



The added value of advance demand information for the planning of a repair kit

Christoph Rippe¹ · Gudrun P. Kiesmüller²

Accepted: 26 June 2022 / Published online: 25 July 2022
© The Author(s) 2022

Abstract

In this paper the repair kit problem is studied, where technicians have to visit several customers to repair broken appliances (such as copiers or heating systems) and they can only take a limited set of parts with them (called the repair kit). In this problem, it has to be decided which spare parts to include in the repair kit. We consider a version of this problem in which partial advance demand information is available. That means we divide the set of parts into two subsets, where the condition of parts in one subset is monitored by sensors. In case an appliance fails and a repair job is requested by a customer the service provider is able to access this sensor data before a technician visits the customer. For this setting, we derive an expression for the job fill rate, which is used as a constraint in the optimization model, where holding and replenishment costs are minimized. We use a greedy heuristic to determine near-optimal repair kits. In a numerical study, we find that integrating advance demand information yields substantial cost savings. In order to find out for which parts having advance demand information is most valuable, we examine the effect of parts' demand probabilities and their prices. We find that monitoring parts that are expensive and likely to fail leads to the largest cost savings. In particular, the price of the monitored parts and the achievable cost savings are strongly correlated.

Keywords Inventory · Repair kit · Spare parts · Advance demand information

Mathematics Subject Classification 90B05

✉ Christoph Rippe
christoph.rippe@ovgu.de

¹ Faculty of Economics and Management, Chair of Operations Management, Otto-von-Guericke-University Magdeburg, P.O. Box 4120, D-39108 Magdeburg, Germany

² TUM School of Management, Technical University of Munich, TUM Campus Heilbronn, Bildungscampus 9, 74076 Heilbronn, Germany

1 Introduction

In some industries it is quite common that companies do not only sell their appliances but also offer after-sales service to support their customers in case of problems with the bought appliance (Cohen et al. 1997). For small and light appliances the customers usually bring or send the malfunctioning appliance to a retailer or to a specific repair shop, where the repair can be performed. However, field service is provided in case the appliance is too large or heavy to be moved easily, as for example copy machines, heating systems, production equipment, and medical systems or white goods like washing machines or refrigerators. This means customers can call a service hotline and request a visit by a service engineer, who should repair the malfunctioning appliance.

In order to offer this service, technicians travel in commercial vans to customer sites where they perform an on-site diagnosis and repair the failed appliance if the necessary spare parts are available in the technician's van. Otherwise, the repair cannot be completed and a technician has to revisit the customer. This is called a broken job and leads to an unsatisfied customer and extra costs for the service provider because a second visit to the customer is necessary. A service technician usually visits three to six customers in one tour (i.e., before the repair van is restocked).

The problem of determining which spare parts to include in the van (or repair kit) and how many units of each part, is called the repair kit problem. We consider a version of this problem where service is measured by the job fill rate, defined as the fraction of jobs that can be completed in the next tour after customer notification. The objective of our repair kit problem is to minimize the expected holding costs for spare parts and handling costs for replenishment orders while satisfying a service level target.

The repair kit problem is studied in the literature under different assumptions. However, most versions assume that the demand for spare parts is not known until the technician visits the customer. In case certain parts within the appliance are equipped with sensors that monitor the condition of these parts, that is no longer true. Whenever an appliance breaks down this sensor data can be accessed by the service provider via remote diagnostics.

In this paper, we extend the standard multi-job repair kit problem by integrating advance demand information (or short ADI) that can be obtained from part-specific sensors. We formulate a model that considers perfectly reliable ADI for some parts and no ADI at all for other parts and show how to compute the job fill rate under these conditions. We compare scenarios in which we include or exclude advance demand information and determine for which spare parts having advance demand information is especially valuable. In particular, we are looking at the impact of a part's demand probability and its price on the costs savings that can be achieved.

The remainder of the paper is organized as follows. In the next section, we discuss the relevant literature and in Sect. 3 a detailed problem and model description is presented. In Sect. 4 we derive a formula for the job-fill rate and provide a heuristic to determine the repair kit. The results of a detailed numerical study are presented in Sect. 5 before we highlight our key findings and give directions for future research in Sect. 6.

2 Literature

There are two streams of literature that are relevant to our research. We first summarize the studies related to the repair kit problem. Then we discuss the contributions on advance demand information for inventory planning. In particular, we confine ourselves to perfect demand information.

Smith et al. (1980), Graves (1982), and Hausman (1982) are the first authors to study the repair kit problem, where a tour consists of only a single job, part demands are independent and at most one unit of a part might be needed for the repair. This problem is extended to dependent demands and multiple units of a part in one job (Mamer and Smith 1982; March and Scudder 1984), spare machines (Mamer and Smith 1985) and budget constraints (Mamer and Shogan 1987). Brumelle and Granot (1993) present a unifying approach for different formulations of the single job repair kit problem.

The multiple job repair kit problem, where a tour can consist of several customers to visit before the repair kit is restocked, is first studied by Heeremans and Gelders (1995). They present a formulation of this more generic form of the repair kit problem and propose a heuristic to determine a solution for this multiple job problem. The authors use the probability that none of the jobs on a tour is broken as a service criterion. Teunter (2006) presents a more general model formulation for the multiple job problem where he assumes that all required parts that are available in the repair kit are left with the customers, regardless of whether the job can be completed (i.e., regardless of whether all required parts are available in the quantities needed.) A cost model as well as a service model with a job fill rate constraint are formulated and a greedy heuristic is developed to solve the problem. Bijvank et al. (2010) discuss the same problem but show how the job fill rate can be computed when spare parts are only taken from the repair kit when all parts necessary to perform the repair are available in the right quantity.

In two recent contributions on the multiple job repair kit problem Saccani et al. (2016) and Prak et al. (2017) introduce replenishment costs to the cost model. Saccani et al. (2016) assume fixed replenishment costs per delivery to the service technician and determine optimal frequencies for these deliveries. Prak et al. (2017) consider non-zero replenishment lead times and part-specific material handling costs. They minimize the sum of these material handling costs and holding costs by determining (s,S) -policies for all parts. All the papers on the repair kit problem discussed so far assume that the service technician is completely unaware of any customer's need for spare parts before the first repair attempt.

Hariharan and Zipkin (1995) consider inventory management in a scenario where customers place advance orders for goods they wish to receive at some distinct point in the future called due date. As opposed to the supply lead time which is defined as the time required to fill a retailer's replenishment order, Hariharan and Zipkin (1995) introduce the demand lead time as the interval between the time an advance customer order is placed with the retailer and the corresponding due date. For a single product continuous review inventory model, where advance orders are arriving according to a Poisson process, they show that the retailer's inventory can be managed using a base stock policy (or an (s, S) -policy if there are fixed order costs) with replenishments

being triggered by the advance orders. The authors find that demand lead time and supply lead time have opposing effects on the system performance. Thus increasing the demand lead time results in a reduced need for safety stock.

A more general model for advance demand information is introduced for a periodic review inventory system by Gallego and Özer (2001). They allow that in period t customers can place orders to be delivered either immediately or in one of N future periods $t + 1$ to $t + N$. For such a situation they show that state-dependent policies with a modified inventory position are optimal. Similar to Hariharan and Zipkin (1995) their study indicates that longer demand lead times lead to improved system performance.

A serial system is investigated by Gallego and Özer (2003) and a divergent system by Özer (2003). While Özer (2003) considers a periodic review inventory system Marklund (2006) studies a divergent system under continuous review. Other extensions include limited production capacities (Özer and Wei 2004), flexible deliveries (Wang and Toktay 2008) and different demand classes (Koçağa and Şen 2007).

The repair kit problem can be considered as a multi-product problem with each set of spare parts required by a customer defining a product. Multi-product problems that integrate ADI have been studied by Lu et al. (2003), Bernstein and DeCroix (2015) and Angelus and Özer (2016). Lu et al. (2003) consider an assembly system under continuous review where advance orders for assembled products are triggering the replenishment of component inventories that are managed using basestock policies. A similar system but with periodic review is studied by Angelus and Özer (2016), who also combine advance demand information with order expediting. In Bernstein and DeCroix (2015) advance demand information is available in aggregated form, either as the total volume of demand across several products or as the demand mix between these products.

Based on our literature review we can conclude, that most research on the repair kit problem disregards advance demand information. Since it is shown (see Hariharan and Zipkin 1995) that advance demand information reduces safety stocks, resulting in lower inventory costs, we also conjecture a positive impact on the safety stocks in case of the repair kit problem. However, since cost savings are usually larger for longer demand lead times, it is unclear how strong the influence is in our setting, where lead times are short. Therefore, we quantify in this contribution the added value of advance demand information for the planning of a repair kit.

Since Angelus and Özer (2016) only consider one final product to be assembled, we have a more complicated situation because repair jobs can be different for different customers. Similar to the repair kit problem, Lu et al. (2003) also consider several final end products. However, they assume that ADI is available for **all** components. If ADI is revealed by sensor technology, it is usually not beneficially from an economic point of view to equip all components with sensors. Therefore, the question arises which parts should be monitored. In this contribution, we will also focus on this question and investigate which factors are relevant in this context.

The only attempt to date that uses ADI for the repair kit problem stems from an initial project (Rippe and Kiesmüller 2022) in which we considered unreliable and non-part-specific ADI when planning the repair kit. In this initial project we focussed on how to cope with the information obtained from a given imprecise and unreliable sensor system that is only able to communicate very few error codes. This contribution,

however, focusses on how to design the sensor system. That means we examine the way in which installing reliable part-specific sensors for different parts impacts the costs incurred by the service technician. For this scenario where some parts are monitored by individual sensors we are also able to derive a closed form expression for the job fill rate.

3 Problem and model description

The model formulated here is an extension of the multiple-job repair kit problem with a job fill rate service constraint that has been previously studied by Teunter (2006), Bijvank et al. (2010), Saccani et al. (2016), and Prak et al. (2017). We first describe the basic problem following the assumptions made in the aforementioned contributions before we discuss the modification with advance demand information. We use the notation summarized in Table 1.

A service technician can hold a number of N different parts in his or her repair kit. For each of these parts, the service technician can place a replenishment order before the start of a repair tour. These orders are delivered instantaneously. Next, the service technician goes on a repair tour and replaces malfunctioning parts in the customers' appliances with spares from the repair kit. After the end of the tour holding and replenishment costs are charged.

The immediate replenishment of spare parts before the start of a repair tour corresponds to an overnight delivery from a central warehouse with ample supply. These overnight deliveries are in line with our experience with home appliance manufacturers who tend to restock their fleet of service technicians using an express parcel service. Similar problems with zero lead time have been studied by Teunter (2006) and Bijvank et al. (2010) for example. In order to determine the order quantities for the different spare parts, we apply part-specific (s_i, S_i) -policies, as suggested by Prak et al. (2017). For single-item problems with fixed order costs (s, S) -policies are optimal. They are also easy to handle and therefore attractive for practitioners.

The demand for spare parts observed by the service technician depends on the number of customers visited in one tour. We model the number of customers served in a repair tour as a random variable denoted by J , where the maximum number of customers that can be visited in any tour is denoted by M . Thus, J can take any value in $\{1, \dots, M\}$. We consider the distribution of the number of customers in one tour to be independent and identically distributed (i.i.d.) across all tours.

Each customer visited by the service technician has a malfunctioning appliance due to a failure of one or multiple components. The repair is completed when all broken components are replaced with spare parts that were brought by the service technician in the repair kit. When a repair job cannot be completed, we assume the service technician still leaves all available parts required by the customer on-site for a second visit. This assumption is in line with Teunter (2006) and Prak et al. (2017). We denote the aggregated demand for part i in the l_1 -th to the l_2 -th customer visit in the same tour (given that the tour contains at least l_2 customer visits) by $D_i(l_1, l_2)$. Following this notation, the demand for part i during the l -th customer visit in a tour is described by the random variable $D_i(l, l)$ (again given there are at least l customers).

Table 1 Notations

Notation	Description
$D_i(l_1, l_2)$	Aggregated demand of the l_1 -th to the l_2 -th customer for part i (given there are at least j_2 customers)
D_i	Total demand for part i in one tour
d_i^{max}	Maximum number of units of part i required by one customer
f_i	Material handling costs incurred per order of part i
$F_i(s_i, S_i)$	Average material handling costs per tour for part i given the policy parameters s_i and S_i
h_i	Holding costs incurred per unit of part i
$H_i(s_i, S_i)$	Average holding costs per tour for part i given the policy parameters s_i and S_i
$C(s, S)$	Expected total costs per tour given $s = (s_1, \dots, s_N)$ and $S = (S_1, \dots, S_N)$
I_i^-	Net inventory level of part i before order placement
I_i^+	Net inventory level of part i after order arrival before the start of a tour
M	Maximum number of customers served in one tour
N	Number of different spare parts
N_1	Number of different NPs
Q_i	Order for part i placed before the start of a tour
J	Number of repair jobs in a tour
J_c	Number of completed repair jobs in a tour
J_l	Binary variable indicating if the number of jobs in the tour is at least l and the l -th job can be completed
p_l	Probability that the l -th job can be completed.
p_l^{AP}	Probability that sufficient units of all APs are available for the l -th job
p_l^{NP}	Probability that sufficient units of all NPs are available for the l -th job
$\pi_i(k)$	Probability that the net inventory level of part i is k (before order arrival for APs, after order arrival for NPs)
s_i	Reorder level of part i
S_i	Order-up-to level of part i
$\gamma(s, S)$	Job fill rate (JFR) given $s = (s_1, \dots, s_N)$ and $S = (S_1, \dots, S_N)$
γ^*	Target job fill rate
$R(a)$	Ratio trading off cost increase and job fill rate increase if $a \in \{s_i, S_i i = 1, \dots, N\}$ is raised by one unit

We assume that the demand for part i is i.i.d. for all customers. Though customers' demand distributions may be heterogeneous in real life, this simplification is common in the repair kit literature (see Teunter 2006; Bijvank et al. 2010; Prak et al. 2017). The distributions of $D_i(l, l)$ are taken to be discrete with the support ranging from 0 to d_i^{max} . Finally, we define D_i as the total demand for part i in one tour. The distribution of D_i can be computed for all $i = 1, 2, \dots, N$ as:

$$P(D_i = d) = \sum_{j=1}^M P(D_i(1, j) = d) \cdot P(J = j); \quad d = 0, \dots, d_i^{max} \cdot M. \quad (1)$$

Costs are charged after finishing the tour. In accordance with Prak et al. (2017), we consider two different part-specific types of costs: holding costs h_i are incurred for each unit of part i , $i = 1, \dots, N$, still in the repair kit after visiting all customers of the tour and material handling costs f_i for each part type i , $i = 1, \dots, N$, ordered to replenish the repair kit for the repair tour regardless of the order quantity. The material handling costs are charged for the order picking process in a central warehouse from which the service technician is replenished. There are no fixed replenishment costs for the actual shipment process to get the parts from the warehouse to the repair kit since it is reasonable to assume that at least one part type needs to be replenished before a tour. (i.e. there is always a shipment).

Contrary to previous contributions related to the repair kit problem we assume that perfectly reliable part-specific advance demand information is available for some of the spare parts carried in the repair kit. Thus, we divide the set of parts into non-advance demand information parts (NP) and advance demand information parts (AP). Without loss of generality, we arrange the parts such that the first N_1 ($N_1 \leq N$) parts are NPs and that for the remaining parts ($N_1 + 1$ to N) advance demand information is available. In this sense, we speak of partial advance demand information.

Focusing first on NPs we need to introduce a tour index t to describe the inventory development of a particular part i , $i = 1, \dots, N_1$ over consecutive tours. Let $I_{i,t}^-$ and $I_{i,t}^+$ define this part's net inventory level before order placement and right after order arrival before the start of tour t . With s_i and S_i being the reorder- and the order-up-to-level for part i the amount ordered $Q_{i,t}$ for this part before the start of tour t is given by

$$Q_{i,t} = \begin{cases} S_i - I_{i,t}^- & \text{if } I_{i,t}^- \leq s_i \\ 0 & \text{if } I_{i,t}^- > s_i. \end{cases} \quad (2)$$

Let $D_{i,t}$ denote the demand for part i , $i = 1, \dots, N_1$ in tour t . This leads to the following recursive relation

$$I_{i,t+1}^+ = \begin{cases} S_i & \text{if } I_{i,t}^+ - D_{i,t} \leq s_i \\ I_{i,t}^+ - D_{i,t} & \text{if } I_{i,t}^+ - D_{i,t} > s_i. \end{cases} \quad (3)$$

Turning to APs we assume that for each AP there is a sensor that is monitoring the condition of this very part. This sensor communicates a binary signal (either up

or down), which we assume to be 100% reliable. In case a customer's appliance experiences a failure the service provider will be able to connect to the appliance via remote diagnostics. This means that given the customer's permission the service provider can read the appliance's error log from afar. In this error log the status of each sensor is listed, giving the service provider valuable insights into what has caused the appliance's failure. Therefore, if a particular sensor detects the failure of a part type, this is essentially advance demand information from the service technician's point of view. We make the assumption that this advance demand information is available to the service technician before he or she is placing the initial replenishment order prior to the start of the repair tour. Thus the order quantity $Q_{i,t}$ of a particular AP i , $i = N_1 + 1, \dots, N$ before the start of a tour t depends not only on the net inventory level $I_{i,t}^-$ but also on the demand $D_{i,t}$ for the upcoming repair tour that the technician is already aware of. With s_i and S_i again being the reorder- and the order-up-to-level we obtain

$$Q_{i,t} = \begin{cases} S_i - (I_{i,t}^- - D_{i,t}) & \text{if } I_{i,t}^- - D_{i,t} \leq s_i \\ 0 & \text{if } I_{i,t}^- - D_{i,t} > s_i. \end{cases} \quad (4)$$

For the net inventory level before order placement this gives us the following recursive equation

$$I_{i,t+1}^- = \begin{cases} S_i & \text{if } I_{i,t}^- - D_{i,t} \leq s_i \\ I_{i,t}^- - D_{i,t} & \text{if } I_{i,t}^- - D_{i,t} > s_i. \end{cases} \quad (5)$$

Please note, that in contrast to the NPs the order quantity already includes the demand of the next tour, which also results in earlier order triggering. Further, for an NP the parameter S_i is a target for the net inventory level before the start of a tour, whereas for an AP it is a target for the net inventory level after all customers have been visited.

Even though the demand for an AP in the next tour is known beforehand and our replenishment lead time is zero, that does not necessarily mean that the corresponding repair job can always be finished successfully. First, it may happen, that additional NP spare parts are needed which are not available. Second, it may be beneficial to allow planned backorders to reduce holding costs if the fixed replenishment costs are high. This would require large ordering batches and reorder levels to be below -1 . However, this will only happen, if the service requirement is not too high.

In order to perform a single repair job, a combination of several NPs and APs might be required. We assume that the failures of different parts are independent of each other. Otherwise, we could simply model them as one part, monitoring them if possible with just one common sensor.

Our objective is to minimize the average sum of holding and material handling costs subject to a service level constraint. To this end our decision variables are the reorder levels $s = (s_1, \dots, s_N)$ and the order-up-to levels $S = (S_1, \dots, S_N)$.

We aim to minimize the expected costs per tour across all parts i , $i = 1, \dots, N$. Let us define $H_i(s_i, S_i)$ and $F_i(s_i, S_i)$ as the holding and the replenishment cost we can expect to incur for part i , $i = 1, \dots, N$ per tour when the policy parameters (s_i, S_i)

are implemented. We obtain

$$H_i(s_i, S_i) = h_i \cdot \mathbb{E}[\max(I_i^+ - D_i, 0)] \quad \text{and} \quad F_i(s_i, S_i) = f_i \cdot P(Q_i > 0). \quad (6)$$

Then the expected total costs per tour can be calculated as

$$C(s, S) = \sum_{i=1}^N H_i(s_i, S_i) + F_i(s_i, S_i) \quad (7)$$

The service level we are interested in is the job fill rate $\gamma(s, S)$ (with $s = (s_1, \dots, s_N)$ and $S = (S_1, \dots, S_N)$), which is defined as the fraction of jobs that can be completed immediately with the spare parts carried in the repair kit. The job fill rate has previously been used by Teunter (2006), Bijvank et al. (2010), and Prak et al. (2017) for similar multi-job problems. It reflects the average customer satisfaction across all customers in a tour, which makes it a good performance indicator from the service provider’s point of view. We assume, that the service technician only visits customers when he or she knows that all APs are available. Otherwise, in case of planned backorders, when it is already clear, that a job cannot be completed successfully, the customer is not visited in the next tour and the repair is planned for a later tour. Such a job still counts as a broken job for the job fill rate, because it can not be completed within the next day. With γ^* being the target service level, we can formulate the optimization problem as follows

$$\begin{aligned} & \min_{s, S} C(s, S) \\ \text{s.t. } & \gamma(s, S) \geq \gamma^* \\ & s \in \mathbb{Z}^N \\ & S \in \mathbb{N}_0^N \end{aligned} \quad (8)$$

4 Analysis

In this section, we first derive a closed-form expression for the job fill rate and the expected holding and material handling costs for a given set of (s,S)-policies for all parts considered. Then we outline a greedy algorithm to determine near-optimal parameters s_i and S_i , ($i \in \{1, \dots, N\}$) for the optimization problem.

4.1 Job fill rate

Let us first define J_c as the number of completed jobs out of the total number of J jobs in a repair tour. In order to derive the job fill rate we can describe J_c as the sum of a number of binary variables J_1, \dots, J_M where $J_l = 1$, ($l \in \{1, \dots, M\}$) if at least l

customers are visited in one tour and the l -th job can be completed. With p_l being the probability that the l -th job can be completed (given $J \geq l$) we obtain

$$\gamma(s, S) = \frac{\mathbb{E}[J_c]}{\mathbb{E}[J]} = \frac{\sum_{l=1}^M \mathbb{E}[J_l]}{\mathbb{E}[J]} = \frac{\sum_{l=1}^M P(J \geq l) \cdot p_l}{\sum_{l=1}^M P(J \geq l)}. \tag{9}$$

The probability p_l , ($l \in \{1, \dots, M\}$) can be split up into the probabilities p_l^{NP} and p_l^{AP} that sufficient units of all NPs and all APs are available to complete the l -th job (given $J \geq l$) with

$$p_l = p_l^{NP} \cdot p_l^{AP}. \tag{10}$$

The probability p_l^{NP} has already been derived by Prak et al. (2017). In a first step, they determine the steady-state distribution of the net inventory level after order arrival before the start of a tour for all parts (see Appendix I). We denote these steady-state probabilities by $\pi_i(k)$ ($i = 1, \dots, N_1$, $k = s_i + 1, \dots, S_i$). Then p_l^{NP} can be determined as follows

$$p_l^{NP} = \begin{cases} \prod_{i=1}^{N_1} \left[\sum_{k=0}^{S_i} \pi_i(k) \cdot P(D_i(l, l) \leq k) \right] & \text{if } l = 1 \\ \prod_{i=1}^{N_1} \left[P(D_i(l, l) = 0) + \sum_{m=1}^{S_i} \sum_{n=0}^{S_i-m} \sum_{k=m+n}^{S_i} P(D_i(l, l) = m) \cdot P(D_i(1, l-1) = n) \cdot \pi_i(k) \right] & \text{if } l \geq 2 \end{cases} \tag{11}$$

This formula reflects that the l -th job can only be completed if sufficient units of all NPs are available for the l -th customer. This is the case if a part is either not required at all by the customer or if the stock at the beginning of the tour minus the quantity requested by the preceding customers still exceeds the demand of customer l .

In order to determine the probability p_l^{AP} we need to consider the net inventory level before order placement and the demand during the course of a tour for all different APs. Be aware that in this case, the ordering decision before the start of a tour depends on the demand of the customers to be served as it is already known due to advance demand information. Comparing Eqs. (3) and (5) we realize that the recursive equation for an NP's net inventory level after order arrival is equivalent to the recursive equation for an AP's inventory level before ordering. Thus, the steady-state distribution of an AP's inventory level **before order placement** corresponds to the steady-state distribution of an NP's net inventory level **after order arrival** before the start of the tour. Because of this property we define $\pi_i(k)$, ($i = N_1 + 1, \dots, N$, $k = s_i + 1, \dots, S_i$) as the steady-state probabilities for the APs' inventory level before order placement. This means that for NPs and APs $\pi_i(k)$ defines the steady-state probabilities of the net inventory level

at different time points, but it is calculated in the exact same way using the formula derived by Prak et al. (2017). (See Appendix I)

For the derivation of the probability that sufficient units of an AP are available for the l -th customer, we have to realize that this probability depends on the total number of customers in the tour $J \geq l$. This means that the l -th customer's chance of being served does not only depend on the preceding customers 1 to $l - 1$ but also on the succeeding customers $l + 1$ to J , because the replenishment decision before the start of the tour has been made taking into account **all** of these customers' advance demand information (4). Therefore, the probability that the l -th job can be completed depends on the tour length $J \geq l$ in case of advance demand information. We can express this probability as follows.

$$\begin{aligned}
 p_l^{AP} &= \sum_{j=l}^M P(J = j | J \geq l) \cdot \prod_{i=N_l+1}^N P(D_i(l, l) \leq \max(I_i^+ - D_i(1, l - 1), 0) | J = j) \\
 &= \sum_{j=l}^M \frac{P(J = j)}{P(J \geq l)} \cdot \prod_{i=N_l+1}^N P(D_i(l, l) \leq \max(I_i^+ - D_i(1, l - 1), 0) | J = j). \tag{12}
 \end{aligned}$$

The term $P(D_i(l, l) \leq \max(I_i^+ - D_i(1, l - 1), 0) | J = j)$ that describes the probability that sufficient units of AP i are available for the l -th out of j customers can now be split into two parts again:

$$\begin{aligned}
 &P(D_i(l, l) \leq \max(I_i^+ - D_i(1, l - 1), 0) | J = j) \\
 &= P(I_i^+ - D_i(1, l) < 0, D_i(l, l) = 0 | J = j) \\
 &\quad + P(I_i^+ - D_i(1, l) \geq 0 | J = j). \tag{13}
 \end{aligned}$$

The first summand describes the probability that the inventory level of part i is negative after visiting customer number l , but not due to the demand of this customer which was zero. Note that the net inventory level I_i^+ depends on the total number of customers j and their demand. Thus, the two events in the first summand are not independent. The second summand is the probability that the net inventory level of part i , after the l -th customer has been visited, is non-negative. To determine both probabilities, summand one and two, we express the inventory level I_i^+ after orders as the inventory level before orders I_i^- plus the ordered quantity Q_i . For APs the ordered quantity depends on the inventory level before orders I_i^- and the total demand $D_i(1, j)$. The total demand in a tour of j customers may take any value between 0 and $j \cdot d_i^{max}$. By conditioning on the values of I_i^- and $D_i(1, j)$ we can determine the exact ordered quantity for all relevant demand-inventory scenarios. That way we reformulate both probabilities $P(I_i^+ - D_i(1, l) \geq 0 | J = j)$ and $P(I_i^+ - D_i(1, l) < 0, D_i(l, l) = 0 | J = j)$ using only the steady-state probabilities $\pi_i(k)$, ($k = s_i + 1, \dots, S_i$) and the probabilities $P(D_i(l_1, l_2) = m)$, ($1 \leq l_1 \leq l_2 \leq M, m = 1, \dots, M \cdot d_i^{max}$) that describe the customer demand. We obtain the results given in the following Eqs. (14) and (15). A detailed derivation

is provided in Appendix II. To cover the cases in which we are either considering the first or the last customer in a tour let us define $D_i(l_1, l_2) := 0$ for $l_1 > l_2$.

$$\begin{aligned}
 &P(I_i^+ - D_i(1, l) \geq 0 | J = j) \\
 &= \sum_{m=1}^{j \cdot d_i^{\max}} \sum_{k=s_i+1}^{\min(s_i+m, S_i)} [P(D_i(1, j) = m) \cdot \pi_i(k)] \\
 &+ \sum_{m=0}^{\min(j \cdot d_i^{\max}, S_i - s_i - 1)} \sum_{k=\max(s_i+m+1, 0)}^{S_i} \sum_{n=0}^{\min(k, m)} [P(D_i(1, l) = n) \cdot P(D_i(l+1, j) = m-n) \cdot \pi_i(k)]
 \end{aligned} \tag{14}$$

$$\begin{aligned}
 &P(I_i^+ - D_i(1, l) < 0, D_i(l, l) = 0 | J = j) \\
 &= \begin{cases} P(D_i(1, 1) = 0) \cdot \left[1 - \sum_{k=\max(0, s_i+1)}^{S_i} \pi_i(k) \right] & \text{if } l = j = 1 \\
 P(D_i(l, l) = 0) \cdot \left[1 - \left(\sum_{m=1}^{(j-1) \cdot d_i^{\max}} \sum_{k=s_i+1}^{\min(s_i+m, S_i)} [\pi_i(k) \cdot P(D_i(1, l-1) + D_i(l+1, j) = m)] \right. \right. \\
 \left. \left. + \sum_{m=0}^{\min((j-1) \cdot d_i^{\max}, S_i - s_i - 1)} \sum_{k=\max(s_i+m+1, 0)}^{S_i} \sum_{n=0}^{\min(k, m)} [P(D_i(l+1, j) = m-n) \right. \right. \\
 \left. \left. \cdot P(D_i(1, l-1) = n) \cdot \pi_i(k)] \right) \right] & \text{else} \end{cases}
 \end{aligned} \tag{15}$$

4.2 Cost function

For all parts $i \in (1, \dots, N)$ we derive part-specific holding costs $H_i(s_i, S_i)$ and material handling costs $F_i(s_i, S_i)$. We consider the expected costs per tour. Both costs are charged at the end of a tour. However, the replenishment costs are determined based on the replenishment that occurs before the start of a repair tour.

4.2.1 Cost function for NPs

In order to determine the holding cost for an NP i , $i = 1, \dots, N_1$, we need to determine the distribution of the net inventory level at the end of the repair tour. We obtain this distribution as the convolution of the steady-state distribution of the net inventory level after order arrival before the start of the tour and the distribution of the demand in an arbitrary tour. This gives us the following result

$$\begin{aligned}
 H_i(s_i, S_i) &= h_i \cdot \mathbb{E} [\max (I_i^+ - D_i, 0)] \\
 &= \sum_{k=s_i+1}^{S_i} \sum_{m=1}^k (k - m) \cdot \pi_i(k) \cdot P (D_i = m)
 \end{aligned} \tag{16}$$

In deriving the material handling cost incurred for a part i , $i = 1, \dots, N_1$, we need to quantify the probability that an order is placed before the start of a tour. In case we consider an NP the ordering decision is entirely based on the net inventory level I_i^- . Unfortunately, at this point we only have the steady-state distribution of the net inventory level after order placement. However, the inventory level before ordering prior to the start of a tour t corresponds to the net inventory level after demand is fulfilled at end of the previous tour $t - 1$. Because of this, we obtain

$$\begin{aligned}
 F_i(s_i, S_i) &= f_i [1 - P(Q_i = 0)] \\
 &= f_i [1 - P(I_i^- > s_i)] \\
 &= f_i [1 - P(I_i^+ - D_i > s_i)] \\
 &= f_i \left[1 - \sum_{k=s_i+1}^{S_i} \sum_{m=0}^{k-(s_i+1)} \pi_i(k) P(D_i = m) \right] \tag{17}
 \end{aligned}$$

4.2.2 Cost function for APs

For the computation of the holding cost incurred for an AP i , $i = N_1 + 1, \dots, N$ again we need to consider the inventory level at the end of the tour after demand fulfillment. Yet this is equivalent to the net inventory level at the very beginning of the next tour before order placement. The steady-state distribution of an AP's net inventory level before order placement, however, is known as it corresponds to the steady-state distribution of an NP's net inventory level after order arrival. This gives us

$$\begin{aligned}
 H_i(s_i, S_i) &= h_i \cdot \mathbb{E}[\max(I_i^+ - D_i, 0)] \\
 &= h_i \cdot \mathbb{E}[\max(I_i^-, 0)] \\
 &= h_i \cdot \sum_{k=s_i+1}^{S_i} k \cdot \pi_i(k) \tag{18}
 \end{aligned}$$

For the material handling costs we derive

$$\begin{aligned}
 F_i(s_i, S_i) &= f_i [1 - P(Q_i = 0)] \\
 &= f_i [1 - P(I_i^- - D_i > s_i)] \\
 &= f_i \left[1 - \sum_{k=s_i+1}^{S_i} \sum_{m=0}^{k-(s_i+1)} \pi_i(k) P(D_i = m) \right] \tag{19}
 \end{aligned}$$

Note that the material handling costs are calculated in the same way for parts with and without advance demand information. Thus, integrating advance demand information for a particular part does not have an impact on the material handling costs at least not if we do not adapt the policy parameters s_i and S_i at the same time.

4.3 Heuristic solution procedure

Since the number of possible combinations of reorder levels $s = (s_1, \dots, s_N)$ and order-up-to-levels $S = (S_1, \dots, S_N)$ is increasing exponentially with the number of different spare parts to be considered, the computational effort would be too high to compute an optimal solution for real-life problems. Therefore, we use a greedy heuristic to determine solutions for optimization problem (8). For repair kit problems that disregard advanced demand information similar approaches have been suggested by Teunter (2006), Bijvank et al. (2010), and Prak et al. (2017).

Our greedy algorithm iteratively increases the different parts' reorder and order-up-to levels until the target job fill rate is achieved. As in Prak et al. (2017), we build upon the EOQ model to obtain an initial solution for the policy parameters. We derive

$$Q_i := S_i - s_i = \max \left(\left\lceil \sqrt{\frac{2f_i \bar{D}_i}{h_i}} \right\rceil, 1 \right), \quad i = 1, 2, \dots, N \quad (20)$$

where \bar{D}_i is the average demand for part i for one tour and set $s_i = -Q_i$ and $S_i = 0$ for all $i \in \{1, \dots, N\}$. Starting from this EOQ-based solution we increase either the reorder or the order-up-to level of a single part in each iteration. In either case, the job fill rate is non-decreasing. The total costs, however, may decrease upon an increase of a part's order-up-to-level. That is due to a decrease in material handling costs. For this reason, selecting the action with the least cost increase to job fill rate increase ratio would favor actions with low job fill rate increases over larger job fill rate increases given the same negative cost increase. To avoid this problem we introduce a new ratio R . Let $\Delta_a C(s, S)$ and $\Delta_a \gamma(s, S)$ denote the cost and job fill rate increase respectively given $a \in \{s_i, S_i | i = 1, \dots, N\}$ is increased by one unit. Then R is defined by

$$R(a) = \Delta_a C(s, S) \cdot (\Delta_a \gamma(s, S) + \epsilon)^{-\text{sgn}(\Delta_a C(s, S))}, \quad (21)$$

with $\epsilon = 10^{-40}$. In each iteration, we select the action a with the least $R(a)$ value. That way we will always choose larger over smaller job fill rate increases given the same (positive or negative) cost increase. We add the small positive constant ϵ to avoid numerical problems in case the job fill rate increase for an action is 0 or close to 0. Once this greedy procedure reaches a solution that satisfies the service constraint, we check if we can further reduce costs without falling below the target job fill rate again. To this end, we try to reduce the reorder levels of all parts in the reverse order in which they were last increased. We refer to this last step as the reduction step. A detailed description of the greedy algorithm is provided in Appendix III.

While we use the same initial solution as Prak et al. (2017) for their non-ADI problem, our algorithm differs from theirs in that we do not increase reorder and order-up-to levels simultaneously in each step. We tested two versions of our algorithm with either simultaneous or independent increments of reorder and order-up-to-levels on the problem instances presented in the following numerical experiment (5.1). Allowing for independent increments led to slightly better solutions that were up to 2.62% cheaper than those generated with the simultaneous-increase approach. We also considered

the impact of different starting solutions. However, we find that the algorithm is rather insensitive to the starting solution for as long as the initial reorder and order-up-to levels are low enough not to preclude potentially good solutions. With EOQ-based initial solutions that does not happen. The reduction step we perform is similar to a procedure suggested by Bijvank et al. (2010) for a repair kit problem without handling costs and ADI.

5 Numerical study

The aim of our numerical study is twofold. First, we quantify the added value of integrating advance demand information and investigate the causes of the cost savings. In a second study, we examine for which spare parts having advance demand information is especially advantageous. In particular, we are looking at the impact of a part's demand probability and its price on the benefits of integrating advance demand information for this part. For both studies, we assume that only one unit of each part can be required by a single customer ($d_i^{\max} = 1$, $i = 1, \dots, N$). In that sense the demand probability for a part i refers to the probability $P(D_i(l, l) = 1)$ that a single customer l requires a unit of this part, where this probability is identical for all customers $l = 1, \dots, J$ in a tour.

5.1 The added value of advance demand information

In this section, we quantify the added value of advance demand information by constructing repair kits with the heuristic solution procedure assuming that advance demand information is either available for some parts or not at all. The difference in the corresponding costs is then the added value of the advance demand information. We measure this added value for examples with 100 different spare parts, which corresponds to the largest problem sizes considered in the experiments conducted by Teunter (2006). However, using the heuristic solution procedure we proposed, we can solve instances with 1000 spare parts in less than 40 minutes (executed in R using a single core of a 1.6 GHz Intel@Core™ i5-8520U processor)

We consider 100 instances that differ with respect to the probability that a spare part is needed for a job and the price of the spare part. The demand probabilities as well as the prices of the spare parts are randomly selected from uniform distributions on $[0.001; 0.2]$ and $[1\text{€}; 500\text{€}]$ respectively. The annual holding cost rate is fixed to 20% of a spare part's price and the material handling costs related to each replenishment are fixed to 1€ for each part. The number of customers to be visited per tour is assumed to follow a discrete uniform distribution between one and six. For all numerical results we have chosen a target job fill rate of 90%, but results are similar for other target service levels.

We compare the situation where none of the parts has advance demand information with several scenarios in which a growing number of the parts can generate advance information. In all cases, the repair kit is determined with the heuristic solution approach and the corresponding costs are computed as in 4.2. Let us denote the




Fig. 1 The relative total cost reduction due to ADI given a target JFR of 90%

reorder and order-up-to-levels obtained when the first x % of the parts provide ADI by s^x and S^x . To determine the benefit of different levels of ADI availability we measure the relative cost reduction

$$\frac{C(s^0, S^0) - C(s^x, S^x)}{C(s^0, S^0)} \cdot 100\% \quad (22)$$

for $x \in \{10, 20, \dots, 90, 100\}$. The results for our 100 problem instances are depicted in Fig. 1.

It is not surprising that the total cost reduction is increasing with the number of parts equipped with technology to deliver advance demand information. With additional information available we can lower the average amount of stock that is required to satisfy the service constraint, which leads to reduced holding costs. These holding cost reductions are opposed by only a small increase in order costs. We find that the total cost reduction we observe is not just due to the parts equipped with sensors to provide ADI. In case less than 50% of the parts are equipped with ADI, these APs only account for 80% to 85% of the cost savings. The other 15% to 20% of the total cost savings can be attributed to NPs, where the safety stock can be reduced as well (see Fig. 2).

With advance information obtained for some parts we are able to increase these parts' availability while reducing their stock levels at the same time compared to the baseline scenario in which there is no advance demand information at all. Due to this increased availability for the APs, we can decrease the availability of some of the NPs and still reach the predefined target service level. Decreasing an NP's availability, however, is achieved by lowering its stocking levels which results in decreasing holding costs. That is why up to 20% of the cost savings can be attributed to the NPs.




Fig. 2 Share of total cost reduction based on APs

5.2 Identification of suitable parts for advance demand information

In a real-life application fitting sensors to all parts in an appliance is rather unrealistic as it may not be technically possible. Further, due to the costs for the installation of sensors, it may also not be feasible from an economic point of view to equip all parts with this technology. Therefore, we examine in the following, for which parts having advance demand information is particularly advantageous regarding the total cost reduction that can be achieved. For this reason, we are looking at the impact that parts' demand probabilities and their prices have on the benefit of having advance demand information for these parts. We would expect the value of having advance demand information to increase with increasing demand frequencies and increasing prices of the monitored parts.

In order to test the first conjecture concerning the impact of the demand probability, we conduct a second experiment. For this experiment, we consider 100 problem instances generated as described in Table 2. Each problem instance consists of 20 parts that differ in their demand probabilities but not in their prices which are identical for all parts. Out of the 20 parts, 5 parts are to be monitored by sensors that provide ADI. The objective is to find out which combinations of APs yield the best results and in how far these results depend on the APs' demand probabilities. In this setting, there are in total $\binom{20}{5} = 15504$ possible combinations of APs. Out of this set, we randomly select 500 different combinations for each problem instance. For each AP combination, we determine the repair kit with the greedy algorithm described in Sect. 4.3. We derive the corresponding total costs and compare them with the costs of a repair kit that is constructed with the same algorithm assuming all parts are NPs. For each AP combination, we determine the potential cost savings as defined in term (22).

To describe the impact of the parts' demand probabilities on the benefit of ADI availability we characterize all AP combinations by the mean demand probability of the respective APs. When we compare the cost savings that can be achieved with

Table 2 Parameter settings

	Conjecture 1 100 instances generated as follows:	Conjecture 2 100 instances generated as follows:
number of parts	20 parts	
available ADI	5 APs, 15 NPs, 500 randomly selected AP combinations	
holding cost rate	0.05/365 (per day)	
number of customers per tour	discrete uniformly distributed between 1 and 6	
target job fill rate	continuous uniformly distributed between 90% and 99%	
material handling cost per order	0.2	
demand probabilities	continuous uniformly distributed between 5% and 25%	15%
price	50	continuous uniformly distributed between 1 and 100




Fig. 3 Impact of the APs' mean demand probability

various AP combinations we find that using sensors for parts with higher demand probabilities pays off in general. However, there is no monotone relationship between the APs' mean demand probability and the costs savings that can be obtained. Out of the 100 instances, we considered the AP combination with the largest mean demand probability led to the largest cost savings in only 19 cases. Yet for 90 instances it is among the 10% most beneficial AP combinations. We find that the correlation coefficient between the APs' mean demand probability and the cost savings ranges from -0.40 to 0.88 with a mean of 0.69 across all instances. For further analysis let us consider the instance with the correlation coefficient closest to the average value of 0.69 . Figure 3 illustrates the relation between the APs' mean demand probabilities and the corresponding cost savings for all AP combinations of this example instance.

We find that the increasing cost savings we observe for increasing mean demand probabilities of the APs are primarily due to decreased holding costs in the case where ADI is available. With the APs' mean demand probability increasing the total amount of stock necessary to satisfy the target job fill rate can generally be decreased. That is why we see a trend for lower holding costs with larger AP mean demand probabilities.




Fig. 4 Impact of the APs' mean price

However, stock levels are managed by the choice of integer reorder and order-up-to levels for all parts. Because of this integrality, exchanging one part from the set of APs with one from the set of NPs that has a slightly higher demand probability may not enable us to decrease the reorder or the order-up-to level of any of the parts. Where this happens this exchange can even cause additional holding costs. That is because with unchanged policy parameters an increased demand probability for a part leads to an increased replenishment frequency for that part. Whenever an AP is replenished its inventory level at the end of the next tour is equal to its order-up-to level (see (5)). Thus, the more often we order an AP i the more often we have S_i units on hand at the end of a tour. This leads to increased stock levels and thus increased holding costs for this AP. This effect is one reason why the cost savings in Fig. 3 are not monotonically increasing in the AP's mean demand probabilities. Another reason is that holding cost reductions are sometimes opposed by increased material handling costs. This occurs when the stock reduction is achieved by lowered order-up-to levels against the same reorder levels.

In the second part of this numerical study, we concentrate on the second conjecture that expensive parts should be equipped with a sensor because this leads to the largest cost savings. To gain insights we conduct a third experiment that again comprises 100 problem instances. These instances are constructed as described in Table 2. Again we consider 20 parts for each instance. This time though the demand probabilities for all parts are identical while they differ in their prices. As for the second experiment, we examine 500 randomly selected AP combinations with 5 APs and 15 NPs for each instance. For each combination, we employ our greedy heuristic to determine the repair kit. Additionally, we construct a repair kit with the same algorithm for an NP-only scenario. Comparing the costs incurred with the AP combinations and in the NP-only scenario, we can calculate the cost savings as described in term (22). We contrast these cost savings achieved with the various AP combinations with the average prices of the respective APs. As can be seen from the results of the exemplary instance depicted in Fig. 4 there is no monotone relation between the average price of the APs and the cost savings that can be achieved. Yet, there is a very strong positive correlation between the average prices of the APs and the cost savings. For all 100 instances, we tested the

correlation coefficient was larger than 0.97. The most valuable combination of APs also led to the highest cost savings in 93 out of the 100 problem instances and for the other 7 instances it was amongst the 5 best combinations. Based on this result, we can assume that given all parts are equally likely to fail, it would be most beneficial to monitor the five most expensive parts as they would form the AP combination with the highest mean price.

6 Summary and outlook

Technical innovations can help to improve inventory planning if more information about future demand can be generated. In the context of spare parts management, this advance demand information can be generated by sensors, which help a service technician decide about the required spare parts for a repair. In this paper, we have extended the repair kit problem to a situation where some parts are equipped with sensors such that perfect advance demand information is obtained. The presented model can be used to quantify the added value of this technology for the repair kit planning problem. We have shown, that with perfect ADI for some parts, safety stocks can be reduced while maintaining the same service level, which can result in large cost savings. These cost savings are not just due to the parts which are monitored. That is because the job fill rate service criterion measures the joint availability of parts such that the increased availability of parts with ADI, may be offset by the decreased availability of other parts without ADI. Further, we find that monitoring combinations of parts with high mean demand probabilities and high mean prices in general leads to the largest cost savings. There is a very strong correlation between the mean price of the parts that are monitored and the cost savings potential compared to a non-ADI scenario. The correlation between the monitored parts' mean demand probabilities and the cost savings that can be achieved is still strong on average across all problem instances that we considered. However, we did observe instances for which monitoring the parts most likely to fail led to sub-optimal cost savings. Summing up, our model can be used to quantify the cost savings obtained by sensor technology and to support decision making, when parts have to be selected for redesign and configuration with sensors.

Funding Open Access funding enabled and organized by Projekt DEAL. The authors did not receive support from any organization for the submitted work.

Data availability The data used for our numerical study has been randomly generated as described in Sect. 5.

Code Availability All code written for the numerical study is provided in two files named “Numerical Study.R” and “Functions.R”.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give

appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix I - Derivation of $\pi_i(k)$ according to Prak et al. (2017)

We follow (Prak et al. 2017) in deriving the steady-state probabilities $\pi_i(k)$, $k = s_i + 1, \dots, S_i$ of the net inventory level of an NP i before the start of the tour but after order arrival (if an order has been placed at all). Let us define the time between two consecutive order arrivals as an order cycle. In a first step the probability that the net inventory level reaches a certain level $k = s_i + 1, \dots, S_i$ at some point within one order cycle has to be determined. We denote this probability by $\Pi_i(k)$, $k = s_i + 1, \dots, S_i$. Then we can use this probability to derive the probability that the net inventory level is in a particular state k at an arbitrary point in time within the order cycle. This latter probability is our steady-state probability $\pi_i(k)$.

The probabilities $\Pi_i(k)$, $k = s_i + 1, \dots, S_i$ can be calculated recursively starting from $k = S_i$ with the following formula

$$\Pi_i(k) = \begin{cases} 1 & \text{if } k = S_i \\ \sum_{l=k+1}^{S_i} \Pi_i(l) \cdot \frac{P(D_i=l-k)}{1-P(D_i=0)} & \text{if } k < s_i, \end{cases} \tag{23}$$

At the beginning of each cycle, the net inventory level is always S_i . Thus, S_i is reached with probability one. The following levels $k = S_i - 1, \dots, s_i + 1$ are then reached if the net inventory level drops to this level from one of the previously reached levels due to the observed customer demand. The denominator $1 - P(D_i = 0)$ accounts for the possibility that no demand for part i might occur for an arbitrary number of tours. Using $\Pi_i(k)$, we can calculate $\pi_i(k)$ as follows

$$\pi_i(k) = \frac{\Pi_i(k)}{\sum_{l=s_i+1}^{S_i} \Pi_i(l)}. \tag{24}$$

Appendix II - Derivation of $P(D_i(l, l) \leq \max(I_i^+ - D_i(1, l - 1), 0) | J = j)$

We can determine the probability $P(D_i(l, l) \leq \max(I_i^+ - D_i(1, l - 1), 0) | J = j)$ that sufficient units of part i are available for the l -th out of j jobs in total as the sum of $P(I_i^+ - D_i(1, l) \geq 0 | J = j)$ and $P(I_i^+ - D_i(1, l) < 0, D_i(l, l) = 0 | J = j)$. $P(I_i^+ - D_i(1, l) \geq 0 | J = j)$ is the probability that the net inventory of part i after visiting the l -th customer is not negative. In this case, the l -th job must be

complete. Even if there are already backorders for part i before the l -th customer is visited, this job can be completed if part i is not required by the l -th customer. $P(I_i^+ - D_i(1, l) < 0, D_i(l, l) = 0 | J = j)$ describes the probability of this event.

Let us start by determining the first summand $P(I_i^+ - D_i(1, l) \geq 0 | J = j)$. We can express this probability by considering the net inventory level before ordering I_i^- , the order quantity Q_i and the demand D_i . Keep in mind that as explained in Sect. 4.1 for an AP i $\pi_i(k)$, $k = s_i + 1, \dots, S_i$ defines the steady-state probabilities of the net inventory level before order placement. Using this we obtain

$$\begin{aligned}
 &P(I_i^+ - D_i(1, l) \geq 0 | J = j) \\
 &= P(I_i^- + Q_i - D_i(1, l) \geq 0 | J = j) \\
 &= \sum_{m=0}^{j \cdot d_i^{\max}} P(I_i^- + Q_i - D_i(1, l) \geq 0 | D_i(1, j) = m) \cdot P(D_i(1, j) = m) \\
 &= \sum_{m=0}^{j \cdot d_i^{\max}} \sum_{k=s_i+1}^{S_i} P(k + Q_i - D_i(1, l) \geq 0 | D_i(1, j) = m, I_i^- = k) \cdot P(D_i(1, j) = m) \cdot \pi_i(k)
 \end{aligned} \tag{25}$$

Since Q_i depends on the values of $D_i(1, j)$ and I_i^- we can split (25) into two parts. The first part considers all relevant events in which an order was placed and the second part considers the relevant events in which no order was placed.

$$\begin{aligned}
 &P(I_i^+ - D_i(1, l) \geq 0 | J = j) \\
 &= \sum_{m=1}^{j \cdot d_i^{\max}} \sum_{k=s_i+1}^{\min(s_i+m, S_i)} [P(D_i(1, l) \leq m + S_i | D_i(1, j) = m) \cdot P(D_i(1, j) = m) \cdot \pi_i(k)] \\
 &+ \sum_{m=0}^{\min((j-1) \cdot d_i^{\max}, S_i - s_i - 1)} \sum_{k=\max(s_i+m+1, 0)}^{S_i} [P(D_i(1, l) \leq k | D_i(1, j) = m) \cdot P(D_i(1, j) = m) \cdot \pi_i(k)] \\
 &= \sum_{m=1}^{j \cdot d_i^{\max}} \sum_{k=s+1}^{\min(s_i+m, S_i)} [P(D_i(1, j) = m) \cdot \pi_i(k)] \\
 &+ \sum_{m=0}^{\min((j-1) \cdot d_i^{\max}, S_i - s_i - 1)} \sum_{k=\max(s_i+m+1, 0)}^{S_i} \sum_{n=0}^{\min(k, m)} [P(D_i(1, l) = n) \cdot P(D_i(l + 1, j) = m - n) \cdot \pi_i(k)]
 \end{aligned} \tag{26}$$

Let us now determine the probability $P(I_i^+ - D_i(1, l) < 0, D_i(l, l) = 0 | J = j)$ that the net inventory level is negative after performing the l -th job but not because of the demand of the l -th customer which was zero. For this purpose, we must first differentiate between the case in which we have only one job in the entire repair tour and all cases with more than one job in the tour. If we assume the total number of jobs

J in the tour is one, then l has to be one as well. We obtain the following result

$$P(I_i^+ - D_i(1, 1) < 0, D_i(1, 1) = 0 \mid J = 1) = P(D_i(1, 1) = 0) \cdot P(I_i^+ < 0) \tag{27}$$

Since we only have one customer on the tour who does not require part i it is safe to say that this part has not been replenished before the start of the tour, as the order quantity is based on the demand in the upcoming tour that is known beforehand. That means that the net inventory levels I_i^- and I_i^+ are identical. Thus we get for $J = 1$.

$$\begin{aligned} &P(I_i^+ - D_i(1, 1) < 0, D_i(1, 1) = 0 \mid J = 1) \\ &= P(D_i(1, 1) = 0) \cdot P(I_i^- < 0) \\ &= P(D_i(1, 1) = 0) \cdot \left[1 - \sum_{k=\max(0, s_i+1)}^{S_i} \pi_i(k) \right] \end{aligned} \tag{28}$$

In case we have more than just one customer we get the following result

$$\begin{aligned} &P(I_i^+ - D_i(1, l) < 0, D_i(l, l) = 0 \mid J = j) \\ &= P(I_i^+ - D_i(1, l - 1) < 0, D_i(l, l) = 0 \mid J = j). \\ &= P(D_i(l, l) = 0) \cdot [1 - P(I_i^+ - D_i(1, l - 1) \geq 0 \mid D_i(l, l) = 0, J = j)] \end{aligned} \tag{29}$$

Note that we defined $D(l_1, l_2) := 0$ for $l_1 > l_2$. Using this definition the above result holds even when we are considering the first customer in a tour ($l=1$). At this point, we can determine $P(I_i^+ - D_i(1, l - 1) \geq 0 \mid D_i(l, l) = 0, J = j)$ analogously to $P(I_i^+ - D_i(1, l) \geq 0 \mid J = j)$ (see (2) and (26)). Thus, we get

$$\begin{aligned} &P(I_i^+ - D_i(1, l) < 0, D_i(l, l) = 0 \mid J = j) \\ &= \begin{cases} P(D_i(1, 1) = 0) \cdot \left[1 - \sum_{k=\max(0, s_i+1)}^{S_i} \pi_i(k) \right] & \text{if } l = j = 1 \\ P(D_i(l, l) = 0) \cdot \left[1 - \left(\sum_{m=1}^{(j-1) \cdot d_i^{\max}} \sum_{k=s_i+1}^{\min(s_i+m, S_i)} [\pi_i(k) \cdot P(D_i(1, l - 1) + D_i(l + 1, j) = m)] \right. \right. \\ \quad \left. \left. + \sum_{m=0}^{\min((j-1) \cdot d_i^{\max}, S_i - s_i - 1)} \sum_{k=\max(s_i+m+1, 0)}^{S_i} \sum_{n=0}^{\min(k, m)} [P(D_i(l + 1, j) = m - n) \right. \right. \\ \quad \left. \left. \cdot P(D_i(1, l - 1) = n) \cdot \pi_i(k)] \right) \right] & \text{else} \end{cases} \end{aligned} \tag{30}$$

Appendix III - Greedy Heuristic

Algorithm 1 Greedy heuristic

Input: unit holding cost h_i , $\forall i \in \{1, \dots, N\}$, material handling cost f_i , $\forall i \in \{1, \dots, N\}$, customer distribution $P(J = j)$, $\forall j \in \{1, \dots, N\}$, Demand distribution $P(D_i(l, l) = d)$, $\forall d \in \{0, \dots, d_i^{\max}\}$, $i \in \{1, \dots, N\}$ and $P(D_i = d)$, $\forall d \in \{0, \dots, d_i^{\max} \cdot M\}$, $i \in \{1, \dots, N\}$, target job fill rate γ^*

Output: reorder and order-up-to levels s and S

$$\bar{D}_i \leftarrow \sum_{d=0}^{M \cdot d_i^{\max}} d \sum_{j=1}^M P(D_i(1, j) = d) \cdot P(J = j), \forall i \in \{1, \dots, N\}$$

$$Q_i \leftarrow \max \left(\left\lfloor \sqrt{\frac{2f_i \bar{D}_i}{h_i}} \right\rfloor, 1 \right), \forall i \in \{1, \dots, N\}$$

$$s \leftarrow -Q, \quad S \leftarrow 0^N$$

order $\leftarrow \emptyset$

while $\gamma(s, S) < \gamma^*$ **do**

$$R(s_i) \leftarrow \Delta_{s_i} C(s, S) \cdot (\Delta_{s_i} \gamma(s, S) + \epsilon)^{-\text{sgn}(\Delta_{s_i} C(s, S))}, \forall i \in \{1, \dots, N\}$$

$$R(S_i) \leftarrow \Delta_{S_i} C(s, S) \cdot (\Delta_{S_i} \gamma(s, S) + \epsilon)^{-\text{sgn}(\Delta_{S_i} C(s, S))}, \forall i \in \{1, \dots, N\}$$

if $\min_i R(s_i) < \min_i R(S_i)$ **then**

$$i^* \leftarrow \arg \min_{i \in \{1, \dots, N\}} R(s_i)$$

$$s_{i^*} \leftarrow s_{i^*} + 1$$

order $\leftarrow (i^*, \text{order})$

else

$$i^* \leftarrow \arg \min_{i \in \{1, \dots, N\}} R(S_i)$$

$$S_{i^*} \leftarrow S_{i^*} + 1$$

end if

end while

for $j = 1$ to $\text{length}(\text{order})$ **do**

if $\text{order}[j] \neq \text{order}[k]$, $\forall k \in \{1, \dots, j-1\}$ **then**

$$i \leftarrow \text{order}[j]$$

$$s^{\text{temp}} \leftarrow s$$

$$s_i^{\text{temp}} \leftarrow s_i + 1$$

if $\gamma(s^{\text{temp}}, S) > \gamma^*$ **then**

$$s \leftarrow s^{\text{temp}}$$

end if

end if

end for

References

- Angelus A, Özer Ö (2016) Knowledge you can act on: Optimal policies for assembly systems with expediting and advance demand information. *Oper Res* 64:1338–1371. <https://doi.org/10.1287/opre.2016.1541>
- Bernstein F, DeCroix GA (2015) Advance demand information in a multiproduct system. *Manuf Serv Oper Manag* 17:52–65. <https://doi.org/10.1287/msom.2014.0502>
- Bijvank M, Koole G, Vis IFA (2010) Optimising a general repair kit problem with a service constraint. *Eur J Oper Res* 204:76–85. <https://doi.org/10.1016/j.ejor.2009.10.004>
- Brumelle S, Granot D (1993) The repair kit problem revisited. *Oper Res* 41:994–1006. <https://doi.org/10.1287/opre.41.5.994>

- Cohen MA, Zheng Y-S, Agrawal V (1997) Service parts logistics: a benchmark analysis. *IIE Trans* 29:627–639. <https://doi.org/10.1080/07408179708966373>
- Gallego G, Özer Ö (2001) Integrating replenishment decisions with advance demand information. *Manage Sci* 47:1344–1360. <https://doi.org/10.1287/mnsc.47.10.1344.10261>
- Gallego G, Özer Ö (2003) Optimal replenishment policies for multiechelon inventory problems under advance demand information. *Manuf Serv Oper Manag* 5:157–175. <https://doi.org/10.1287/msom.5.2.157.16072>
- Graves SC (1982) Note-a multiple-item inventory model with a job completion criterion. *Manage Sci* 28:1334–1337. <https://doi.org/10.1287/mnsc.28.11.1334>
- Hariharan R, Zipkin P (1995) Customer-order information, leadtimes, and inventories. *Manage Sci* 41:1599–1607. <https://doi.org/10.1287/mnsc.41.10.1599>
- Hausman WH (1982) On optimal repair kits under a job completion criterion. *Manage Sci* 28:1350–1351
- Heeremans D, Gelders L (1995) Multiple period repair kit problem with a job completion criterion: A case study. *Eur J Oper Res* 81:239–248. [https://doi.org/10.1016/0377-2217\(94\)00177-E](https://doi.org/10.1016/0377-2217(94)00177-E)
- Koçağa YL, Şen A (2007) Spare parts inventory management with demand lead times and rationing. *IIE Trans* 39:879–898. <https://doi.org/10.1080/07408170601013646>
- Lu Y, Song J-S, Yao DD (2003) Order fill rate, leadtime variability, and advance demand information in an assemble-to-order system. *Oper Res* 51:292–308. <https://doi.org/10.1287/opre.51.2.292.12781>
- Mamer JW, Shogan AW (1987) A constrained capital budgeting problem with applications to repair kit selection. *Manage Sci* 33:800–806. <https://doi.org/10.1287/mnsc.33.6.800>
- Mamer JW, Smith SA (1982) Optimizing field repair kits based on job completion rate. *Manage Sci* 28:1328–1333. <https://doi.org/10.1287/mnsc.28.11.1328>
- Mamer JW, Smith SA (1985) Job completion inventory systems: Optimal policies for repair kits and spare machines. *Manage Sci* 31:703–718. <https://doi.org/10.1287/mnsc.31.6.703>
- March ST, Scudder GD (1984) Note-on “optimizing field repair kits based on job completion rate”. *Manage Sci* 30:1025–1028. <https://doi.org/10.1287/mnsc.30.8.1025>
- Marklund J (2006) Controlling inventories in divergent supply chains with advance-order information. *Oper Res* 54:988–1010. <https://doi.org/10.1287/opre.1060.0319>
- Özer Ö (2003) Replenishment strategies for distribution systems under advance demand information. *Manage Sci* 49:255–272. <https://doi.org/10.1287/mnsc.49.3.255.12738>
- Özer Ö, Wei W (2004) Inventory control with limited capacity and advance demand information. *Oper Res* 52:988–1000. <https://doi.org/10.1287/opre.1040.0126>
- Prak D, Saccani N, Syntetos A, Teunter R, Visintin F (2017) The repair kit problem with positive replenishment lead times and fixed ordering costs. *Eur J Oper Res* 261:893–902. <https://doi.org/10.1016/j.ejor.2017.02.019>
- Rippe C, Kiesmüller GP (2022) The repair kit problem with imperfect advance demand information. *Eur J Oper Res*. <https://doi.org/10.1016/j.ejor.2022.04.019>
- Saccani N, Visintin F, Mansini R, Colombo M (2016) Improving spare parts management for field services: a model and a case study for the repair kit problem. *IMA J Manag Math* 28:185–204. <https://doi.org/10.1093/imaman/dpw023>
- Smith SA, Chambers JC, Shlifer E (1980) Optimal inventories based on job completion rate for repairs requiring multiple items. *Manage Sci* 26:849–852. <https://doi.org/10.1287/mnsc.26.8.849>
- Teunter RH (2006) The multiple-job repair kit problem. *Eur J Oper Res* 175:1103–1116. <https://doi.org/10.1016/j.ejor.2005.05.023>
- Wang T, Toktay BL (2008) Inventory management with advance demand information and flexible delivery. *Manage Sci* 54:716–732. <https://doi.org/10.1287/mnsc.1070.0831>