCroix (2015) consider a multi-product system where the decision maker receives imperfect signals for the demand volume (i.e., signals for the total aggregate demand) or mixed demand (i.e., signals revealing information about the market shares for each product).

Some papers in this stream assume multiple customer classes. Liberopoulos and Koukoumialos (2008) investigate how the uncertainty in ADI affects the performance of a make-to-stock supplier. They investigate the uncertainty in ADI by assuming two customer classes, where one class does not provide any ADI and one class provides reservations on a requested due date. Tan (2008) considers a demand forecasting problem in a make-to-stock system. Tan et al. (2009) consider an inventory problem with two customer classes having different priorities. Available stock is reserved for the future demand of preferred customers at the expense of losing the current orders of other customers. Gayon et al. (2009) study an inventory–production system with multiple customer classes where customers provide imperfect ADI on the due date of their orders.

Among the papers in the first literature stream, only Topan et al. (2018) and Zhu et al. (2020) specifically focus on spare parts inventory, like we do in our paper. Topan et al. (2018) focus on a spare parts management problem with imperfect ADI and the option of returning inventory. Zhu et al. (2020) assume a single-item, periodic-review setting for a spare parts management problem under imperfect ADI. The source of ADI is the spare part demand forecast based on the planned maintenance tasks. In our study, we consider the failure signals as the source of imperfect ADI.

Among all studies in the first literature stream, only Song and Zipkin (2012), Gayon et al. (2009), Benjaafar et al. (2011), Topan et al. (2018), and Zhu et al. (2020) assume that demands are lost or satisfied via emergency shipments in stockout situations similar to our paper, while others assume unmet demand is backlogged. Among the studies in this smaller group, our paper comes closest to Topan et al. (2018), because they also consider a single-item, single-location, periodic-review, infinite-horizon inventory control problem with an imperfect signal generation mechanism for future demands similar to our problem. To be specific, Topan et al. (2018) assume signals are generated for a fraction of all demands, signals can be false, and the actual demand occurs a stochastic time after the signal was generated. They derive the structure of the optimal ordering and return policy, and they show the value of the imperfect ADI in a computational experiment. In our paper, we assume a simpler model, but explicitly characterize how the optimal costs behave as a function of the precision, sensitivity and the demand lead time of the imperfect demand signals. This leads to new, crisp insights. In particular, none of the studies in the first literature stream has found the theoretical result that the optimal costs only depend on the sensitivity and the lead time demand through their product.

The second stream of literature is about condition monitoring in spare parts management. The condition of a component can be used to predict when the component fails. In that way, also advance demand information is obtained.

Within this second stream, Deshpande et al. (2006) use the part-age information to model the degradation of aircraft spare parts at the U.S. Coast Guard. Elwany and Gebraeel (2008) develop a sensor-driven decision model for making joint component replacement and spare parts inventory decisions. Similarly, Li and Ryan (2011) exploit the real-time condition monitoring information for the inventory control of spare parts. They assume a Wiener process as the degradation model. Lin et al. (2017) consider a single critical component of multiple installed machines in the field. The installed components follow a Markov degradation process and that information is used for optimizing the spare parts inventory. Eruguz et al. (2018) study an integrated maintenance and spare part optimization problem for moving assets where the degradation level of a single critical component is observable. They model the degradation of the component by a continuous-time Markov chain. Basten and Ryan (2019) consider two classes of spare parts demand (i.e., one for planned maintenance and the other for corrective maintenance) and investigate the benefit of delaying planned maintenance when there is perfect ADI for the number of spare parts needed for planned maintenance. These studies all assume that the condition of the components can be observed perfectly in a regular manner over time, while we consider signals that randomly arrive over time and are imperfectly generated about the condition of the components. Hence, different from our paper, the papers above implicitly assume that both precision and sensitivity are perfect but only the moments that the failures occur are uncertain.

Within the second literature stream, there are also studies where the information collected on the component condition can be imperfect. Karabağ et al. (2020) study an integrated maintenance and spare parts selection decision for a multi-component system, where a single sensor gives imperfect information about the condition of the system. Yan et al. (2022) introduce a remaining-useful life prediction method, and use it for making joint replacement and spare parts ordering decisions with a fixed lead time. Different from our paper, at most one spare part can be stored. More recently, Shi et al. (2023) consider imperfect IoT-enabled condition predictions to jointly optimize condition-based maintenance and spare parts inventory decisions. Rippe and Kiesmüller (2023) consider failure codes provided by customers as the source of imperfect ADI in a repair kit problem setting with multiple components. The problem in our paper is simpler because it does not include maintenance decisions and it only includes a single spare part. As a result, also in comparison to the literature in the second stream, we have more crisp results on how the optimal costs behave as a function of the precision, sensitivity and demand lead time of the imperfect demand signals.

3. Model description

In this section, we provide the detailed description of the model to address the problem introduced in Section 1. We consider a setting where a significant number of technical systems is supported by a local warehouse that keeps spare parts on stock. These spare parts are needed to execute maintenance actions. The technical systems are operated during a time horizon that is assumed to be infinite. The local warehouse is part of a service network consisting of a central warehouse and multiple local warehouses.

We focus on a single critical component that is part of all technical systems. All components are identical and have the same failure behavior. We assume that this component has a generally distributed lifetime. A certain amount of time before a failure (i.e., the end of component lifetime), a signal is generated by a predictive model. After this signal generation, the component still functions, but it is known that the component may fail soon. The time from the signal generation until the end of lifetime is referred to as demand lead time and denoted by D. The demand lead time is relatively short compared to the lifetime of the component (Fig. 1). For ease of exposition, we assume that the demand lead time is deterministic. However, all results in this paper can be generalized to a setting with a stochastic demand lead time; see Remark 1.

We assume that the technical systems operate continuously. Therefore, interrupting their operation for a preventive replacement is equally expensive as a corrective replacement. Hence, we assume that replacements are only executed when a component fails. The technical systems are at close distance from the local warehouse and we assume that a spare part is provided to a technical system within such a short time that it does not cause extra downtime of the technical system.

Spare parts of the critical component are kept on stock at the local warehouse. The local warehouse is replenished periodically (e.g., every week). Hence, we divide the time horizon in periods of length one, and the periods are numbered as $0, 1, \ldots$. The beginning of a period t is called time t. The local service point is replenished by the central warehouse, which is assumed to have ample stock. The corresponding replenishment lead time is short and is assumed to be 0. Hence, ready-for-use parts are ordered at the beginning of each period t, and they arrive immediately.

The generated signals for upcoming failures are subject to false positives and false negatives. The predictive model that generates the signals is based on sensor data and data that is collected via the control system software, and possibly other sources. Such a predictive model can be tuned to be more sensitive to certain data patterns, making it more likely to trigger signals with more false positives (i.e., low precision). On the other hand, making the sensor less sensitive to such patterns leads to more false negatives (i.e., low sensitivity). Precision and sensitivity are two related but different metrics. To be specific, suppose that the prediction model is run for a long time, and let TP denote the number of correctly signaled failures, FP the number of incorrect failure signals (i.e., a signal is generated, but no failure is observed at the end of the demand lead time), and FN the number of cases where no signal is generated for a failure. The precision is given by TP/(TP+FP), while sensitivity is given by TP/(TP+FN). In words, precision captures the correctness of the generated signals, while sensitivity captures the ability to detect upcoming failures. Let $p \in [0, 1]$ denote the precision, and let $q \in [0, 1]$ denote the sensitivity. If $p = 1$ and $q = 1$, we have perfect signals. If $p = 0$ or $q = 0$, the signals are useless. A binary classifier with false negatives and false positives is a common signal-generation mechanism in the maintenance literature; see e.g., Berrade et al. (2013), MacPherson and Glazebrook (2014), Zhang et al. (2021), and Akcay (2022).

If the demand lead time D is at least one period, there is at least one replenishment moment during the demand lead time and at the last replenishment moment a part can be ordered via a regular replenishment (and that part can be used for a corrective replacement as soon as the failure occurs). This implies that having a value of D that is larger than 1 is equally good as having D equal to 1. Therefore, without loss of generality, we limit ourselves to values of $D \in [0, 1]$. If $D = 1$, there is always precisely one replenishment moment during the demand lead time of a failing component and that moment is used to order a spare part in order to replace the component as soon as the failure occurs. If $D = 0$, the signal and the corresponding failure of a

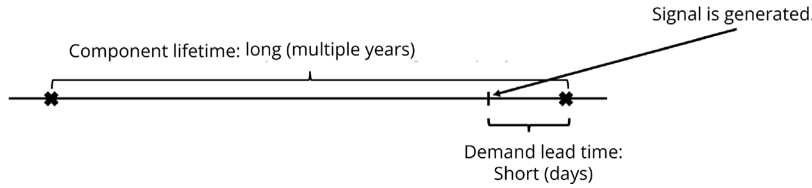

Fig. 1. Generation of predictions.

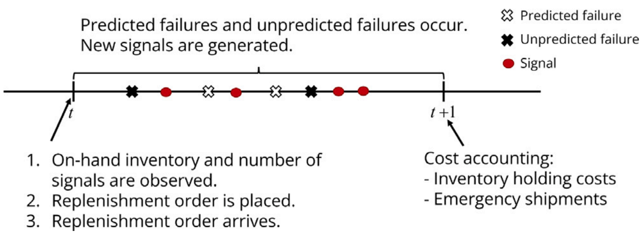

Fig. 2. Order of events in period t.

component occur at the same time, and hence the generated signals are useless.

In each period t, we have the following order of events. First, at the beginning of the period, based on the active number of signals (collected during the time interval $(t-D,t)$) and the on-hand spare parts stock, a replenishment order for ready-for-use spare parts is placed and delivered. Next, during the whole period, replacements of failed parts are executed. Replacements are also executed for failures of systems for which no signals were generated. The replacements can be executed without any delay as long as spare parts are on stock. When the on-hand stock is 0 and a failure occurs, a spare part is delivered from the central warehouse via an emergency shipment, and the failed part is replaced immediately when the emergency shipment has arrived. This leads to a short delay for the execution of the replacement and hence to a downtime cost for the involved system. In addition, we have an extra cost for the emergency shipment. We refer to this procedure as the *emergency procedure* and we denote the corresponding costs by c_{em}. Because this cost factor c_{em} includes the cost of a downtime, it has a large value in many applications in practice. For every failure, a spare part is sent to the local warehouse via a regular replenishment or via an emergency shipment. Hence, the unit costs for the delivered parts are constant under any reasonable policy and therefore these costs are excluded in our model (they are also excluded in the cost factor c_{em}). For parts that are on stock at the end of a period, we have inventory holding costs c_{h} per part. The order of events in each period t is summarized in Fig. 2.

The moments at which failures of technical systems occur are independent of each other and the spare parts provisioning. Hence, the process of the failures per technical system is a renewal process with renewal intervals that follows the general lifetime distribution. The overall failure process, i.e., the process that describes the failures for all technical systems together, is the merger of all these renewal processes. We assume that the number of technical systems is such that the overall failure process behaves as a Poisson process with a constant rate (this will hold when the number of technical systems is sufficiently large; how large that number has to be depends on the lifetime distribution, e.g., it holds for any number when we have an exponential lifetime distribution). Notice that this assumption is often made for spare parts inventory models. Further, we would like to note that the age per component can be included in the data used by the prediction model.

Let λ (>0) be the constant rate of the overall failure process. Correct signals will be generated for a fraction q of these failures. That means that correct signals (true positives) occur according to a Poisson process

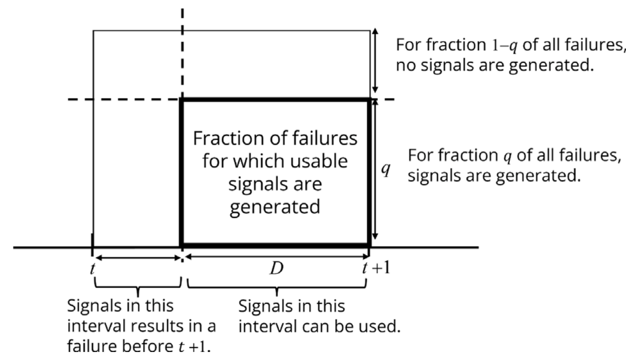

Fig. 3. Fraction of failures for which usable signals are generated.

with rate $q\lambda$. Let $\hat{\lambda}$ be the rate of the Poisson process with which signals arrive; i.e., $\hat{\lambda}$ is the rate with which true positives and false positives arrive. The rate with which true positives arrive is $p\hat{\lambda}$. Since we already know the true positives arrive at rate $q\lambda$, it follows that $\hat{\lambda} = \frac{q\lambda}{p}$ (we take this rate equal to ∞ when $p=0$).

At the beginning of each period, we have to decide how much parts must be replenished. The objective is to minimize the long-run average costs per period, which consist of inventory holding costs and emergency procedure costs. The minimal costs are denoted by $\tilde{C}(p,q,D)$. We include the precision p, sensitivity q, and demand lead time D as parameters, because later we are interested in how these minimal costs behave as a function of p, q, and D.

4. Reduction of the main problem

In this section, we reduce the main problem with a cost function having three parameters (i.e. p, q and D) into a problem with a cost function that has two parameters (i.e. p and r, where $r=qD$). During a period t, a signal that occurs in the first $1-D$ time units will result in a failure before the end of the period (in case of a true positive) or it vanishes before the end of the period (in case of a false positive). For the signals that occur in the last D time units, the failure occurs in period $t+1$. For a fraction q of all upcoming failures a signal is generated, and the required spare parts can be ordered at the beginning of period $t+1$. Overall, for all failures occurring in period $t+1$, qD is the fraction for which signals are generated in period t and $1-qD$ is the fraction for which no signals are generated (see Fig. 3).

The above reasoning shows that the number of predicted failures in a given period $t+1$ is Poisson distributed with rate $qD\lambda$, and the number of unpredicted failures in that period $t+1$ is Poisson distributed with rate $(1-qD)\lambda$. This latter amount is denoted by X^{u}. For the predicted failures, the corresponding number of signals in the preceding period t is Poisson distributed with rate $qD\lambda/p$. These signals are all active at the beginning of period $t+1$. We denote this Poisson distributed amount by X^{s}.

Let us now consider the dynamics in period $t+1$ and how the demand behaves in that period. At the beginning of that period, we have a number of active signals that is a realization of X^{s}. Let us denote this amount by a. For a given a, the number of predicted failures in period $t+1$ is Binomially distributed with a trials and success probability

p. This Binomially distributed amount is denoted by $X^{\mathrm{p}}(a)$. The number of unpredicted failures in period $t+1$ is given by X^{u}. Hence, the total demand in period $t+1$ equals $X^{\mathrm{p}}(a) + X^{\mathrm{u}}$. The parameters of the distributions of $X^{\mathrm{p}}(a)$ and X^{u} only depend on q and D via their product qD. Because q and D play no role in other aspects of our inventory model, this leads to the following theorem.

Theorem 1. *For a given $p \in [0,1]$, the optimal policy and optimal costs $\tilde{C}(p,q,D)$ only depend on the sensitivity q and demand lead time D through their product.*

This theorem implies that a sensitivity $q = \alpha \in [0,1]$ and demand lead time $D = \beta \in [0,1]$ lead to the same optimal policy and optimal costs as a sensitivity $q = \beta$ and demand lead time $D = \alpha$ for a given $p \in [0,1]$. That is, it is equally important to have a high value for the sensitivity q as having a high value for the demand lead time D. Based on Theorem 1, the optimal costs function $\tilde{C}(p,q,D)$ is simplified to the function $C(p,r)$ with $r = qD$ (notice that r represents the fraction of failures for which usable signals are generated).

Remark 1. As stated in Section 3, we have assumed a deterministic demand lead time D, but all results in this paper can be generalized to a setting with a stochastic lead time. That is what we explain in this remark.

Consider a stochastic delay time $D^{\mathrm{st}} \in [0,1]$ with mean $E\{D^{\mathrm{st}}\} = D$. Consider an arbitrary failure in period $t+1$. Given that failures arrive according to a Poisson process, this failure occurs at time $t+1+U$, where U is uniformly distributed on $[0,1]$. A corresponding signal is generated at time $t+1+U-D^{\mathrm{st}}$ with probability q. This signal arrives in period t and can be incorporated for the replenishment decision at time $t+1$ if and only if $t+1+U-D^{\mathrm{st}} \le t+1$. The latter inequality is satisfied with probability

$$P\{D^{\mathrm{st}} \ge U\} = \int_0^1 (1 - F_{D^{\mathrm{st}}}(u))du = E\{D^{\mathrm{st}}\} = D.$$

where $F_{D^{\mathrm{st}}}$ is the distribution function of D^{st}. Hence, for each failure in period $t+1$, we have a usable signal at time $t+1$ with probability qD, and no usable signal is generated with probability $1-qD$. This implies that we get exactly the same dynamics as under a deterministic demand lead time D. And thus the reduction as described in this section and all results in the rest of the paper also hold under the stochastic demand lead time D^{st}.

5. MDP formulation

In this section, we provide the MDP formulation for the reduced problem presented in Section 4. We need this MDP formulation because decisions in subsequent periods depend on each other. This can be seen as follows. At the beginning of a period, the stock will be increased to a certain amount, and with that amount the demands X^{u} and $X^{\mathrm{p}}(a)$ have to be covered. The larger the number of signals, the larger the level to which the inventory position will be increased, but if this level is chosen relatively high and the number of realized demands from the active signals is low, then a relatively large stock is left at the end of the period, and that may lead to a larger stock than desired in the next period. That means that a simple myopic policy that minimizes the costs in the current period will not be optimal.

For the MDP, the state at the beginning of a period is described by (y,a), where y is the on-hand stock and a is the number of active signals. The state space is given by $\mathcal{S} = \{(y,a)|y,a \in \mathbb{N}_0\}$. At the beginning of a period, based on the state (y,a), the on-hand stock is increased by a replenishment order. We describe the action by the level $z \ge y$ to which the inventory position is increased. The replenishment order arrives immediately, and thus the on-hand stock becomes also immediately equal to z. Given action z, the direct expected costs are

$$d(z,a) = \sum_{x=0}^{z} (z-x)P\{X^{\mathrm{u}}+X^{\mathrm{p}}(a) = x\}c_{\mathrm{h}} + \sum_{x=z+1}^{\infty} (x-z)P\{X^{\mathrm{u}}+X^{\mathrm{p}}(a) = x\}c_{\mathrm{em}}.$$

If the total demand is x, then the on-hand stock at the beginning of the next period is $(z-x)^+$. The number of active signals $\hat{a}$ at the beginning of the next period is a realization of X^{s}. This results in the following formulas for the n-period costs $V_n(y,a)$:

$$V_{n+1}(y,a) = \min_{z\ge y} \hat{V}_{n+1}(z,a), \quad (y,a) \in \mathcal{S}, \tag{1}$$

where

$$\hat{V}_{n+1}(z,a) = d(z,a) + \sum_{\hat{a}=0}^{\infty} P\{X^{\mathrm{s}} = \hat{a}\}\Bigg(P\{X^{\mathrm{u}} + X^{\mathrm{p}}(a) \ge z\}V_n(0,\hat{a}) + \sum_{x=0}^{z-1} P\{X^{\mathrm{u}} + X^{\mathrm{p}}(a) = x\}V_n(z-x,\hat{a})\Bigg)$$

and $V_0(y,a) = 0$ for all $(y,a) \in \mathcal{S}$. Appendix B presents how the contributions of the inventory holding costs and emergency procedure costs to $V_n(y,a)$ are calculated by using the MDP formulation. The optimal costs $C(p,r)$ are obtained by $C(p,r) = \lim_{n\to\infty} \frac{V_n(0,0)}{n}$. In order to see the effect of p and r on the optimal costs, we compare $C(p,r)$ with respect to the worst case situation where $p=0$ and $r=0$. For this purpose, we define $\hat{C}(p,r) = \frac{C(p,r)}{C(0,0)}$. Then, $\hat{C}(0,0) = 1$ and $\hat{C}(p,r)$ denotes how close we are to the worst case situation at each point (p,r). For example, $\hat{C}(p,r) = 0.8$ means that we have 80% of the costs associated with the worst case situation.

6. Special cases

In this section, we study two special cases of our model and provide analytical and numerical results. For this purpose, we first define a *base instance* with parameters $\lambda = 0.2$ (failures/demands per week), $c_{\mathrm{h}} = 1$ (Euro per part per week), and $c_{\mathrm{em}} = 10^4$ (Euro per application of the emergency procedure). The value for λ is a common value for the demand rate at a local stockpoint for one component. The value for c_{h} can be chosen w.l.o.g.; a value of 1 Euro per part per week (hence, 50 Euro per part per year) corresponds to a part with a price of 250–500 Euro. The value of c_{em} includes the costs for executing an emergency shipment and the costs for extra downtime of the involved system while it is waiting for the delivery of the part. Normally an emergency shipment takes multiple hours and then the downtime costs can be significant. Hence, 10,000 Euro as costs for the application of an emergency procedure is a quite common number in practice.

6.1. Special case 1: Perfect precision ($p = 1$)

Consider the special case with precision $p = 1$. Then every active signal at the beginning of a period results in an actual failure and it is optimal to take one part on stock per active signal. The number of unpredicted demands X^{u} is Poisson distributed with rate $(1-r)\lambda$. The optimal amount of stock for the unpredicted failures is like the optimal base stock level in a basic model with only unpredicted failures. This basic model is described in Appendix A. In this case, we have a basic model instance with a Poisson demand process with rate $(1-r)\lambda$, and with cost parameters c_{h} and c_{em} for inventory holding and the emergency procedure. The optimal base stock level is denoted by $S^*((1-r)\lambda)$, and this denotes the optimal stock for the unpredicted failures. For the predicted and unpredicted failures together, it is optimal to increase the on-hand stock to $z^*(y,a) = a + S^*((1-r)\lambda)$, when being in state (y,a) at the beginning of a period. If this rule is followed in every period, then the on-hand stock y at the beginning of a period will never exceed $S^*((1-r)\lambda)$ and hence is never larger than $a + S^*((1-r)\lambda)$. This leads to part (a) of the following lemma. Parts (b) and (c) of this lemma follow directly from Lemma A1.

Lemma 1. *For precision $p = 1$, it holds that:*

(1) It is optimal to increase the on-hand stock to $z^*(y,a) = a + S^*((1-r)\lambda)$ at the beginning of each period when being in state (y,a). The base stock level $S^*((1-r)\lambda)$ is non-increasing as a function of r;

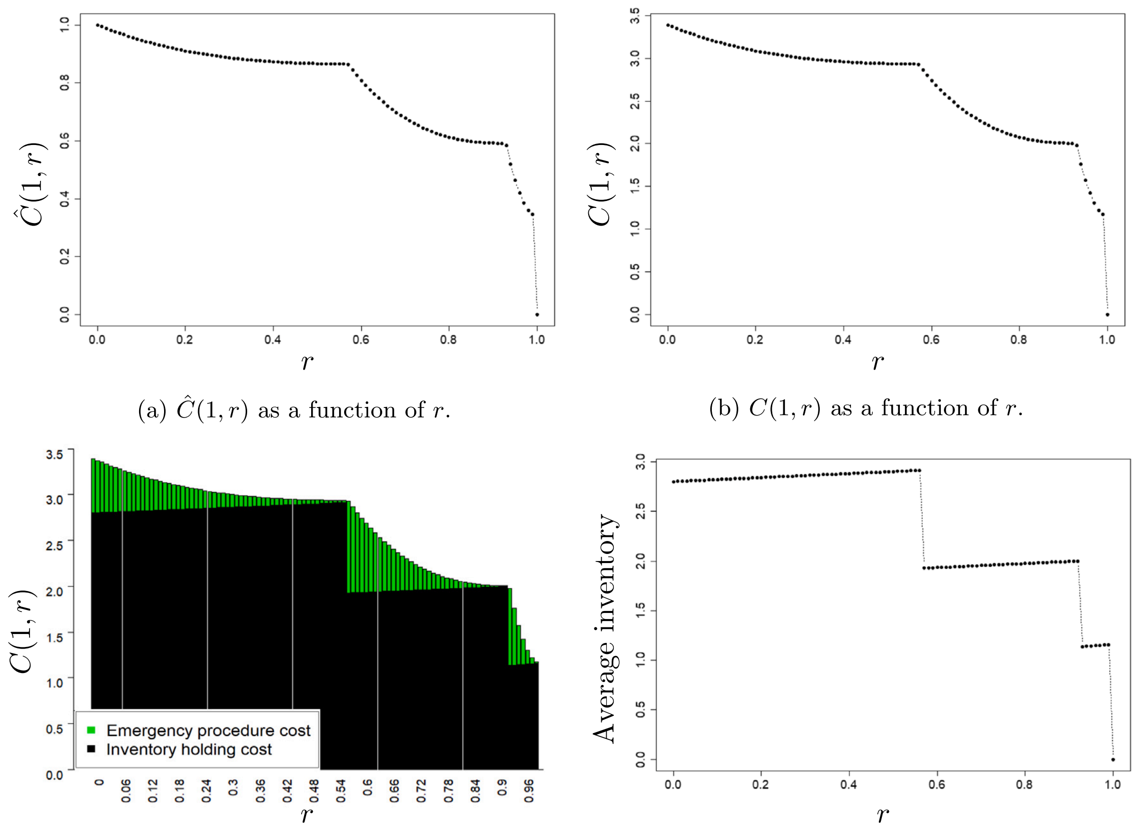

(a) $\hat{C}(1, r)$ as a function of r.

(b) $C(1, r)$ as a function of r.

(c) Share of inventory holding and emergency procedure costs in $C(1, r)$ as a function of r.

(d) Average on-hand inventory as a function of r.

Fig. 4. Change in costs and inventory levels for the base instance with $p = 1$.

(2) The base stock level $S^*((1-r)\lambda)$ equals 0 if and only if $(1-r)\lambda \le ln(1 + (c_h/c_{em}))$;

(3) If $(1-r)\lambda \le ln(1 + (c_h/c_{em}))$, then the optimal costs are equal to $(1-r)\lambda c_{em}$.

In the best case, $r = 1$. Then all failures are predicted and for each failure a part can be ordered at the first order moment after a signal occurs. In that case, the optimal costs are equal to zero and no parts have to be kept on stock for unpredicted failures. These observations lead to the following corollary.

Corollary 1. *For precision $p = 1$, it holds that: If in addition $r = 1$, then the optimal costs are equal to 0 and $S^*((1-r)\lambda) = 0$.*

Next, we investigate the behavior of the optimal costs and the average on-hand inventory (at the end of a period) for the base instance. In Fig. 4(a-b), we show how $\hat{C}(1, r)$ and $C(1, r)$ change as a function of r. We see that costs are non-increasing and piecewise convex functions as a function of r. We also see an inverse Pareto principle: 70% perfectness for r leads to only a 35% reduction in optimal costs. In Fig. 4(c), we illustrate the share of emergency procedure costs and inventory holding costs in $C(1, r)$ as a function of r. We see that the emergency procedure costs decrease as r increases until a certain point. This can be explained by the average on-hand inventory which is depicted in Fig. 4(d). For low values of r, the emergency procedure costs are relatively high. As r increases, the emergency procedure costs decrease until the point where the base stock level $S^*((1-r)\lambda)$ is decreased from 3 to 2. At that point, the average on-hand inventory decreases with a large jump, and the emergency procedure costs increase. After that point, a similar behavior is obtained until a second jump point, and that behavior is also obtained in the interval between that second jump point and $r = 1$.

In Fig. 5, we provide the optimal order-up-to levels $z^*(y, a)$ for the base instance with two different values of r. The optimal order-up-to level $z^*(y, a)$ is increasing as a function of a for a given y, and the other way around. We also see that, in all states, $z^*(y, a)$ is smaller for $r = 0.6$ than for $r = 0.5$ (which is in line with Lemma 1(a)).

6.2. Special case 2: Perfect sensitivity and perfect timing of predictions ($r = 1$)

In this special case, we assume both a perfect sensitivity and a perfect demand lead time (i.e. $r = 1$). We again investigate the behavior of the optimal costs and the average on-hand inventory for the base instance. In Fig. 6(a-b), we observe that $\hat{C}(p, 1)$ and $C(p, 1)$ are non-increasing in p.

We further note that the decreasing behavior of the cost functions is different than what we observed in Fig. 4(a-b). Specifically, the optimal costs decrease fastly for low values of p, while they decrease slowly for large values of p. Further, we now observe a Pareto principle: 30% perfectness for the precision p brings 70% reduction for the optimal costs.

We illustrate the share of emergency procedure costs and inventory holding costs in $C(p, 1)$ as a function of p in Fig. 6(c) and the average on-hand inventory in Fig. 6(d). For very low values of p, the emergency procedure costs are relatively high. After that, they quickly decrease to zero. The inventory holding costs and the average on-hand inventory are first non-decreasing for very low values of p and they decrease on the rest of the interval [0,1].

In Fig. 7, we provide the optimal order-up-to levels $z^*(y, a)$ with respect to on-hand inventory levels y and the number of active signals a. We see that even though we have a relatively low level of p (i.e. $p =$

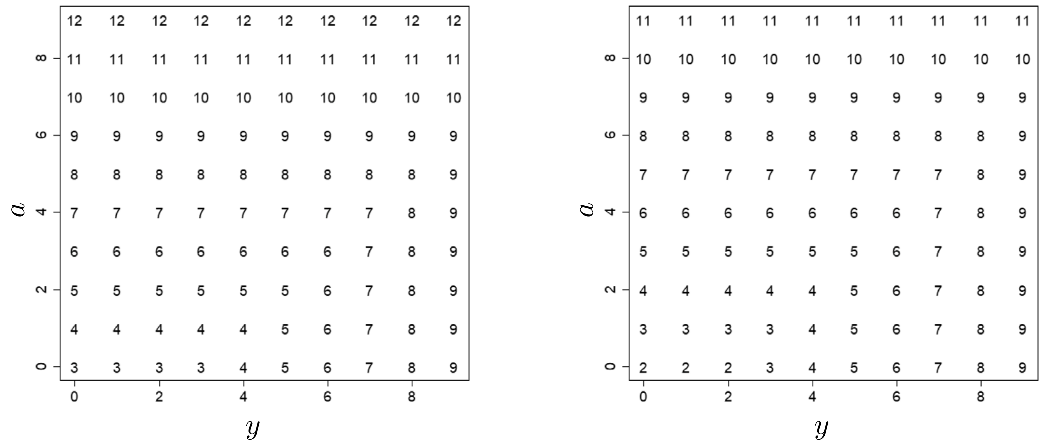

(a) For $p = 1$ and $r = 0.5$. (b) For $p = 1$ and $r = 0.6$.

Fig. 5. Optimal actions $z^*(y, a)$ for the base instance.

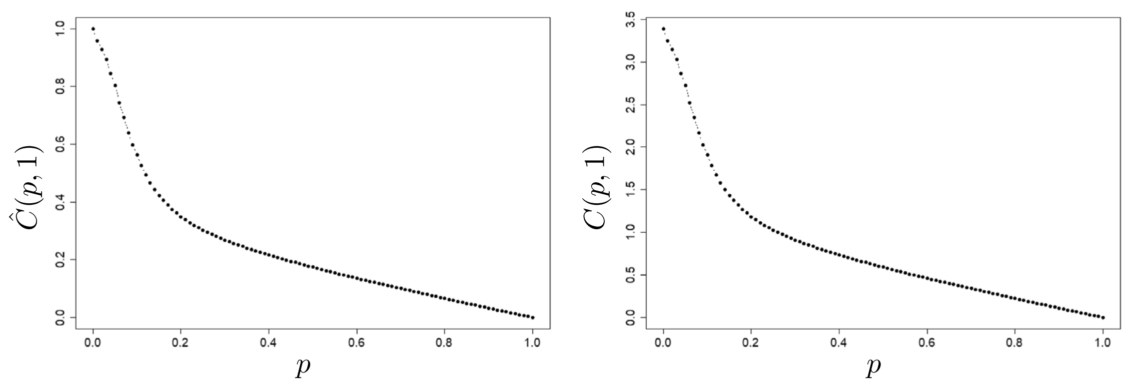

(a) $\hat{C}(p, 1)$ as a function of p. (b) $C(p, 1)$ as a function of p.

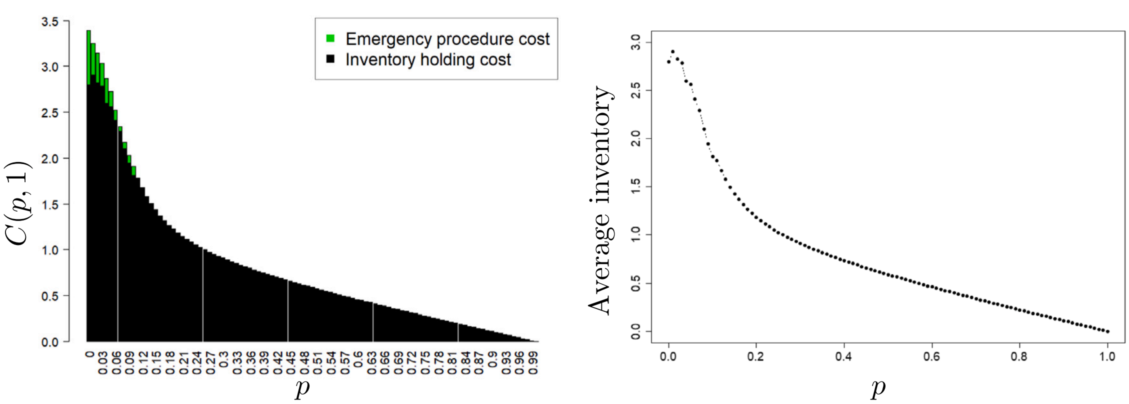

(c) Share of inventory holding and emergency procedure costs in $C(p, 1)$ as a function of p. (d) Average on-hand inventory as a function of p.

Fig. 6. Change in costs and inventory levels for the base instance with $r = 1$.

Table 1
$\hat{C}(p,r)$(%) for the base instance.

$r \backslash p$	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
0.1	100.0	97.7	96.1	95.6	95.3	95.1	95.0	94.9	94.8	94.8	94.7
0.2	100.0	96.2	94.3	93.1	92.4	91.9	91.7	91.4	91.3	91.2	91.1
0.3	100.0	94.3	93.5	91.9	90.8	90.2	89.7	89.3	89.1	88.8	88.7
0.4	100.0	93.3	91.1	90.8	90.2	89.3	88.7	88.2	87.8	87.5	87.3
0.5	100.0	92.2	89.4	88.6	88.3	88.2	88.1	87.8	87.4	87.0	86.7
0.6	100.0	89.0	85.4	83.3	82.2	81.7	81.4	81.2	81.1	81.0	80.9
0.7	100.0	85.2	78.1	74.2	71.9	70.5	69.6	68.9	68.4	68.1	67.8
0.8	100.0	82.1	74.1	70.0	67.4	65.4	64.0	63.0	62.3	61.7	61.3
0.9	100.0	75.8	67.0	64.3	63.6	63.4	62.3	61.3	60.5	59.8	59.2
1	100.0	56.3	34.9	26.9	21.6	17.4	13.6	10.0	6.6	3.2	0.0

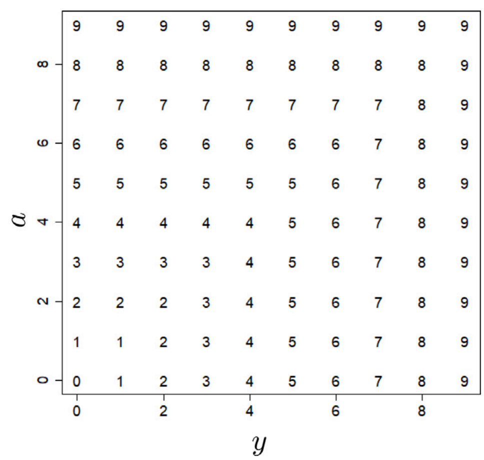

Fig. 7. Optimal actions $z^*(y,a)$ for the base instance with $p = 0.5$ and $r = 1$.

0.5), it holds that $z^*(y,a) = \min\{a,y\}$ at all points in this figure. That means that for all active signals at the beginning of a period a spare part is taken on stock, which explains that we have zero emergency procedure costs in this case (see Fig. 6(c)).

7. Computational experiments

In this section, we provide the results of our computational experiments for varying values of p and r. We also perform a sensitivity analysis on the parameters of the base instance.

7.1. Computational experiments for the base instance for a general p and r

The goal of this section is to generate further insights on the effect of p and r on $\hat{C}(p,r)$, the average on-hand inventory, and the average number of emergency procedure applications (per week) for the base instance. In Table 1, we observe how $\hat{C}(p,r)$ changes with respect to p and r for the base instance. For a constant $r > 0$, $\hat{C}(p,r)$ decreases in p. We observe the Pareto principle in each row of Table 1. For example, for $r = 1$, $\hat{C}(p,r)$ decreases 65% when p is only 20% of the perfect level. This effect can be seen more clearly in Fig. 8(a). We also see that the larger r, the stronger $\hat{C}(p,r)$ decreases as a function of p. On the other hand, for a constant $p > 0$, we see the inverse Pareto principle. For example, for $p = 1$, in order to achieve an about 40% decrease in $\hat{C}(p,r)$ (i.e., $\hat{C}(p,r)$ equal to 59.2%), the value of r should be 90% of the perfect level. This behavior can also be observed in Fig. 8(b). Table 1 and Fig. 8 also show that a large value for r is needed to obtain a significant reduction in optimal costs, while for p a moderate value suffices.

In Table 2, we show how the average on-hand inventory behaves as a function of p and r. In the worst case scenario, the average on-hand inventory is 2.80 units. It is a non-monotonic function of p for a fixed r and a non-monotonic function of r for a fixed p. In general, the average on-hand inventory is non-increasing as a function of p and r. Notice that very low values for the average on-hand inventory (≤ 0.5, say) are only obtained for $r = 1$ and $p \geq 0.6$.

In Table 3, we show how the average number of emergency procedure applications behaves as a function of p and r. In the worst case, the average number of emergency procedure applications is 0.59×10^{-4}. Due to the relatively high cost of an emergency procedure, the average number of emergency procedure applications is in general low under the optimal policy. Similar to the average on-hand inventory, the average number of emergency procedure applications is a non-monotonic function of p and r. We see that the average number of emergency procedure applications can be less than or equal to 0.1×10^{-4} for $p \geq 0.1$ with $r = 1$ and for $p \geq 0.6$ with $r \geq 0.8$. Again, in Table 3, we see that focusing on having large values of r is more crucial than focusing on having large values of p for a low average number of emergency procedure applications.

Finally, in Fig. 9, we show the values of the optimal order-up-to level $z^*(y,a)$ as a function of the state variables a and y for $p = 0.8$ and $r = 0.8$. The value of $z^*(y,a)$ is at least 2 in all states (y,a). In fact, we observe that $z^*(y,a) = \min\{2 + a, y\}$ at all points in this figure. That means that in all states one spare part is taken on stock for each active signal and that a stock of (at least) 2 is kept for the unpredicted failures.

7.2. Sensitivity analysis

In this section, we perform a sensitivity analysis for the parameters λ and c_{em} by considering $\lambda \in \{0.1, 0.2, 0.5\}$ (in demands per week) and $c_{\text{em}} \in \{10^2, 10^4, 10^6\}$ (in Euro per emergency procedure application). Because c_{em} consists of a downtime cost and the cost of an emergency shipment, $c_{\text{em}} = 10^2$ is a low value. Some of the insights that we obtain below become different when you have lower values for c_{em}, which may occur for another application than spare parts for technical systems with high system availability requirements; see Remark 2.

In Table 4, we show how $\hat{C}(p,r)$ changes for varying values of λ and c_{em} as a function of r and p. Similar to our earlier observations for the base instance, for all considered values of λ and c_{em}, we see that the Pareto principle holds for an optimal costs reduction in terms of p and an inverse Pareto principle holds in terms of r. Observations regarding the Pareto and inverse Pareto principles are insensitive to the failure rate and the cost of an emergency procedure.

In Table 5, we do a sensitivity analysis for $C(p,r)$ for varying values of λ and c_{em}. We see that $C(p,r)$ is non-decreasing both in the failure rate and in the emergency procedure costs for any (p,r) pair. Similarly, for any (p,r) pair, the average on-hand inventory is non-decreasing both in λ and in c_{em} (see Table 6). The average on-hand inventory levels are very low (≤ 0.5, say) for only a few cases: $r = 1$ and p is moderate to large. The average number of emergency procedure applications in the worst-case scenario is non-monotonic in λ and c_{em}. However, we cannot directly observe a monotonic behavior for the average on-hand inventory (see Table 6) and the average number of emergency procedure applications (see Table 7) for varying values of λ and c_{em}.

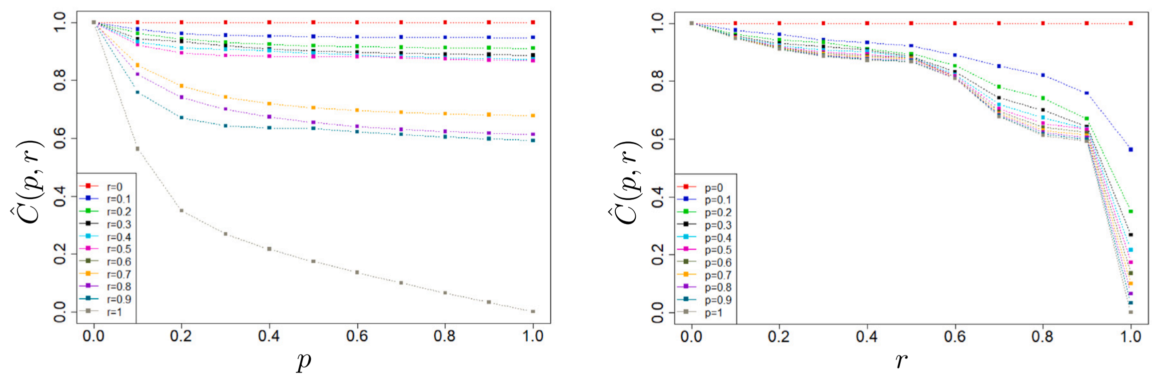

(a) Change in $\hat{C}(p,r)$ as a function of p for varying values of r.

(b) Change in $\hat{C}(p,r)$ as a function of r for varying values of p.

Fig. 8. Combined effect of p and r on $\hat{C}(p,r)$.

Table 2
Average on-hand inventory for the base instance.

$r \backslash p$	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
0	2.80	2.80	2.80	2.80	2.80	2.80	2.80	2.80	2.80	2.80	2.80
0.1	2.80	2.98	2.89	2.86	2.85	2.84	2.83	2.83	2.83	2.82	2.82
0.2	2.80	2.86	2.98	2.92	2.90	2.88	2.87	2.86	2.85	2.84	2.84
0.3	2.80	2.92	2.84	2.98	2.94	2.92	2.90	2.89	2.88	2.87	2.86
0.4	2.80	2.99	2.86	2.83	2.98	2.95	2.93	2.91	2.90	2.89	2.88
0.5	2.80	2.88	2.89	2.85	2.83	2.82	2.82	2.93	2.92	2.91	2.90
0.6	2.80	2.61	2.38	2.20	2.11	2.05	2.01	1.98	1.96	1.95	1.93
0.7	2.80	2.72	2.35	2.26	2.16	2.09	2.04	2.01	1.99	1.97	1.95
0.8	2.80	2.36	2.40	2.24	2.20	2.13	2.08	2.04	2.01	1.99	1.97
0.9	2.80	2.45	2.10	1.96	1.90	1.87	2.07	2.04	2.04	2.01	2.00
1	2.80	1.81	1.18	0.91	0.73	0.59	0.46	0.34	0.22	0.11	0.00

Table 3
Average number of emergency shipments ($\times 10^{-4}$) for the base instance.

$r \backslash p$	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
0	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.59
0.1	0.59	0.33	0.37	0.38	0.38	0.39	0.39	0.39	0.39	0.39	0.39
0.2	0.59	0.40	0.22	0.23	0.23	0.24	0.24	0.24	0.25	0.25	0.25
0.3	0.59	0.28	0.33	0.14	0.14	0.14	0.14	0.14	0.15	0.15	0.15
0.4	0.59	0.17	0.23	0.24	0.08	0.08	0.08	0.08	0.08	0.08	0.08
0.5	0.59	0.24	0.14	0.15	0.16	0.17	0.17	0.04	0.05	0.04	0.04
0.6	0.59	0.40	0.52	0.63	0.68	0.72	0.75	0.77	0.79	0.80	0.81
0.7	0.59	0.18	0.30	0.26	0.28	0.30	0.31	0.33	0.33	0.34	0.34
0.8	0.59	0.42	0.11	0.14	0.09	0.09	0.09	0.10	0.10	0.10	0.10
0.9	0.59	0.12	0.17	0.22	0.25	0.28	0.04	0.04	0.01	0.01	0.01
1	0.59	0.10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Remark 2. We provide additional computational results in Appendix C for $\lambda = 0.2$ and $c_{em} \in \{1, 10\}$. For $c_{em} = 10$, we still observe a reverse Pareto principle for the product of sensitivity and demand lead time under a perfect precision but it is weaker than for higher c_{em} values. A 75% increase in the perfectness of sensitivity and demand lead time brings a 50% decrease in costs under perfect precision. For $c_{em} = 1$, we see a linear behavior. We no longer observe the Pareto principle for the precision under a perfect sensitivity and a perfect demand lead time when $c_{em} = \{1, 10\}$ and $\lambda = 0.2$. A 75% increase in the perfectness of p brings 15% and 72% reduction in costs under a perfect sensitivity and a perfect demand lead time for $c_{em} = 1$ and $c_{em} = 10$, respectively.

The different insights are obtained because of a different balance between the inventory holding costs and the emergence procedure costs that is obtained under an optimal policy for low values of c_{em}. For high values of c_{em}, the optimal inventory levels are such that the application of the emergency procedure is avoided for sufficiently high values of r and p. For low values of c_{em}, the optimal inventory levels are much lower, and hence the emergency procedure is much less avoided for many values of p and r. This leads to a larger share of the emergency procedure costs in the total costs and to a different behavior of the total costs as a function of p and r, respectively.

8. Conclusion

We have studied the spare parts inventory problem of a single critical component that is kept on stock in a single local stockpoint. For upcoming failures of the component in the supported technical systems, signals are generated. For these signals, we distinguish the factors precision, sensitivity, and demand lead time, and we investigated how the average inventory and costs for inventory holding and emergency procedure applications depend on these three factors under optimal inventory control of the spare parts stock. This optimal control is obtained via a Markov decision process. Our investigation gives directions for the trade-off between precision, sensitivity, and demand lead time for developers of these signals. Even when the predictive models perform close to the ideal case, decision-makers should keep inventory on-hand. We found that for a given precision level the optimal inventory control and optimal costs only depend on the sensitivity and the demand lead

Table 4
$\hat{C}(p,r)$ (%) for varying values of λ and c_{em} as a function of r and p.

$r\backslash p$	$\lambda = 0.1$ and $c_{em} = 10^2$					$\lambda = 0.2$ and $c_{em} = 10^2$					$\lambda = 0.5$ and $c_{em} = 10^2$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
0.25	100.0	91.7	89.2	88.3	87.8	100.0	99.6	99.2	99.1	99.1	100.0	92.5	89.8	88.7	88.1
0.5	100.0	88.7	83.3	81.3	80.3	100.0	82.4	77.5	75.6	74.7	100.0	86.3	81.2	77.5	75.2
0.75	100.0	81.1	80.0	78.6	77.1	100.0	76.0	67.6	63.7	61.6	100.0	79.4	69.1	65.5	64.1
1	100.0	58.8	35.9	17.1	0.0	100.0	53.0	30.7	14.6	0.0	100.0	61.5	33.7	15.8	0.0

$r\backslash p$	$\lambda = 0.1$ and $c_{em} = 10^4$					$\lambda = 0.2$ and $c_{em} = 10^4$					$\lambda = 0.5$ and $c_{em} = 10^4$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
0.25	100.0	89.5	88.8	88.7	88.6	100.0	93.0	90.9	90.1	89.7	100.0	92.2	89.8	88.8	88.3
0.5	100.0	77.3	74.7	73.8	73.3	100.0	88.9	88.2	87.6	86.7	100.0	86.3	83.9	83.5	82.1
0.75	100.0	72.9	70.6	69.0	68.2	100.0	73.6	67.3	65.0	63.8	100.0	76.1	70.2	66.4	64.0
1	100.0	27.8	17.0	8.1	0.0	100.0	30.3	17.4	8.3	0.0	100.0	38.3	19.5	9.1	0.0

$r\backslash p$	$\lambda = 0.1$ and $c_{em} = 10^6$					$\lambda = 0.2$ and $c_{em} = 10^6$					$\lambda = 0.5$ and $c_{em} = 10^6$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
0.25	100.0	99.3	99.3	99.3	99.1	100.0	92.1	91.1	90.8	90.7	100.0	91.0	89.6	88.6	88.1
0.5	100.0	83.3	81.4	80.8	80.5	100.0	86.3	83.4	82.2	81.6	100.0	83.3	79.5	77.6	76.7
0.75	100.0	78.0	76.9	75.8	75.2	100.0	72.6	68.0	66.4	65.7	100.0	73.2	67.1	64.1	62.6
1	100.0	20.6	12.5	6.0	0.0	100.0	21.0	12.1	5.7	0.0	100.0	27.2	13.8	6.5	0.0

Table 5
$C(p,r)$ for varying values of λ and c_{em} as a function of r and p.

	$\lambda = 0.1$ and $c_{em} = 10^2$					$\lambda = 0.2$ and $c_{em} = 10^2$					$\lambda = 0.5$ and $c_{em} = 10^2$				
$r\backslash p$	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	1.39	1.39	1.39	1.39	1.39	1.92	1.92	1.92	1.92	1.92	2.70	2.70	2.70	2.70	2.70
0.25	1.39	1.27	1.24	1.23	1.22	1.92	1.91	1.91	1.91	1.91	2.70	2.49	2.42	2.39	2.38
0.5	1.39	1.23	1.16	1.13	1.11	1.92	1.58	1.49	1.45	1.44	2.70	2.33	2.19	2.09	2.03
0.75	1.39	1.13	1.11	1.09	1.07	1.92	1.46	1.30	1.23	1.18	2.70	2.14	1.86	1.77	1.73
1	1.39	0.82	0.50	0.24	0.00	1.92	1.02	0.59	0.28	0.00	2.70	1.66	0.91	0.42	0.00

$r\backslash p$	$\lambda = 0.1$ and $c_{em} = 10^4$					$\lambda = 0.2$ and $c_{em} = 10^4$					$\lambda = 0.5$ and $c_{em} = 10^4$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	2.94	2.94	2.94	2.94	2.94	3.39	3.39	3.39	3.39	3.39	4.65	4.65	4.65	4.65	4.65
0.25	2.94	2.63	2.61	2.61	2.60	3.39	3.15	3.08	3.06	3.04	4.65	4.29	4.18	4.13	4.11
0.5	2.94	2.27	2.20	2.17	2.16	3.39	3.02	2.99	2.97	2.94	4.65	4.02	3.90	3.89	3.82
0.75	2.94	2.14	2.07	2.03	2.00	3.39	2.50	2.28	2.20	2.17	4.65	3.54	3.26	3.09	2.98
1	2.94	0.82	0.50	0.24	0.00	3.39	1.03	0.59	0.28	0.00	4.65	1.78	0.91	0.42	0.00

$r\backslash p$	$\lambda = 0.1$ and $c_{em} = 10^6$					$\lambda = 0.2$ and $c_{em} = 10^6$					$\lambda = 0.5$ and $c_{em} = 10^6$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	3.98	3.98	3.98	3.98	3.98	4.88	4.88	4.88	4.88	4.88	6.57	6.57	6.57	6.57	6.57
0.25	3.98	3.95	3.95	3.95	3.94	4.88	4.49	4.45	4.43	4.42	6.57	5.98	5.88	5.82	5.78
0.5	3.98	3.32	3.24	3.22	3.20	4.88	4.21	4.07	4.01	3.98	6.57	5.47	5.22	5.09	5.03
0.75	3.98	3.10	3.06	3.02	2.99	4.88	3.54	3.32	3.24	3.20	6.57	4.80	4.40	4.21	4.11
1	3.98	0.82	0.50	0.24	0.00	4.88	1.03	0.59	0.28	0.00	6.57	1.78	0.91	0.42	0.00

time via their product. This implies that, when developing signals, getting a high value for sensitivity is equally important as getting a high value for the demand lead time. Furthermore, a low value in sensitivity (demand lead time) will decrease the effectiveness of a high value in demand lead time (sensitivity) on cost reduction. Further, we found that both factors need to have a high value in order to get a significant reduction in optimal costs and average inventory in comparison to the situation without signals. For the precision, a significant cost reduction is obtained for high values but also for moderate values. So, it is much better to develop signals with a moderate value for the precision and high values for the sensitivity and demand lead time than the other way around.

For future research, it would be interesting to investigate whether the main insights also hold in settings with relaxed assumptions. It would be relevant to study a setting with a stochastic demand lead time that can be larger than one period. It would also be interesting to investigate how much the results change under continuous review and/or a positive replenishment lead time.

Data availability

No data was used for the research described in the article.

Acknowledgments

The authors are grateful to the reviewers whose comments have helped a lot to improve the paper.

Appendix A. Basic model for unpredicted demand ($q = 0$)

In this appendix, we describe a basic model for unpredicted demand that is used as a building block for the analysis of the general model of

Table 6
Average on-hand inventory for varying values of λ and c_{em} as a function of r and p.

$r \backslash p$	$\lambda = 0.1$ and $c_{em} = 10^2$					$\lambda = 0.2$ and $c_{em} = 10^2$					$\lambda = 0.5$ and $c_{em} = 10^2$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	0.90	0.90	0.90	0.90	0.90	1.80	1.80	1.80	1.80	1.80	2.50	2.50	2.50	2.50	2.50
0.25	0.90	1.02	0.98	0.96	0.95	1.80	1.82	1.81	1.81	1.80	2.50	1.94	1.79	1.71	1.67
0.5	0.90	1.12	1.04	1.01	1.00	1.80	1.21	1.09	1.02	0.99	2.50	1.82	1.95	1.88	1.81
0.75	0.90	0.96	0.92	1.06	1.04	1.80	1.35	1.16	1.12	1.08	2.50	1.65	1.45	1.24	1.13
1	0.90	0.81	0.50	0.24	0.00	1.80	1.02	0.59	0.28	0.00	2.50	1.61	0.91	0.42	0.00

$r \backslash p$	$\lambda = 0.1$ and $c_{em} = 10^4$					$\lambda = 0.2$ and $c_{em} = 10^4$					$\lambda = 0.5$ and $c_{em} = 10^4$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	2.90	2.90	2.90	2.90	2.90	2.80	2.80	2.80	2.80	2.80	4.50	4.50	4.50	4.50	4.50
0.25	2.90	2.00	1.95	1.94	1.93	2.80	2.98	2.90	2.87	2.85	4.50	3.90	3.75	3.67	3.63
0.5	2.90	2.08	2.00	1.97	1.95	2.80	2.86	2.82	2.93	2.90	4.50	3.78	3.61	3.79	3.75
0.75	2.90	1.94	2.04	2.00	1.98	2.80	2.28	2.11	2.01	1.96	4.50	3.16	3.08	3.01	2.89
1	2.90	0.82	0.50	0.24	0.00	2.80	1.03	0.59	0.28	0.00	4.50	1.78	0.91	0.42	0.00

$r \backslash p$	$\lambda = 0.1$ and $c_{em} = 10^6$					$\lambda = 0.2$ and $c_{em} = 10^6$					$\lambda = 0.5$ and $c_{em} = 10^6$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	3.90	3.90	3.90	3.90	3.90	4.80	4.80	4.80	4.80	4.80	6.50	6.50	6.50	6.50	6.50
0.25	3.90	3.91	3.90	3.90	3.93	4.80	4.00	3.90	3.87	3.85	6.50	5.61	5.72	5.67	5.63
0.5	3.90	3.10	3.00	2.97	2.95	4.80	4.13	4.00	3.93	3.90	6.50	5.21	4.98	4.83	4.75
0.75	3.90	2.94	3.04	3.00	2.98	4.80	3.37	3.10	3.00	2.95	6.50	4.53	4.24	4.00	3.88
1	3.90	0.82	0.50	0.24	0.00	4.80	1.03	0.59	0.28	0.00	6.50	1.78	0.91	0.42	0.00

Table 7
Average number of emergency procedure applications for varying values of λ and c_{em} as a function of r and p.

$r \backslash p$	$\lambda = 0.1$ and $c_{em} = 10^2$ $(\times 10^{-2})$					$\lambda = 0.2$ and $c_{em} = 10^2$ $(\times 10^{-2})$					$\lambda = 0.5$ and $c_{em} = 10^2$ $(\times 10^{-2})$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	0.48	0.48	0.48	0.48	0.48	0.12	0.12	0.12	0.12	0.12	0.19	0.19	0.19	0.19	0.19
0.25	0.48	0.25	0.26	0.27	0.27	0.12	0.09	0.10	0.10	0.10	0.19	0.55	0.63	0.68	0.71
0.5	0.48	0.11	0.12	0.12	0.12	0.12	0.37	0.40	0.43	0.45	0.19	0.50	0.24	0.21	0.22
0.75	0.48	0.17	0.19	0.03	0.04	0.12	0.11	0.14	0.11	0.11	0.19	0.49	0.42	0.53	0.60
1	0.48	0.01	0.00	0.00	0.00	0.12	0.00	0.00	0.00	0.00	0.19	0.04	0.00	0.00	0.00

$r \backslash p$	$\lambda = 0.1$ and $c_{em} = 10^4$ $(\times 10^{-4})$					$\lambda = 0.2$ and $c_{em} = 10^4$ $(\times 10^{-4})$					$\lambda = 0.5$ and $c_{em} = 10^4$ $(\times 10^{-4})$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	0.04	0.04	0.04	0.04	0.04	0.59	0.59	0.59	0.59	0.59	0.15	0.15	0.15	0.15	0.15
0.25	0.04	0.63	0.66	0.67	0.68	0.59	0.17	0.18	0.19	0.19	0.15	0.39	0.43	0.46	0.48
0.5	0.04	0.19	0.19	0.20	0.20	0.59	0.15	0.17	0.04	0.04	0.15	0.23	0.29	0.10	0.07
0.75	0.04	0.20	0.03	0.03	0.03	0.59	0.22	0.17	0.19	0.20	0.15	0.38	0.18	0.08	0.09
1	0.04	0.00	0.00	0.00	0.00	0.59	0.00	0.00	0.00	0.00	0.15	0.00	0.00	0.00	0.00

$r \backslash p$	$\lambda = 0.1$ and $c_{em} = 10^6$ $(\times 10^{-6})$					$\lambda = 0.2$ and $c_{em} = 10^6$ $(\times 10^{-6})$					$\lambda = 0.5$ and $c_{em} = 10^6$ $(\times 10^{-6})$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.07	0.07	0.07	0.07	0.07
0.25	0.08	0.05	0.05	0.05	0.02	0.08	0.50	0.55	0.56	0.57	0.07	0.37	0.16	0.15	0.16
0.5	0.08	0.22	0.24	0.25	0.25	0.08	0.07	0.07	0.08	0.08	0.07	0.26	0.24	0.26	0.28
0.75	0.08	0.16	0.02	0.02	0.02	0.08	0.17	0.22	0.24	0.25	0.07	0.28	0.17	0.21	0.23
1	0.08	0.00	0.00	0.00	0.00	0.08	0.00	0.00	0.00	0.00	0.07	0.00	0.00	0.00	0.00

Section 3. For this basic model, everything is the same as for the model of Section 3, but we assume that no signals are generated (i.e., $q = 0$) and that failures occur according to a Poisson process with rate μ. We introduce a new parameter for this rate because we will use this basic model for different demand rates. For the emergency procedure costs and the inventory holding costs, we still use the cost parameters c_{em} and c_h.

For this basic model, the demand per period is denoted by X, which is Poisson distributed with rate μ. Hence,

$$P\{X = x\} = \frac{\mu^x}{x!}e^{-\mu}, \quad x \in \mathbb{N}_0.$$

Consider the replenishment decision at the beginning of period 0. The initial inventory level is 0. Suppose that the on-hand inventory is increased to S (by ordering S units), then the expected costs in period 0 are equal to

$$\tilde{G}(\mu, S) = \sum_{x=0}^{S}(S - x)P\{X = x\}c_h + \sum_{x=S+1}^{\infty}(x - S)P\{X = x\}c_{em}.$$

This function $\tilde{G}(\mu, S)$ is similar to the cost function for a newsvendor problem. It is convex as a function of S, and is minimized at the lowest S for which

$$P\{X \leq S\} \geq \frac{c_{em}}{c_{em} + c_h}.$$

This optimal S is denoted by $S^*(\mu)$. Let the corresponding minimal costs for period 0 be denoted by $G(\mu) = \tilde{G}(\mu, S^*(\mu))$.

Let us now look at the whole time horizon. It is not possible to get strictly lower expected costs per period than $\tilde{G}(\mu, S^*(\mu))$. By following a base stock policy with base stock level $S^*(\mu)$ (i.e., by increasing the on-hand inventory at the beginning of each period to $S^*(\mu)$), we get expected costs $\tilde{G}(\mu, S^*(\mu))$ in each period, and thus the resulting

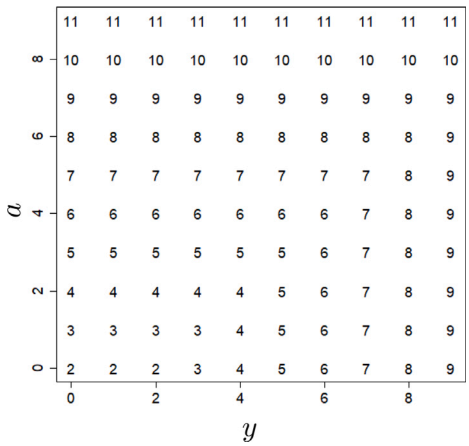

Fig. 9. Optimal action $z^*(y,a)$ for the base instance, where $p = 0.8$ and $r = 0.8$.

average costs per period are also equal to $\tilde{G}(\mu, S^*(\mu))$. Therefore, this base stock policy is optimal, and the corresponding minimal costs are equal to $G(\mu) = \tilde{G}(\mu, S^*(\mu))$. The following results hold for $S^*(\mu)$ and $G(\mu)$.

Lemma A1.

(a) $S^*(\mu)$ *is non-decreasing as a function of* μ.
(b) $S^*(\mu) = 0$ *if and only if* $\mu \leq ln(1 + (c_h/c_{em}))$.
(c) If $\mu \leq ln(1 + (c_h/c_{em}))$, *then* $G(\mu) = \mu c_{em}$.

Proof. First, we prove Lemma A1(a). A rate μ leads to the optimal inventory level $S^*(\mu)$. Let us assume another rate $\mu + \delta$, where $\delta > 0$ is a small increment. Then $X_{\mu+\delta}$ stochastically dominates X_μ, where they respectively represent the random variables of a Poisson distribution with rate $\mu + \delta$ and μ. This results in $S^*(\mu)$ being the smallest value for which $P\{X_\mu \leq S^*(\mu)\} \geq \frac{c_{em}}{c_{em}+c_h}$ and $S^*(\mu+\delta)$ being the smallest value for which $P\{X_{\mu+\delta} \leq S^*(\mu+\delta)\} \geq \frac{c_{em}}{c_{em}+c_h}$. Because $X_{\mu+\delta}$ stochastically dominates X_μ, it holds that $P\{X_\mu \leq S\} \geq P\{X_{\mu+\delta} \leq S\}$ for all S. Therefore $P\{X_\mu \leq S^*(\mu+\delta)\} \geq P\{X_{\mu+\delta} \leq S^*(\mu+\delta)\} \geq \frac{c_{em}}{c_{em}+c_h}$ and hence $S^*(\mu) \leq S^*(\mu+\delta)$. This proves that $S^*(\mu)$ is non-decreasing as a function of μ.

Next we prove Lemma A1(b). It holds that $S^*(\mu) = 0$ if and only if

$$P\{X \leq 0\} \geq \frac{c_{em}}{c_{em}+c_h} \Leftrightarrow e^{-\mu} \geq \frac{c_{em}}{c_{em}+c_h}$$
$$\Leftrightarrow -\mu \leq ln\left(\frac{c_{em}}{c_{em}+c_h}\right) \Leftrightarrow \mu \leq ln\left(1 + \frac{c_h}{c_{em}}\right).$$

Lemma A1(c) follows directly from the observation that all failures are solved by applying the emergency procedure if no parts are kept on stock. □

Appendix B. Value iteration algorithm

In this section, we show how the inventory holding and emergency procedure costs under the optimal policy can be calculated when solving the MDP formulation. Algorithm 1 provides a pseudocode of our algorithm. We let $d^h(z,a) = \sum_{x=0}^{z}(z-x)P(X^u + X^p(a) = x)c_h$ denotes the direct expected costs of inventory holding and $d^{em}(z,a) = \sum_{x=z+1}^{\infty}(z-x)P(X^u + X^p(a) = x)c_{em}$ denotes the direct expected costs for emergency procedure applications. The direct expected costs $d(z,a)$ are equal to the sum of the direct expected costs of inventory holding and the direct expected costs for emergency procedure applications, i.e., $d(z,a) = d^h(z,a) + d^{em}(z,a)$. By splitting direct expected costs into two, we can calculate the costs contribution of each into the total costs separately. For this purpose, we introduce $\hat{V}^h_{n+1}(z,a)$ and $\hat{V}^{em}_{n+1}(z,a)$ denoting the so-called value functions for inventory holding and emergency procedure costs, respectively. We calculate these functions by a value iteration algorithm. At each iteration, we calculate $\hat{V}^h_{n+1}(z,a)$ and $\hat{V}^{em}_{n+1}(z,a)$ independently, then we sum them up to update the value of $\hat{V}_{n+1}(z,a)$ for all (z,a), and then we calculate the optimal value of $V_{n+1}(y,a)$ for all (y,a). The value iteration algorithm stops when the long-run average cost per period (i.e., value of $V_n(y,a)/n$) converges to a constant at some sufficiently large value of n. The convergence is checked by comparing the deviation of the average cost per-period in two subsequent steps of the algorithm to a small number ϵ (=0.01).

Algorithm 1 Value Iteration Algorithm

Initialize $V_0(y,a) \leftarrow 0$, $V^h_0(y,a) \leftarrow 0$, $V^{em}_0(y,a) \leftarrow 0 \; \forall (y,a) \in S$, $n = 0$, Stop=False

while Stop=False **do**

  for $\forall (y,a) \in S$ **do**

    for $\forall z \geq y$ **do**

$$\hat{V}^h_{n+1}(z,a) = d^h(z,a) + \sum_{\hat{a}=0}^{\infty} P\{X^s = \hat{a}\}\Bigg(P\{X^u + X^p(a) \geq z\}V^h_n(0,\hat{a}) + \sum_{x=0}^{z-1} P\{X^u + X^p(a) = x\}V^h_n(z-x,\hat{a})\Bigg)$$

$$\hat{V}^{em}_{n+1}(z,a) = d^{em}(z,a) + \sum_{\hat{a}=0}^{\infty} P\{X^s = \hat{a}\}\Bigg(P\{X^u + X^p(a) \geq z\}V^{em}_n(0,\hat{a}) + \sum_{x=0}^{z-1} P\{X^u + X^p(a) = x\}V^{em}_n(z-x,\hat{a})\Bigg)$$

$$\hat{V}_{n+1}(z,a) = \hat{V}^h_{n+1}(z,a) + \hat{V}^{em}_{n+1}(z,a)$$

    end for

    $V_{n+1}(y,a) \leftarrow \min_{z \geq y}\{\hat{V}(z,a)\}$

    $z^*(y,a) \leftarrow \text{argmin}_{z \geq y}\{\hat{V}(z,a)\}$

    $V^h_{n+1}(y,a) = \hat{V}^h_{n+1}(z^*(y,a),a)$

    $V^{em}_{n+1}(y,a) = \hat{V}^{em}_{n+1}(z^*(y,a),a)$

  end for

  if $n > 0$ and $\max_{\forall (y,a) \in S}\{|\frac{V_{n+1}(y,a)}{n+1} - \frac{V_n(y,a)}{n}|\} \leq \epsilon$ **then**
    Stop=True

  $n = n + 1$

end while

Finally, we define the costs of inventory holding and the emergency procedure costs as

$$C^h(p,r) = \lim_{n\to\infty} \frac{V^h_n(0,0)}{n}, \quad C^{em}(p,r) = \lim_{n\to\infty} \frac{V^{em}_n(0,0)}{n}.$$

Dividing $C^h(p,r)$ by c_h gives the average on-hand inventory level for a given (p,r), and dividing $C^{em}(p,r)$ by c_{em} results in the average number of emergency procedure applications per week.

Appendix C. Computational experiment referred to in Remark 2

In this appendix, we provide computational results for problem instances with low ratios for c_{em}/c_h (i.e., we take $c_h = 1$ and $c_{em} \in$

Table C.8
$\hat{C}(p,r)$ (%) for $\lambda = 0.2$ and $c_{em} \in \{1, 10\}$ as a function of r and p.

$r \backslash p$	$c_{em} = 1$					$c_{em} = 10$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
0.25	100.0	100.0	100.0	100.0	75.0	100.0	99.0	99.0	98.0	98.0
0.5	100.0	100.0	100.0	95.0	50.0	100.0	99.0	96.0	95.0	95.0
0.75	100.0	100.0	100.0	90.0	25.0	100.0	88.1	73.3	60.4	49.5
1	100.0	100.0	95.0	85.0	0.0	100.0	81.2	57.4	27.7	0.0

Table C.9
Average on-hand inventory for $\lambda = 0.2$ and $c_{em} \in \{1, 10\}$ as a function of r and p.

$r \backslash p$	$c_{em} = 1$					$c_{em} = 10$				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	0.00	0.00	0.00	0.00	0.00	0.82	0.82	0.82	0.82	0.82
0.25	0.00	0.00	0.00	0.00	0.00	0.82	0.82	0.83	0.82	0.82
0.5	0.00	0.00	0.00	0.01	0.00	0.82	0.83	0.84	0.83	0.82
0.75	0.00	0.00	0.01	0.01	0.00	0.82	0.64	0.46	0.21	0.00
1	0.00	0.00	0.01	0.02	0.00	0.82	0.67	0.57	0.28	0.00

Table C.10
Average number of emergency procedure applications per period for $\lambda = 0.2$ and $c_{em} \in \{1, 10\}$ as a function of r and p.

$r \backslash p$	$c_{em} = 1$					$c_{em} = 10$ ($\times 10^{-1}$)				
	0	0.25	0.5	0.75	1	0	0.25	0.5	0.75	1
0	0.20	0.20	0.20	0.20	0.20	0.19	0.19	0.19	0.19	0.19
0.25	0.20	0.20	0.20	0.19	0.15	0.19	0.18	0.17	0.17	0.17
0.5	0.20	0.20	0.19	0.18	0.10	0.19	0.17	0.13	0.13	0.14
0.75	0.20	0.20	0.19	0.17	0.05	0.19	0.25	0.28	0.40	0.50
1	0.20	0.20	0.19	0.16	0.00	0.19	0.16	0.01	0.00	0.00

$\{1, 10\}$). For these problem instances, we fix the arrival rate of failures per period at $\lambda = 0.2$. In Tables C.8–C.10, we provide the values for $\hat{C}(p,r)$ (%), the average on-hand inventory and the average number of emergency procedure applications per period as a function of r and p.

References

Akcay, A., 2022. An alert-assisted inspection policy for a production process with imperfect condition signals. European J. Oper. Res. 298 (2), 510–525.

Barraza-Barraza, D., Limón-Robles, J., Beruvides, M.G., 2014. Opportunities and challenges in condition-based maintenance research. In: Industrial and Systems Engineering Research Conference.

Basten, R.J., Ryan, J.K., 2019. The value of maintenance delay flexibility for improved spare parts inventory management. European J. Oper. Res. 278 (2), 646–657.

Benjaafar, S., Cooper, W.L., Mardan, S., 2011. Production-inventory systems with imperfect advance demand information and updating. Nav. Res. Logist. 58 (2), 88–106.

Bernstein, F., DeCroix, G.A., 2015. Advance demand information in a multiproduct system. Manuf. Serv. Oper. Manag. 17 (1), 52–65.

Berrade, M.D., Scarf, P.A., Cavalcante, C.A.V., Dwight, R.A., 2013. Imperfect inspection and replacement of a system with a defective state: A cost and reliability analysis. Reliab. Eng. Syst. Saf. 120, 80–87.

Deshpande, V., Iyer, A.V., Cho, R., 2006. Efficient supply chain management at the U.S. coast guard using part-age dependent supply replenishment policies. Oper. Res. 54 (6), 1028–1040.

Elwany, A.H., Gebraeel, N.Z., 2008. Sensor-driven prognostic models for equipment replacement and spare parts inventory. IIE Trans. 40 (7), 629–639.

Eruguz, A.S., Tan, T., van Houtum, G.J., 2018. Integrated maintenance and spare part optimization for moving assets. IISE Trans. 50 (3), 230–245.

Gayon, J.P., Benjaafar, S., de Véricourt, F., 2009. Using imperfect advance demand information in production-inventory systems with multiple customer classes. Manuf. Serv. Oper. Manag. 11 (1), 128–143.

Karabağ, O., Eruguz, A.S., Basten, R., 2020. Integrated optimization of maintenance interventions and spare part selection for a partially observable multi-component system. Reliab. Eng. Syst. Saf. 200, 106955.

Li, R., Ryan, J.K., 2011. A Bayesian inventory model using real-time condition monitoring information. Prod. Oper. Manage. 20 (5), 754–771.

Liberopoulos, G., Koukoumialos, S., 2008. On the effect of variability and uncertainty in advance demand information on the performance of a make-to-stock supplier. MIBES Trans. Int. J. 2 (1), 95–114.

Lin, X., Basten, R.J.I., Kranenburg, A., van Houtum, G.J., 2017. Condition based spare parts supply. Reliab. Eng. Syst. Saf. 168, 240–248.

MacPherson, A.J., Glazebrook, K.D., 2014. How and when to deploy error prone sensors in support of the maintenance of two-phase systems with ageing. IEEE Trans. Reliab. 63 (1), 118–133.

Rippe, C., Kiesmüller, G.P., 2023. The repair kit problem with imperfect advance demand information. European J. Oper. Res. 304 (2), 558–576.

Shi, J., Rozas, H., Yildirim, M., Gebraeel, N., 2023. A stochastic programming model for jointly optimizing maintenance and spare parts inventory for IoT applications. IISE Trans. 55 (4), 419–431.

Song, J.S., Zipkin, P.H., 2012. Newsvendor problems with sequentially revealed demand information. Nav. Res. Logist. 59 (8), 601–612.

Tan, T., 2008. Using imperfect advance demand information in forecasting. IMA J. Manag. Math. 19 (2), 163–173.

Tan, T., Güllü, R., Erkip, N., 2007. Modelling imperfect advance demand information and analysis of optimal inventory policies. European J. Oper. Res. 177 (2), 897–923.

Tan, T., Güllü, R., Erkip, N., 2009. Using imperfect advance demand information in ordering and rationing decisions. Int. J. Prod. Econ. 121 (2), 665–677.

Thonemann, U., 2002. Improving supply-chain performance by sharing advance demand information. European J. Oper. Res. 142 (1), 81–107.

Topan, E., Tan, T., van Houtum, G.J., Dekker, R., 2018. Using imperfect advance demand information in lost-sales inventory systems with the option of returning inventory. IISE Trans. 50 (3), 246–264.

Tortorella, G.L., Fogliatto, F.S., Cauchick-Miguel, P.A., Kurnia, S., Jurburg, D., 2021. Integration of industry 4.0 technologies into total productive maintenance practices. Int. J. Prod. Econ. 240, 108224.

Tortorella, G.L., Saurin, T.A., Hines, P., Antony, J., Samson, D., 2023. Myths and facts of industry 4.0. Int. J. Prod. Econ. 255, 108660.

van Donselaar, K., Rock Kopczak, L., Wouters, M., 2001. The use of advance demand information in a project-based supply chain. European J. Oper. Res. 130 (3), 519–538.

Yan, T., Lei, Y., Li, N., Si, X., Pintelon, L., Dewil, R., 2022. Online joint replacement-order optimization driven by a nonlinear ensemble remaining useful life prediction method. Mech. Syst. Signal Process. 173, 109053.

Zhang, F., Shen, J., Liao, H., Ma, Y., 2021. Optimal preventive maintenance policy for a system subject to two-phase imperfect inspections. Reliab. Eng. Syst. Saf. 205, 107254.

Zhu, S., van Jaarsveld, W., Dekker, R., 2020. Spare parts inventory control based on maintenance planning. Reliab. Eng. Syst. Saf. 193, 106600.