Three extensions to the repair kit problem

Schriftliche Promotionsleistung

zur Erlangung des akademischen Grades

Doctor rerum politicarum

vorgelegt und angenommen

an der Fakultät für Wirtschaftswissenschaft

der Otto-von-Guericke-Universität Magdeburg

Verfasser:	Christoph Rippe
Geburtsdatum und -ort:	19.04.1989, Minden
Arbeit eingereicht am:	03.06.2022

Gutachter der schriftlichen Promotionsleistung:

Prof. Dr. Gudrun P. Kiesmüller

Prof. Dr. Sven Müller

Prof. Dr. Marlin Ulmer

Datum der Disputation: 10.11.2022

Contents

* The page numbers are adopted from the original publications.

List of figures

II The Repair Kit Problem with Fixed Delivery Costs*

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

IV The repair kit problem with imperfect advance demand information*

* The page numbers are adopted from the original publications.

List of tables

* The page numbers are adopted from the original publications.

I

Introduction

For many original equipment manufacturers (OEMs) after-sales service has become crucial to their overall performance. That is on one hand because good after-sales service increases customer satisfaction and satisfied customers are more likely to purchase further equipment from an OEM and recommend the OEM's products to others. Thus, after-sales service is directly linked to product sales. On the other hand and at least equally important after-sales service can be a significant source of revenue and profit itself. A survey conducted by Deloitte (Koudal (2006)) among 80 OEMs from Europe, North America and the Asia-Pacific region finds that on average 26% of their total revenues are from services including spare part sales. In a second survey by Bain & Company (Straehle et al. (2015)) among 45 European OEMs the average contribution of services towards total revenue is 22%. Both studies find that profit margins are higher in the service business than for original equipment sales.

While after-sales service comprises anything from the initial set-up of an appliance to its disposal at the end of the product lifetime (see Dombrowski et al. (2020) and Durugbo (2020)), maintenance certainly is a key aspect from the customer's point of view. In this work, the focus is on corrective rather than preventive maintenance for durable goods such as printers, washing machines, or heating systems. That means customers experience failures of their appliances and request immediate repairs. As the malfunctioning appliances are typically too heavy or bulky to be transported to a repair shop, repairs must be carried out at the customers' locations. To perform these on-site repair jobs OEMs or third-party service providers employ traveling repair technicians. The technicians' job is to visit the customers in need of help, to identify the sources of their appliances' malfunctions and to repair them instantaneously. Wherever the malfunction is due to a failure of one or multiple parts the repair technician has to replace all failed parts with spare parts to complete the repair job. For this purpose, the repair technician is equipped with a set of different spare parts that are carried in the technician's company van. This set is called the repair kit and the question of how to manage its contents is the repair kit problem.

This work comprises three articles specified in Table 1 that extend the research on the repair kit problem in two distinct ways. The remainder of this chapter relates the three papers to the relevant literature and summarizes the main findings. The papers themselves are then presented in the following chapters II, III, and IV.

Table 1: Extensions to the repair kit problem

	Rippe (2022)	Rippe & Kiesmüller (2023a)	Rippe & Kiesmüller (2023b)
Title	The Repair Kit Problem with Fixed Delivery Costs	The added value of advance demand information for the planning of a repair kit	The repair kit problem with imperfect advance demand information
Chapter	Chapter II	Chapter III	Chapter IV

1. Contribution

The articles presented in this work contribute to 3 different streams of literature. That is first and foremost research on the repair kit problem but also research on the stochastic joint replenishment problem and on the use of advance demand information (ADI).

As for other spare part management problems the repair kit problem trades off inventory holding costs for spare parts on stock against either the service level that can be achieved with them or penalty costs that are incurred when required spare parts are unavailable. The repair kit problem is similar to system-oriented spare part problems (see e.g. Basten & van Houtum (2014)) in so far that the service level considered depends on the availability of spares for all critical parts of an appliance. However, the repair kit problem differs from these problems in that when an appliance fails, the cause of this failure is not immediately apparent. Instead, the repair technician must first visit the customer to diagnose the failure.

The repair kit problem was first studied by Smith et al. (1980) as a single-job problem meaning that the technician may be restocked after every customer visit. They suggest using the probability that all parts necessary to complete a repair job are available as a service criterion which they term the job fill rate. They aim to minimize the sum of inventory holding costs and penalty costs incurred for each failed repair attempt. Contrary to this unconstrained problem formulation, Graves (1982) and Hausman (1982) study a repair kit problem with a service constraint. That is they minimize holding costs subject to a job fill rate service constraint. Mamer & Smith (1982) and March & Scudder (1984) generalize the previous approaches by allowing for more complex failure scenarios where the demand for different spare parts need not be independent. Further extensions of the single-job repair kit problem by Mamer & Smith (1985) and Mamer & Shogan (1987) consider additional spare machines and capacity constraints. Brumelle & Granot (1993) show that optimal repair kits for different relative weights of holding and penalty costs form a monotone sequence.

Heeremans & Gelders (1995) are the first to assume that several customers are visited in a repair tour before the service technician gets a chance to restock the repair kit. For this multi-job problem, they measure service by the probability that all jobs in the tour can be completed. Teunter (2006) and Bijvank et al. (2010) study similar multi-job problems but return to the job fill rate service level that is used for the single-job problems. Since the chance to complete a repair decreases the more customers the technician has already seen, they consider an average job fill rate across all customers visited. While papers up to this point exclusively trade off holding costs for spare parts against service levels or penalty costs, two recent publications by Saccani et al. (2017) and Prak et al. (2017) also take costs incurred for the replenishment of the repair kit into account. Prak et al. (2017) consider fixed material handling costs per spare part that is ordered. They suggest managing the contents of the repair kit with individual (s,S)-policies for all spare parts. Saccani et al. (2017) present a problem formulation with both fixed order costs per ordered part and fixed delivery costs per shipment of multiple parts to the repair technician. By explicitly considering delivery costs their paper is the first to treat the repair kit replenishment frequency as a decision variable. However, they determine the repair kit that would have performed best given the exact spare part demand experienced in the past. Thus, they examine a deterministic-demand version of the repair kit problem.

As can be seen in the overview in Table 2 the articles presented in this work all contribute to the research on the more complex multi-job repair kit problem. Rippe (2022) returns to the wide-scale cost scenario with fixed order and fixed delivery costs studied by Saccani et al. (2017). However, demand for spare parts is modeled as a random variable, which is more realistic for real-life problems and thus a standard assumption made for all repair kit problems except for that by Saccani et al. (2017). Rippe & Kiesmüller (2023a) and Rippe & Kiesmüller (2023b) extend the repair kit problems of Prak et al. (2017) and Teunter (2006) by utilizing information on the condition of the customers' failed appliances.

The repair kit literature so far assumes that the repair technician is entirely unaware of a customer's needs for spare parts before the first on-site visit. While this may be true for older machinery, state-of-the-art appliances are often equipped with sensors that monitor the condition the appliance is in. In case a failure occurs these sensors generate additional information and help to detect the cause of the error. This sensor information is available to the service technician before a first customer visit. It is, therefore, advance demand information for spare parts. Rippe & Kiesmüller (2023a) and Rippe & Kiesmüller (2023b) consider different types of ADI for the replenishment of the repair kit, which stem from different types of sensor systems. While Rippe & Kiesmüller (2023a) assume that reliable, part-specific ADI is available for only a fraction of an appliance's critical parts, Rippe & Kiesmüller (2023b) consider ADI that is uncertain and not specific to individual parts.

Table 2: Literature on the repair kit problem

	Problem Specification										Solution	
	single job	multiple jobs	fixed order costs	fixed delivery costs	deterministic demand	stochastic demand	no ADI	ADI	unconstrained	service constraint	optimal	approximate
Smith et al. (1980)	✓					✓	✓		✓		✓	
Graves (1982)	✓					✓	✓			✓	✓	
Hausman (1982)	✓					✓	✓			✓	✓	
Mamer & Smith (1982)	✓					✓	✓		✓		✓	
March & Scudder (1984)	✓					✓	✓		✓		✓	
Mamer & Smith, 1985	✓					✓	✓			✓		✓
Mamer & Shogan, 1987	✓					✓	✓			✓		✓
Brumelle & Granot (1993)	✓					✓	✓		✓		✓	
Heeremans & Gelders (1995)		✓				✓	✓			✓		✓
Teunter (2006)		✓				✓	✓		✓	✓		✓
Bijvank et al. (2010)		✓				✓	✓		✓	✓		✓
Saccani et al. (2017)		✓	✓	✓	✓		✓		✓		✓	
Prak et al. (2017)		✓	✓			✓	✓			✓		✓
Rippe (2022)		✓	✓	✓		✓	✓			✓		✓
Rippe & Kiesmüller (2023a)		✓	✓			✓		✓		✓		✓
Rippe & Kiesmüller (2023b)		✓				✓		✓	✓			✓

Because the repair kit problem is extended to integrate fixed delivery costs for coordinated shipments of multiple parts, Rippe (2022) can be seen as a contribution to the research on the stochastic-demand joint replenishment problem (SJRP). The joint replenishment problem in general is the problem of managing the inventories of several products that are sourced from one particular supplier, such that replenishments of different products can be delivered at once. For each shipment fixed delivery costs are charged and for each product included in that shipment fixed order costs are incurred that account for product-specific material handling operations. These two types of costs often termed major and minor order costs are traded off against inventory holding costs in any case and stock-outs in case product demand is stochastic. The repair kit problem with fixed delivery costs stands out in the SJRP-literature for two reasons. First, time is not continuous. Second, shortage costs or service targets in the SJRP are typically considered per product, whereas the technician in the repair kit problem relies on the joint availability of all parts needed for a repair job. Since the structure of the optimal policy for the SJRP is unknown, there is a large variety of policies suggested for different versions of the problem. Multi-product deliveries may be triggered when the inventory level of a single product falls below a product-specific reorder level (e.g Balintfy (1964), Federgruen et al. (1984), Kiesmüller (2010)) or when the aggregated inventory across all products undercuts a total stock reorder level (e.g. Renberg & Planche (1967), Nielsen & Larsen (2005)). Alternatively, deliveries may be dispatched periodically as in Atkins & Iyogun (1988) and Viswanathan (1997). In Rippe (2022) the repair kit is replenished using independent, periodic review (s,S)-policies for all parts with common review points for parts sourced from the same supplier. Viswanathan (1997) has already shown that this type of policy works well for a standard SJRP. For the repair kit problem, this policy is a good fit because it allows the explicit calculation of the job fill rate service level, which would not be possible for more complex policies. More detailed reviews of the extensive literature on the (stochastic) joint replenishment problem and further extensions can be found in Khouja & Goyal (2008) and Bastos et al. (2017).

The papers Rippe & Kiesmüller (2023a) and Rippe & Kiesmüller (2023b) that consider information obtained from different types of sensors also complement the literature on inventory models with ADI. ADI can arise in all sorts of indications, from explicit orders

placed in advance by a customer to click-stream data that may or may not turn into actual demand. Most contributions to the ADI literature, however, study single product problems. For a review of these single-item problems, the reader is referred to the overview by Karaesmen (2013). The contributions closest to Rippe & Kiesmüller (2023a) and Rippe & Kiesmüller (2023b) are the multi-item ADI-problems studied by Lu et al. (2003), Angelus & Özer (2016), Thonemann (2002), Bernstein & DeCroix (2015) and Chen et al. (2017). Lu et al. (2003) and Angelus & Özer (2016) both consider perfect advance orders for products manufactured in assembly systems. A repair job may be understood as a product assembled from the set of spare parts needed to complete that repair job. The ADI in Lu et al. (2003) and Angelus & Özer (2016), however, is for final products and thus all components required for a product are known in advance, whereas in Rippe & Kiesmüller (2023a) ADI is part-specific and it is only available for a subset of all parts that parts may be needed.
Thonemann (2002),Bernstein & DeCroix (2015) and Chen et al. (2017) consider imperfect ADI as in Rippe & Kiesmüller (2023b). In Bernstein & DeCroix (2015) ADI comes as information on either the total demand across several products or the relative market shares of these products. Thus, in any case, ADI is aggregated across all customers in the entire planning period. In Chen et al. (2017) and Thonemann (2002) advance information on future orders is available for individual customers. Contrary to the multi-item demand observed in the repair kit problem, though, both studies assume that customers request only one out of several items at once. Still, the replenishment policy that is shown to be optimal by Chen et al. (2017) in a single-item case is very useful for a heuristic solution procedure presented in Rippe & Kiesmüller (2023b).

2. Main Results

The main value of the research presented in this thesis is to show how the total costs incurred by a technician performing on-site repairs can be reduced for three different real-life scenarios. To this end, the articles in chapters II, III, and IV demonstrate how costs and service levels can be determined given a repair kit replenishment policy and how to derive good replenishment policies using novel heuristic solution procedures.
In Rippe (2022), presented in chapter II, the standard multi-job, stochastic-demand repair

kit problem is extended to integrate fixed delivery costs for each shipment from a supplier to the repair kit. While the repair kit may be restocked by different suppliers, each spare part is only sourced from one supplier. Therefore the inventories of all parts sourced from the same supplier are reviewed periodically at the same time intervals to allow for coordinated deliveries of multiple items. As mentioned before individual periodic (s,S)-policies are applied, similar to Viswanathan (1997). For this setting, closed-form expressions are derived for the expected total costs per repair tour and the job fill rate service level, the latter depending on the joint availability of all parts sourced from all suppliers. In a numerical experiment reorder and order-up-to levels are determined for different review period lengths using the same greedy heuristic adapted from Prak et al. (2017). It can be shown that longer review periods often decrease delivery costs significantly, whereas safety stock increases necessary to counteract the increased uncertainty that comes with longer review periods are modest. Thus, in the absence of any ADI replenishing the repair technician less frequently can help to reduce costs.

In Rippe & Kiesmüller (2023a), reproduced in chapter III, a repair kit problem is studied in which the appliances in need of repair are equipped with sensors that monitor some parts' conditions individually. A closed-form job fill rate is presented that takes into account both the monitored parts for which ADI is available and the parts that are not monitored. Further, a new greedy heuristic is proposed that allows for individual increments of all decision variables until the target job fill rate is reached. Comparing appliances with different shares of monitored and unmonitored parts, it can be shown that advance demand information helps to reduce safety stocks of both monitored and unmonitored parts and thus inventory holding costs. Numerical experiments reveal that monitoring frequently required parts with high values is most beneficial in general. However, there is no monotone relation between the parts' values or their demand frequencies and the cost savings that can be achieved when they are monitored.

In Rippe & Kiesmüller (2023b), presented in chapter IV, uncertain ADI is used for the management of the repair kit. This uncertain ADI comes as error codes triggered by sensors that monitor indicators such as heat pressure or vibration which can be affected by the conditions of several parts. Given this type of ADI, the repair kit problem is formulated as a Markov decision process. Optimal policies, however, can only be derived for small-scale problems, because the state space increases exponentially in the number of spare parts in the repair kit. To obtain solutions for real-world problems two heuristics are suggested that decompose the original problem in different ways. Both heuristics are shown to produce near-optimal results on small instances. Yet, the non-greedy heuristic that optimizes parts' inventories individually performs slightly better than the greedy heuristic that considers all parts at once, which is surprising given that most contributions to the repair kit problem suggest greedy heuristics. Extensive numerical studies show that using imperfect ADI for the replenishment of the repair kit helps to reduce costs compared to a situation where this information is not available or not used. The value of the ADI, though, increases the more precise it gets.

References

Angelus, A., Ö. Özer. 2016. Knowledge you can act on: Optimal policies for assembly systems with expediting and advance demand information. Operations Research, 64, 1338–1371.

Atkins, D. R., P. O. Iyogun. 1988. Periodic versus "can-order" policies for coordinated multi-item inventory systems. Management Science, 34, 791–796.

Balintfy, J. L.. 1964. On a basic class of multi-item inventory problems. Management Science, 10, 287–297.

Basten, R., G. van Houtum. 2014. System-oriented inventory models for spare parts. Surveys in Operations Research and Management Science, 19, 34–55.

Bastos, L. S. L., M. L. Mendes, D. R. L. Nunes, A. C. S. Melo, M. P. Carneiro. 2017. A systematic literature review on the joint replenishment problem solutions: 2006-2015. Production, 27.

Bernstein, F., G. A. DeCroix. 2015. Advance demand information in a multiproduct system. Manufacturing & Service Operations Management, 17, 52–65.

Bijvank, M., G. Koole, I. F. A. Vis. 2010. Optimising a general repair kit problem with a service constraint. European Journal of Operational Research, 204, 76–85.

Brumelle, S., D. Granot. 1993. The repair kit problem revisited. Operations Research, 41, 994–1006.

Chen, L., Şafak Yücel, K. Zhu. 2017. Inventory management in a closed-loop supply chain with advance demand information. Operations Research Letters, 45, 175 – 180.

Dombrowski, U., S. Fochler, C. Malorny (Eds.). 2020. After Sales Service - Zukunftsfähig und prozessorientiert gestalten. Springer.

Durugbo, C. M.. 2020. After-sales services and aftermarket support: a systematic review, theory and future research directions. International Journal of Production Research, 58, 1857–1892.

Federgruen, A., H. Groenevelt, H. C. Tijms. 1984. Coordinated replenishments in a multi-item inventory system with compound poisson demands. Management Science, 30, 344–357.

Graves, S. C.. 1982. A multiple-item inventory model with a job completion criterion. Management Science, 28, 1334–1337.

Hausman, W. H.. 1982. On optimal repair kits under a job completion criterion. Management Science, 28, 1350–1351.

Heeremans, D., L. Gelders. 1995. Multiple period repair kit problem with a job completion criterion: A case study. European Journal of Operational Research, 81, 239–248.

Karaesmen, F.. 2013. Value of advance demand information in production and inventory systems with shared resources. In J. MacGregor Smith, & B. Tan (Eds.), Handbook of Stochastic Models and Analysis of Manufacturing System Operations (pp. 139–165). Springer.

Khouja, M., S. Goyal. 2008. A review of the joint replenishment problem literature: 1989–2005. European Journal of Operational Research, 186, 1–16.

Kiesmüller, G. P.. 2010. Multi-item inventory control with full truckloads: A comparison of aggregate and individual order triggering. European Journal of Operational Research, 200, 54–62.

Koudal, P.. 2006. The Service Revolution in Global Manufacturing Industries. Deloitte Research.

Lu, Y., J.-S. Song, D. D. Yao. 2003. Order fill rate, leadtime variability, and advance demand information in an assemble-to-order system. Operations Research, 51, 292–308.

Mamer, J. W., A. W. Shogan. 1987. A constrained capital budgeting problem with applications to repair kit selection. Management Science, 33, 800–806.

Mamer, J. W., S. A. Smith. 1982. Optimizing field repair kits based on job completion rate. Management Science, 28, 1328–1333.

Mamer, J. W., S. A. Smith. 1985. Job completion inventory systems: Optimal policies for repair kits and spare machines. Management Science, 31, 703–718.

March, S. T., G. D. Scudder. 1984. On "optimizing field repair kits based on job completion rate". Management Science, 30, 1025–1028.

Nielsen, C., C. Larsen. 2005. An analytical study of the q(s,s) policy applied to the joint replenishment problem. European Journal of Operational Research, 163, 721–732. Supply Chain Management and Advanced Planning.

Prak, D., N. Saccani, A. Syntetos, R. Teunter, F. Visintin. 2017. The repair kit problem with positive replenishment lead times and fixed ordering costs. European Journal of Operational Research, 261, 893–902.

Renberg, B., R. Planche. 1967. Un modèle pour la gestion simultanée de n articles d'un stock. Revue Française d'Informatique et de Recherche Opérationelle, 1, 47–59.

Rippe, C.. 2022. The repair kit problem with fixed delivery costs. In N. Trautmann, & M. Gnägi (Eds.), Operations Research Proceedings 2021: Selected Papers of the Annual International Conference of the German Operations Research Society (GOR), Bern, Switzerland, September 1-3, 2021 (pp. 382–388). Springer.

Rippe, C., G. P. Kiesmüller. 2023b. The repair kit problem with imperfect advance demand information. European Journal of Operational Research, 304, 558–576.

Rippe, C., G. P. Kiesmüller. 2023a. The added value of advance demand information for the planning of a repair kit. Central European Journal of Operations Research, 31, 311–335.

Saccani, N., F. Visintin, R. Mansini, M. Colombo. 2017. Improving spare parts management for field services: a model and a case study for the repair kit problem. IMA Journal of Management Mathematics, 28, 185–204.

Smith, S. A., J. C. Chambers, E. Shlifer. 1980. Optimal inventories based on job completion rate for repairs requiring multiple items. Management Science, 26, 849–852.

Straehle, O., P. Roth, C. Herr. 2015. Winning in industrial service: The hallmarks of a service champion. Bain & Company, Inc.

Teunter, R. H.. 2006. The multiple-job repair kit problem. European Journal of Operational Research, 175, 1103–1116.

Thonemann, U. W.. 2002. Improving supply-chain performance by sharing advance demand information. European Journal of Operational Research, 142, 81 – 107.

Viswanathan, S.. 1997. Note. periodic review (s, s) policies for joint replenishment inventory systems. Management Science, 43, 1447–1454.

II
The Repair Kit Problem with Fixed Delivery Costs

First published as *Rippe, C. (2022) The repair kit problem with fixed delivery costs. in N. Trautmann & M. Gnägi (Eds.), Operations Research Proceedings 2021: Selected Papers of the Annual International Conference of the German Operations Research Society (GOR), Bern, Switzerland, September 1-3 2021, 382-388.* https://doi.org/10.1007/978-3-031-08623-6_56 by Springer Nature. Reproduced with permission from Springer Nature.

The Repair Kit Problem with Fixed Delivery Costs

Christoph Rippe(✉)

Otto von Guericke University Magdeburg, Magdeburg, Germany
christoph.rippe@ovgu.de

Abstract. The repair kit problem is the problem of managing the spare parts inventory of a field service technician. Contrary to most previous contributions we acknowledge that fixed costs per delivery to the technician make up an important share of the total field repair costs. Thus, we treat the replenishment frequency as a decision variable and suggest to manage the content of the repair kit using individual (R, s, S)-policies for each spare part with common review periods R for all parts sourced from the same supplier. We derive a closed-form expression for the job-fill-rate service level and suggest a heuristic to determine the length of the review period(s) as well as the reorder and order-up-to-levels for spare parts carried in the repair kit. Using a numerical experiment we show that lowering the replenishment frequency can lead to a substantial cost reduction. That is because delivery cost reductions outweigh the costs for additional safety stock.

Keywords: Repair kit problem · Inventory management · Stochastic models

1 Introduction

Manufacturers who offer on-site repair services to their customers must provide their service technicians with a set of spare parts called a repair kit. The multiple-job repair kit problem (RKP) is the problem of managing the content of a repair kit that can only be restocked after several customers have been visited in a tour. This problem was first studied by [1–3], who trade off holding costs for spare parts stocked in the repair kit against the service level that can be achieved or against penalty costs incurred for failed repair attempts. [4] consider additional part-specific fixed order costs that are associated with material handling activities in a warehouse. Thus they apply individual (s, S)-policies rather than base stock policies to manage the repair kit. Even though papers on the RKP agree that repair kits are restocked from only one or very few suppliers (e.g. a regional and a national warehouse), the opportunity to reduce delivery costs for shipments from the supplier to the service technicians by means of coordinated replenishment has been largely overlooked. So far the only contribution that considers fixed costs per delivery is a deterministic-demand RKP studied

N. Trautmann and M. Gnägi (Eds.): OR 2021, LNOR, pp. 382–388, 2022.
https://doi.org/10.1007/978-3-031-08623-6_56

by [5]. In this paper, our contribution is to integrate fixed delivery costs into the more realistic stochastic-demand RKP studied by [1–4]. Thus, our contribution is at the intersection of RKP and stochastic joint replenishment problem (SJRP). What distinguishes our problem from papers on the SJRP is that we consider a job-fill-rate service level that requires a more complex analysis than the part-fill-rates used for backorder cost calculations in the SJRP literature (for a review see [8]). We suggest to apply (R, s, S)-policies as introduced by [7] for our RKP and derive closed-form expressions for expected costs and the job-fill-rate service constraint for given policy parameters. Further, we introduce a heuristic solution procedure for our RKP and present a numerical experiment that demonstrates the cost savings potential of longer replenishment cycles. Our work can be seen as an extension to [4], that allows for cycle times different from one. For this reason, we follow their notation wherever possible.

2 Problem Description

There are N different parts in the repair kit. The number of repair jobs per repair tour and the number of units of each part $i = 1, \ldots, N$ required for a single job are stochastic. Let $p_c(l), c_{\min} \leq l \leq c_{\max}$ denote the probability for l customers in one tour and $p_i(k), d_{\min} \leq k \leq d_{\max}$ define the probability that k units of part i are required for a job. Demands for different parts are independent. By $P_c^t(l)$ and $P_i^{J|j}(k)$ we define the probability for l customers in t tours and the probability for a demand of k items of part i in j jobs. We obtain

$$P_c^t(l) = \sum_{\substack{l_1,\ldots,l_t \\ l_1+\cdots+l_t=l}} \prod_{r=1}^{t} p_c(l_r) \quad (1) \qquad P_i^{J|j}(k) = \sum_{\substack{k_1,\ldots,k_j \\ k_1+\cdots+k_j=k}} \prod_{m=1}^{j} p_i(k_m). \quad (1)$$

The repair kit is replenished from G suppliers. Each part $i = 1, \ldots, N$ is sourced from exactly one warehouse with $w(i) = g \in \{1, \ldots, G\}$ defining this warehouse. To manage the content of the repair kit we suggest periodic (s, S)-policies. That means the inventory positions (IPs) of all parts sourced from one supplier g are reviewed every R_g tours and if the IP of a part i is at or below a reorder level s_i at review time it is raised to an order-up-to level S_i. All orders

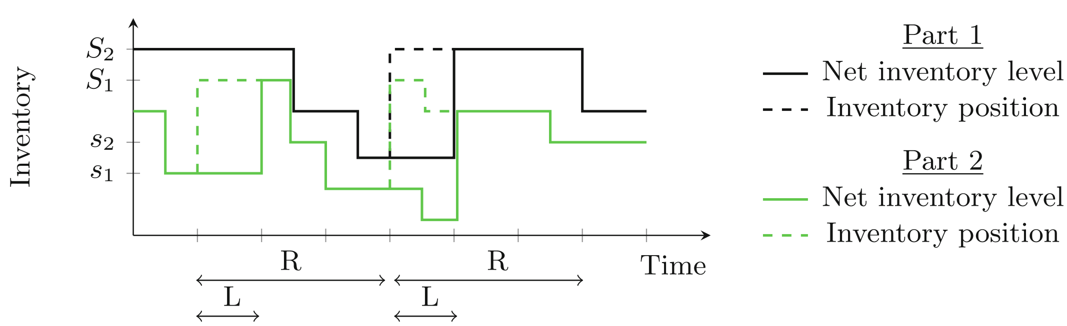

Fig. 1. Exemplary inventory development of two parts from the same supplier

from the supplier are delivered together after a lead time of L_g tours. See Fig. 1. for an example.

We define this policy by (R, s, S) with $R = (R_1, \ldots, R_G)$, $s = (s_1, \ldots, s_N)$ and $S = (S_1, \ldots, S_N)$. Let $\pi_i^{IP}(k)$ define the steady state distribution of part i's IP after order placement at review time. For a detailed derivation of $\pi_i^{IP}(k)$ we refer to [4]. Let us refer to the point in time $L_{w(i)}$ tours after review time as potential delivery time. The net inventory level (IL) at potential delivery time corresponds to the IP at review time minus demand during the lead time. With a review period longer than one tour we need to explicitly consider the IL after each tour within one review period. For a part i with $w(i) = g$ let us denote the steady state distribution of the IL $r = 0, \ldots, R_g - 1$ tours after the last potential delivery time by $\pi_i^{IL|r}(k)$. We obtain

$$\pi_i^{IL|r}(k) = \sum_{l=\max(s_i+1,k)}^{S_i} \pi_i^{IP}(l) \times \sum_{j=(L_g+r)\cdot c_{\min}}^{(L_g+r)c_{\max}} P_i^{J|j}(k)\ P_c^{L_g+r}(j). \tag{2}$$

With h_i defining the unit holding costs for part i per tour we can derive the expected holding costs per tour as

$$EHC = \sum_{i=1}^{N} \frac{h_i \sum\limits_{r=0}^{R_{w(i)}-1} \sum\limits_{k=1}^{S_i} k\ \pi_i^{IL|r}(k)}{R_{w(i)}}. \tag{3}$$

We incur fixed order costs f_i for every order of part i and fixed delivery costs of F_g for every shipment from supplier g. Let P_i^o and $P_i^{o|j}$ denote the probability that an order for part i is placed at review time, regardless of the number of customers during the last review period and given that j customers have been visited respectively.

$$P_i^{o|j} = \sum_{l=s_i+1}^{S_i} \pi_i^{IP}(l) \times \sum_{k=l-s_i}^{j\ d_{\max}} P_i^{J|j}(k) \tag{4}$$

$$P_i^o = \sum_{j=R_g c_{\min}}^{R_g c_{\max}} P_i^{o|j}\ p_c^t(j) \tag{5}$$

Whenever at least one part sourced from supplier g needs to be replenished at review time a delivery from that supplier to the service technician is initiated. The chance P_g^D that a delivery from supplier g is triggered at review time is

$$P_g^D = \sum_{j=R_g c_{\min}}^{R_g c_{\max}} \left(1 - \prod_{i|w(i)=g} \left(1 - P_i^{o|j}\right)\right)\ P_c^{R_g}(j). \tag{6}$$

Using (5)–(6) we can determine the expected fixed order costs and the expected delivery costs per tour as follows

$$EOC = \sum_{i=1}^{N} \frac{f_i P_i^o}{R_{w(i)}} \tag{7}$$

$$EDC = \sum_{g=1}^{G} \frac{F_g P_g^D}{R_g}. \tag{8}$$

We aim to minimize the sum of holding, order, and delivery costs per tour subject to a job-fill-rate (JFR) constraint. The JFR is defined as the average probability that a repair job can be completed with the spare parts available in the repair kit. As shown earlier the availability of different parts varies significantly across tours within a review period. Let us again use the potential delivery time as a reference point. Then the joint availability of multiple parts depends on the combination of the number of tours that elapsed since the last potential delivery times for each part. The number of possible combinations is given by the least common multiple of the different review periods $R_g, g = 1, \ldots, G$. Thus, we need to consider the average JFR across an observation period of $\text{lcm}(R_1, \ldots, R_G)$ tours to cover all possible cases. Let us assume synchronized potential deliveries at the start of the observation period. Then the chance that job $j = 1, \ldots, c_{\max}$ in tour $T = 0, \ldots \text{lcm}(R_1, \ldots, R_G) - 1$ is completed is

$$p_j^{c,T} = \begin{cases} \prod_{i=1}^{N} \left[p_i(0) + \sum_{l=1}^{S_i} \pi_i^{IL|T \bmod R_{w(i)}}(l) \sum_{k=0}^{l} p_i(k) \right], & j = 1 \\ \prod_{i=1}^{N} \left[p_i(0) + \sum_{l=1}^{S_i} \sum_{n=0}^{S_i - l} \sum_{m=l+n}^{S_i} p_i(l) P_i^{J|j-1}(n) \pi_i^{IL|T \bmod R_{w(i)}}(m) \right], & j > 1. \end{cases} \tag{9}$$

Dividing the expected number of completed jobs by the expected number of jobs during the observation period we obtain the JFR.

$$JFR = \frac{\sum_{T=0}^{\text{lcm}(R_1,\ldots,R_G)-1} \sum_{j=1}^{J} p_j^{c,T} \sum_{l=\max(j, c_{\min})}^{c_{\max}} p_c(l)}{\text{lcm}(R_1, \ldots, R_G) \sum_{l=c_{\min}}^{c_{\max}} l \, p_c(l)}. \tag{10}$$

3 Heuristic Solution Approach

The multi-job RKP presented in Sect. 2 is an integer optimization problem with $2 \cdot N + G$ decision variables. Just like other multi-job problems [1–4] our problem cannot be solved to optimality for real-life-sized repair kits. Even for a fixed combination of review periods $(R_1, \ldots, R_G)$ optimal reorder and order-up-to-levels can only be determined for very small numbers of parts N. For larger problems, the solution must be determined heuristically. To this end, we propose to adapt [4]'s job heuristic (JH) to account for delivery costs and review periods different from one. Their JH calculates reorder and order-up-to-levels jointly. That means in a first step the differences $Q_i := S_i - s_i, i = 1, \ldots, N$ are determined. Given these fixed differences a greedy algorithm is applied to iteratively increase stock levels until the target JFR is reached. Differing from [4] we suggest to determine the quantities Q_i in the following way: Let us first assume constant and continuous demand for each spare part. For a given review period of R_g the time between to consecutive orders for each part with $w(i) = g$ must be $m_i \cdot R_g$, where m_i is an integer multiplier. Under these conditions [6] derives the optimal integer multiplier as

$$m_i^* = \left\lceil \frac{1}{2}\sqrt{1 + \frac{8f_i}{h_i D_i R_g^2}} - \frac{1}{2} \right\rceil, \tag{11}$$

where D_i is the demand for part i per time unit. Because demand is stochastic in our case we consider $\mathrm{E}[D_i]$. We set the quantity Q_i equal to the expected demand during $m_i^* \cdot R_g$ tours, but at least 1 unit.

$$Q_i = \max(\lfloor m_i^* R_g \mathrm{E}[D_i] \rceil, 1) \tag{12}$$

With the quantities Q_i calculated as described above we apply [4]'s greedy algorithm to determine s and S, obviously using the formulas for the JFR and the expected holding costs derived in Sect. 2. That way we obtain a heuristic solution for a given combination of review periods. For small numbers of suppliers, we can repeat this procedure for all reasonable combinations of review periods, to identify the best one. For instances with 500 parts and 3 suppliers we were able to run this procedure in less than 30 min on a Mac Pro 7.1 with an Intel 24-Core Xeon W-processor.

4 Numerical Experiment

The point of this experiment is to compare a situation in which orders are placed after every tour (as considered by [4]) to situations in which the review period is longer than just one day. We assume there is only one supplier from which all spare parts are sourced. Orders are delivered after a lead time of 2 tours. We consider a repair kit that consists of 500 parts. At most one unit of each part may be required by one customer. The unit demand probabilities for each part are drawn from a continuous uniform distribution on $[0, 0.01408]$. The number of customers per tour ranges from 1 to 3 with the following probabilities: 1: 25%, 2: 70%, 3: 5%. The value of each part is drawn from a log-normal distribution with a mean of 55€ and a standard deviation of 154€. Both, demand and value scenario have been designed to resemble the characteristics of the dataset described by [4]. The fixed order costs are set to 1€ for each part and each order. Using the algorithm described in Sect. 3, we determine reorder and order-up-to levels for all review periods from 1 (daily) to 5 (weekly) given different combinations of fixed delivery cost, holding cost rate, and target job fill rate.

Table 1 shows the review period lengths that yield the least total costs (in brackets) and the corresponding relative cost decreases compared to situations with daily reviews. For most combinations, weekly reviews perform best. Only in case holding cost rates are high and delivery costs are low, it can be beneficial to review the inventory more frequently. In these cases, longer review periods would lead to an increase in holding costs that outweighs any further savings on delivery costs. Across all scenarios, the average possible cost savings are 31.55%, which shows that considering the review period length can lead to substantial benefits. These savings can be attributed to lower delivery costs with longer review periods. This is offset by only a slight increase in holding costs. We find that it takes surprisingly little additional safety stock to counteract the increased demand uncertainty caused by an extended review period. The average total number and the average total value of all units in the repair kit increase by less than 20%, when reviews are conducted weekly rather than daily in all cases.

Table 1. Relative cost savings in % comparing the best review period length (in brackets) to the daily review option

	Fixed delivery costs								
	5			10			20		
	Holding cost rate (in % per year)								
	5	10	20	5	10	20	5	10	20
JFR = 0.8	33 (5)	26 (5)	16 (5)	47 (5)	40 (5)	29 (5)	59 (5)	53 (5)	44 (5)
JFR = 0.95	26 (5)	17 (5)	9 (4)	40 (5)	31 (5)	20 (5)	54 (5)	45 (5)	34 (5)
JFR = 0.99	22 (5)	12 (4)	5 (3)	36 (5)	24 (5)	14 (4)	50 (5)	40 (5)	27 (5)

5 Conclusions

We presented an extension to the multi-job RKP that integrates fixed costs per delivery and treats the review period length as a decision variable. With a numerical experiment, we could demonstrate that in many cases weekly replenishments should be favored over daily replenishments of a service technician. That way delivery costs can be reduced significantly against only a small increase in safety stocks. The latter increase is so moderate that most likely a van suitable for a daily restocked repair kit will also fit the weekly replenished kit.

References

1. Heeremans, D., Gelders, L.: Multiple period repair kit problem with a job completion criterion: a case study. Eur. J. Oper. Res. **81**, 239–248 (1995)

2. Teunter, R.H.: The multiple-job repair kit problem. Eur. J. Oper. Res. **175**, 1103–1116 (2006)
3. Bijvank, M., Koole, G., Vis, I.F.A.: Optimising a general repair kit problem with a service constraint. Eur. J. Oper. Res. **204**, 76–85 (2010)
4. Prak, D., Saccani, N., Syntetos, A., Teunter, R., Visintin, F.: The repair kit problem with positive replenishment lead times and fixed ordering costs. Eur. J. Oper. Res. **261**, 893–902 (2017)
5. Saccani, N., Visintin, F., Mansini, R., Colombo, M.: Improving spare parts management for field services: a model and a case study for the repair kit problem. IMA J. Manage. Math. **28**, 185–204 (2016)
6. Wildeman, R.E., Frenk, J.B.G., Dekker, R.: An efficient optimal solution method for the joint replenishment problem. Eur. J. Oper. Res. **99**, 433–44 (1997)
7. Viswanathan, S.: Note. Periodic review (s, S) policies for joint replenishment inventory systems. Manage. Sci. **43**(10), 1447–1454 (1997)
8. Khouja, M., Goyal, S.: A review of the joint replenishment problem literature: 1989–2005. Eur. J. Oper. Res. **186**, 1–16 (2008)

III

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

First published as *Rippe, C., Kiesmüller, G.P. (2023). The added value of advance demand information for the planning of a repair kit. Central European Journal of Operations Research, 31, 311-335. https://doi.org/10.1007/s10100-022-00809-1* by Springer Nature under under a Creative Commons Attribution 4.0 International License (CC BY 4.0).

Central European Journal of Operations Research (2023) 31:311–335
https://doi.org/10.1007/s10100-022-00809-1

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

Christoph Rippe[1] · **Gudrun P. Kiesmüller**[2]

Accepted: 26 June 2022 / Published online: 25 July 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 $\{1, \ldots, M\}$. 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 $s = (s_1, \ldots, s_N)$ and $S = (S_1, \ldots, S_N)$
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 $s = (s_1, \ldots, s_N)$ and $S = (S_1, \ldots, S_N)$
γ^*	Target job fill rate
$R(a)$	Ratio trading off cost increase and job fill rate increase if $a \in \{s_i, S_i \mid i = 1, \ldots, N\}$ 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 $i = 1, 2, \ldots, N$ as:

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

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 $i,\ i = 1, \ldots, N$, still in the repair kit after visiting all customers of the tour and material handling costs f_i for each part type $i,\ i = 1, \ldots, N$, 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 $i,\ i = 1, \ldots, N_1$ 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

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

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

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

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 $i, i = N_1 + 1, \ldots, N$ 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

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

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

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

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 $s = (s_1, \ldots, s_N)$ and the order-up-to levels $S = (S_1, \ldots, S_N)$.

We aim to minimize the expected costs per tour across all parts $i, i = 1, \ldots, N$. 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 $i, i = 1, \ldots, N$ per tour when the policy parameters (s_i, S_i)

are implemented. We obtain

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

Then the expected total costs per tour can be calculated as

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

The service level we are interested in is the job fill rate $\gamma(s, S)$ (with $s = (s_1, \ldots, s_N)$ and $S = (S_1, \ldots, S_N)$), 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

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

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, $(i \in \{1, \ldots, N\})$ 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 $J_1, \ldots, J_M$ where $J_l = 1,\ (l \in \{1, \ldots, M\})$ 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, $(l \in \{1, \ldots, M\})$ 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 $\pi_i(k)$ $(i = 1, \ldots, N_1,\ k = s_i + 1, \ldots, S_i)$. Then p_l^{NP} can be determined as follows

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

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 $\pi_i(k)$, $(i = N_1 + 1, \ldots, N,\ k = s_i + 1, \ldots, S_i)$ 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.

$$\begin{aligned} p_l^{AP} &= \sum_{j=l}^{M} P(J=j|J\geq l)\cdot \prod_{i=N_1+1}^{N} P\left(D_i(l,l)\leq \max\left(I_i^+ - D_i(1,l-1),0\right)|J=j\right) \\ &= \sum_{j=l}^{M} \frac{P(J=j)}{P(J\geq l)}\cdot \prod_{i=N_1+1}^{N} P\left(D_i(l,l)\leq \max\left(I_i^+ - D_i(1,l-1),0\right)|J=j\right). \end{aligned} \tag{12}$$

The term $P\left(D_i(l,l)\leq \max\left(I_i^+ - D_i(1,l-1),0\right)|J=j\right)$ 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:

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

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 $\pi_i(k)$, ($k=s_i+1,\ldots,S_i$) and the probabilities $P(D_i(l_1,l_2)=m)$, $(1\leq l_1\leq l_2\leq M,\ m=1,\ldots,M\cdot d_i^{max})$ 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 $l_1 > l_2$.

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

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

4.2 Cost function

For all parts $i \in (1, \ldots, N)$ 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, $i = 1, \ldots, N_1$, 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

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

In deriving the material handling cost incurred for a part $i, \ i = 1, \ldots, N_1$, 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

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

4.2.2 Cost function for APs

For the computation of the holding cost incurred for an AP $i, i = N_1 + 1, \ldots, N$ 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

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

For the material handling costs we derive

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

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 $s = (s_1, \ldots, s_N)$ and order-up-to-levels $S = (S_1, \ldots, S_N)$ 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

$$Q_i := S_i - s_i = \max\left(\left\lfloor\sqrt{\frac{2f_i\overline{D_i}}{h_i}}\right\rceil, 1\right), \quad i = 1, 2, \ldots, N \tag{20}$$

where $\overline{D}_i$ is the average demand for part i for one tour and set $s_i = -Q_i$ and $S_i = 0$ for all $i \in \{1, \ldots, N\}$. 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 $a \in \{s_i, S_i | i = 1, \ldots, N\}$ is increased by one unit. Then R is defined by

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

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 ($d_i^{\max} = 1,\ i = 1, \ldots, N$). 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 $l = 1, \ldots, J$ 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 ;500] 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

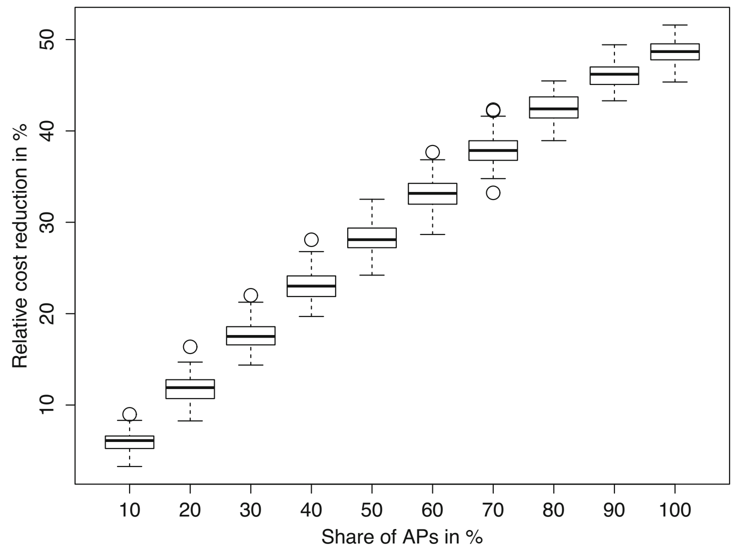

Fig. 1 The relative total cost reduction due to ADI given a target JFR of 90%

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

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

for $x \in \{10, 20, \ldots, 90, 100\}$. 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.

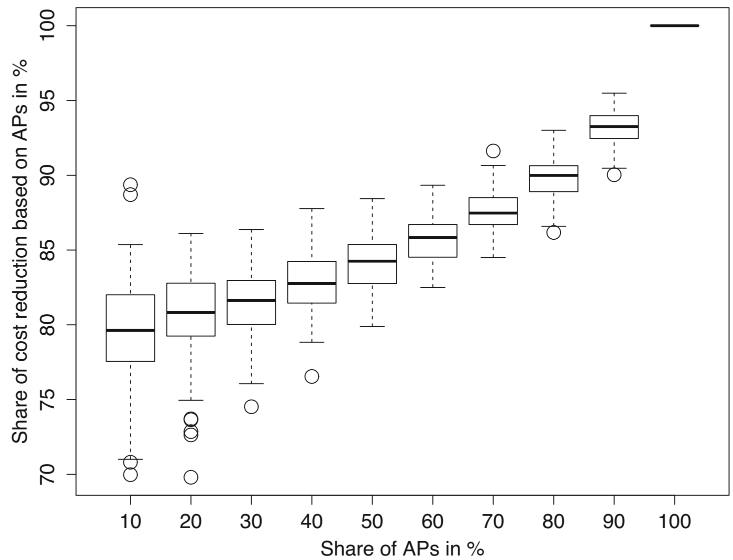

Fig. 2 Share of total cost reduction based on APs

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 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

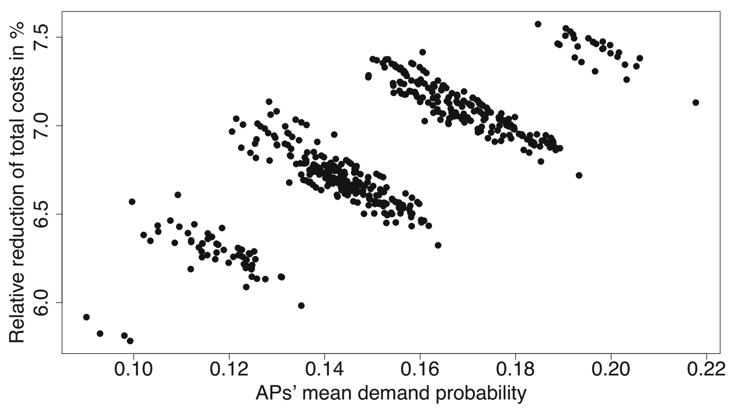

Fig. 3 Impact of the APs' mean demand probability

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.

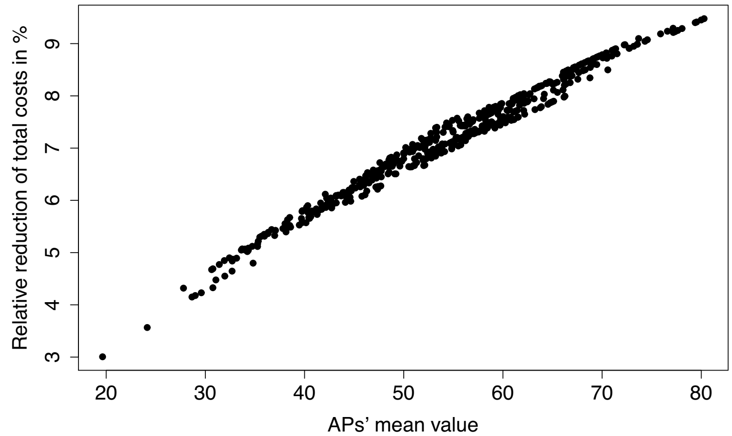

Fig. 4 Impact of the APs' mean price

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 availibility 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.

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 $\pi_i(k),\ k = s_i + 1, \ldots, S_i$ 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 $k = s_i + 1, \ldots, S_i$ at some point within one order cycle has to be determined. We denote this probability by $\Pi_i(k),\ k = s_i + 1, \ldots, S_i$. 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 $\Pi_i(k),\ k = s_i + 1, \ldots, S_i$ can be calculated recursively starting from $k = S_i$ with the following formula

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

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 $k = S_i - 1, \ldots, s_i + 1$ 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 $P\left(D_i(l,l) \leq \max\left(I_i^+ - D_i(1,l-1), 0\right) | J = j\right)$

We can determine the probability $P\left(D_i(l,l) \leq \max\left(I_i^+ - D_i(1,l-1), 0\right) | J = j\right)$ that sufficient units of part i are available for the l-th out of j jobs in total as the sum of $P\left(I_i^+ - D_i(1,l) \geq 0 | J = j\right)$ and $P(I_i^+ - D_i(1,l) < 0, D_i(l,l) = 0 | J = j)$. $P\left(I_i^+ - D_i(1,l) \geq 0 | J = j\right)$ 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 $P\left(I_i^+ - D_i(1,l) \geq 0 | J = j\right)$. 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 $\pi_i(k),\ k = s_i + 1, \ldots, S_i$ defines the steady-state probabilities of the net inventory level before order placement. Using this we obtain

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

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.

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

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 $J = 1$.

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

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

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

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 $P\left(I_i^+ - D_i(1,l) \geq 0 \mid J = j\right)$ (see (2) and (26)). Thus, we get

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

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 $P(D_i(l,l)=d)$, $\forall d \in \left\{0,\dots,d_i^{\max}\right\}$, $i \in \{1,\dots,N\}$ and $P(D_i=d)$, $\forall d \in \left\{0,\dots,d_i^{\max}\cdot M\right\}$, $i \in \{1,\dots,N\}$, target job fill rate γ^*

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

$$\overline{D}_i \leftarrow \sum_{d=0}^{M\cdot d_i^{\max}} d \sum_{j=1}^{M} P(D_i(1,j)=d)\cdot P(J=j),\ \forall i \in \{1,\dots,N\}$$
$$Q_i \leftarrow \max\left(\left\lfloor \sqrt{\frac{2 f_i \overline{D}_i}{h_i}} \right\rceil, 1\right),\ \forall i \in \{1,\dots,N\}$$

$s \leftarrow -Q, \qquad S \leftarrow 0^N$
order $\leftarrow \emptyset$
while $\gamma(s,S) < \gamma^*$ **do**
  $R(s_i) \leftarrow \Delta_{s_i} C(s,S)\cdot\left(\Delta_{s_i}\gamma(s,S)+\epsilon\right)^{-\mathrm{sgn}(\Delta_{s_i} C(s,S))},\ \forall i \in \{1,\dots,N\}$
  $R(S_i) \leftarrow \Delta_{S_i} C(s,S)\cdot\left(\Delta_{S_i}\gamma(s,S)+\epsilon\right)^{-\mathrm{sgn}(\Delta_{S_i} C(s,S))},\ \forall i \in \{1,\dots,N\}$
  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 length(order) **do**
  if order[j] $\neq$order[k], $\forall k \in \{1,\dots,j-1\}$ **then**
    $i \leftarrow$ 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
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

IV
The repair kit problem with imperfect advance demand information

First published as *Rippe, C., Kiesmüller, G.P. (2023). The repair kit problem with imperfect advance demand information. European Journal of Operational Research, 304, 558-576* *https://doi.org/10.1016/j.ejor.2022.04.019* by Elsevier.

European Journal of Operational Research 304 (2023) 558–576

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Production, Manufacturing, Transportation and Logistics

The repair kit problem with imperfect advance demand information

Christoph Rippe[a,*], Gudrun P. Kiesmüller[b]

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

[b] *Technical University of Munich, TUM School of Management, TUM Campus Heilbronn, Bildungscampus 9, 74076 Heilbronn, Germany*

ARTICLE INFO

Article history:
Received 22 December 2020
Accepted 14 April 2022
Available online 21 April 2022

Keywords:
Inventory
Repair kit
Spare parts
Advance demand information
Multiple jobs

ABSTRACT

To promote customer loyalty and generate revenue in after-sales service, many companies provide field repair services to their customers. Service technicians who perform these repair jobs typically carry a set of spare parts called a repair kit in their company van. The repair kit problem aims to determine which parts to include in this kit and in what quantity. Currently, many appliances are equipped with sensors that monitor their different functionalities. If an appliance breaks down because one of these functionalities is disturbed, then the respective sensor triggers a failure code that describes the appliance's condition. From the service technician's perspective, this failure code serves as potentially imperfect advance demand information for spare parts. In this paper, we present an extension of the repair kit problem that uses this information for the replenishment decision. We formulate this repair kit problem with advance demand information as a Markov decision process and propose two heuristic solution procedures. Our first heuristic is far-sighted and optimizes the inventory of all parts individually, while our second heuristic is a myopic greedy algorithm that considers all parts at once. We conduct an extensive numerical study to evaluate the performance of both heuristics and to identify which heuristic performs best under which circumstances. Comparing both heuristics to a state-of-the-art algorithm that disregards any available advance demand information, we find that utilizing this information yields substantial cost savings.

1. Introduction

Spare parts provisioning and technical assistance are key features of after-sales service. Since after-sales service accounts for 40–50% of all profits of original equipment manufacturers (Holmströem, Cheikhrouhou, Farine, & Främling, 2011), optimizing maintenance operations has a major impact on the overall company performance. For heavy and bulky appliances, such as printers, heating systems, medical systems, or white goods such as washing machines or refrigerators, services are typically performed on-site. That means that customers with malfunctioning appliances are visited by service technicians who attempt to repair the appliance using a set of spare parts that are carried in a company van. This set of spare parts is commonly referred to as the repair kit. The repair kit problem is to determine which spare parts to include in the repair kit and if so, in which quantity. If the technician visits several customers in a repair tour before restocking, which is usual for appliances with reasonable complexity, we speak of a multi-job repair kit problem.

To date, almost all contributions regarding the repair kit problem assume that the service technician has no information about the condition of a customer's failed appliance prior to a first visit to the customer. This assumption, however, is often unrealistic because many modern appliances are fitted with sensors that display failure codes if they detect an error. When a customer contacts the service provider to request a repair job, he/she is likely to disclose this information. Thus, the failure code serves as advance demand information (ADI) for spare parts from the technician's point of view. In particular, this ADI can be uncertain, because the fault discovered by the sensor system (e.g., abnormal temperatures, pressure, or vibration) may originate from malfunctions of different components.

We study a version of the multi-job repair kit problem that incorporates this uncertain ADI. We assume that the service technician adjusts the content of the repair kit on a daily basis to match the presumed needs of the customers taking into account the failure codes they disclose. Thus, the repair kit problem with ADI becomes a dynamic problem. Similar to most previous contributions, we assume the technician may receive instantaneous re-

* Corresponding author.
E-mail addresses: christoph.rippe@ovgu.de (C. Rippe), gudrun.kiesmueller@tum.de (G.P. Kiesmüller).

https://doi.org/10.1016/j.ejor.2022.04.019
0377-2217/

plenishments from a warehouse with ample supply, which in practice corresponds to overnight deliveries of goods. Further, we allow for the technician to return spare parts to that warehouse at additional costs in order to clear excess stock that may have been built up due to the uncertain nature of the failure codes. When deciding about replenishments and returns, the service technician must trade-off expected inventory holding costs against penalty costs incurred for unsuccessful repair attempts.

The contributions of this paper are as follows: First, this is the first time that the repair kit problem with imperfect ADI is studied. We formulate the problem as a Markov decision process (MDP) and show that the state space of this MDP is increasing exponentially in the number of parts, which makes it impossible to solve for reasonably sized problem instances. Second, we decompose the multi-part multi-tour problem in two different ways, by part and by tour. For the single-part multi-tour problem we can show the structure of the optimal policy and for the multi-part single-tour problem we demonstrate that the cost function is submodular. Motivated by these structural results for the two decompositions we develop two different heuristic solution approaches, a part and a greedy heuristic, for the original repair kit problem with ADI. Third, we conduct an extensive numerical study to compare the performance of both heuristics under various circumstances. Additionally, we benchmark both ADI heuristics against the state-of-the-art job heuristic suggested by Teunter (2006) that disregards ADI. In this way, we demonstrate the potential for cost savings that can be achieved if information available about the condition of malfunctioning appliances is employed by the service technician. Fourth, we show how to adjust our greedy heuristic to determine solutions for repair kit problems with an additional capacity constraint.

The remainder of this paper is structured as follows. In Section 2, we review the relevant literature. In Sections 3 and 4, we present the problem formulation and the heuristic solution methods. The numerical study used to test these heuristics is described in Section 5. In Section 6, we consider the impact of an additional capacity constraint. Finally, in Section 7, we summarize our main findings.

2. Background

The repair kit problem with ADI is at the interface of two streams of literature, namely, the contributions on the repair kit problem and the (imperfect) advance demand information literature. In this section, we will provide a brief overview of papers from both fields and show how our work fits into the context of the problems studied before.

The repair kit problem was first introduced by Smith, Chambers, & Shlifer (1980) as the problem of determining the cost-minimizing set of spare parts to be carried by a service technician who performs repair jobs in the field. After every repair job, the service technician's repair kit is restocked making this a so-called single-job problem. To derive the optimal repair kit, Smith et al. (1980) employ an unconstrained cost model that weighs inventory holding costs for parts held in the repair kit against penalty costs incurred for failed repair attempts. While several parts might be needed to complete a repair job, only one unit per part may be needed and demands for different parts are independent. Adopting these assumptions, Graves (1982) and Hausman (1982) consider a service model with a job fill rate constraint that ensures that the chance to complete a repair job is above a given threshold.

Further extensions to the single-job repair kit problem incorporate dependent demands and demands for multiple units by one customer (Mamer & Smith, 1982; March & Scudder, 1984), spare machines (Mamer & Smith, 1985) and budget constraints (Mamer & Shogan, 1987). Brumelle & Granot (1993) derive a monotone sequence of optimal repair kits for different parametrizations of their cost model's objective function.

Heeremans & Gelders (1995) are the first ones to study a multiple-job repair kit problem, where the service technician performs several repair jobs before restocking. In contrast to the single-job papers, however, they consider a tour fill rate instead of a job fill rate meaning that they regard the probability that **all** jobs in a tour can be completed rather than the chance that an arbitrary job is completed. Teunter (2006) is the first one to derive a job fill rate formula for the multiple-job repair kit problem. He develops a very general formulation for both, a cost and a service model. Nevertheless, for the calculation of the job fill rate, he requires independent demands for different parts and at most one unit of each part may be needed for each job. He also assumes that parts are left with the first customer who needs them regardless of whether this customer's job can be completed. To determine the repair kit, two greedy heuristics are proposed, a job heuristic that makes use of the job fill rate and a simpler part heuristic that is based on a part fill rate. Following Teunter (2006), three further extensions to the multiple-job repair kit problem have been presented by Bijvank, Koole, & Vis (2010), Saccani, Visintin, Mansini, & Colombo (2017), and Prak, Saccani, Syntetos, Teunter, & Visintin (2017). Bijvank et al. (2010) assume a repair process in which available spare parts are not left with customers for a second visit if their repair jobs cannot be completed at the first try. Saccani et al. (2017) solve an integer linear model to determine the repair kit that would have performed best given past demand data and assess its future performance based on simulation. Prak et al. (2017) consider a non-zero replenishment lead time for spare parts and manage the contents of the repair kit using individual (s,S)-policies for the different spare parts.

All repair kit papers discussed up to this point assume that the condition of the malfunctioning appliance in need of a repair is unknown to the service technician before his/her first visit to the customer. All customers are taken to be homogeneous and historical demand data aggregated across all past customers are used to decide which parts to put into the repair kit. In contrast to this, we include information about the source of the failure in our model, which is obtained through sensors and translated into a failure code. That means, that not only historical data are used to estimate spare part requirements, but also real-time information about the condition of the failed system. As a consequence, the customers' spare part demand in the model cannot be assumed to be identically distributed, and the inventory levels in the repair kit have to be adapted dynamically. The only contribution so far that considers ADI for the repair kit problem is Rippe & Kiesmüller (2020), who assume part-specific and perfectly reliable sensors. This paper, however, assumes that sensors monitor functionalities rather than individual parts, which implies that ADI obtained from these sensors is imperfect. We extend the model described in the seminal multi-job, repair kit paper by Teunter (2006), adopting most of his assumptions and adjusting it to incorporate imperfect ADI. As a consequence, the inventory policy also has to be adapted and therefore, we allow spare parts to be returned to the warehouse, which is a completely new aspect of the repair kit problem. Since the companies we have spoken with only apply overnight deliveries, we have decided to assume a lead time of zero because this reflects the situation of the companies. Table 1 illustrates that our study is the first to integrate uncertain real-time information (ADI) and returns in a repair kit model.

As mentioned above, our research is also related to the stream of literature discussing inventory models with advance demand information. In this stream of literature, information on the customer demand is revealed to a decision-maker before the actual demand is due, which is called advance demand information. One possibility to model the advance demand information is the so-called

Table 1
Literature overview part 1: repair kit problem.

	single item	multiple item	single job	multiple jobs	zero lead time	positive lead time	with returns	without returns	imperfect ADI	perfect ADI	no ADI
Smith et al. (1980)		✓	✓		✓			✓			✓
Graves (1982)		✓	✓		✓			✓			✓
Hausman (1982)		✓	✓		✓			✓			✓
Mamer & Smith (1982)		✓	✓		✓			✓			✓
March & Scudder (1984)		✓	✓		✓			✓			✓
Mamer & Smith (1985)		✓	✓		✓			✓			✓
Mamer & Shogan (1987)		✓	✓		✓			✓			✓
Brumelle & Granot (1993)		✓	✓		✓			✓			✓
Heeremans & Gelders (1995)		✓		✓	✓			✓			✓
Teunter (2006)		✓		✓	✓			✓			✓
Bijvank et al. (2010)		✓		✓	✓			✓			✓
Saccani et al. (2017)		✓		✓	✓			✓			✓
Prak et al. (2017)		✓		✓		✓		✓			✓
Rippe & Kiesmüller (2020)		✓		✓		✓		✓		✓	
our paper		✓		✓	✓		✓		✓		

"demand lead time" introduced by Hariharan & Zipkin (1995) and defined as the time span between information disclosure and the demand due date. For a continuous review inventory model where customers place firm orders to be delivered at a given due date, they demonstrate that demand lead time and supply lead time have opposing effects on the system performance. Gallego & Özer (2001) consider a similar problem with firm advance orders in discrete-time and find that the optimal replenishment policies are state-dependent where the state reflects the currently known orders for future time periods.

While the two seminal papers mentioned above study stochastic inventory models with perfect advance demand information, there also exists another stream of literature where information about demand is imperfect. Since the ADI available in our case is imperfect we limit the further discussion of ADI literature to publications on imperfect ADI. For further research on models incorporating perfect advance demand information, we refer to the detailed review of Karaesmen (2013).

In several contributions, the added value of imperfect advance demand information in production/inventory systems is considered. Inventory replenishment orders are placed at a production facility with limited capacity, and it has to be decided when to place an order and how many units to keep as safety stock. For continuous review production/inventory systems, Liberopoulos, Chronis, & Koukoumialos (2003), Liberopoulos & Koukoumialos (2008) and Claudio & Krishnamurthy (2009) assume that customers place advance orders that may be canceled at a later point in time. Gayon, Benjaafar, & de Véricourt (2009) consider a joint production and inventory allocation problem with different customer classes, where some classes order in advance. However, the demand due dates are subject to changes, and again, orders might be canceled altogether. Similar problems are studied by Benjaafar, Cooper, & Mardan (2011) and Kim, Ahn, & Righter (2009), who assume that a supplier obtains status updates on a customer's product in a stochastic multi-stage production process in order to deliver a particular component exactly when it is needed. In a simulation study, Fischer, Benzaman, Diegel, Gimenez, & Claudio (2020) compare the impact of different types of uncertainties (cancelations, changing due dates, and order quantities) for a three-stage production system with exponential processing times. Benbitour & Sahin (2015) conduct a similar study for a single-stage production system with limited manufacturing capacities in discrete time. Another contribution on a periodic review problem with capacitated resources is examined by Gao, Xu, & Ball (2012) who model both demand quantity and due date with Markov chains. All these studies have in common that they only consider a single item. Additionally, the replenishment lead time is load dependent due to the fixed production capacity.

In our model, we assume that replenishment orders are placed at a warehouse that has ample supply of parts, and therefore, the replenishment lead time can be assumed to be constant. In the following contributions, a similar assumption is made such that they are more closely related to our problem. DeCroix & Mookerjee (1997) and Zhu & Thonemann (2004) study problems where retailers may purchase potentially imperfect ADI in the form of customer demand forecasts before making their replenishment decision. van Donselaar, Rock Kopczak, & Wouters (2001) and Thonemann (2002) consider project environments in the construction industry where manufacturers receive tentative orders from installers bidding for different projects that are confirmed and specified if a bid is accepted. Tan, Güllü, & Erkip (2007) and Huang & Van Mieghem (2014) both examine situations in which prospective customers identified by sales representatives or clickstream data, respectively, are tracked until they either place an order or leave the system. A two-period model that uses imperfect ADI for replenishment and rationing decisions with different priority customer classes is developed by Tan, Güllü, & Erkip (2009).

The papers closest to ours are those that either use imperfect ADI for the inventory management of spare parts (Topan, Tan, van Houtum, & Dekker, 2018; Zhu, Jaarsveld, & Dekker, 2020) or consider multi-product systems (Bernstein & DeCroix, 2015; Chen, Şafak Yücel, & Zhu, 2017; Thonemann, 2002). In Topan et al. (2018), ADI comes as a warning signal for the imminent failure of a particular component that may or may not occur in the following periods. In the event that expected failures do not happen, Topan et al. (2018) allow for returns of excess stock to an upstream supplier. In Zhu et al. (2020), the maintenance plan serves as ADI for spare parts, in so far that the number of inspections of the system containing a certain part is known in advance. Whether that part actually needs to be replaced is only revealed upon inspection, but the probability for a spare part demand can be derived from historic data. While Topan et al. (2018) and Zhu et al. (2020) consider, similar as in our problem setting, an inventory model in the context of maintenance planning, Chen et al. (2017) study a video rental system, and Bernstein & DeCroix (2015) a production planning problem. The latter two contributions have in common that they also allow multiple items. However, Bernstein & DeCroix (2015) examine a single period model and decide about production capacities. In Chen et al. (2017) an item-specific ADI is modeled and cus-

Table 2
Literature overview part 2: imperfect advance demand information.

	single item	multiple items	simple jobs	complex jobs	load-independent lead time	load-dependent lead time	continuous review	periodic review	with returns	without returns
Liberopoulos et al. (2003)	✓		✓			✓	✓			✓
Liberopoulos & Koukoumialos (2008)	✓		✓		✓	✓	✓			✓
Claudio & Krishnamurthy (2009)	✓		✓			✓	✓			✓
Gayon et al. (2009)	✓		✓			✓	✓			✓
Benjaafar et al. (2011)	✓		✓			✓	✓			✓
Kim et al. (2009)	✓		✓			✓	✓			✓
Fischer et al. (2020)	✓		✓			✓	✓			✓
Benbitour & Sahin (2015)	✓		✓			✓		✓		✓
Gao et al. (2012)	✓		✓			✓		✓		✓
DeCroix & Mookerjee (1997)	✓		✓		✓			✓		✓
Zhu & Thonemann (2004)	✓		✓		✓			✓		✓
van Donselaar et al. (2001)	✓		✓		✓			✓		✓
Thonemann (2002)		✓	✓		✓			✓	✓	
Tan et al. (2007)	✓		✓		✓			✓		✓
Huang & Van Mieghem (2014)	✓		✓		✓			✓		✓
Tan et al. (2009)	✓		✓		✓			✓		✓
Topan et al. (2018)	✓		✓		✓			✓	✓	
Zhu et al. (2020)	✓		✓		✓			✓		✓
Chen et al. (2017)	✓	✓	✓		✓			✓	✓	
Bernstein & DeCroix (2015)		✓	✓		✓			✓		✓
our paper		✓		✓	✓			✓	✓	

tomers keep online queues that indicate the items they intend to rent in the future.

Our contribution belongs to the class of multi-item inventory models under periodic review with imperfect ADI and with constant lead time. As illustrated in Table 2, there are only a few studies where inventory models with multiple items are investigated. However, our multi-item inventory model differs substantially from these contributions, because of the demand process and how we model ADI. Since in Thonemann (2002) and Chen et al. (2017) the customer is only allowed to select one of the multiple items, we call this a simple job, while we consider complex jobs where several items may be demanded by one customer. This demand model is induced by the characteristics of the repair jobs, because for one repair job several items may be needed. Thus, to measure service, we use the job fill rate and not an item-specific service level as in Thonemann (2002). A pure cost model is considered in Bernstein & DeCroix (2015), who also use a different model for the advance demand information. They only allow one demand signal, either for the total demand volume across products or for the demand mix between products. In our model, each customer observes a different failure code and thus communicates an individual demand signal. Additionally, we consider a multi-period model while Bernstein & DeCroix (2015) study a single-period model. Then, due to the longer time horizon and the low demand rates, it can also be economically beneficial to return items to the warehouse.

To summarize, we combine the overviews in Tables 1 and 2 and conclude, that our study is the first contribution where the repair kit problem is extended to allow for imperfect ADI. Additionally, we extend the literature on stochastic periodic inventory models with imperfect ADI by simultaneously considering several customers' demand signals for complex repair jobs that may require multiple spare parts. All these different problem aspects require a new model and novel solution approaches, which are presented in the following.

3. Problem description

For each repair tour, we consider the following sequence of events. First, customers are allocated to tours based on geographical locations. The information available on the customers' failure codes is passed on to the technician who adjusts the repair kit to the customers' presumed needs. Next, the technician visits his/her customers one by one, attempting to repair their appliances with spare parts carried in the repair kit. Costs attributed to the repair tour are charged after the last visit to a customer.

The number of customers in a repair tour is modeled with a random variable C. We denote the minimum and the maximum number of customers in a tour by $C_{\min}$ and $C_{\max}$. The probability that $c \in \{C_{\min}, \ldots, C_{\max}\}$ customers need to be visited in one repair tour is given by $P(C = c)$. Usually, 4 to 6 customers can be visited in one tour.

Each of the C customers' broken appliances may display one out of M different failure codes that signal different types of errors. If communicated to the service technician before the first repair attempt, these failure codes are essentially ADI for spare parts. However, not every error can be detected by a sensor and not every customer is aware of the failure code that is displayed. In these cases, a dummy failure code is assigned to the customer in question. In this way, the total number of failure codes is increased to $M + 1$. Let $F = (F_1, \ldots, F_C)$ denote the vector of failure codes observed by the C customers in the tour. The probability to observe a specific failure code $P(F_j = f_j)$ $\left(f_j = 1, \ldots, M + 1\right)$ at customer j is identical and independent for all customers $j = 1, \ldots, C$.

The service technician can carry N different spare parts on his van, which might be required to perform a repair job. For ease of calculations, we make the slightly simplifying assumption that only one unit of each part may be required by each customer. Let $D_{j,n} \in \{0, 1\}$ $(j = 1, \ldots, C,\ n = 1, \ldots, N)$ denote the jth customer's demand for part n. The conditional demand probabilities $P(D_{j,n} = 1 | F_j = f_j)$ $\left(f_j = 1, \ldots, M + 1,\ j = 1, \ldots, C\right)$ for each part $n = 1, \ldots, N$ given the jth customer's failure code f_j are known to the service technician. In practice, these conditional demand probabilities can be determined on the basis of previous experience with the different failure codes. For that purpose, the service technician may either use his own experience or the results obtained from a durability test conducted by the original equipment manufacturer. The conditional demand probability for any part n is the same for all customers, and the demands of different customers for the same part as well as the requirements of one customer for different parts are independent. Under these conditions and given the failure code vector f, the aggregated demand $D_n = \sum_{j=1}^{C} D_{j,n}$ for a part $n = 1, \ldots, N$ follows a Poisson binomial distribution, which is a generalization of the standard binomial distribution to situations with different probabilities of success in each trial. In our case, a repair job corresponds to a trial and a success to a demand for part n. The different probabilities of success are the conditional demand probabilities $P(D_{j,n} = 1 | F_j = f_j)$. Let $D = (D_1, \ldots, D_N)$ denote the total demand vector across all parts.

Given the number of customers in the upcoming tour and their failure codes, the service technician modifies the content of the repair kit to meet the presumed needs of the customers. All modifications come into effect immediately. We denote the content of the repair kit before ordering with $X = (X_1, \ldots, X_N)$, and after the arrival of orders and the return of units but before the start of the tour by $Y = (Y_1, \ldots, Y_N)$. For each part n the technician can update the inventory level to any value $Y_n \geq 0$, where $Y_n > X_n$ corresponds to a replenishment and $Y_n < X_n$ to a return. In this way, the technician can react to both, anticipated stockouts and redundant stock. In case $Y_n = X_n$, the inventory level of part n is not altered at all.

Next, the service technician visits the customers one by one to perform on-site repairs. To fix the malfunctioning appliances, the technician identifies broken components and replaces them with spare parts from the repair kit. If all required spare parts are available, we call the repair job complete and broken otherwise. Following Teunter (2006), we assume that all spare parts are left with the first customer who needs them regardless of whether that customer's repair job can be completed.

Once the service technician returns from the repair tour, costs incurred in or attributed to that tour are determined. We consider three different types of costs. First, part-specific return costs r_n are incurred for each unit of any part $n = 1, \ldots, N$ returned from the repair kit to the supplier. These return costs arise for essentially wasteful material handling operations, such as testing, labeling, and storing the returned part, that do not contribute to the fulfillment of customer requests. Unlike other handling operations, these actions could have been avoided if the service technician had not built up excess stock. Second, the service technician incurs part-specific holding costs h_n for each unit of part n held in the repair kit after the last customer has been visited for all parts $n \in \{1, \ldots, N\}$. Third and last, the service technician incurs penalty costs P for each broken job. These penalty costs include the actual costs for a second visit to the repair site and a penalty for the loss of customer goodwill.

The repair kit problem with ADI as formulated above can be modeled as a Markov decision process (MDP) with states (X, C, F) defined by the repair kit X before orders and returns, the number of customers C to be served in the upcoming tour and their error codes F. If we have a maximum of $C_{\max}$ customers, it makes sense to assume that the inventory level X_n of any part $n = 1, \ldots, N$ lies between $-C_{\max}$, which means that all customers required the part that was not available at all, and $C_{\max}$, which means that we expected every customer to require the part but none of them actually did. That means that $X \in \{-C_{\max}, \ldots, C_{\max}\}^N$. As defined before, the number of customers to be served ranges from $C_{\min}$ to $C_{\max}$ and each customer may communicate one out of $M+1$ failure codes. Thus, for a given number of customers C, we obtain $F \in \{1, \ldots, M+1\}^C$.

Our overall objective is to find a stationary policy π^* that minimizes the long-run average cost per tour over an infinite time horizon. The existence of such an optimal stationary policy is guaranteed for average-cost MDPs with finite state and action spaces, which holds in our case.

To derive the long-run average costs, let us first consider a single tour. Given a state (x, c, f) and new inventory levels y, the calculation of return and expected holding costs for the upcoming tour is straightforward. To compute the expected penalty costs, we need to determine the probability that a randomly selected job can be completed, which is determined by the job fill rate (JFR) introduced by Teunter (2006). However, since ADI and repair kits differ from tour to tour in our model, we do not derive a JFR across all tours but only for the next tour ahead. This next-tour-JFR incorporates the available ADI and can be calculated as follows

$$JFR(c, f, y) = \frac{1}{c} \sum_{j=1}^{c} \left(\prod_{n=0}^{N} \left(1 - P\left(D_{j,n} = 1 \middle| F_j = f_j \right) \right) + P\left(D_{j,n} = 1 \middle| F_j = f_j \right) \cdot P\left(\sum_{l=1}^{j-1} D_{l,n} < y_n \middle| F = f \right) \right). \quad (1)$$

For any combination of customer $j = 1, \ldots, C$ and part $n = 1, \ldots, N$, the first summand describes the probability that no unit of part n is required, while the second summand describes the probability that part n is required and still available for the jth customer's repair job. Let $V_1(x, c, f, y)$ denote the expected single-tour costs given a state (x, c, f) and post-modification inventory levels y. Using the JFR-formula we obtain

$$V_1(x, c, f, y) = \sum_{n=1}^{N} \left[r_n \cdot (x_n - y_n)^+ + h_n \mathrm{E}\left[(y_n - D_n)^+ | F = f \right] \right] + P \cdot c \cdot [1 - JFR(c, f, y)]. \quad (2)$$

Let π be a stationary policy that maps each state (x, c, f) to a new repair kit y. Then, $V_1(x, c, f, \pi)$ denotes the expected single-tour costs if the action in state (x, c, f) is determined by policy π. In a similar way, we define $V_t(x, c, f, \pi)$ as the total expected cost in t consecutive tours when the initial state is (x, c, f) and the action at the beginning of each tour is determined by policy π. For any $t > 1$, we can calculate the expected costs in t tours recursively. Let $\mathfrak{D}(f) = \{d | P(D = d | F = f) > 0\}$ define the set of possible demand events given ADI-vector f and $\mathfrak{F}(c) = \{f = (f_1, \ldots, f_c) | f_i \in \{1, \ldots, M+1\}, i \in \{1, \ldots c\}\}$ describe the set of all possible ADI-vectors for c customers; then, $V_t(x, c, f, \pi)$ can be derived as follows

$$V_t(x, c, f, \pi) = V_1(x, c, f, \pi) + \sum_{d \in \mathfrak{D}(f)} \sum_{c'=C_{\min}}^{C_{\max}} \sum_{f' \in \mathfrak{F}(c')} \left(V_{t-1}\left(x + \pi(x, c, f) - d, c', f', \pi \right) \cdot P\left(F' = f' | C' = c' \right) \cdot P\left(C' = c' \right) \cdot P(D = d | F = f) \right). \quad \forall t > 1 \quad (3)$$

The long-run average cost per tour when the system is controlled by policy π and the initial state is (x, c, f) is then given by the limit

$$\lim_{t \to \infty} \frac{V_t(x, c, f, \pi)}{t}. \quad (4)$$

Since the Markov chain that describes the content of the repair kit and the available ADI at the beginning of a tour for any given policy π is unichain (all states can be reached from each other), the long-run average cost does not depend on the initial state. Thus, we can express the optimization problem as

$$\min_{\pi} \lim_{t\to\infty} \frac{V_t(\pi)}{t}. \text{ Table 3} \tag{5}$$

Table 3
Notations.

Input parameters	
$C_{\min}$	Minimum number of jobs in a tour
$C_{\max}$	Maximum number of jobs in a tour
M	Number of different genuine failure codes
N	Number of different spare parts
P	Penalty cost per broken job
h_n	Unit holding cost for spare part n
r_n	Unit return cost for spare part n
$P(C=c)$	Probability that the number of jobs in a tour is c
$P(F_j=f_j)$	Probability that costumer j communicates failure code f_j
$P(D_{j,n}=1\|F_j=f_j)$	Conditional demand probability for part n given failure code f_j
Stochastic variables	
C	Number of jobs in a tour
$F=(F_1,\ldots,F_C)$	Failure codes communicated by the customers to be served
$D_{j,n}$	Demand of customer j for spare part n
D_n	Total Demand for part n in a tour
$D=(D_1,\ldots,D_N)$	Total demand vector
$X=(X_1,\ldots X_N)$	Net inventory level of the spare parts in the repair kit before orders or returns
Other notations	
$Y=(Y_1,\ldots,Y_N)$	Adjusted content of the repair kit after orders and returns before the start of the tour
$JFR(c,f,y)$	Job fill rate for the next tour given the information c,f and the inventory y
$V_t(i,c,f,\pi)$	Total expected cost in t tours starting in state (i,c,f) and applying policy π

4. Heuristic solution methods

To demonstrate the need for heuristic solution methods to the MDP described above, let us first determine the size of its state space. As discussed in the previous section, each part's inventory level before orders or returns can take any value between $-C_{\max}$ and $C_{\max}$. With N different parts, there are $(2C_{\max}+1)^N$ states that the repair kit can be in before orders and returns. Considering that each of the customers in the following tour observes one out $M+1$ failure codes, we obtain a state space of size $(2\cdot C_{\max}+1)^N \cdot \sum_{c=C_{\min}}^{C_{\max}} (M+1)^c$. Even for a moderate size example with a fixed number of 6 customers, 20 parts, and 10 error codes, the size of the state space is $1.9\cdot 10^{28}$. Because of this enormous number of possible states, the repair kit problem with advance demand information can only be solved to optimality for very small problem instances. The methods that can be employed to determine the optimal policy for an MDP, such as value or policy iteration, require the iterative update of value function estimates for all states. Even a single update, however, can be computationally intractable for a large state space. To overcome this difficulty, we develop two different heuristics that determine solutions to the repair kit problem with ADI by decomposing it into either single-part or single-tour problems.

4.1. Part heuristic

One of the main reasons why the repair kit problem with ADI cannot simply be solved using value iteration is that we have to optimize the replenishments of all spare parts simultaneously. That is why the size of the state space increases exponentially in the number of spare parts. The first heuristic we propose is a part heuristic that addresses this issue by decomposing the repair kit problem with ADI into several smaller problems for each of the spare parts considered. These individual replenishment problems have much smaller state spaces than the joint optimization problem.

With regard to holding and return costs, the decomposition of the repair kit problem into single-part problems is straightforward as these types of costs are part-specific in the original problem anyway. Instead of penalty cost incurred for an incomplete job, however, we need to consider part-specific penalty costs incurred for every stock-out. Naturally, the penalty costs for an incomplete repair job should be attributed to the spare parts whose stockouts have caused the broken job. Thus, with only one part missing, the penalty cost P incurred for the broken job should be attributed entirely to that missing part. While several parts might be required to perform a repair job with a well-equipped repair kit, it is very likely that only one of them is missing if the job cannot be completed. Due to this property, we assume that the penalty costs per missing part required in the individual optimization problems are equal to the penalty costs per broken job considered in the original joint optimization problem. In this way, we can decompose the original MDP considering N parts simultaneously into N MDPs where only one part each is optimized.

For the N individual MDPs, the state spaces that need to be considered are significantly smaller than the state space of the joint optimization MDP due to two reasons. First and most obviously, a state no longer has to reflect the inventory position of all spare parts at the same time. A second aspect that helps us to reduce the state space even further is that we do not need to consider the individual customers' demands $D_{j,n},\ j=1,\ldots,C$ for a part n anymore but only the aggregate demand D_n across all customers in the tour. That is because the single-part MDPs do not require the calculation of a job fill rate that explicitly considers individual customers' demands (see equation (1)). As a consequence, it is possible to prove the optimal policy structure for a single part as we will show below.

Instead of using the vector F to describe the available ADI, it is sufficient to keep track of the numbers of customers that have disclosed the different types of failure codes. Let us denote the distribution of failure codes among customers by $B=(B_1,\ldots,B_{M+1})$ with

$$B_l = \left|\left\{ j\in\{1,\ldots,C\} \middle| F_j = l\right\}\right| \quad l=1,\ldots,M+1. \tag{6}$$

Given the number of customers C in the upcoming tour, the new information vector B follows a multinomial distribution with C trials and the set of different failure codes $\{1,\ldots,M+1\}$ describing the possible outcomes in each trial. Thus, we obtain

$$P(B=b|C=c) = \begin{cases} \frac{c!}{b_1!\cdot\ldots\cdot b_{M+1}!}\cdot\prod_{l=1}^{M+1} P(F_j=l)^{b_l} & \text{if } \sum_{l=1}^{M+1} b_l = c \\ 0 & \text{otherwise,} \end{cases} \tag{7}$$

for any $j\in\{1,\ldots,c\}$ since the probability to observe a particular failure code is identical for all customers. The conditional distribution of the aggregated demand D_n given the information vector B still follows a Poisson binomial distribution. However, instead of listing the different success probabilities for a sequence of trials, we count the number of trials that are conducted with the different probabilities. Using the new information vector B, the state of the MDP specific to part n can be described by (X_n, C, B).

For each part n, the part-specific long-run average costs per tour can be derived similarly to the joint approach considering all parts at once. Given a particular state (x_n, c, b) and a part-specific policy π_n, we define the single-tour expected costs $V_{n,1}$ attributed to part n by

$$V_{n,1}(x_n, c, b, \pi_n) = r_n \cdot (x_n - \pi_n(x_n, c, b))^+ + h_n \cdot \mathrm{E}\Big[(\pi_n(x_n, c, b) - D_n)^+ | B = b\Big] + P \cdot \mathrm{E}\Big[(D_n - \pi_n(x_n, c, b))^+ | B = b\Big]. \quad (8)$$

Let $\mathfrak{D}_n(b) = \{d_n | P(D_n = d_n | B = b) > 0\}$ define the set of all demand events with positive probabilities for part n given the information vector b. Further, let $\mathfrak{B}(c) = \left\{b \in \{1, \ldots, c\}^{M+1} \middle| \sum_{l=1}^{M+1} b_l = c\right\}$ describe the set of all condensed information vectors for exactly c customers. Then, again given an initial state (x_n, c, b) and a policy π_n, the total expected cost over t tours $V_{n,t}$ attributable to part n can be defined recursively by

$$V_{n,t}(x_n, c, b, \pi_n) = V_{n,1}(x_n, c, b, \pi_n) + \sum_{d \in \mathfrak{D}_n(b)} \sum_{c'=C_{\min}}^{C_{\max}} \sum_{b' \in \mathfrak{B}(c')} \Big(V_{n,t-1}\big(\pi_n(x_n, c, b) - d_n, c', b', \pi_n\big) \cdot P\big(B' = b' | C' = c'\big) \cdot P\big(C' = c'\big) \cdot P(D_n = d_n | B = b)\Big). \quad (9)$$

The part-specific average costs per tour can then be expressed as

$$\lim_{t \to \infty} \frac{V_{n,t}(\pi_n)}{t}, \quad (10)$$

where we can disregard the initial state as it does not impact the system in the long run. For the single-part average cost minimization problem we can characterize the optimal policy π_n^* as follows.

Theorem 1. *The optimal replenishment policy π_n^* that minimizes the average costs per tour for part n is a state-dependent (L, U)-policy. That means for any combination (c, b) of customers c to be visited and condensed information vectors b, there are two thresholds $L_n(c, b)$ and $U_n(c, b)$ such that*

$$\pi_n^*(x_n, c, b) = \begin{cases} L_n(c, b) & \text{if} \quad x_n < L_n(c, b) \\ x_n & \text{if} \quad L_n(c, b) \le x_n \le U_n(c, b) \\ U_n(c, b) & \text{if} \quad U_n(c, b) < x_n. \end{cases} \quad (11)$$

Chen et al. (2017) derived a similar policy structure for a finite-time horizon total cost minimization problem in the context of a closed-loop rental system. Although our problem is slightly different, the same line of arguments can be followed to prove Theorem 1 (see Appendix). Even though we know the structure of the optimal policy, we cannot derive the state-dependent thresholds analytically. However, since the state space of the part-specific MDP is much smaller than that of the joint optimization problem we can employ value iteration to determine these thresholds. This way we obtain individually optimal policies π_n^* for each part $n \in \{1, \ldots, N\}$. If we combine these policies π_n^*, we get a (potentially sub-optimal) policy $\pi = \left(\pi_1^*, \ldots, \pi_N^*\right)$ for the original joint optimization problem (5).

4.2. 2-Step greedy algorithm

As a second alternative solution method, we propose a 2-step greedy algorithm that is a combination of two different greedy algorithms. The 2-step greedy algorithm is in several ways the exact opposite of the part heuristic explained in Section 4.1. While the part heuristic optimizes the inventory levels for all spare parts individually, the 2-step greedy algorithm attempts to put together all parts' inventory levels simultaneously. However, the 2-step greedy algorithm is myopic, in so far that we are not looking beyond the current tour, whereas the part heuristic is considering the future implications of today's stocking decisions. Finally, using the part heuristic, we pre-compute actions to be chosen for all possible states before we start the very first repair tour. In contrast to this, the 2-step greedy heuristic uses decision time planning, which means that actions are only determined once a particular state is actually encountered.

Immediately before the start of each tour, we optimize the content of the repair kit regarding only the immediate return costs as well as the expected penalty and holding cost incurred for the upcoming tour. Thus, we are solely considering the single-tour cost function $V_1(x, c, f, y)$.

Even for this single-tour problem, the optimal solution can only be found for small-scale problem instances. That is because we still need to determine the JFR which depends on the inventory levels of **all** parts at the same time. For that reason, we cannot optimize the inventory levels y_n of the different parts $n \in \{1, \ldots, N\}$ individually. To derive near-optimal solutions we propose a greedy heuristic, which is a common approach in the repair kit literature (see, e.g., Bijvank et al., 2010; Teunter, 2006, and Prak et al., 2017). This greedy heuristic iteratively adds units of different parts to the repair kit or removes them from it. This increase or decrease of one part's inventory level has an impact on the further change in costs that is realized with additional changes to the repair kit. In particular let us consider the way in which an increase of one part's inventory level y_n affects the marginal change of costs we induce by adding yet another unit of either the exact same part n or one of the other parts $m \in \{1, \ldots, N\} \setminus \{n\}$.

Surprisingly, we find that when we successively add two units of one particular part, the marginal cost reduction achieved with the first unit can be smaller than that obtained with the second. That means the cost function does not need to be convex in the post-modification inventory level y_n of any part $n \in \{1, \ldots, N\}$. Consider the following minimal counterexample: We examine a single-period problem with 2 parts, 2 jobs, and 2 failure codes. The conditional demand probabilities are defined in Table 4 and the customers' failure codes are given by $f = (1, 2)$. We increase y_1 while the remaining repair kit is defined by $x_1 = x_2 = y_2 = 0$. Figure 1 shows that the expected number of failed repair attempts is not convex in y_1. The marginal penalty cost reduction of the first unit added is less than that of the second unit because the first unit is likely to be used for a job that cannot be completed whereas adding a second unit guarantees the completion of the second job. Thus the penalty costs are not convex y_1, which means that the total costs need not be convex in y_1. This property is unique to repair kit problems with customers that are heterogeneous with respect to their demand probabilities. It shows that

Table 4
Conditional demand probabilities.

	failure code 1	failure code 2
part 1	0.75	0.75
part 2	1	0

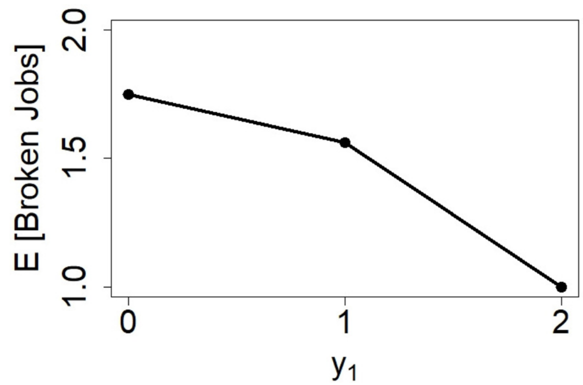

Fig. 1. Expected number of failed repair attempts.

even for given inventory levels of all other parts optimizing a single part's inventory level y_n is not straightforward.

To characterize how the increase of one part's inventory level affects the marginal cost reductions that can be achieved with an additional increase of another parts' inventory level we build upon the concept of submodularity, which is defined as follows:

Definition 1. Let S define a partially ordered set of vectors. For $s, t \in S$ let $s \wedge t$ and $s \vee t$ define the component-wise minimum and maximum respectively. A partially ordered set S that contains $s \wedge t$ and $s \vee t$, $\forall s, t \in S$ is called a lattice. A function $f : S \to \mathbb{R}$ defined on a lattice S is called submodular if

$$f(s) + f(t) \geq f(s \wedge t) + f(s \vee t), \quad \forall s, t \in S. \tag{12}$$

From Corollary 2.6.1 of Topkis (1998) we can derive that a function f defined on a sublattice of $\mathbb{Z}^N$ is submodular if and only if it has decreasing differences that is if

$$\Delta_{s_k}\Delta_{s_l} f := f(s + e_k + e_l) - f(s + e_l) - f(s + e_k) + f(s) \leq 0 \qquad \forall s \in S,\ \forall k \neq l \in \{1, \dots, N\}, \tag{13}$$

where e_j is the jth standard unit vector. Using this more intuitive characterization we can prove that the cost function is submodular in the inventory levels. In fact, we can show by induction that not only the single-period cost function is submodular but also the t-tour cost function. We obtain the following stronger result. (A proof is provided in the Appendix.)

Theorem 2. *Let $V_t^*(x, c, f, y)$ define the t-tour cost function that gives the expected costs if y is implemented in state (x, c, f) and an optimal policy is followed thereafter for $t-1$ tours. Thus, $V_1^*(x, c, f, y) = V_1(x, c, f, y)$ and for $t \geq 2$ the cost function is defined recursively by*

$$V_t^*(x, c, f, y) = V_1(x, c, f, y) + \sum_{d \in \mathfrak{D}(f)} \sum_{c'=C_{\min}}^{C_{\max}} \sum_{f' \in \mathfrak{F}(c')} \left(\min_{y'} V_{t-1}^*\left(y - d, c', f', y'\right) \cdot P\left(F' = f' | C' = c'\right) \cdot P\left(C' = c'\right) \cdot P(D = d | F = f) \right) \tag{14}$$

Let $\mathfrak{X} := \{-C_{\max}, \dots, C_{\max}\}^N$ and $\mathfrak{Y} := \{0, \dots, C_{\max}\}^N$ denote the sets of all possible inventory level combinations before and after orders. Then $V_t^*(x, c, f, y)$ is submodular in x and y on $\mathfrak{X} \times \mathfrak{Y}$ for any $c \in \{C_{\min}, \dots, C_{\max}\}$ and $f \in \mathfrak{F}(c)$.

The submodularity of the cost function essentially means that increasing one part's inventory level by one unit makes an increase of the other parts' inventory levels more favorable (or at least less unfavorable.) The same is true for the impact of the reduction of a parts' inventory level on the marginal effect of the reduction of other part's inventory levels. Because of this property, it makes sense to consider the reduction and the increase of several parts' inventory levels consecutively. For that reason, our greedy heuristic consists of two distinct steps. First, a reduction heuristic is used to strip the current repair kit of any units unlikely to contribute to a fulfilled repair request in the upcoming tour. Then, we continue with a re-equipment algorithm that raises the inventory levels of the parts likely to be needed.

The reduction heuristic (Algorithm 1) successively identifies spare parts for which the isolated removal of a unit currently held in the repair kit results in a reduction of the expected costs. In each iteration, the unit that yields the largest expected cost reduction is returned from the repair kit to the supplier until no further cost reduction can be achieved.

Once we have run the reduction heuristic, we take the resulting repair kit as an initial solution for the second greedy algorithm (Algorithm 2) that identifies the spare parts that should be reordered as well as the respective reorder quantities. In each step of this re-equipment algorithm, we raise the inventory level of exactly one part by one unit. The part selected is the part with the highest job fill rate increase to holding cost increase ratio. Though this might seem counter-intuitive, we also consider those parts for which the reduction algorithm suggests that we reduce the inventory level. This is because the first algorithm might suggest returning units of a part that are likely to be needed for jobs in the next tour if due to the lack of some other parts in the repair kit, the

Algorithm 1 Reduction heuristic.

Input: current repair kit x, number of customers c, error codes f, conditional demand probabilities $P(D_{j,n} = 1 | F_j = f_j)\ \forall n \in \{1, \dots, N\}$, $f_j \in \{1, \dots, M+1\}$, unit holding costs $h_n\ \forall n \in \{1, \dots, N\}$, unit return cost $r_n\ \forall n \in \{1, \dots, N\}$, penalty costs P

Output: new inventory levels y

$y \leftarrow x$
stop $\leftarrow 0$
while stop $= 0$ **do**
  for $n \in \{1, \dots, N\}$ with $y_n > 0$ **do**
    $y^{temp} \leftarrow y$
    $y_n^{temp} \leftarrow y_n^{temp} - 1$
    $\Delta_n \leftarrow V_1(x, c, f, y) - V_1\left(x, c, f, y^{temp}\right)$
  end for
  $n^* \leftarrow \arg\max_{n: y_n > 0} \Delta_n$
  if $\Delta_{n^*} > 0$ **then**
    $y_n \leftarrow y_n - 1$
  else
    stop $\leftarrow 1$
  end if
  if $\{n | y_n > 0\} = \emptyset$ **then**
    stop $\leftarrow 1$
  end if
end while

Algorithm 2 Re-equipment heuristic.

Input: old repair kit x, current repair kit y as determined by reduction heuristic, number of customers c, error codes f, conditional demand probabilities $P(D_{j,n} = 1 | F_j = f_j)\ \forall n \in \{1, \dots, N\}$, $f_j \in \{1, \dots, M+1\}$, unit holding costs $h_n\ \forall n \in \{1, \dots, N\}$, unit return cost $r_n\ \forall n \in \{1, \dots, N\}$, penalty costs P

Output: new inventory levels y^{opt}

$y^{opt} \leftarrow y$
stop $\leftarrow 0$
while stop$=0$ **do**
  for $n \in \{1, \dots, N\}$ with $y_n < c$ **do**
    $y^{temp} \leftarrow y$
    $y_n^{temp} \leftarrow y_n + 1$
    $\Delta_n \leftarrow \frac{JFR(c,f,y^{temp}) - JFR(c,f,y)}{h_n\left(\mathrm{E}\left[\left(y_n^{temp} - D_n\right)^+ \middle| F=f\right] - \mathrm{E}\left[(y_n - D_n)^+ \middle| F=f\right]\right) - r_n \cdot \mathbb{1}_{\{y_n < x_n\}}}$
  end for
  $n^* \leftarrow \arg\max_{n: y_n < c} \Delta_n$
  $y_{n^*} \leftarrow y_{n^*} + 1$
  if $V_1(x, c, f, y) < V_1(x, c, f, y^{opt})$ **then**
    $y^{opt} \leftarrow y$
  end if
  if $y_n = c\ \forall n \in \{1, \dots, N\}$ **then**
    stop$\leftarrow 1$
  end if
end while

chance that these jobs can be completed is very small. In this case, the reduction algorithm trades a small increase in penalty costs for a larger decrease in holding costs. While this reduction makes sense when the inventory levels of all parts cannot be increased (which is the situation the reduction algorithm is based on), we are likely to obtain a more favorable repair kit if we raise the inventory levels of all parts towards the expected demands. That is why we allow the re-equipment algorithm to consider all parts. If we increase a part's inventory level with the re-equipment algorithm after we have first reduced it with the reduction algorithm, we have to bear in mind that this is reducing the return costs we incur. To take this effect into account, we subtract any potential return cost reductions from the holding cost increase considered when deciding upon the spare part to be selected in each iteration.

5. Numerical study

In this section, we examine the performance of our heuristic solution procedures using a small and a large-scale numerical experiment. For both experiments, the problem instances we examine are designed such that the heuristic solutions achieve realistic JFR service levels above 70%. First, we evaluate the solutions found with part and greedy heuristic for small problem instances. For these instances, it is possible to determine the exact expected costs per tour for any given solution as well as the optimal solution. We benchmark the solutions found with part and greedy algorithm against the optimal solution on the one hand and against a solution determined with the job heuristic by Teunter (2006) that disregards advance demand information.

Second, we study the performance of the heuristic solution procedures for real-life-sized repair kits in a simulation study. That means we compare the solutions obtained with the different heuristics in various settings that involve a larger number of failure codes and parts considered for the repair kit. As the state space grows exponentially in the number of parts in the repair kit we cannot calculate the exact costs incurred for a given policy or even determine the optimal solution for large-scale problem instances. For this reason, we compare the solutions provided by the ADI heuristics to each other and to the non-ADI benchmark. In particular, we look in detail at the composition of the total costs incurred with the different heuristics and characterize their mechanics.

5.1. Small-scale experiment

For this experiment, we examine problem instances with 6 different parts, 3 different failure codes, and at most 3 customers per tour, because this enables us to determine the optimal solution. That way we can determine the optimality gap of the solutions obtained with our novel solution procedures. Additionally, we aim to quantify the value of considering advance demand information for inventory management. To this end, we compare the performance of ADI-sensitive repair kits updated on a daily basis to the performance of a one-fits-all repair kit that does not take ADI into account. This one-fits-all repair kit is determined with the state-of-the-art job heuristic suggested by Teunter (2006) for repair kit problems where advance demand information is either not available or not used. In particular, our problem differs from that studied by Teunter in three aspects: The availability of advance demand information, the stochasticity of customer numbers, and the time at which holding costs are incurred. Since we want to study the impact of the availability of advance demand information exclusively, we have to adjust Teunter's problem and thus his solution procedure with respect to the latter two characteristics. To this end, we first extend Teunter's job fill rate formula to allow for stochastic numbers of customers. That means we are using the following average JFR formula

$$\mathrm{JFR} = \frac{\sum_{c=C_{\min}}^{C_{\max}} c \cdot P(C=c) \cdot \mathrm{JFR}(c)}{\sum_{c=C_{\min}}^{C_{\max}} c \cdot P(C=c)}, \tag{15}$$

where $\mathrm{JFR}(c)$ is Teunter's JFR formula for a fixed number of c customers. Second, when we apply Teunter's job heuristic we charge holding costs only for units left in the repair kit at the end of the tour. With these changes in place the difference between our heuristics and the benchmark heuristic is solely due to the use of available ADI. While our ADI heuristics consider conditional demand probabilities for each part given each failure code, the benchmark heuristic uses average demand probabilities across all failure codes. Let us denote the average demand probabilities by p_n $n \in \{1, \dots, N\}$. Using the law of total probability, we obtain

$$p_n = \sum_{l=1}^{M+1} P\big(D_{j,n} = 1 | F_j = l\big) \cdot P\big(F_j = l\big) \quad \text{for any customer } j. \tag{16}$$

Based on these average demand probabilities the slightly adjusted job heuristic computes expected penalty and holding costs and weights them against each other to determine ADI-independent order-up-to levels for each spare part. The combination of these order-up-to levels characterizes the one-fits-all repair kit, that we compare to ADI-sensitive repair kits. Note that returns and thus return costs need not be considered for a one-fits-all repair kit that does not change from tour to tour.

For the small-scale experiment, we consider 10 different demand and holding cost scenarios that are combined with three different values for both return and holding cost parameters. Altogether that gives us 90 different instances. For each scenario, the repair kit can be equipped with 6 different parts. The number of customers C follows a discrete uniform distribution on $\{1, 2, 3\}$. Each customer communicates one out of three failure codes where one of them is a dummy failure code. The probability that no failure code is communicated is assumed to be 30% for all scenarios. The remaining 70% are divided randomly among the two genuine failure codes. That is the chance to observe the first genuine failure code is drawn from a continuous uniform distribution on [0,0.7]. The conditional demand probabilities for the different failure codes are again drawn from continuous uniform distributions. For the genuine failure codes that is the continuous uniform distribution on [0.1,0.5] and for the dummy failure code the support is $[0.2, 0.3]$. Finally different unit holding costs for each part and each scenario are drawn from a continuous uniform distribution on [0,1]. Each of these 10 different demand and holding cost scenarios are combined with unit return costs $r = r_n$ for all $n \in \{1, \dots, N\}$ of $0.5, 1$ and 2 and penalty cost rates of $5, 10$ and 15, to give us in total 90 different problem instances.

For each of the problem instances, we derive the optimal solution and compare it to the solutions found with the ADI heuristics and the non-ADI benchmark. Table 5 shows the relative cost differences between the heuristic solutions and the optimal solution. The given figures are averages across all 10 different demand and holding cost scenarios. We find that the solutions obtained with the part heuristic are very close to the optimal solutions with an optimality gap of just 0.27% on average and at most 1.68% across all instances. In particular, for instances with larger penalty costs of 10 and 15 this gap is always below 0.5%. In all but two instances, the part-heuristic outperforms the greedy heuristic. Nonetheless, the solutions found with the greedy heuristic are still close to the optimal solutions. The optimality gap is just 1.73% on average and at most 6.13% across all instances.

Both ADI heuristics provide better solutions than the non-ADI benchmark in 85 out of the 90 instances. Across these 85 instances the average relative difference between costs incurred with

Table 5
Relative cost difference with the optimal solution on average across all scenarios (in %).

	r= 0.5			r= 1			r= 2		
	P= 5	P= 10	P= 15	P= 5	P= 10	P= 15	P= 5	P= 10	P= 15
part heuristic	0.65	0.09	0.08	0.58	0.14	0.07	0.62	0.11	0.08
greedy heuristic	1.40	0.99	1.10	2.36	2.11	2.31	2.12	1.52	1.70
non-ADI heuristic	10.59	20.36	21.76	9.34	17.63	18.40	9.09	16.96	17.70

Table 6
Run times on average across all scenarios (in s).

	r= 0.5			r= 1			r= 2		
	P= 5	P= 10	P= 15	P= 5	P= 10	P= 15	P= 5	P= 10	P= 15
optimal solution	694.29	669.05	634.69	833.25	782.40	777.09	1089.49	993.71	971.76
part heuristic	34.79	35.29	34.81	35.00	35.29	35.40	35.70	35.43	35.43
greedy heuristic	96.33	98.53	95.05	92.57	95.91	100.07	95.98	95.53	95.98
non-ADI heuristic	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.07	0.06

the ADI-sensitive repair kit and the ADI-insensitive one-fits-all repair kit is 13.90% and 12.64% for part and greedy heuristic respectively. This difference is the cost savings potential that can be realized with ADI heuristics in case that advance demand information is available. The potential cost savings are smallest for instances with low penalty costs that induce low service levels. That is because in case service levels may be low, the ability to react to customers' failure codes becomes less valuable. Not surprisingly the 5 instances for which the one-fits-all repair kit found with the adjusted job heuristic performs better than at least one of the ADI heuristics' solutions are all characterized by low penalty costs. In these cases, the ADI heuristics increase the service levels above the optimal levels at the price of additional holding and return costs which does not happen with the non-ADI benchmark. Nevertheless, the costs incurred with the ADI solutions are within 5% of the costs of the non-ADI repair kit in all of these 5 atypical instances.

Table 6 shows the average run times of the different solution procedures on a Mac Pro-7.1 with an Intel 24-Core Xeon W-processor. To evaluate the performance of a policy we must explicitly derive decisions for all states and under these circumstances the part heuristic appears to be quicker than the greedy algorithm. Note however that with the greedy algorithm we can determine the decision for one state independent of the decisions for all other states. Thus, with the greedy algorithm we can determine the decision to be taken in a single state (e.g. the state currently encountered) in no time, while the part heuristic requires the solution of an MDP for each part up front.

5.2. Large-scale experiment

In this section, we examine the performance of our ADI heuristics for real-life-sized problem instances with up to 50 parts, 10 failure codes, and up to 6 customers per tour. The numerical experiment can be divided into a basic experiment and three subsequent experiments that examine the sensitivity of the initial results to different factors of influence. Namely, these factors are the stochasticity of customer numbers, the share of customers that provide ADI and the precision of the ADI available. Because of the size of the state space, we cannot evaluate the average costs per tour for a given policy let alone minimize them for large problem instances. However, we can still compare the performance of the ADI heuristics and benchmark them against the ADI-independent repair kit found with Teunter's job heuristic on the basis of estimates of the average costs per tour. To obtain these estimates we apply a sequential sampling procedure following Law & Kelton (2000). Each sample corresponds to a simulation run that comprises 110 repair tours. The initial repair kit before the very first replenishment decision at the start of tour 1 is empty. We treat the first 10 tours as a warm-up period. Thus, the average costs across the remaining 100 tours are an unbiased estimation of the actual average costs per tour. Our overall estimation of the actual costs per tour is given by $S(L) = \frac{1}{L}\sum_{l=1}^{L} S_l$ with S_l denoting the estimate from the lth simulation run. For each problem instance, we iteratively increase the number of simulation runs L until the half-width of the 95% confidence interval of the average costs per tour is within 2% of the actual value for each heuristic. At least 10 simulation runs are performed in any case.

5.2.1. Basic experiment

For the basic experiment, we run the heuristic solution procedures on 900 different problem instances. To create these problem instances, we combine 100 randomly generated demand and holding cost scenarios with 3 different return and penalty cost rates each. For each scenario, we consider 50 different spare parts and a fixed number of 6 customers in each repair tour. The impact of stochastic numbers of customers will be discussed separately in a subsequent experiment. The holding costs h_n for the different parts $n = 1, \ldots N$ are drawn from a uniform distribution on [0,1]. This choice of model parameters is inspired by the set of problem instances generated by Teunter (2006) for his second numerical experiment. Additionally, our demand scenarios, which are explained in the following, are designed to match the mean demand probabilities generated by Teunter on average across all parts.

To define a demand scenario, we must specify the number of different failure codes M, their respective probabilities, and the conditional demand probabilities for each part given each failure code. For each scenario, we distinguish between 9 distinct failure codes and 1 dummy failure code that indicates that no ADI was communicated by the customer. We assume the chance that an arbitrary customer does not provide ADI is 30%. To divide the remaining 70% among the 9 genuine failure codes, we draw 9 random numbers $R_1, \ldots, R_9$ from a uniform distribution on [0,1]. The probability that an arbitrary customer j communicates the genuine failure code i is then generated as

$$P(F_j = i) = 0.7 \cdot \frac{R_i}{\sum_{i=1}^{9} R_i}. \tag{17}$$

The different parts' conditional demand probabilities given the dummy failure code are drawn from a uniform distribution on [0.01,0.3]. For the genuine failure codes, the conditional demand probabilities have to reflect the additional information that is available. To represent the impact of this additional information, we divide the parts into three different sets for each failure code. Parts in the first set are deemed likely to have triggered the given failure

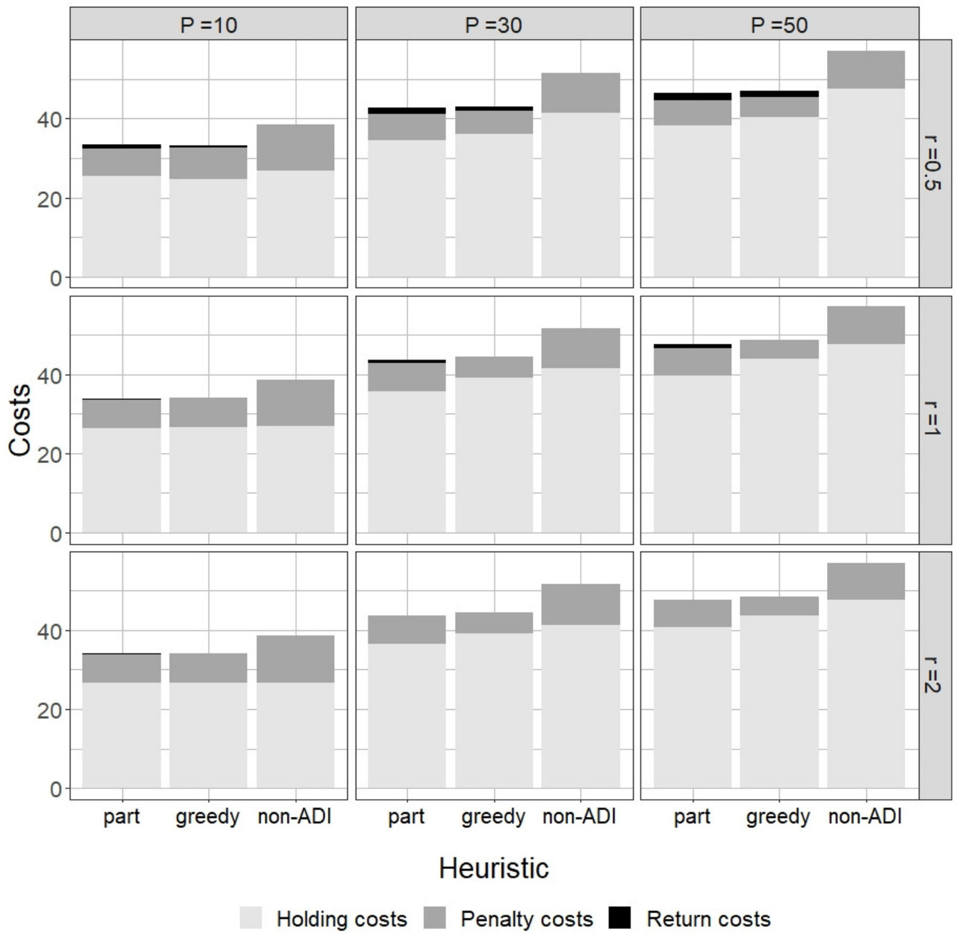

Fig. 2. Costs per tour on average across all 100 scenarios.

Table 7
Demand characteristics for genuine failure codes.

	number of parts	conditional demand probabilities
Set 1	3	drawn from uniform distribution on [0.5,0.9]
Set 2	15	drawn from uniform distribution on [0,0.2]
Set 3	32	0

code. The second set comprises parts that either may have caused the failure in an unlikely event or that could have been damaged by whatever set off the failure code (e.g., electronics damaged by leakage). Parts in the third set are assumed to have no connection to the given failure code at all. A fixed number of spare parts is randomly assigned to each set. The number of parts in each set and the construction of their respective conditional demand probabilities are described in Table 7.

The values considered for the return cost rate $r = r_n$ for all $n \in \{1, \ldots, N\}$ are 0.5, 1, and 2. For the penalty costs, we consider values of 10, 30, and 50.

Applying the simulation-based evaluation procedure described above, we obtain the expected costs per tour for each heuristic and each problem instance. Average values across all 100 demand and holding cost scenarios are displayed in Fig. 2 for each combination of return and penalty cost rate.

Even though the two ADI heuristics manage the content of the repair kit in completely different ways, the resulting replenishment or return decisions and thus the average costs per tour are surprisingly similar. The costs obtained with the part heuristic are within −5% to 5% of the respective values determined with the greedy heuristic for all instances. Table 8 shows the number of scenarios for which we observe the least costs per tour with the part and greedy heuristic. The figures in brackets denote the number of cases for which the differences between both heuristics are significant to an approximate confidence level of 95%. As in the small experiment, the part heuristic outperforms the greedy heuristic in the majority of cases. However, the greedy heuristic provides the more favorable solution for one of 6 instances in this experiment compared to one out of 45 in the small-scale experiment. This suggests that the greedy heuristic which considers all parts simultaneously becomes more competitive when the number of different parts is increased. We find that the greedy-heuristic performs best for instances with small penalty cost rates whereas the part heuristic yields better solutions for instances with large penalty cost rates. The superiority of the part heuristic for instances with large penalty cost rates can be explained by the myopic nature of the greedy heuristic. For large penalty cost rates, the total amount of stock held to avoid broken jobs is increased, which in turn increases the chance that a unit remains in the repair kit over several tours before it is sold. In all of these tours, we incur holding costs for the unsold units. The part heuristic considers this effect but the greedy heuristic does not as it does not look beyond the next tour.

Unlike in the small-scale experiment, the ADI heuristics outperform the non-ADI benchmark in all instances of the large-scale experiment. The reduction of the average cost incurred per tour that can be realized in comparison to this benchmark ranges from 8% to 22% for both ADI heuristics with means of 15.50% and 14.36%

Table 8
Number of scenarios in which the respective heuristic performs best.

	r= 0.5			r= 1			r= 2		
	P= 10	P= 30	P= 50	P= 10	P= 30	P= 50	P= 10	P= 30	P= 50
part heuristic	20 (2)	100 (68)	100 (66)	73 (34)	100 (96)	100 (100)	57 (21)	100 (85)	100 (86)
greedy heuristic	80 (43)	0 (0)	0 (0)	27 (9)	0 (0)	0 (0)	43 (10)	0 (0)	0 (0)

Table 9
MAD between successive start-of-tour stock levels on average across all parts and all scenarios.

	r= 0.5			r= 1			r= 2		
	P= 10	P= 30	P= 50	P= 10	P= 30	P= 50	P= 10	P= 30	P= 50
part heuristic	0.43	0.50	0.53	0.37	0.42	0.44	0.36	0.39	0.40
greedy heuristic	0.42	0.50	0.53	0.38	0.42	0.43	0.38	0.42	0.43

for part and greedy heuristic, respectively. As can be seen in Fig. 2, these cost reductions are due to savings on holding and penalty costs, whereas additional return costs that are not incurred with the benchmark heuristic counteract these savings.

For both ADI heuristics, the composition of the cost reductions depends on the penalty cost rate. For the smallest rate of P= 10, more than 85% of the mean cost reduction can be attributed to penalty cost savings. For a penalty cost rate of P= 50, on the other hand, savings on holdings costs account for more than 50% of the total cost reduction. In the first case, the ADI heuristics increase the average JFR and thus decrease penalty costs at comparable holding costs. In the second case, the high penalty cost rate induces the benchmark heuristic to aim for JFRs that are already very high. Thus, there is little room for a reduction of penalty costs. Instead, the ADI heuristics decrease the total cost per tour primarily by reducing the holding costs against a similar JFR.

The additional return costs that are incurred with the ADI heuristics have only a minor impact on the total costs per tour. For both the part and greedy heuristic, the return costs account for less than 5% of the total costs in each of the 900 problem instances. This shows that the return option is not used excessively in any case. Even for the smallest rate considered of r= 0.5, the average number of units returned per tour is just 2.77 and 2.14 for the part and greedy heuristic, respectively.

Looking at the composition of the repair kits over time, we find that the main advantage of the ADI heuristics is the opportunity to start each tour with a different stock level for each part. To quantify the changing of the start-of-tour stock levels, we derived the mean absolute difference (MAD) between the stock levels of a part at the beginning of successive tours (see Table 9). An MAD of 0.5, for example, means that the stock level of a part is changed by one unit on average every second tour. The stock level may be increased or decreased. A decrease does not necessarily coincide with the return of a unit. It might also correspond to the decision not to replenish a unit that was taken out of the kit in the previous tour. Comparing the values from Table 9 to the corresponding value of 0 for the non-ADI benchmark heuristic, it is easy to see how the flexibility to adjust the repair kit on a tour-to-tour basis results in lower holding and penalty costs.

5.2.2. Stochastic number of customers

In the basic experiment, we have a fixed number of customers who communicate rather precise ADI for the majority of failures (70%). To examine the sensitivity of the results obtained for this setting with respect to stochastic customer numbers, different shares of customers providing ADI, and less precise conditional demand probabilities given a failure code, we conduct three subsequent experiments. For each of these experiments, we adopt the cost settings from the initial experiment. Only the demand settings are adjusted in a different way for each of the subsequent experiments.

Table 10
Number of scenarios in which the respective heuristic performs best.

	part heuristic	greedy heuristic
fixed number of customers	779	121
stochastic number of customers	804	96

For the first additional experiment, we compare a setting with a fixed number of customers per tour (setting 1) to an equivalent setting with stochastic numbers of customers (setting 2). The fixed number of customers for the first setting is 5, while the number of customers to be served in the second setting is drawn from a discrete uniform distribution on $\{4, 5, 6\}$. Thus, the average number of customers to be served is the same in either case. For both settings, we rerun the ADI heuristics as well as the non-ADI benchmark heuristic on all 900 problem instances generated for the basic experiment.

We find that whether a fixed or a stochastic number of customers is considered does not have an impact on which of the ADI heuristics performs best. Table 10 shows the number of instances (out of 900) for which we obtain the least costs with the part and greedy heuristic. For both settings, the part heuristic equally seems to outperform the greedy heuristic for most instances. However, on average across all instances, the costs obtained with the greedy heuristic are just 1.64% (1.84%) above the costs obtained with the part heuristic for setting 1(2).

If we compare the costs obtained with the ADI heuristics to those obtained with the benchmark heuristic, we see that the benefit of using ADI for the replenishment decision is larger when the number of customers is stochastic. Figure 3 shows boxplots of the relative cost reduction that can be obtained with the ADI heuristics in comparison to the non-ADI benchmark for both settings. For both ADI heuristics, the potential for costs savings is approximately 7% higher for stochastic numbers of customers than it is for a fixed number. That is because the varying number of customers is known before the start of a repair tour and therefore serves as additional ADI for the part and greedy heuristic.

5.2.3. Share of customers providing ADI

So far we have assumed that the share of customers who communicate ADI to the service technician is 70%. This ADI-share depends on the number of potential error causes that can be detected by the sensors that monitor an appliance's condition. In some cases, monitoring 70% of the system can be technically infeasible, while in other cases, it may only be achieved with expensive investments in additional sensors. For that reason, we conduct a

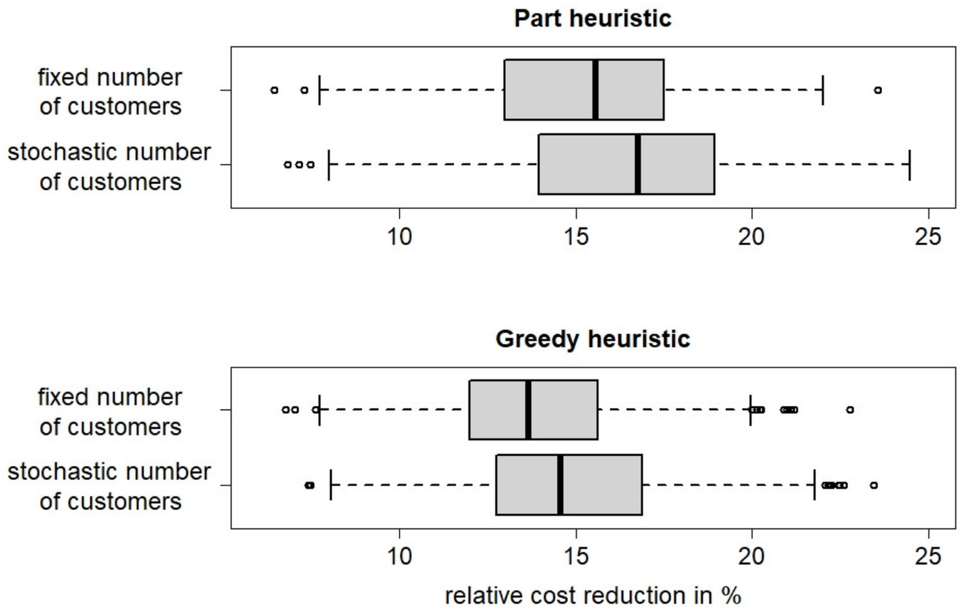

Fig. 3. relative cost reduction compared to the benchmark heuristic in %.

Table 11
Number of scenarios in which the respective heuristic performs best.

	part heuristic	greedy heuristic
ADI share of 30%	594	306
ADI share of 50%	676	224
ADI share of 70%	750	150

second subsequent study to determine the impact of the ADI share on the performance of the ADI heuristics.

We compare three different settings with ADI shares of 30%, 50%, and 70%. In any case, the respective ADI share α is divided amongst 9 genuine failure codes following the procedure described for the basic experiment. For each setting, we use the same random numbers such that the relative importance of the genuine failure codes remains the same across all settings.

To exclusively study the effect of the ADI share, we must ensure that the mean demand probability for each part across all failure codes (including the dummy failure code) is the same for each setting. Otherwise, we would be measuring the effect of different mean demand probabilities. Since the different ADI shares have an impact on a part's mean demand probability, we must offset this impact by altering the conditional demand probabilities accordingly. To be precise, it is enough to alter only the conditional demand probability given the dummy failure code. For the genuine failure codes, we use the same conditional demand probabilities for all settings. To create 100 different demand scenarios, we simply adopt the conditional demand probabilities generated for the first experiment, changing only the values for the dummy failure codes to achieve identical mean demand probabilities for each setting. We return to a fixed number of customers per tour of 6 and combine each demand scenario with 3 different return and penalty costs rates each (0.5, 1, 2 and 10, 30, 50). Thus, we again obtain 900 test instances.

As can be seen in Table 11, the number of instances for which the greedy heuristic outperforms the part heuristic is increasing with a decreasing ADI share. That is because the part heuristic's key advantage is that it implicitly anticipates the spare part demand of future customers. If, however, the share of future customers who communicate ADI is decreasing, this property becomes less important. Especially when a low penalty cost rate (P= 10) leads to a relatively small amount of spare parts stocked in the repair kit, the ADI share has an effect on which heuristic performs best. For an ADI share of just 30%, the greedy algorithm performs better in over 90% of the instances with a penalty cost rate P= 10, whereas for an ADI share of 70%, however, that is true in only 50% of those instances.

When we compare the ADI heuristics to the non-ADI benchmark, we find that the gap between them widens with an increasing ADI share. That is not surprising considering that the benchmark does not use ADI. However, as can be seen in Fig. 4, the relative cost reduction that is achievable by using ADI heuristics increases more than just linearly in the ADI share. It can also be seen that even if only a small number of customers transfers information about failure codes, there is still a relative cost reduction of approximately 4% per tour on average across all instances. Thus, over a whole year and across all technicians, large savings can be gathered.

5.2.4. Precision of the ADI

The precision of advance demand information depends on the number of sensors in an appliance, the number of different parts potentially needed when a sensor triggers a failure code, and the conditional demand probabilities for these parts. In this section, we determine the ADI heuristics' solutions for different ADI precision levels. We benchmark their performances against that of the ADI-insensitive one-fits-all repair kit. First, we consider different appliances with an identical number of sensors, but different accuracy levels given each failure code. Second, we compare different monitoring systems for one particular appliance that involve different numbers of sensors.

For the first part of this experiment, again, let us consider appliances with 9 sensors, such that each customer communicates either one out of 9 genuine failure codes or a dummy failure code. To represent the impact of a failure code, we have divided the spare parts into three sets for parts likely, unlikely, and impossible to be needed given that failure code. The precision of the ADI essentially depends on the number of parts in each of these sets and their respective conditional demand probabilities. To characterize differ-

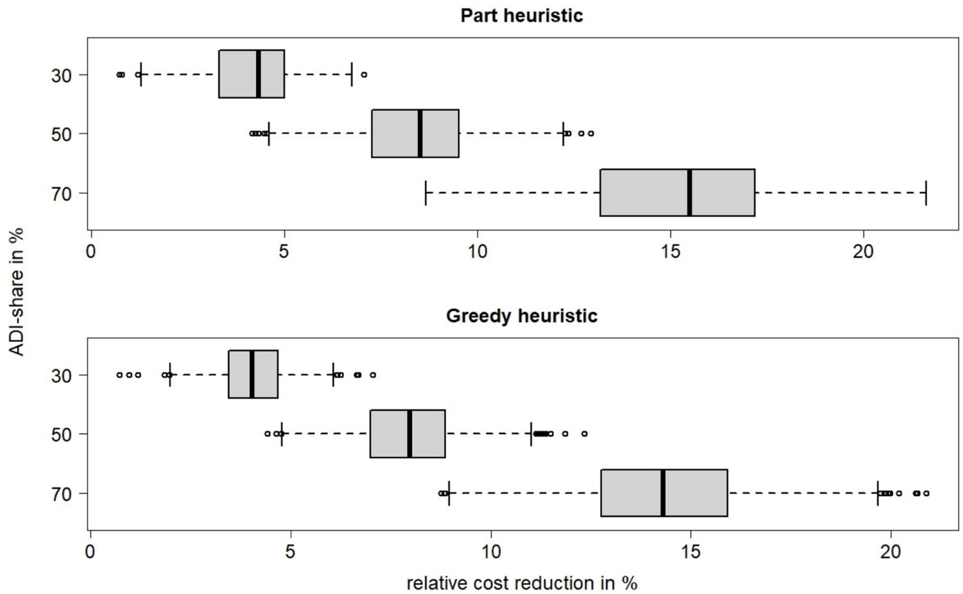

Fig. 4. Relative cost reduction compared to the benchmark heuristic in %.

Table 12
Demand characteristics for genuine failure codes.

setting	set	number of parts	conditional demand probabilities
high precision ADI	Set 1	3	drawn from a uniform distribution on [0.5,0.9]
	Set 2	15	drawn from a uniform distribution on [0,0.2]
	Set 3	32	0
medium precision ADI	Set 1	6	drawn from a uniform distribution on [0.2,0.5]
	Set 2	15	drawn from a uniform distribution on [0,0.2]
	Set 3	29	0
low precision ADI	Set 1	9	drawn from a uniform distribution on [0.1667,0.3]
	Set 2	15	drawn from a uniform distribution on [0,0.2]
	Set 3	26	0

Table 13
Number of scenarios in which the respective heuristic performs best.

	part heuristic	greedy heuristic
high precision ADI	750	150
medium precision ADI	788	112
low precision ADI	739	161

ent appliances, we vary the level of ADI precision by altering the number of spare parts in each of these sets as well as the conditional demand probabilities. We consider three different settings where the information given any genuine failure code is generated as described in Table 12. The settings are designed such that the expected demand probability for any spare part given a genuine failure code is the same for each setting.

Assumptions regarding costs and the number of customers in a tour are adopted from the basic experiment. The same is true for the failure code distribution and the conditional demand probabilities for all parts given the dummy failure code. For each precision setting, we run both ADI heuristics and the benchmark heuristic on 900 test instances.

As can be seen in Table 13, the precision of the ADI does not have a substantial effect on which of the ADI heuristics performs best. For all three settings, the part heuristic outperforms the greedy heuristic in approximately 5 out of 6 instances. Comparing the results obtained with the ADI heuristics on the one hand and the non-ADI benchmark heuristic on the other hand, we find that the differences increase the more precise the ADI is (see Fig. 5). Both the part and greedy heuristic attempt to identify the parts that are required for the upcoming repair tour based on the available failure codes, which is, however, particularly difficult when a single failure code can be triggered by the malfunction of a number of different parts. Even if the largest cost reduction can be observed in the case of the most precise information, there is still considerable benefit if the information is less precise.

Let us now consider the impact of different monitoring systems for one particular appliance. This appliance is built in a modular design and consists of 8 components that work independently of each other. That is malfunctions of one component do not affect the other components. For this appliance, we consider three different monitoring systems displayed in Fig. 6. The most precise system monitors all 8 components individually. The second best system monitors 4 sub-assemblies that are built from two components each. In technical terms, this could mean that we are measuring current flow through two electrical circuits connected in series. In case a failure code is triggered we know that one of the components is malfunctioning but not which one. The least precise system monitors two assemblies that comprise 2 sub-assemblies or 4 components each. Depending on the type of monitoring system one out of 8, 4, or just 2 different failure codes describe the state of a malfunctioning appliance. For each monitoring system, the service technician is aware of the probabilities with which the different failure codes are triggered and the respective conditional demand probabilities for each part in each component. In particular, we assume that the technician has no information about the correlation of demands for different parts given any failure code.

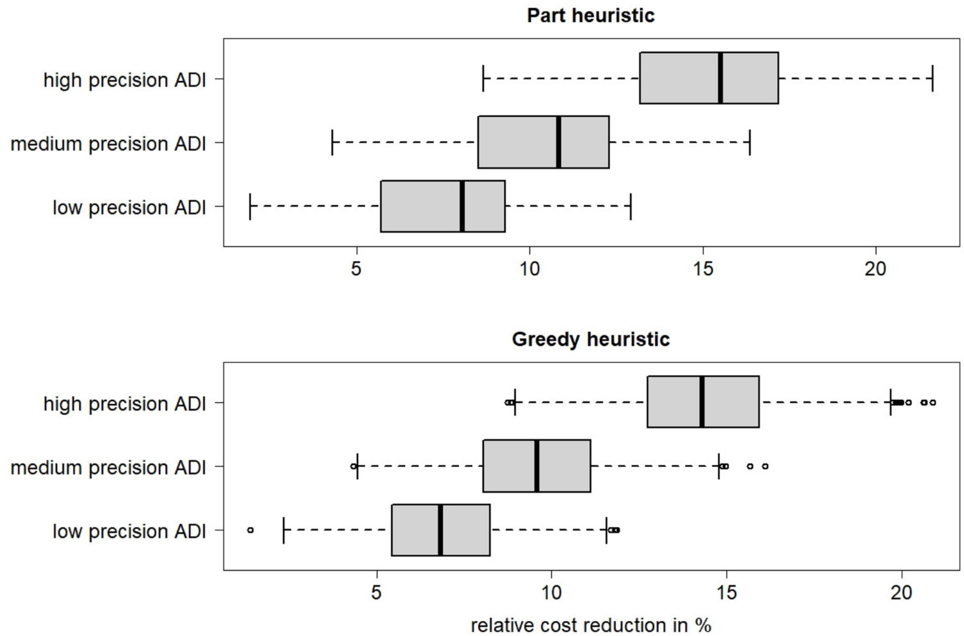

Fig. 5. Relative cost reduction compared to the benchmark heuristic in %.

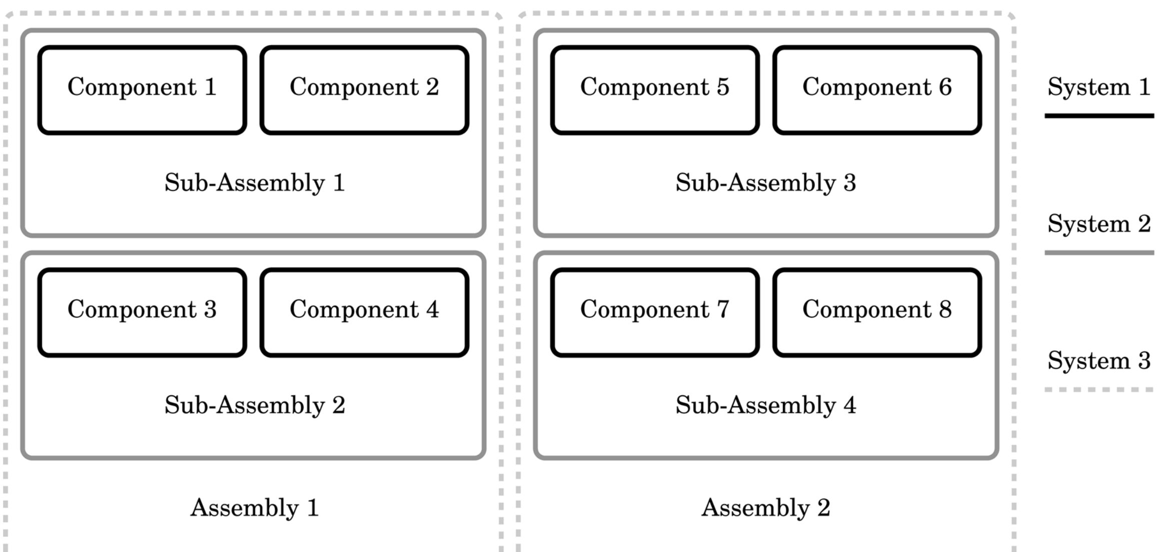

Fig. 6. Design of different monitoring systems.

We compare the performance of the ADI heuristics when using information of different accuracy from each of the three monitoring systems. To this end again we consider 900 problem instances that are comprised of 100 demand and holding cost scenarios combined with three different values each for return and penalty cost rate. The demand scenarios are constructed in the following way. We assume that all 8 components in the appliance are monitored and that all customers communicate the available information. Thus, ADI is available for each repair job and there is no need for a dummy failure code. Only one component fails at a time. To generate failure probabilities for the different components we draw 8 random numbers $R_1, \ldots, R_8$ from a uniform distribution on [0,1] and normalize their sum. That way the chance that component i in the jth customer's appliance has failed is given by

$$P(F_j^{CO} = i) = \frac{R_i}{\sum_{i=1}^{8} R_i} \quad i \in \{1, \ldots, 8\} \tag{18}$$

All components are built from parts that are specific to them. For each component, there are 6 critical parts, 2 parts that are likely to be needed given a failure of that component and four parts that are less likely to be needed. The conditional demand probabilities for these parts given a component failure are drawn from uniform distributions on [0.5,0.9] and [0,0.2], respectively. Thus we need to consider 48 parts in total for the repair kit. Depending on the type of monitoring system, the information available to the technician is either on a component, a sub-assembly, or an assembly level. The failure probabilities for sub-assemblies and assemblies can be derived by adding up the failure probabilities of all the components they contain. The conditional demand probability for any part given the failure of an assembly or sub-assembly is simply the weighted mean of the conditional demand probabilities for this part across all of the components in that assembly or sub-assembly. Similar to the basic experiment the holding cost rates

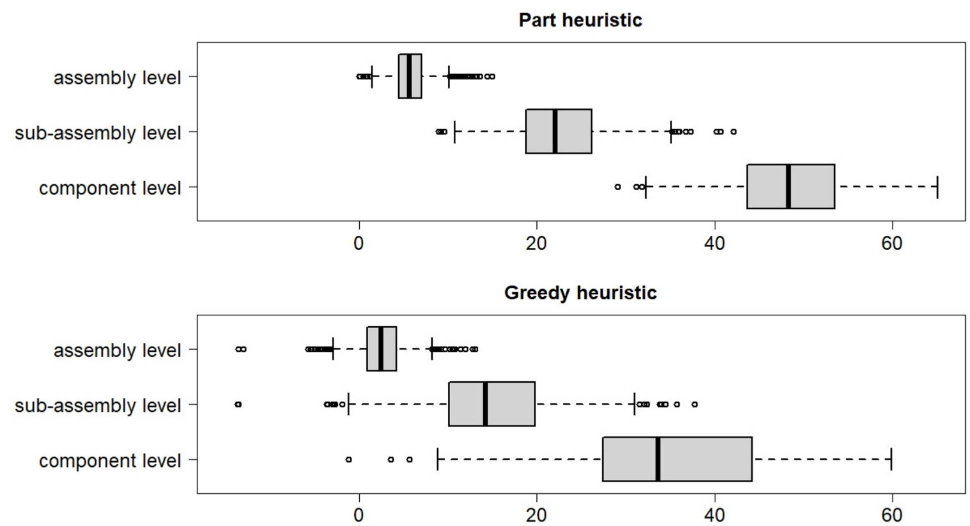

Fig. 7. Relative cost reduction compared to the benchmark heuristic in %.

for all 48 parts are drawn from a uniform distribution on [0,1]. Each demand and holding cost scenario is combined with return and penalty cost rates of 0.5, 1 and 3 and 10, 30 and 50, respectively. Again, we assume a fixed number of 6 customers in every repair tour.

We find that the part heuristic outperforms the greedy heuristic for all three monitoring systems on each of the 900 problem instances. The gap between solutions of both ADI heuristics can be explained with the modular design of the appliance. Contrary to previous experiments each part affects only one component. That means each part can only be needed upon the failure of one particular component. If, however, no failure of this component is signaled in the upcoming tour it makes sense to return spare parts for this component from the repair kit. The part heuristic is more likely to do that as it anticipates holding costs incurred in future tours.

Figure 7 shows the relative cost differences between costs incurred with the ADI heuristics assuming different monitoring systems and the non-ADI benchmark solution. For the part heuristic, we find that for each instance the usefulness of information is increasing the more precise it gets. Using the information from system 3 results in solutions that are on average 5.92% cheaper than the one-fits-all repair kits. Using information on a sub-assembly level (system 2) rather than on an assembly level (system 3) reduces the costs by another 17.72%. Finally, we can save an average of 33.77% of the total costs if we apply the part heuristic to information on a component level (system 1) rather than on a sub-assembly level (system 2). Not only do we save costs with increasing precision of the ADI, the percentage of cost savings actually increases from step to step. For the greedy heuristic, the picture is less clear. That is because there are instances for which the solutions derived using less precise information actually perform better. As already seen in the small-scale experiment the non-ADI benchmark outperforms the greedy heuristic in a number of cases. This happens when the greedy heuristic reacts to the available ADI by stocking up on parts that are then neither used nor returned. Thus, compared to the benchmark solution the increase in holding costs is larger than the reduction of penalty costs. In the majority of instances, however, the greedy heuristic provides a cheaper solution than the non-ADI benchmark and the potential for cost savings increases significantly with the precision of the available ADI.

6. Extension to problems with a capacity constraint

The literature on stochastic-demand multi-job repair kit problems so far only provides heuristic solution procedures for problems that disregard weight or volume limitations of the repair kit. In practice though, these limitations might restrict the applicability of the solution. For our repair kit problem with ADI, we find that a constraint on the total weight or volume of the units in the repair kit would rule out the use of the part heuristic. That is because the total capacity limit cannot easily be decomposed into capacity limits for each part. The greedy heuristic, however, can quickly be adjusted to handle a capacity constraint. To that end, we suggest running the re-equipment heuristic as described in Section 4 until the capacity constraint is violated. Once that happens, we successively reduce units of parts with the least JFR-contribution per unit of weight or volume until the constraint is satisfied again and resume the re-equipment procedure. To avoid infinite inclusion-removal loops, however, we exclude all parts once removed in the additional reduction step from further inclusion in the continued re-equipment procedure. A detailed description of the adjusted greedy heuristic is provided in the Appendix.

To demonstrate the impact of a capacity constraint and to test the performance of the adjusted greedy heuristic we return to the instances generated for the small-scale experiment. We add unit weights w_n for each part $n = 1, \dots, N$ representing either physical weights or volumes and consider different capacity limits. The 6 different parts in each instance are assigned weights of 0.1 (3), 0.5 (2), and 2 (1). That is we consider 3 small or light parts, 2 mid-sized parts, and 1 heavy or bulky part. The capacity constraints examined are expressed as multiples of the average weighted demand per tour (AWD) calculated as

$$AWD = \sum_{n=1}^{N} w_n p_n \cdot \sum_{c=C_{\min}}^{C_{\max}} c \cdot P(C = c), \tag{19}$$

where p_n is the mean demand probability as derived in Equation 16. For each instance, we determine the optimal and the heuristic solution for capacity limits of 2 to 5 times the respective AWD. Figure 8a shows the mean total costs across all instances given the different capacity limits. We find that the adjusted greedy heuristic generates close-to-optimal solutions on the constrained problem instances. The average optimality gap across

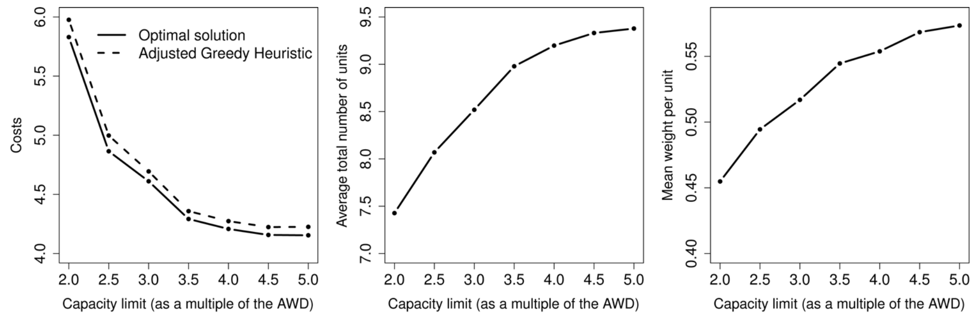

(a) Impact on total costs (b) Impact on total number of units (c) Impact on mean unit weight

Fig. 8. Impact of capacity constraints.

all 90 instances and all capacity limits is just 1.85%, with a maximum deviation of 15.18% from the optimal solution. The optimal solutions with low capacity limits, however, differ significantly from the optimal solutions for the unconstrained problem. With tighter capacity limits both, the total number of units and their mean weight per unit, are decreasing on average across the 90 problem instances as shown in Figs. 8b and c. That means that a number of particularly large or heavy parts are excluded from the repair kit due to the capacity constraint. This effect causes lower JFR-levels and thus increases penalty costs with decreasing capacity limits.

7. Conclusion

Today, many appliances from printers to washing machines are equipped with sensors that trigger failure codes in the event of faults. Wherever these sensors are monitoring indicators that are influenced by the condition of several different parts, the failure codes displayed by the appliance in case of an error is essentially imperfect ADI. In this paper, we extended previously studied problem formulations to account for the uncertain information on an appliance's condition that is communicated by customers to the service technician before a first repair attempt. The resulting repair kit problem with imperfect ADI is a sequential multi-part decision problem that can be modeled as a Markov decision process. We presented two heuristics to solve this problem, a part heuristic that splits the original MDP into substantially smaller MDPs for each part and a myopic greedy heuristic that decomposes the original problem into single-tour problems. For small-scale problems, we found that both heuristics provide solutions that are very close to the optimal solutions. Surprisingly though for most of the problem instances we examined, small-scale as well as large-scale, the part heuristic outperformed the greedy heuristic. The only exceptions to this rule seem to be instances in which the penalty cost rates are low compared to the inventory holding costs. That is a remarkable result since the solution procedures suggested in the literature on repair kit problems without ADI have always been greedy algorithms that handle multiple parts at the same time. Furthermore, we compared both ADI heuristics to a benchmark heuristic that determines a one-fits-all repair kit considering only average demand probabilities. We found that using advance demand information to modify the contents of a repair kit on a daily basis yields substantial advantages over a replenishment process that is defined by a single one-fits-all repair kit for most of the examined instances. Just how much can be saved depends on the share of customers that provide ADI, and the precision of the ADI.

For appliances already in the field, we recommend that all information available on failure codes is used for the replenishment decision. On the one hand, that means that failure codes must be communicated to the service technician before the first repair attempt. On the other hand, it means that we require statistics on how often a failure code occurs and which parts are likely to be needed in that case. Given this information, the service technician can use the ADI heuristics to reduce the average costs per tour.

For new appliances to be developed with service operations in mind, we recommend installing sensor systems that can detect as many potential malfunctions as precisely as possible. Our rationale is that the average costs incurred per repair tour are decreasing the more that customers are communicating failure codes, and the easier it is to deduce the actual cause of a malfunction from the failure code. The marginal cost reduction achieved with a small increase in the share of customers who communicate ADI is even increasing with the share of customers that already provide ADI. In particular, it pays off to design an appliance in such a way that it can be built from independent components. If this is possible monitoring these components individually rather than assemblies or sub-assemblies comprised of several components helps a lot to increase the ADI precision and thus reduce costs.

Finally, we could show that a slightly adjusted greedy heuristic provides excellent solutions when a capacity constraint is introduced to the repair kit problem with ADI. However, as the capacity constraint gets tighter the average costs per tour increase significantly.

While we determined heuristic solutions to the repair kit problem with ADI by decomposing it by part or by tour, one possible direction for further research could be to tackle this problem with the methods of approximate stochastic programming. By using a clever value function approximation, it could be possible to circumvent the need to evaluate each individual state in each step of a value iteration, a computationally intractable task.

Declaration of Competing Interest

None

Appendix A

Proof of Theorem 1. For the proof of Theorem 1 we first consider the following finite-time total cost version of our optimization problem for part $n \in \{1 \ldots, N\}$:

$$\overline{V}_{n,t}(x_{n,t}, c_t, b_t) = \begin{cases} \min\limits_{0 \le y_{n,t} \le C_{\max}} V_{n,1}\left(x_{n,t}, c_t, b_t, y_{n,t}\right) + \mathrm{E}\left[\overline{V}_{n,t-1}(y_{n,t} - D_{n,t}, C_{t-1}, B_{t-1})\right], & 1 \le t \le T \\ 0, & t = 0 \end{cases} \tag{A.1}$$

This problem is similar to the problem studied by Chen et al. (2017) with additional order costs and different style ADI being the only differences. Following the line of argumentation in Chen et al. (2017) we show by induction over t that $\overline{V}_{n,t}$ is convex in $x_{n,t}$ and that the optimal policy is a state and time-dependent (L,U)-policy.

For the initial step we can easily see that the claim is true for t=0. Let us now assume that it holds for any $t-1 \geq 0$. To show that the statement must be true for t we reformulate $\overline{V}_{n,t}$ as

$$\overline{V}_{n,t}(x_{n,t}, c_t, b_t) = \min_{0 \leq y_{n,t} \leq C_{\max}} \left\{ \min_{y_{n,t} \geq x_{n,t}} f_1(c_t, b_t, y_{n,t}), \min_{y_{n,t} \leq x_{n,t}} f_2(x_{n,t}, c_t, b_t, y_{n,t}) \right\}, \quad \text{with} \tag{A.2}$$

$$f_1(c_t, b_t, y_{n,t}) = h_n \cdot \mathrm{E}[(y_{n,t} - D_{n,t})^+] + P \cdot \mathrm{E}[(D_{n,t} - y_{n,t})^+] + \mathrm{E}\Big[\overline{V}_{n,t-1}(y_{n,t} - D_{n,t}, C_{t-1}, B_{t-1})\Big], \quad \text{and} \tag{A.3}$$

$$f_2(x_{n,t}, c_t, b_t, y_{n,t}) = r_n \cdot (x_{n,t} - y_{n,t}) + f_1(c_t, b_t, y_{n,t}). \tag{A.4}$$

We see that $f_1(c_t, b_t, y_{n,t}) \leq (\geq) f_2(x_{n,t}, c_t, b_t, y_{n,t})$ for $y_{n,t} < (>) x_{n,t}$ and $f_1(c_t, b_t, y_{n,t}) = f_2(x_{n,t}, c_t, b_t, y_{n,t})$ for $y_{n,t} = x_{n,t}$. Following the induction hypothesis we also know that f_1 and f_2 are convex in $y_{n,t}$. Next, we define the state and time dependent thresholds $L_{n,t}(c_t, b_t)$ and $U_{n,t}(c_t, b_t)$ by

$$L_{n,t}(c_t, b_t) = \underset{0 \leq y_{n,t} \leq C_{\max}}{\arg\min}\ f_1(c_t, b_t, y_{n,t}) \tag{A.5}$$

$$U_{n,t}(c_t, b_t) = \underset{0 \leq y_{n,t} \leq C_{\max}}{\arg\min}\ f_2(x_{n,t}, c_t, b_t, y_{n,t}) \quad \text{for any } x_{n,t} \tag{A.6}$$

Note that, while f_2 depends on x_t, $U_{n,t}$ does not. By the construction of the objective functions f_1 and f_2 we get $L_{n,t}(c_t, b_t) \leq U_{n,t}(c_t, b_t)$. Because of the relative positions of f_1 and f_2 and their convexity in $y_{n,t}$, the minimal costs in t tours is given by

$$\overline{V}_{n,t}(x_{n,t}, c_t, b_t) = \begin{cases} V_{n,1}(x_{n,t}, c_t, b_t, L_{n,t}(c_t, b_t)) + \mathrm{E}\Big[\overline{V}_{n,t-1}(L_{n,t}(c_t, b_t) - D_{n,t}, C_{t-1}, B_{t-1})\Big], & x_{n,t} \leq L_{n,t}(c_t, b_t) \\ V_{n,1}(x_{n,t}, c_t, b_t, x_{n,t}) + \mathrm{E}\Big[\overline{V}_{n,t-1}(x_{n,t} - D_{n,t}, C_{t-1}, B_{t-1})\Big], & L_{n,t}(c_t, b_t) < x_{n,t} < U_{n,t}(c_t, b_t) \\ V_{n,1}(x_{n,t}, c_t, b_t, U_{n,t}(c_t, b_t)) + \mathrm{E}\Big[\overline{V}_{n,t-1}(U_{n,t}(c_t, b_t) - D_{n,t}, C_{t-1}, B_{t-1})\Big], & x_{n,t} \geq U_{n,t}(c_t, b_t), \end{cases} \tag{A.7}$$

which means that we have a time and state dependent (L,U)-policy. Using the expression in Equation (A.7) we see that the difference $V_{n,t}(x_{n,t}+1, c_t, b_t) - V_{n,t}(x_{n,t}, c_t, b_t)$ is non-decreasing in $x_{n,t}$. Thus $V_{n,t}(x_{n,t}, c_t, b_t)$ is convex in $x_{n,t}$, which completes the proof for the finite-time total cost optimization problem. Since the (L,U)-policy minimizes the t-period total cost function it must also minimize the t-period average costs. This does not change for $t \to \infty$. However, for an infinite-time average cost MDP with a finite state and action space, there must be a stationary solution. This shows that for $t \to \infty$ $L_{n,t}(c_t, b_t)$ and $U_{n,t}(c_t, b_t)$ converge against state dependent but time independent thresholds $L_n(c_t, b_t)$ and $U_n(c_t, b_t)$ which completes the proof of Theorem 1. □

Proof of Theorem 2. We prove Theorem 2 by induction. For $t=1$ we obtain the following differences.

$$\Delta_{x_k}\Delta_{x_l}V_1^* = 0 \quad \forall k \neq l \in \{1, \ldots, N\} \tag{A.8}$$

$$\Delta_{y_k}\Delta_{y_l}V_1^* = -P\sum_{j=1}^{c}\left[\prod_{n\in\{1,\ldots,N\}\setminus\{k,l\}}\left(\left(1 - P\Big(D_{j,n} = 1\Big|F_j = f_j\Big)\right) + P\Big(D_{j,n} = 1\Big|F_j = f_j\Big) \cdot P\left(\sum_{m=1}^{j-1} D_{m,n} < y_n\Big|F = f\right)\right) \cdot \prod_{n\in\{k,l\}} P\Big(D_{j,n} = 1\Big|F_j = f_j\Big) \cdot P\left(\sum_{m=1}^{j-1} D_{m,n} = y_n\Big|F = f\right)\right] \leq 0 \quad \forall k \neq l \in \{1, \ldots, N\} \tag{A.9}$$

Algorithm 3 Adjusted re-equipment heuristic.

Require: old repair kit x, current repair kit y as determined by reduction heuristic, number of customers c, error codes f, conditional demand probabilities $P(D_{j,n} = 1|F_j = f_j)\ \forall n \in \{1, \ldots, N\}$, $f_j \in \{1, \ldots, M+1\}$, unit holding costs $h_n\ \forall n \in \{1, \ldots, N\}$, unit return cost $r_n\ \forall n \in \{1, \ldots, N\}$, penalty costs P, capacity limit (cap)

Ensure: new inventory levels y^{opt}

```
y^opt ← y
stop ← 0
α ← 0^N
while stop= 0 do
    for n ∈ {1,...,N} with y_n < c ∧ α_n ≠ 1 do
        y^temp ← y
        y_n^temp ← y_n + 1
        Δ_n ← (JFR(c,f,y^temp) − JFR(c,f,y)) / (h_n (E[(y_n^temp − D_n)^+ | F=f] − E[(y_n − D_n)^+ | F=f]) − r_n · 1_{y_n < x_n})
    end for
    n* ← arg max_{n: y_n < c} Δ_n
    y_n* ← y_n* + 1
    while Σ_{n=1}^{N} w_n · y_n > cap do
        for n ∈ {1,...,N} with y_n ≥ 1 do
            y^temp ← y
            y_n^temp ← y_n − 1
            Δ_n ← (JFR(c,f,y) − JFR(c,f,y^temp)) / w_n
        end for
        n* ← arg min_{n: y_n ≥ 1} Δ_n
        y_n* ← y_n* − 1
        α_n* ← 1
    end while
    if V_1(x,c,f,y) < V_1(x,c,f,y^opt) then
        y^opt ← y
    end if
    if y_n = c ∨ α_n = 1 ∀n ∈ {1,...,N} then
        stop ← 1
    end if
end while
```

$$\Delta_{x_k}\Delta_{y_l}V_1^* = 0 \quad \forall k \neq l \in \{1, \ldots, N\} \tag{A.10}$$

$$\Delta_{x_k}\Delta_{y_k}V_1^* = -r_k\mathbb{1}_{\{x_k = y_k\}} \leq 0 \quad \forall k \in \{1, \ldots, N\} \tag{A.11}$$

Since all differences are non-increasing by Corollary 2.6.1 of Topkis (1998) the single period cost function V_1^* is submodular in x and y on $\mathfrak{X} \times \mathfrak{Y}$. Because the inventory x before orders does not directly impact the stocking decision in the following tour we obtain the following results for $t > 1$.

$$\Delta_{x_k}\Delta_{x_l}V_t^* = \Delta_{x_k}\Delta_{x_l}V_1^* = 0 \quad \forall k \neq l \in \{1, \ldots, N\} \tag{A.12}$$

$$\Delta_{x_k}\Delta_{y_l}V_t^* = \Delta_{x_k}\Delta_{y_l}V_1^* = 0 \quad \forall k \neq l \in \{1, \ldots, N\} \tag{A.13}$$

$$\Delta_{x_k}\Delta_{y_k}V_t^* = \Delta_{x_k}\Delta_{y_k}V_1^* = -r_k\mathbb{1}_{\{x_k = y_k\}} \leq 0 \quad \forall k \in \{1, \ldots, N\} \tag{A.14}$$

According to the induction hypothesis V_{t-1}^* is submodular on $\mathfrak{X} \times \mathfrak{Y}$. Then by Theorem 2.7.6 of Topkis (1998) $\min_y V_{t-1}^*(x, c, f, y)$ is submodular in x on $\mathfrak{X}$. Thus, $\min_{y'} V_{t-1}^*\big(y - d, c', f', y'\big)$ is submodular in y on $\mathfrak{Y}$ for any d,c' and f'. Because we know already that

V_1^* is submodular on $\mathfrak{Y}$ and the expected future costs are just a conical combination of functions shown to be submodular on $\mathfrak{Y}$ we can conclude that V_t^* is submodular in y on $\mathfrak{Y}$. □

References

Benbitour, M. H., & Sahin, E. (2015). Evaluation of the impact of uncertain advance demand information on production/inventory systems. *IFAC-PapersOnLine, 48*(3), 1738–1743. 15th IFAC Symposium on Information Control Problems in Manufacturing

Benjaafar, S., Cooper, W. L., & Mardan, S. (2011). Production-inventory systems with imperfect advance demand information and updating. *Naval Research Logistics (NRL), 58*(2), 88–106.

Bernstein, F., & DeCroix, G. A. (2015). Advance demand information in a multiproduct system. *Manufacturing & Service Operations Management, 17*(1), 52–65.

Bijvank, M., Koole, G., & Vis, I. F. A. (2010). Optimising a general repair kit problem with a service constraint. *European Journal of Operational Research, 204*(1), 76–85.

Brumelle, S., & Granot, D. (1993). The repair kit problem revisited. *Operations Research, 41*(5), 994–1006.

Chen, L., Şafak Yücel, & Zhu, K. (2017). Inventory management in a closed-loop supply chain with advance demand information. *Operations Research Letters, 45*(2), 175–180.

Claudio, D., & Krishnamurthy, A. (2009). Kanban-based pull systems with advance demand information. *International Journal of Production Research, 47*(12), 3139–3160.

DeCroix, G. A., & Mookerjee, V. S. (1997). Purchasing demand information in a stochastic-demand inventory system. *European Journal of Operational Research, 102*(1), 36–57.

Fischer, S., Benzaman, B., Diegel, E., Gimenez, M. S., & Claudio, D. (2020). Effects of different types of imperfect advance demand information in production systems. *Journal of Simulation*, 1–13.

Gallego, G., & Özer, Ö. (2001). Integrating replenishment decisions with advance demand information. *Management Science, 47*(10), 1344–1360.

Gao, L., Xu, S. H., & Ball, M. O. (2012). Managing an available-to-promise assembly system with dynamic short-term pseudo-order forecast. *Management Science, 58*(4), 770–790.

Gayon, J.-P., Benjaafar, S., & de Véricourt, F. (2009). Using imperfect advance demand information in production-inventory systems with multiple customer classes. *Manufacturing & Service Operations Management, 11*(1), 128–143.

Graves, S. C. (1982). A multiple-item inventory model with a job completion criterion. *Management Science, 28*(11), 1334–1337.

Hariharan, R., & Zipkin, P. (1995). Customer-order information, leadtimes, and inventories. *Management Science, 41*(10), 1599–1607.

Hausman, W. H. (1982). On optimal repair kits under a job completion criterion. *Management Science, 28*(11), 1350–1351.

Heeremans, D., & Gelders, L. (1995). Multiple period repair kit problem with a job completion criterion: A case study. *European Journal of Operational Research, 81*(2), 239–248.

Holmströem, J., Cheikhrouhou, N., Farine, G., & Främling, K. (2011). Product centric organization of after-sales supply chain planning and control. In J. C. Fransoo, T. Wäfler, & J. R. Wilson (Eds.), *Behavioral operations in planning and scheduling* (pp. 187–198). Springer.

Huang, T., & Van Mieghem, J. A. (2014). Clickstream data and inventory management: Model and empirical analysis. *Production and Operations Management, 23*(3), 333–347.

Karaesmen, F. (2013). Value of advance demand information in production and inventory systems with shared resources. In J. MacGregor Smith, & B. Tan (Eds.), *Handbook of stochastic models and analysis of manufacturing system operations* (pp. 139–165). Springer.

Kim, J.-H., Ahn, H.-S., & Righter, R. (2009). Optimal production policies with multistage stochastic demand lead times. *Probability in the Engineering and Informational Sciences, 23*, 515–543.

Law, A. M., & Kelton, W. D. (2000). *Simulation modeling and analysis*. McGraw-Hill.

Liberopoulos, G., Chronis, A., & Koukoumialos, S. (2003). Base stock policies with some unreliable advance demand information. In *Proceedings of the 4th aegean international conference on analysis of manufacturing systems* (pp. 1–4).

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 Transactions International Journal, 2*, 95–114.

Mamer, J. W., & Shogan, A. W. (1987). A constrained capital budgeting problem with applications to repair kit selection. *Management Science, 33*(6), 800–806.

Mamer, J. W., & Smith, S. A. (1982). Optimizing field repair kits based on job completion rate. *Management Science, 28*(11), 1328–1333.

Mamer, J. W., & Smith, S. A. (1985). Job completion inventory systems: Optimal policies for repair kits and spare machines. *Management Science, 31*(6), 703–718.

March, S. T., & Scudder, G. D. (1984). On “optimizing field repair kits based on job completion rate”. *Management Science, 30*(8), 1025–1028.

Prak, D., Saccani, N., Syntetos, A., Teunter, R., & Visintin, F. (2017). The repair kit problem with positive replenishment lead times and fixed ordering costs. *European Journal of Operational Research, 261*(3), 893–902.

Rippe, C., & Kiesmüller, G. P. (2020). *The added value of advance demand information for the planning of a repair kit*. Unpublished Results

Saccani, N., Visintin, F., Mansini, R., & Colombo, M. (2017). Improving spare parts management for field services: A model and a case study for the repair kit problem. *IMA Journal of Management Mathematics, 28*, 185–204.

Smith, S. A., Chambers, J. C., & Shlifer, E. (1980). Optimal inventories based on job completion rate for repairs requiring multiple items. *Management Science, 26*(8), 849–852.

Tan, T., Güllü, R., & Erkip, N. (2007). Modelling imperfect advance demand information and analysis of optimal inventory policies. *European Journal of Operational Research, 177*(2), 897–923.

Tan, T., Güllü, R., & Erkip, N. (2009). Using imperfect advance demand information in ordering and rationing decisions. *International Journal of Production Economics, 121*(2), 665–677.

Teunter, R. H. (2006). The multiple-job repair kit problem. *European Journal of Operational Research, 175*(2), 1103–1116.

Thonemann, U. W. (2002). Improving supply-chain performance by sharing advance demand information. *European Journal of Operational Research, 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 Transactions, 50*(3), 246–264.

Topkis, D. M. (1998). *Supermodularity and complementarity*. Princeton University Press.

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

Zhu, K., & Thonemann, U. W. (2004). Modeling the benefits of sharing future demand information. *Operations Research, 52*(1), 136–147.

Zhu, S., Jaarsveld, W., & Dekker, R. (2020). Spare parts inventory control based on maintenance planning. *Reliability Engineering and System Safety, 193*.