PhD-FDEF-2021-021
The Faculty of Law, Economics and Finance

DISSERTATION

Defended on October 29, 2021, in Luxembourg
to obtain the degree of

Docteur de l'université du Luxembourg

en gestion

by

Melvin Drent

Born on October 3, 1993, in Geleen, The Netherlands

Stochastic Models of Critical Operations

Dissertation Defense Committee:

Supervisor:	prof. dr. Joachim Arts	Université du Luxembourg
Chair:	prof. dr. Benny Mantin	Université du Luxembourg
Vice-chair:	prof. dr. Nils Löhndorf	Université du Luxembourg
Members:	prof. dr. Yossi Sheffi	Massachussetts Institute of Technology
	prof. dr. Robert Boute	Katholieke Universiteit Leuven Vlerick Business School

Stochastic Models of Critical Operations

This research has been funded by the National Research Fund of Luxembourg through AFR grant 12451704.

Acknowledgments

While this PhD thesis has only one author, it is in fact the product of countless interactions with colleagues, friends, and family members. It fills me with pride and gratitude to thank those who have supported and encouraged me over the past years.

First and foremost, I am deeply indebted to my advisor, Joachim Arts, for his guidance, mentorship, and friendship over the past six years. I am grateful for his support and time that made this thesis possible. He has been instrumental in shaping my ideas about research, academic life, and life in general. His positive attitude towards research problems and extensive technical knowledge is inspiring. He has become a dear friend of mine and I look forward to our future collaborations.

Next I wish to thank my committee chair, Benny Mantin, for his continuous support from the moment I joined the LCL and for his enthusiasm to share his experience and advice whenever I needed it. I am grateful to him for giving me the freedom to pursue my own research ideas and for always showing an active interest in my work. Last year we also started a research project together and I was amazed by how quick he finds interesting research directions. I hope we continue our collaboration in the future.

I am grateful to Yossi Sheffi for his time to share his valuable advice and for being part of my committee. He was very approachable during my time at MIT and he helped me in my academic job search. His constructive feedback on my research, particularly during the stimulating discussions after the CTL seminar in the fall of 2019, was very helpful.

I would like to thank Collin Drent and Stella Kapodistria for our collaboration on Chapter 5 of this thesis. My collaboration with Collin extends far beyond the contents of this chapter; he has been my (intellectual) sparring partner for as far as I can remember. He also performed a very thorough proofreading of all chapters. I also thank Stella for our time in Athens where she proved to be an excellent guide and travel mate.

I am grateful to the maintenance development department of the Dutch Railways for our collaboration that resulted in Chapter 2 of this thesis.

I thank Poulad Moradi for his contributions to Chapter 4 of this thesis. It is great to see that he started his own PhD research. Chapter 4 also benefited from the master thesis project of Ranit Sinha.

Further I wish to thank Nils Löhndorf. Besides being part of my committee, I also enjoyed the many conversations we had about a wide range of topics.

My time at the LCL would not have been as enjoyable as it has been without the many coffee breaks, lunches, dinners (of which the ones at Himalaya I remember most), and other countless interactions with LCL staff members and students. In particular I thank Nicole Perez-Becker, You Wu, Sarah van der Auweraer, Roozbeh Qorbanian, Bikey Kleiford Seranilla, Anne Lange, Jackie Brown, Weichun Chen, Çağıl Koçyiğit, Steffen Klosterhalfen, Ganesh Balasubramanian, Mélanie Winter, and Carla Rosen-Vacher; they all have been instrumental in the completion of this thesis.

I am grateful to the National Research Fund of Luxembourg for funding my research through an AFR grant.

My visit to MIT in the fall of 2019 was truly inspiring. I wish to thank Josué Velázquez Martínez for hosting me and all members of the CTL at large for making my time at MIT into an amazing experience. During my time there, I also had the opportunity to interact with Sean Willems. He was always very sharp during our meetings. I am also grateful for his help with my academic job search. My visit to MIT was partially funded by a Prins Bernhard Fellowship supported by the Jan de Ruijsscher / Pia Huisman fund.

I am indebted to Robert Boute for being on my committee and to Tim Huh for serving as an advisor to the committee. They both provided valuable feedback on my work.

During my time as a PhD student I had the privilege to visit many conferences and interact with many great scholars. Their advice has been invaluable. In particular I would like to thank Loe Schlicher, Jasper Veldman, and Chaaben Kouki for our stimulating interactions.

I am thankful for having friends that provide me with the necessary distractions, help me to keep things in perspective, and visit me wherever I am, being it in Luxembourg, Thiaumont or Boston.

My parents and two brothers have always been there for me. I cannot even articulate how grateful I am for their unconditional, selfless, and continuous support, encouragement, and advice. Levi has always kept me grounded, and Collin has literally been the co-author of my life. Without them I would not be where I am today.

My final word of gratitude is to Tim. Thank you for always being so positive and for supporting me through thick and thin. I look forward to our life together.

Contents

Chapter 1

Introduction

> "It can scarcely be denied that the supreme goal of all theory is to make the irreducible basic elements as simple and as few as possible without having to surrender the adequate representation of a single datum of experience."
>
> Albert Einstein

At its core, the field of operations management is concerned with the improvement of operational decision making in its broadest sense. Researchers in this field are typically interested in the effects of operational decisions on important performance measures. Quantifying those effects through real-life experimentation is very expensive or risky – not to mention the fact that the consequences of a given operational policy can only be quantified accurately after a prolonged period of operating that specific policy. For these reasons we generally develop a quantitative representation of the specific operation of interest. These so-called mathematical models allow us to parsimoniously study and evaluate the performance of operational policies without the necessity to undertake expensive and time-consuming real-life interventions. Needless to say, such models need to be sufficiently complex to do justice to real-life operations yet simple enough to tractably generate meaningful insights and advance decision making.

Many operations are characterized by uncertainty. Examples include supermarkets that observe fluctuations in daily demands for nearly all their products, airlines that sometimes must perform costly emergency repairs of their aircraft, and hospitals that

may have to delay surgeries when critical equipment abruptly ceases to function. While the inclusion of uncertainty generally leads to a more complex mathematical model, ignoring uncertainty altogether would render the practical implications of the model useless. If uncertainty is too important to be ignored, the implementation of an operational policy based on a deterministic model likely leads to poor operational performance. Mathematical models that take into account explicitly the presence of uncertainty are called stochastic models. The development and analysis of stochastic models of several critical operations is the overarching topic of this thesis.

1.1. Contributions of the thesis

This thesis comprises four self-contained chapters. We can group these chapters into two categories depending on the stream of literature they contribute to. The first three chapters focus on various stochastic inventory systems with two supply modes. The fourth chapter studies condition based maintenance of stochastically deteriorating systems. Below we present a brief overview of the contents of each chapter. The contributions and positioning relative to the existing literature are described in more detail in each chapter separately.

1.1.1 Stochastic inventory systems with two supply modes

Companies can often replenish their inventories through a regular supply mode and an expedited supply mode, the latter having a shorter lead time than the former. The flexibility of the expedited mode – albeit at the expense of a premium in terms of a cost or additional resource usage – can help companies to cope with uncertainties. The first part of this thesis studies such inventory systems with two supply modes, also called dual-sourcing inventory systems. Since it is well-known that optimal policies for dual-sourcing inventory systems have complex structures (e.g., Whittemore and Saunders, 1977; Feng et al., 2006), we mostly focus on the development and application of heuristic policies.

Chapter 2 studies a two-echelon distribution network for repairable spare parts consisting of one central warehouse and multiple local warehouses. Each warehouse keeps multiple types of repairable parts to maintain several types of capital goods, such as aircraft, rolling stock, and lithography systems. The repair shop at the central warehouse has two repair options for each failed part: a regular repair option and an expedited repair option. The latter is faster than the former but it comes with a higher repair workload. In the design of these spare parts inventory systems, companies need to decide on stocking levels and expedite thresholds such that stock investments are

minimized while satisfying asset availability and repair shop workload constraints. We use queueing theory to model the dynamics of the dual-index policy in the repair shop, and we rely on Dantzig-Wolfe decomposition to develop an effective and efficient solution algorithm for the decision problem. A key insight of this chapter is that anticipating expediting decisions that will be made later can lead to substantial reductions in stock investments required to meet customer service levels. Based on a case study at Netherlands Railways, we show how managers can significantly reduce the investment in repairable spare parts when dynamic expediting policies are leveraged to prioritize repair of parts whose inventory is critically low. The contents of Chapter 2 are based on Drent and Arts (2020).

Chapter 3 studies the inbound transport and inventory management decision making for a company that sells an assortment of products sourced from outside suppliers. The inbound transport is outsourced to a third party logistics provider that offers two distinct transport modes for each product. These modes differ in terms of their carbon emissions, speed, and costs. The company needs to decide periodically how much it wants to ship with each transport mode such that total inventory costs are minimized while keeping the total carbon emissions from transportation for the entire assortment below a certain target level. Such assortment-wide constraints will be increasingly prevalent, either voluntarily or enforced by government regulation. Assuming that shipment decisions are governed by the dual-index policy of Veeraraghavan and Scheller-Wolf (2008), we formulate the decision problem as a mixed integer linear program that we solve through Dantzig-Wolfe decomposition. We benchmark our decision model against two state-of-the art approaches in a large testbed based on real-life carbon emissions data. Relative to our decision model, the first benchmark lacks the flexibility to dynamically ship products with two transport modes while the second benchmark makes transport decisions for each product individually rather than holistically for the entire assortment. Our computational experiment shows that our decision model can significantly outperform both benchmarks for realistic targets for carbon emission reduction and furthermore indicates that dynamic mode selection, as opposed to static and blanket mode selection, has great potential to efficiently curb carbon emissions from transportation at relatively little additional costs. The contents of Chapter 3 are based on Drent et al. (2021b).

The first two chapters apply well-known dual-sourcing inventory policies to multi-item settings and subsequently focus on generating valuable managerial insights. By contrast, Chapter 4 is more fundamental as it focuses on devising a new dual-sourcing inventory policy. In particular, we study the canonical single-echelon single-item inventory system with two suppliers under periodic review facing stochastic demand where excess demand is backlogged. The expedited supplier has a shorter lead time than the regular supplier but charges a higher price. We introduce the Projected Expedited Inventory Position (PEIP) policy to control this inventory system. Under

this policy, the expedited supplier is operated according to an order-up-to rule that keeps the expedited inventory position at (or above) a certain target level. Regular orders placed in the past can cause the expedited inventory position to exceed this target level. The PEIP policy accounts explicitly for this overshoot by placing regular orders such that the projected expedited inventory position is kept at a target level. We show that the relative difference between the long run average cost per period of the PEIP policy and the optimal policy converges to zero when both the shortage cost and the cost premium for expedited units become large, with their ratio held constant. A corollary of this result is that several existing heuristics are also asymptotically optimal in this non-trivial regime. We show through an extensive numerical investigation that the PEIP policy outperforms the current best performing heuristic policies. The contents of Chapter 4 are based on Drent and Arts (2021).

1.1.2 Condition based maintenance of stochastically deteriorating systems

Companies responsible for maintaining capital goods often have access to degradation data of critical components. They can leverage this degradation data to replace components preventively before they fail. We refer to this as condition based maintenance. Since replacing preventively is cheaper than replacing correctively, the decision maker needs to decide when she should intervene and replace the component by a new one, thereby trading off costly premature interventions with costly tardy replacements. This decision problem has been studied since the sixties, but almost exclusively under the assumptions that the parameters of the underlying degradation process are known to the decision maker and that these parameters are the same from component to component (e.g., Kolesar, 1966; Alaswad and Xiang, 2017). Recently built capital goods often have integrated sensor technology that allows degradation data of a critical component to be gathered at hardly any additional cost. This data can be relayed in real-time to the decision maker through the Internet-of-Things. An important implication of this is that it allows us to relax the above-mentioned assumptions. Indeed, real-time degradation data allows us to learn degradation behavior on the individual component level and tailor our decision making accordingly in an integrated fashion.

Learning from real-time degradation behavior of stochastically deteriorating components to tailor our decision making is the topic of the second part of this thesis. In particular, Chapter 5 studies a single component that deteriorates according to compound Poisson degradation, but the parameters of the Poisson process as well as the compounding distribution vary from one component to the next, that is, the population of components is heterogeneous. These parameters cannot be

observed directly and therefore they need to be learned from the degradation signal that is relayed through an online sensor in real-time. We model this situation as a partially observable Markov decision process. The entire past degradation path of a component is relevant state information in this setting, which unfortunately leads to tractability issues. By using concepts from Bayesian statistics, we show that the high dimensional state space can be collapsed to only 3 dimensions while retaining all relevant information. This collapse gives insight into how all relevant information in a real-time degradation signal can be parsimoniously represented, and it makes the model both tractable numerically and amenable to structural analysis. We characterize the optimal replacement policy as a state dependent control limit. The control limit increases with age but may decrease as a result of other information in the degradation signal. Based on a case study on interventional X-ray machines from Philips Healthcare, we show that integration of learning and decision making leads to cost reductions of 10.50% relative to approaches that do not learn from the real-time signal and 4.28% relative to approaches that separate learning and decision making. The contents of Chapter 5 are based on Drent et al. (2021a).

1.2. Organization of the thesis

Table 1.1 gives an overview of the thesis based on the literature streams we contribute to and the scope of analysis as discussed in the previous section. Table 1.2 gives an overview of the thesis based on the main methodology that we use in each chapter. Each chapter is self-contained and can be read independently.

Table 1.1 Navigating the thesis based on related literature stream and scope of analysis.

	Topic		Scope	
Chapter	Dual-sourcing	Maintenance	Single-item	Multi-item
2	x			x
3	x			x
4	x		x	
5		x	x	

Table 1.2 Navigating the thesis based on main methodology.

Chapter	Markov decision theory	Asymptotics	Column generation	Queueing theory	Bayesian statistics
2			x	x	
3			x		
4		x			
5	x				x

Chapter 2

Expediting in two-echelon spare parts inventory systems

2.1. Introduction

For many industries and service organizations, the availability of capital goods such as rolling stock, manufacturing equipment and aircraft is of crucial importance for their operations. To ensure high availability of these capital goods, companies stock critical components and replace a defective component with a ready-for-use spare component after failure. Since many critical components represent a significant financial investment, defective components are usually repaired and put back on stock rather than discarded. Consequently, the availability of capital goods largely depends on the design of the underlying spare parts inventory system for repairing and supplying these so-called repairable components.

Spare parts inventory systems for capital goods often have a two-echelon structure, in which many different types of components are stocked (Cohen et al., 1997). In this chapter, we study such a multi-item two-echelon spare parts inventory system. The system consists of a set of local warehouses, i.e. operating sites, that are supported by one central warehouse. Each local warehouse maintains field inventories for spare components and sends defective components to the repair shop. The repair shop repairs these defective components and sends them to the central warehouse which replenishes the local warehouses. It is obvious that next to the inventory levels of spare components, the repair operations at the repair shop affect the availability of capital goods at the operating sites. Hence, the determination of spare components inventory levels and the design of the repair operations in the repair shop are two key

aspects in the design of these two-echelon spare parts inventory systems.

In the capital goods industry, it is common practice to acquire spare components together with the acquisition of the capital good because, at that time, it is possible to negotiate reasonable prices. The determination of spare components inventory levels is therefore closely related to what is known in literature as the initial spare parts supply problem (e.g., Van Houtum and Kranenburg, 2015). With respect to the repair operations at the repair shop, companies often have the flexibility to expedite the repair of defective components. Although expediting comes at an extra price, either because internal repair resources are limited or because an external repair shop charges a higher price, the possibility to expedite can significantly reduce the required initial financial investment in spare parts. Indeed, expediting the repair of defective components more often implies that a smaller initial financial investment in spare parts is required to ensure the same availability of capital goods as in spare parts inventory systems where no repair flexibility is incorporated.

Hence, in the design of the spare parts inventory systems sketched in the last two paragraphs, decision makers face two major questions:

1. How many spare parts of each repairable type should the company initially purchase and place at each warehouse?
2. When should the repair of a detective part of a given repairable type be expedited?

The objective of this chapter is to present a tractable optimization model that assists decision makers in answering these two questions. These questions are faced among others by Netherlands Railways (NS), the principal Dutch passenger railway operator. Our collaboration with their maintenance department led to the present work. To establish the practical value of our optimization model, we report on a case study on their data.

As is often the case in practice, we consider a setting with several capital good types (e.g. regional trains and inter-city trains; wide-body aircraft and narrow-body aircraft) and where repairables may use different repair resources (e.g. electronic and mechanical). Because not all repairs can be expedited, many companies, including NS, use agreements between the repair shop manager and the inventory manager that determine how much of the total workload can be expedited per repair resource. Hence, the objective of our optimization model is to minimize the total investment costs in spare parts while

- not exceeding a given maximum total mean number of backorders over all local warehouses for each capital good type, and

- keeping the fraction of repairs that are expedited per repair resource below a given target level. (We will also consider an alternate setting in which expediting repairs comes at additional costs rather than being constrained).

Because we consider critical components, a backorder for a spare part implies that the affected capital good becomes inoperable. Since failures of components typically occur very infrequently, a common assumption in the spare parts literature is that the probability that two or more backorders are from the same capital good at any point in time is negligible (e.g., Muckstadt, 2005; Sherbrooke, 2004). Under that assumption the average availability of a capital good type is the number of capital goods of that type minus the expected number of backorders of parts in that capital good type. As such, the first constraint of our optimization model guarantees a certain availability of each capital good type throughout the geographical region covered by the local warehouses.

In this chapter, we provide a mathematical model for the decision problem described above. We assume that each local warehouse is replenished by an $(S-1, S)$ base stock policy. This means that each defective part is replaced with a ready-for-use item and is sent to the repair shop at the central warehouse immediately after the defect occurs. This replenishment policy is common in practice and is considered as well-suited for spare parts inventory control (Van Houtum and Kranenburg, 2015). The central warehouse operates under an (S, T) policy similar to Song and Zipkin (2009), which keeps the usual inventory position at constant level S, just as in a standard base stock policy. In addition, expedite threshold T triggers expedited repairs when outstanding orders in the repair pipeline are too far away. This dynamic repair policy thus takes into account real-time information about the repair pipeline of the repair shop, which can be obtained through modern tracking technologies. We assume that unsatisfied demand is backordered at all warehouses. Furthermore, we assume deterministic lead times for the replenishments of the local warehouses as well as for both repair options at the central warehouse.

The main contributions of this chapter are summarized as follows:

1. We are the first to integrate stocking and expedited repair decisions in multi-item two-echelon spare parts inventory systems, where parts belong to different capital good types and where parts that use the same repair resource compete for expedited repair.

2. We provide a tractable optimization model that yields a tight lower bound on the optimal solution and near optimal feasible solutions. We show that our formulation allows us to decompose the non-linear non-convex integer programming problem into sub-problems per repairable type and subsequently use column generation algorithms. For the resulting sub-problem, whose state

space has dimensions equal to the number of locations plus one, we provide an efficient solution algorithm that searches over only two dimensions and where each instance involves independent Newsvendor type problems.

3. As an alternative solution approach, we provide a greedy heuristic that yields excellent results. Different from most literature on greedy heuristics in spare parts inventory systems, our greedy heuristic does not only decide upon stocking levels given a certain target service level, but also on expedite thresholds such that the fraction of the total demand that receives expedited repair per repair resource remains below a certain target level.

4. Based on a case study at NS, we present insights that will help managers to understand how a dynamic repair policy can be leveraged to reduce the total investment costs in spare parts while meeting availability targets.

In his seminal paper on the METRIC model, Sherbrooke (1968) already argued that in practice, parts in short supply should be scheduled into repair first. Though, he and most contributions on the METRIC model assume that the repair lead times of each part are i.i.d. distributed, meaning that no scheduling or prioritization in repairs is possible. As a direct consequence, performance obtained in practice (either investments in stock or availability) is better than theory predicts (e.g., Rustenburg, 2000; Rustenburg et al., 2001), though exact percentages are lacking. In this chapter, we are the first to relax this assumption by explicitly incorporating the possibility to change the repair lead time of a part based on the current state of the system, thereby actually scheduling parts in short supply into repair first. We show that effective usage of this possibility may lead to reductions in stock investments of up to 19.61 percent compared to static repair lead times. We show that such reductions remain attainable even when we incur additional costs for using this flexibility.

The remainder of this chapter is organized as follows. Section 2.2 reviews related literature. In Section 2.3, we provide a formal description of the model. In Section 2.4, we present an exact evaluation procedure for a given control policy as well as the mathematical formulation of our decision problem. Section 2.5 presents two solution approaches to solve this decision problem. We show in Section 2.6 that these solution approaches also apply to an alternate setting in which expediting has an additional cost. Section 2.7 provides managerial insights based on a case study at NS and evaluates the performance of both solution approaches in a large test bed. Finally, some concluding remarks are presented in Section 2.8.

2.2. Literature review

Although spare parts inventory systems have been studied extensively in a variety of settings, our review involves literature with similar modeling assumptions or similar solution approaches as those used in this chapter. For an extensive discussion of the existing literature in the broad field of spare parts inventory management, we refer the reader to Basten and Van Houtum (2014), Van Houtum and Kranenburg (2015) and Muckstadt (2005).

This chapter contributes to the classical research line of multi-item spare parts inventory systems that started in 1968 with the seminal paper of Sherbrooke on the METRIC model. This model assumes that demand follows a Poisson process and that all warehouses operate under base stock policies. Via an approximative evaluation method, expected backorders at all local warehouses are determined for a given control policy. Since then, many extensions have been made to the METRIC model: While some researchers have focused on deriving exact steady state distributions (e.g., Graves, 1985; Simon, 1971), others have extended the model itself by integrating hierarchical or indentured parts structures (e.g., Muckstadt, 1973), by allowing for part failures that lead to downtime after a delay (e.g., Bitton et al., 2019), or by including emergency shipments (e.g., Alfredsson and Verrijdt, 1999; Lee, 1987; Howard et al., 2015). The exact evaluation procedure of our chapter shows similarities with Graves (1985). The main difference is that he considers only one supply mode at the central warehouse, whereas we consider both a regular and an expedited supply mode.

The system studied in this chapter extends previous research which examined inventory models with multiple supply modes. We refer to Svoboda et al. (2021) for an extensive discussion of such inventory models, here we discuss only the important and more relevant results. Since optimal control policies for inventory systems with expediting have complex structures (e.g., Feng et al., 2006; Whittemore and Saunders, 1977), most recent papers study relatively simple heuristic policies and aim at finding (near) optimal parameters.

For single-echelon inventory systems under periodic review, an often studied heuristic policy is the dual-index policy, in which two different inventory positions are kept track off: The inventory position including arrivals within the expedited lead time and the inventory position including arrivals within the regular lead time (e.g., Arts et al., 2011; Sheopuri et al., 2010; Sun and Van Mieghem, 2019; Veeraraghavan and Scheller-Wolf, 2008). Moinzadeh and Schmidt (1991) consider a similar policy for single-echelon inventory systems facing Poisson under continuous review. They focus on obtaining performance measures for a given dual-index policy when both the expedited and regular lead time are deterministic. Song and Zipkin (2009) reinterpret

and extend the work of Moinzadeh and Schmidt (1991) by showing that the same inventory system with a dual-index policy and stochastic lead times is a special type of product form queueing network with one or more overflow bypasses. The dual-index policy in the setting of Moinzadeh and Schmidt (1991) and Song and Zipkin (2009) is in fact optimal for the special case where the regular repair lead time has a shifted exponential distribution and the base stock level for the regular inventory position is fixed (Arts et al., 2016). The policy that we consider for the central warehouse is equivalent to the dual-index policy of Song and Zipkin (2009). The methods of Song and Zipkin (2009) have been incorporated in a two-echelon spare parts inventory system before, albeit to decide upon emergency shipments from a so-called support warehouse to the local warehouses (Howard et al., 2015).

Literature on multiple supply modes in multi-echelon distribution systems is relatively scarce. Building upon the dual-index policy of Moinzadeh and Schmidt (1991), Moinzadeh and Aggarwal (1997) consider a two-echelon distribution system facing Poisson demand under continuous review in which all warehouses have the option to replenish their inventory through an expedited or regular supply channel. Similar to Moinzadeh and Schmidt (1991), they assume deterministic lead times for both types of shipments to all warehouses. Moinzadeh and Aggarwal (1997) describe a procedure to find optimal policy parameters and show that this system substantially improves its single-sourcing counterpart. By contrast to their paper, we consider a system where only the central warehouse has two supply modes. Yet, we impose no limitations on the lead times of those supply modes. For more variations of multi-echelon distribution systems with multiple supply modes under different cost structures and control policies, see Aggarwal and Moinzadeh (1994), Alvarez and Van der Heijden (2014), Dada (1992) and Minner et al. (2003).

A different category of models that shares some aspects of multiple supply modes are multi-echelon distribution systems with lateral transshipments, in which, next to regular replenishments from the central warehouse, transshipments from other warehouses with sufficient inventory are allowed in case of (imminent) stock-outs. Such models have been studied extensively, particularly in spare parts inventory systems (e.g., Axsäter, 1990; Grahovac and Chakravarty, 2001; Kranenburg and Van Houtum, 2009). For a general overview of the existing literature and important results in this field, we refer to Paterson et al. (2011).

Within the stream of literature focusing on inventory systems for repairable items, many contributions have been made on either expediting the repair or prioritizing the scheduling of repairs in the repair shop. As deriving structural properties of optimal policies is known to be complex when the number of different repairable types increases (Tiemessen and Van Houtum, 2013), most contributions in this area resort to heuristic priority rules. We distinguish two categories of such heuristic

priority rules. Under *static* priority rules, the priority of a repairable depends on its type only. Although these type of priority rules are relatively simple, several studies have shown that such rules outperform simple first come first serve rules in terms of investment costs (e.g., Adan et al., 2009; Sleptchenko et al., 2005). Under more sophisticated *dynamic* priority rules, the priority of a repairable also depends on the current state of the system. The expediting policy in our model falls into this latter category as it essentially changes the repair lead time of a part based on the current state of the repair pipeline. In a recent contribution, Arts et al. (2016) study an expediting policy similar to the present model, albeit in a single-echelon single-item setting under fluctuating demand. They remark that this expediting policy does not suffer from the tractability issues that other dynamic priority rules suffer from, while still providing the lead time flexibility inherent to this category of heuristic priority rules. Loeffen (2012) shows that this expediting policy can also be implemented in elaborate simulation models of the repair shop and yield similar performance.

Few researchers have considered dynamic repair priority rules in multi-echelon inventory systems for repairable items. Pyke (1990) jointly addresses dynamic repair and inventory allocation decisions in a two-echelon system very similar to the one we study. He sketches a mathematical formulation of the problem to emphasize its complexity and computational intractability and subsequently resorts to simulation experiments. More recently, Caggiano et al. (2006) consider a similar problem related to dynamic repair and inventory allocation decisions. Different from the present work, their model is a finite-horizon, periodic-review model involving only one repair resource focusing on operational decisions for repairable spare parts in the exploitation phase of capital goods.

On the analysis side, we use two techniques that are widely used in the context of multi-item spare parts inventory optimization. The first technique, decomposition and column generation, is appropriate for problems that have a complicated aggregation constraint that links the different repairable types. Decomposing this problem leads to relatively simple sub-problems per repairable type. This technique has been used extensively in recent contributions on spare part inventory optimization (e.g., Alvarez et al., 2013, 2015; Arts, 2017; Kranenburg and Van Houtum, 2007; Topan et al., 2017; Wong et al., 2007). Most contributions only consider an aggregated service level constraint that links the different repairable types. In this chapter, repairable types are not only linked through such a service level constraint, but also through the maximally allowed mean fraction of expedited repairs over all repairable types that use the same repair resource. Arts (2017) considers a similar optimization model with linking constraints on both expedited repairs and service levels. The major difference between our work and Arts (2017) is that we consider a two-echelon spare parts inventory system. For an extensive discussion on decomposition and column generation, we refer to Dantzig and Wolfe (1960) and Lübbecke and Desrosiers (2005).

The second technique, a greedy method, is a search algorithm that iteratively selects the alternative that has the highest ratio of improvement in performance over cost increase until a feasible solution is obtained. A greedy method is quick, intuitive, easy to implement and provides satisfactory results. Although the technique has been applied in many papers on multi-item spare parts inventory optimization (e.g., Cohen et al., 1990; Kranenburg and Van Houtum, 2009; Topan et al., 2017; Wong et al., 2007), none have proposed a greedy method on both stocking and expediting decisions that yields good results.

2.3. Model description

In this section, we first provide a brief description of the two-echelon spare parts inventory system and introduce the notation that we use throughout this chapter. We then describe the policy we propose to control the system.

2.3.1 Description and notation

We consider a two-echelon spare parts inventory system consisting of a central warehouse and multiple local warehouses. Let the non-empty set of local warehouses be denoted by N_l. The set of all warehouses is denoted by N, i.e. $N = \{0\} \cup N_l$. Hence, the central warehouse has index zero while the local warehouses are numbered as $n = 1, 2, \ldots, |N_l|$. Each local warehouse n is responsible for serving an operating site consisting of a number of capital goods, which may be of the same or different type. Let C denote the non-empty set of capital good types. Each capital good type $c \in C$ consists of a number of critical components that fail infrequently and independently. These critical components are crucial for operating the capital good, i.e. the capital good is down if one of these components fails. The components are at such levels in the material breakdown structure of the capital good that they can be replaced as a whole by spare parts. Component types are also called Stock Keeping Units (SKUs). Let M denote the non-empty set of critical SKUs that occur in the configurations of the different capital good types. The SKUs are numbered as $m = 1, 2, \ldots, |M|$ and each part of SKU $m \in M$ has an acquisition cost c_a^m. The set of SKUs that occur in the configuration of capital good type $c \in C$ is denoted by M_c^C. There is a set of repair resources, denoted by R, that are used to repair failed parts in the repair shop (at the central warehouse). The SKUs that use repair resource $r \in R$ in their repair are contained in the set M_r^R. We assume that M_c^C and M_r^R partition M, i.e. $\cup_{c \in C} M_c^C = \cup_{r \in R} M_r^R = M$ and $\cap_{c \in C} M_c^C = \cap_{r \in R} M_r^R = \emptyset$. This assumption is common in practice and simplifies notation considerably; it is however not essential to our analysis. We briefly show in Appendix 2.D how this assumption can be readily

relaxed along similar lines as is done in Kranenburg and Van Houtum (2007).

Demand for SKU $m \in M$ at local warehouse $n \in N_l$ is a Poisson process with rate $\lambda_{m,n}$. This demand model is common in literature and accurate in practice for spare parts (e.g., Graves, 1985). When a demand for SKU m occurs at local warehouse n, it will be filled from stock, or backordered if the stock is depleted. In the latter case, the capital good remains down until a spare part becomes available at the local warehouse. The failed part is shipped to the repair shop at the central warehouse, where all failed parts are immediately sent into regular or expedited repair, where the corresponding resource $r \in R$ is used for repair. At the same time, the central warehouse ships a spare part to the local warehouse from its inventory, if it has an available spare part. Otherwise, the replenishment order is backordered at the central warehouse until a part is repaired and becomes available. Upon completion of repair a part is put back on stock at the central warehouse.

The order and shipment time for a spare part of SKU m from the central warehouse to local warehouse n is fixed and denoted by $t_{m,n}$. Note that $t_{m,n}$ excludes any waiting time at the central depot when a spare part is not available. For returned failed parts at the repair shop, it takes either $t_{m,0}^{reg}$ time units, in case of the regular repair, or $t_{m,0}^{exp}$ time units, in case of the expedited repair, until the part is returned to the spare parts stock at the central warehouse. We assume that both repair times are fixed, with $t_{m,0}^{reg} > t_{m,0}^{exp} > 0$. Figure 2.1 provides a graphical representation of the system under consideration and notation is summarized in Table 2.1 (including notation introduced later).

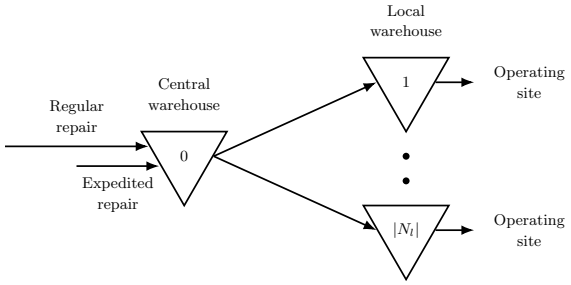

Figure 2.1 Two-echelon spare parts inventory system with expediting.

2.3.2 Control policy

Each failed part at a local warehouse results in an immediate replenishment order at the central warehouse. This implies that the inventory positions of a given SKU $m \in M$ remain constant at all local warehouses. Hence, we have base stock control at

Table 2.1 Overview of notation.

Notation	Description
Sets	
N	Set of all warehouses.
$N_l \subset N$	Set of local warehouses.
M	Set of all SKUs.
C	Set of capital good types.
R	Set of repair resources.
$M_r^R \subseteq M$	Set of SKUs that use repair resource $r \in R$ in the repair of failed parts.
$M_c^C \subseteq M$	Set of SKUs that occur in the configuration of capital good type $c \in C$.
Input parameters	
$\lambda_{m,n}$	Demand intensity for SKU $m \in M$ at warehouse $n \in N$.
δ_m^r	Fraction of demands over all parts of SKUs $k \in M_r^R$ that are from SKU m, i.e. $\frac{\lambda_{m,0}}{\sum_{k \in M_r^R} \lambda_{k,0}}$.
$t_{m,n}$	Lead time from the central warehouse to local warehouse $n \in N_l$ of SKU $m \in M$.
$t_{m,0}^{reg}$	Regular repair lead time of SKU $m \in M$.
$t_{m,0}^{exp}$	Expedited repair lead time of SKU $m \in M$, also denoted by $t_{m,0}^2$.
$t_{m,0}^1$	Additional regular repair lead time of SKU $m \in M$.
c_a^m	Acquisition cost for SKU $m \in M$.
τ	Useful lifespan of each SKU $m \in M$.
κ	Expediting cost multiplier for each SKU $m \in M$.
c_e^m	Cost for expediting the repair of SKU $m \in M$, i.e. $\kappa \cdot c_a^m$.
c_d^m	Depreciation cost rate for SKU $m \in M$, i.e. c_a^m/τ.
$\mathcal{B}_c^{max}$	The maximally allowed mean number of backorders over all SKUs $m \in M_c^C$ for capital good type $c \in C$.
$\mathcal{E}_r^{max}$	The maximally allowed mean fraction of expedited repairs over all SKUs $m \in M_r^R$ that use repair resource $r \in R$ during their repair.
Decision variables	
$S_{m,n}$	Base stock level of SKU $m \in M$ at warehouse $n \in N$.
$\mathbf{S}_m$	The vector $(S_{m,0}, S_{m,1}, \ldots, S_{m,\lvert N_l \rvert})$.
$\mathbf{S}$	The base stock levels matrix $[S_{m,n}]$.
T_m	Expedite threshold of SKU $m \in M$.
T	The vector $(T_1, T_2, \ldots, T_{\lvert M \rvert})$.
State variables	
$X_{m,0}$	Number of outstanding repairs of SKU $m \in M$ at the central warehouse.
$X_{m,0}^2$	Number of outstanding repairs of SKU $m \in M$ at the central warehouse that will be repaired within $t_{m,0}^2$ time units.
$X_{m,0}^1$	Number of outstanding repairs of SKU $m \in M$ at the central warehouse that will not be repaired within $t_{m,0}^2$ time units, i.e. $X_{m,0}$-$X_{m,0}^2$.
$X_{m,n}$	Number of outstanding orders of SKU $m \in M$ at local warehouse $n \in N_l$.
Output of model	
$EBO_{m,n}(\mathbf{S}_m, T_m)$	Mean number of backorders for SKU m at local warehouse $n \in N_l$ under a given control policy $(\mathbf{S}_m, T_m)$, i.e. $\sum_{x=S_{m,n}+1}^{\infty}(x - S_{m,n})\mathbb{P}\{X_{m,n} = x\}$.
$EBO_c(\mathbf{S}, T)$	Aggregate mean number of backorders for capital good type $c \in C$ under a given control policy $(\mathbf{S}, T)$, i.e. $\sum_{m \in M_c^C} \sum_{n \in N_l} EBO_{m,n}(\mathbf{S}_m, T_m)$.
$EXP_m(T_m)$	Fraction of failed parts of SKU $m \in M$ that utilize the expedited repair option under a given expedite threshold (T_m), i.e. $\mathbb{P}\{X_{m,0}^1 = T_m\}$.
$EXP_r(T)$	Aggregate mean fraction of failed parts over all SKUs $m \in M_r^R$ that utilize the expedited repair option under a given expedite threshold vector T, i.e. $\sum_{m \in M_r^R} \delta_m^r EXP_m(T_m)$.
$C(\mathbf{S})$	The total investment costs in spare parts under a given base stock levels matrix $\mathbf{S}$, i.e. $\sum_{m \in M} \sum_{n \in N} c_a^m S_{m,n}$.
$C_d(\mathbf{S})$	The total depreciation cost rate in spare parts under a given base stock levels matrix $\mathbf{S}$, i.e. $\sum_{m \in M} \sum_{n \in N} c_d^m S_{m,n}$.
$C_e(T)$	The total repair expediting cost rate under a given expedite threshold vector T, i.e. $\sum_{m \in M} c_e^m \lambda_{m,0} EXP_m(T_m)$.
C_P^{UB} (C_P^{LB})	Upper (lower) bound for the optimal solution to problem (P).
C_{BM}^{LB}	Lower bound for the optimal solution of a benchmark instance BM.

each local warehouse $n \in N_l$ for each SKU $m \in M$ and we denote the corresponding base stock levels by $S_{m,n}$.

The central warehouse is controlled by a dual-index policy. This policy has two parameters for each SKU $m \in M$, integers $S_{m,0}$ and $S'_{m,0}$, with $S_{m,0} \geq S'_{m,0}$. Let $t^1_{m,0} = t^{reg}_{m,0} - t^{exp}_{m,0}$, i.e. the additional regular lead time, and $t^2_{m,0} = t^{exp}_{m,0}$. We define two inventory positions for each SKU m: $IP^1_{m,0}$ and $IP^2_{m,0}$. $IP^1_{m,0}$ is the usual local inventory position and includes net inventory $IN_{m,0}$ (on-hand stock $OH_{m,0}$ minus any backorders $BO_{m,0}$) plus all parts in repair $X_{m,0}$. $IP^2_{m,0}$ is similar but only includes those parts in repair $X^2_{m,0}$ that will be repaired and returned to on-hand stock within $t^2_{m,0}$ time units. Hence, the number of parts in repair that will not be repaired and returned to on-hand stock within $t^2_{m,0}$ time units $X^1_{m,0}$ is equal to $X_{m,0} - X^2_{m,0}$. Figure 2.2 provides a graphical representation of the two different inventory positions at the central warehouse.

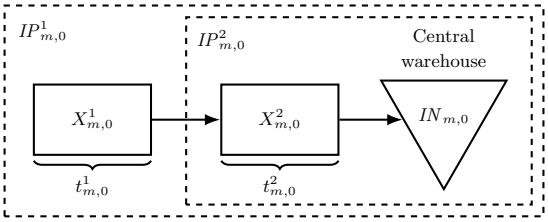

Figure 2.2 Inventory positions at the central warehouse.

The dual-index policy works as follows: Keep $IP^1_{m,0}$ at constant level $S_{m,0}$ (as in standard base stock control) and also $IP^2_{m,0} \geq S'_{m,0}$. Thus upon the demand of a part and the return of a failed part of SKU $m \in M$, we first examine $IP^2_{m,0}$. If $IP^2_{m,0}$ (after the failed part is returned, but before deciding upon the repair option) is already $S'_{m,0}$ or greater, we send it into regular repair. However, if a regular repair would leave $IP^2_{m,0} < S'_{m,0}$, then we use the expedited repair option. Note that $IP^2_{m,0} = IN_{m,0} + X^2_{m,0} = IP^1_{m,0} - X^1_{m,0} = S_{m,0} - X^1_{m,0}$, and thus, equivalently, the dual-index policy keeps $S_{m,0} - X^1_{m,0} \geq S'_{m,0}$. Hence, defining expedite threshold $T_m = S_{m,0} - S'_{m,0}$ $\forall m \in M$, the dual-index policy sends failed parts into regular repair as long as $X^1_{m,0} \leq T_m$ (cf. Song and Zipkin, 2009).

Let $\mathbf{S}_m = (S_{m,0}, S_{m,1}, \ldots, S_{m,|N_l|})$, $m \in M$, denote the vector of base stock levels for SKU m. Then, a control policy $(\mathbf{S}, T)$ is denoted by base stock levels matrix $\mathbf{S}$ and a vector $T = (T_1, T_2, \ldots, T_{|M|})$ containing the expedite thresholds of each SKU $m \in M$.

2.4. Performance evaluation and problem formulation

In this section, we provide an exact evaluation procedure for a given control policy $(\mathbf{S}, T)$, and we present the mathematical formulation of the decision problem.

2.4.1 Exact evaluation of a given control policy

The evaluation of a given control policy $(\mathbf{S}, T)$ can be done per SKU. Consider therefore some SKU $m \in M$ that has base stock vector $\mathbf{S}_m$ and expedite threshold T_m. We first consider the performance of SKU m at the central warehouse, and subsequently link this to its performance at all local warehouses.

Key in evaluating the performance of the central warehouse for SKU m is to obtain the distribution of the number of parts in repair $X_{m,0}$. Since each failure of SKU m results in an immediate replenishment request for SKU m at the central warehouse, the demand process for parts of SKU m as seen by the central warehouse is a Poisson process with constant rate $\lambda_{m,0} = \sum_{n \in N_l} \lambda_{m,n}$. Each replenishment request for SKU m is accompanied by a failed part that goes into repair. Hence, failed parts of SKU m enter the repair pipeline according to a Poisson process with constant rate $\lambda_{m,0}$. The fraction of demands for SKU m over demands from all SKUs that use the same repair resource $r \in R$ as SKU m uses, is then given by $\delta_m^r = \frac{\lambda_{m,0}}{\sum_{k \in M_r^R} \lambda_{k,0}}$. Now, under the dual-index policy described in the previous section, the repair pipeline of each SKU m can be seen as an open queueing network with outside Poisson arrivals at constant rate $\lambda_{m,0}$ and two $\bullet/D/\infty$ queues that cause delays of t_m^1 and t_m^2 time units (see Figure 2.3).

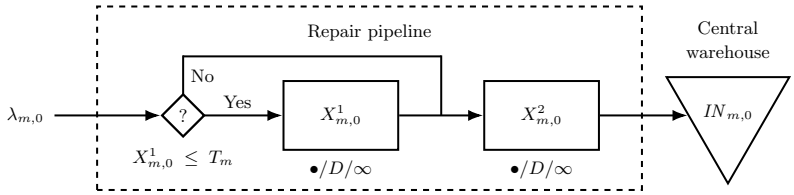

Figure 2.3 Repair pipeline as an open queueing network.

Observe that a normal repair first passes through queue 1, where it remains t_m^1 units, and then passes through queue 2, where it remains t_m^2 units. Failed parts that receive expedited repair bypass the first part of the repair pipeline and only pass through queue 2. After a part completes queue 2, it arrives to inventory at the central warehouse. In effect, the dual-index policy directs failed parts of SKU m into normal repair as long as the number of repairs at queue 1, i.e. $X_{m,0}^1$, is not greater than the expedite threshold T_m. When an arriving failed part would overflow T_m, the failed

part bypasses queue 1 and goes directly to queue 2, that is, the failed part goes into expedited repair. In the queueing literature, this is sometimes referred to as *jump over blocking* or as an *overflow bypass* (e.g., Lam, 1977; Song and Zipkin, 2009).

The distribution of the number of parts in repair $X_{m,0}$, follows from the joint distribution of $(X^1_{m,0}, X^2_{m,0})$. Let $p_m(x_1, x_2) = \mathbb{P}\{X^1_{m,0} = x_1, X^2_{m,0} = x_2\}$ denote the steady-state joint distribution of $(X^1_{m,0}, X^2_{m,0})$. Let $\phi^1_m(x_1)$ and $\phi^2_m(x_2)$ denote the Poisson probabilities $e^{\lambda_{m,0}\cdot t^1_m}(\lambda_{m,0} \cdot t^1_m)^{x_1}/x_1!$ and $e^{\lambda_{m,0}\cdot t^2_m}(\lambda_{m,0} \cdot t^2_m)^{x_2}/x_2!$, respectively. The support of $(X^1_{m,0}, X^2_{m,0})$ is denoted by $\mathscr{X}(T_m) = \{(x_1, x_2) \in \mathbb{N}^2_0 : x_1 \leq T_m\}$. Then, as shown by Lam (1977) and Song and Zipkin (2009), the joint distribution of $(X^1_{m,0}, X^2_{m,0})$ has product-form

$$p_m(x_1, x_2) = \frac{\phi^1_m(x_1)\phi^2_m(x_2)}{\sum_{x_1 \leq T_m} \phi^1_m(x_1)}, \qquad (x_1, x_2) \in \mathscr{X}(T_m).$$

Letting $p_m(x) = \mathbb{P}\{X_{m,0} = x\}$ denote the equilibrium probability of the number of parts in repair $X_{m,0}$ and $\phi_m(x)$ denote the Poisson probability $e^{\lambda_{m,0}\cdot t^{reg}_m}(\lambda_{m,0} \cdot t^{reg}_m)^x/x!$, we obtain

$$\begin{aligned} p_m(x) &= \sum_{i=0}^{x} p_m(i, x-i) \\ &= \left(\sum_{x_1 \leq T_m} \phi^1_m(x_1) \right)^{-1} \begin{cases} \phi_m(x) & x \leq T_m, \\ \sum\limits_{i=0}^{T_m} \phi^1_m(i)\phi^2_m(x-i) & x > T_m. \end{cases} \end{aligned}$$

The main performance measures of the central warehouse are now easily obtained. In particular, the number of backorders $BO_{m,0}$ for SKU m is equal to $(X_{m,0} - S_{m,0})^+$, where $x^+ = \max(0, x)$. For its probability distribution, we have

$$\mathbb{P}\{BO_{m,0} = x\} = \begin{cases} \sum\limits_{j=0}^{S_{m,0}} p_m(j), & x = 0, \\ p_m(S_{m,0} + x), & x > 0. \end{cases}$$

In addition, letting $\rho_m = \lambda_{m,0} \cdot t^1_{m,0}$, the fraction of failed parts of SKU m that utilize the expedite repair option is given by

$$EXP_m(T_m) = \mathbb{P}\{X^1_{m,0} = T_m\} = \frac{\rho^{T_m}_m}{T_m!} \left(\sum_{i=0}^{T_m} \frac{\rho^i_m}{i!} \right)^{-1},$$

which is the (Erlang) blocking probability of an $M/G/c/c$ queue, where the numbers

of parallel servers c is equal to expedite threshold T_m (e.g., Gross et al., 2008).

Key in evaluating the performance of each local warehouse $n \in N_l$ for SKU m is to obtain the distribution of orders outstanding for each local warehouse. Therefore we need to determine the distribution of backorders for a SKU m at the central warehouse that belong to local warehouse n.

Simon (1971) shows that when outstanding orders at the central warehouse are filled on a first-come first-served bases, then each backorder at the central warehouse belongs to local warehouse n with probability $\frac{\lambda_{m,n}}{\lambda_{m,0}}$, independently across backorders. Let $BO^n_{m,0}$ denote the number of backorders of local warehouse n in the backorder queue of SKU m at the central warehouse. Then, by conditioning on the number of backorders of SKU m at the central warehouse and using Simon's result that the conditional distribution of $BO^n_{m,0}$ is a binomial distribution, we obtain the following probability distribution for this number of backorders

$$\begin{aligned}\mathbb{P}\{BO^n_{m,0} = x\} &= \sum_{y=x}^{\infty} \mathbb{P}\{BO^n_{m,0} = x | BO_{m,0} = y\}\mathbb{P}\{BO_{m,0} = y\} \\ &= \sum_{y=x}^{\infty} \binom{y}{x} \left(\frac{\lambda_{m,n}}{\lambda_{m,0}}\right)^x \left(1 - \frac{\lambda_{m,n}}{\lambda_{m,0}}\right)^{y-x} \mathbb{P}\{BO_{m,0} = y\}. \qquad (2.1)\end{aligned}$$

Now, we determine the distribution of the outstanding orders of SKU m at each local warehouse n. The outstanding orders at any time t consists of demand that occurred in the interval $(t - t_{m,n}, t)$ (notation $D_{m,n}(t - t_{m,n}, t)$) and backorders at the central warehouse that belong to local warehouse n at time $t - t_{m,n}$ (notation $BO^n_{m,0}(t)$), i.e. $X_{m,n}(t) = D_{m,n}(t - t_{m,n}, t) + BO^n_{m,0}(t)$. Since the Poisson process has independent increments, $D_{m,n}(t - t_{m,n}, t)$ and $BO^n_{m,0}(t)$ are independent random variables so that in stationary state

$$X_{m,n} = D_{m,n} + BO^n_{m,0},$$

where $D_{m,n}$ has a Poisson distribution with mean $\lambda_{m,n}t_{m,n}$ and the distribution of $BO^n_{m,0}$ is given in (2.1). Therefore the stationary distribution of $X_{m,n}$ is obtained by convolution.

From this point, the main performance measure of each local warehouse is easily obtained. For SKU m, the number of backorders $BO_{m,n}$ at local warehouse n is equal to $(X_{m,n} - S_{m,n})^+$. Its probability distribution is then obtained in a similar way as the probability distribution of the number of backorders at the central warehouse. In particular, the mean number of backorders for SKU m at local warehouse n is given

by

$$EBO_{m,n}(\mathbf{S}_m, T_m) = \sum_{x=S_{m,n}+1}^{\infty} (x - S_{m,n})\mathbb{P}\{X_{m,n} = x\}. \tag{2.2}$$

2.4.2 Problem formulation

For a given control policy $(\mathbf{S}, T)$, we define the total investment costs in spare parts as

$$C(\mathbf{S}) = \sum_{m\in M}\sum_{n\in N} c_a^m S_{m,n},$$

the aggregate mean number of backorders for capital good type c as

$$EBO_c(\mathbf{S}, T) = \sum_{m\in M_c^C}\sum_{n\in N_l} EBO_{m,n}(\mathbf{S}_m, T_m),$$

and the aggregate mean fraction of failed parts of all SKUs $m \in M_r^R$ that utilize the expedited repair option using repair resource $r \in R$ as

$$EXP_r(T) = \sum_{m\in M_r^R} \delta_m^r EXP_m(T_m).$$

The objective of our decision problem is to minimize the total investment costs in spare parts while keeping the mean number of aggregate backorders for each capital good type $c \in C$ below $\mathcal{B}_c^{max}$ and keeping the fraction of repairs that are expedited per repair resource $r \in R$ below $\mathcal{E}_r^{max}$. Combining the aforementioned results in the following mathematical formulation of our decision problem which we call problem (P):

$$\begin{array}{lllll} (P) & \min & C(\mathbf{S}) & & \\ & \text{subject to} & EBO_c(\mathbf{S}, T) \le \mathcal{B}_c^{max}, & \forall c \in C \\ & & EXP_r(T) \le \mathcal{E}_r^{max}, & \forall r \in R \\ & & \mathbf{S} \in \mathscr{S}, \quad T \in \mathbb{N}_0^{|M|}, & \end{array}$$

where $\mathscr{S} = \{\mathbf{S} : S_{m,n} \in \mathbb{N}_0,\ \forall m \in M \text{ and } \forall n \in N\}$. Let $(\mathbf{S}^*, T^*)$ denote an optimal solution to problem (P) and let C_P be the corresponding optimal cost.

We note that problem (P) can be considered as a non-linear non-convex knapsack problem with multiple constraints, where more than one copy of each item can be

selected. It is well-known that even the simplest type of knapsack problems belongs to the class of $\mathcal{NP}$-hard problems (Kellerer et al., 2004). As our knapsack problem is more complex, it is very likely that also for problem (P) no polynomial time optimization algorithm exists.

2.5. Optimization of base stock levels and expedite thresholds

The focus of this section is on finding the optimal base stock levels and expedite thresholds. We first present a decomposition and column generation (DCG) algorithm to construct a lower bound for problem (P). We then show how the sub-problem of this algorithm can be solved efficiently. We continue with showing how to find a good feasible solution for problem (P). We conclude this section with devising an alternative solution approach in which we greedily optimize the expedite thresholds and base stock levels.

2.5.1 Constructing lower bounds

We first reformulate problem (P) as a partitioning problem so that we can apply the technique of column generation (also known as Dantzig-Wolfe decomposition). This technique was pioneered by Dantzig and Wolfe (1960) and a thorough modern treatment is given by Lübbecke and Desrosiers (2005). Thus we obtain an integer linear program for which we relax the integrality constraints. We refer to this problem as the master problem (MP). Let K_m be the set of all policies k for SKU $m \in M$ that are feasible for problem (P). Each policy $k \in K_m$ has base stock vector $\mathbf{S}_m^k := (S_{m,0}^k, S_{m,1}^k, \ldots, S_{m,|N_l|}^k)$ and expedite threshold T_m^k. Let $x_m^k \in \{0,1\}$, $m \in M$, $k \in K_m$, denote the decision variable indicating whether policy k is chosen ($x_m^k = 1$) for SKU m or not ($x_m^k = 0$). Then, by relaxing the integrality constraint on x_m^k, the

master problem (MP) is defined as follows:

$$(MP) \quad \min \quad \sum_{m\in M}\sum_{n\in N}\sum_{k\in K_m} c_a^m S_{m,n}^k x_m^k \tag{2.3}$$

$$\text{subject to} \sum_{m\in M_c^C}\sum_{n\in N_l}\sum_{k\in K_m} EBO_{m,n}(\mathbf{S}_m^k, T_m^k)x_m^k \leq \mathcal{B}_c^{max}, \forall c \in C \tag{2.4}$$

$$\sum_{m\in M_r^R}\sum_{k\in K_m} \delta_m^r EXP_m(T_m^k)x_m^k \leq \mathcal{E}_r^{max}, \qquad \forall r \in R \tag{2.5}$$

$$\sum_{k\in K_m} x_m^k = 1, \qquad \forall m \in M \tag{2.6}$$

$$x_m^k \geq 0, \qquad \forall m \in M, \; \forall k \in K_m$$

Let C_P^{LB} denote the optimal cost for master problem (MP). Due to the relaxation of the integrality constraint on x_m^k, an optimal cost C_P^{LB} is also a lower bound on the optimal cost for problem (P), C_P.

Since the set K_m contains an infinite number of policies, a restricted master problem (RMP) is introduced in which, for each SKU $m \in M$, only a small subset of policies $K_m^{res} \subseteq K_m$ is considered. After solving (RMP) to optimality, we are interested in policies $K_m \setminus K_m^{res}$ that will improve the solution of (RMP) if they are added. To check whether such policies exist, we solve, for each SKU m, a column generation sub-problem. To this end, let p_c denote the dual variable of (RMP) corresponding with the expected backorder constraint (2.4) for capital good type $c \in C$, let ρ_r denote the dual variable of (RMP) corresponding with the expected fraction of expedited repairs constraint (2.5) for repair resource $r \in R$ and let υ_m denote the dual variable of (RMP) corresponding to constraint (2.6) that assures that for each SKU $m \in M$ a convex combination of policies is chosen. (The dual variables p_c, ρ_r, and υ_m can also be interpreted as Lagrange multipliers of relaxing the corresponding constraints; see Brooks and Geoffrion (1966) and Lübbecke and Desrosiers (2005).) Then, the column generation sub-problem for SKU $m \in M_r^R \cap M_c^C$ of (RMP) is given by:

$$(SP(m)) \quad \min \sum_{n\in N} c_a^m S_{m,n} - p_c \sum_{n\in N_l} EBO_{m,n}(\mathbf{S}_m, T_m) - \rho_r \delta_m^r EXP_m(T_m) - \upsilon_m \tag{2.7}$$

$$\text{subject to} \quad \mathbf{S}_m \in \mathbb{N}_0^{|N|}, \quad T_m \in \mathbb{N}_0. \tag{2.8}$$

If a feasible solution to $(SP(m))$ exists with a negative objective value, then the objective of (RMP) can be improved by adding this policy to K_m^{res} and solving (RMP) with the larger set K_m^{res}. An optimal solution for (RMP) is also an optimal solution

for (MP) if for none of the SKUs a policy with negative reduced costs exists.

In the next section, we present an exact solution method to solve $(SP(m))$. However, we remark that all policies that yield a negative objective value for $(SP(m))$, can improve the solution of (RMP). Hence, we do not necessarily have to solve $(SP(m))$ to optimality each time we obtain new dual variables from (RMP).

2.5.2 Solving the sub-problem

This section treats an exact solution method for $(SP(m))$. All proofs are in Appendix 2.A. If we fix the control policy parameters at the central warehouse, then this warehouse simply becomes a supplier with a known stochastic lead time from the perspective of each local warehouse. Hence, for fixed T_m and $S_{m,0}$, each local warehouse $n \in N_l$ operates as an independent Newsvendor subsystem, and we can optimize them separately:

Theorem 2.1. *The optimal $S_{m,n}$, $n \in N_l$, for fixed values of $S_{m,0}$ and T_m, $S^*_{m,n}(S_{m,0}, T_m)$, is the smallest $S_{m,n}(S_{m,0}, T_m)$ that satisfies*

$$\mathbb{P}\{X_{m,n}(S_{m,0}, T_m) \leq S_{m,n}(S_{m,0}, T_m)\} \geq \frac{p_c + c_a^m}{p_c}. \tag{2.9}$$

The remaining problem of finding the optimal control policy parameters at the central warehouse is more involved. In fact, it is known that objective function (2.7) is not convex in $S_{m,0}$ for a fixed T_m and corresponding $S^*_{m,n}$, $n \in N_l$ (e.g., Gallego et al., 2007; Rong et al., 2017). Similarly, it can readily be verified that the objective function (2.7) is also not convex in T_m for a fixed $S_{m,0}$ and corresponding $S^*_{m,n}$, $n \in N_l$. Finding the optimal control policy parameters at the central warehouse therefore requires an enumerative search.

To simplify this search, we establish an upper bound on the optimal base stock level at the central warehouse for a given expedite threshold. If the expedite threshold is fixed and the local warehouses carry no inventories, then only the base stock level at the central warehouse can influence the expected backorders at all local warehouses. Hence, the following lemma shows that for fixed T_m and $S_{m,n} = 0 \ \forall n \in N_l$, the central warehouse also operates as an independent Newsvendor subsystem:

Lemma 2.1. *The optimal $S_{m,0}$ for fixed T_m and $S_{m,n} = 0$ for all $n \in N_l$, say $\bar{S}_{m,0}(T_m)$, is the smallest $S_{m,0}$ that satisfies*

$$\mathbb{P}\{X_{m,0}(T_m) \leq S_{m,0}\} \geq \frac{p_c + c_a^m}{p_c}. \tag{2.10}$$

Observe that if the local warehouses increase their base stock levels, then the amount of inventory that the central warehouse should carry can only decrease (assuming that the expedite threshold is fixed). It is therefore clear that $\bar{S}_{m,0}(T_m)$ obtained using Lemma 2.1 is in fact an upper bound on $S^*_{m,0}(T_m)$ because it assumes no inventories at the local warehouses. This is formalized in the next two results.

Lemma 2.2. *Let $S^*_{m,n}(S_{m,0}, T_m)$ be the optimal value of $S_{m,n}$, $n \in N_l$, for given values of T_m and $S_{m,0}$. Then $S^*_{m,n}(S_{m,0}, T_m)$ is non-increasing in $S_{m,0}$.*

Theorem 2.2. *$\bar{S}_{m,0}(T_m)$, as specified in Lemma 2.1, is an upper bound for $S^*_{m,0}(T_m)$.*

Based on the results presented above, we propose the following exact solution method to solve $(SP(m))$. We set T_m to 0 and then search over T_m. For each value of T_m, we vary $S_{m,0}$ over $0 \leq S_{m,0} \leq \bar{S}_{m,0}(T_m)$, where $\bar{S}_{m,0}(T_m)$ is determined using Lemma 2.1. For each pair $(S_{m,0}, T_m)$, we optimize $S_{m,n}$ for all $n \in N_l$ using Theorem 2.1. Since the objective function of $(SP(m))$ for fixed values of $S_{m,0}$ and corresponding $S^*_{m,n}(S_{m,0}, T_m)$, $n \in N_l$, is not convex in T_m, we continue the search over T_m by examining a few values beyond the last observed local minimum.

2.5.3 Constructing a good feasible solution

When no more policies can be added to K_m^{res}, then a solution to the final version of problem (RMP) provides a lower bound, C_P^{LB}, on the optimal cost for problem (P), C_P. In case there are no fractional solutions for any x_m^k, $m \in M$, $k \in K_m$, this also is an upper bound, C_P^{UB}, for C_P. If there are fractional solutions for any x_m^k, we solve the final version of problem (RMP) as an integer linear program. Alvarez et al. (2013, 2015) show that this approach yields very good results compared to other methods such as local search algorithms. To speed up the solution process of solving the final version of problem (RMP) as an integer linear program, we use the feasibility pump heuristic of Fischetti et al. (2005), and we stop the solution of the integer linear program as soon as a feasible solution with optimality gap of less than 0.5 percent is found or 1 minute has elapsed (whichever occurs first). This results in a good feasible solution to problem (P). The corresponding cost of this solution is also an upper bound, C_P^{UB}, for C_P.

Pseudo-code of the DCG algorithm as well as the greedy heuristic described in the next section can be found in Appendix 2.B.

2.5.4 A two-step greedy approach

We now describe a greedy heuristic for problem (P). This greedy heuristic consists of two steps that are executed consecutively. In the first step, we determine, independent

of base stock level matrix $\mathbf{S}$, expedite threshold vector T. Subsequently, based on the vector of expedite thresholds T determined in the first step, we find base stock levels matrix $\mathbf{S}$.

Expediting the repair of a given SKU $m \in M$ implies that fewer parts of m are needed to provide the same availability as when no repairs are expedited. Hence, given that repair resources are limited, we want to expedite the repair of expensive parts more often than cheaper parts. In addition, the cost benefit of expediting the repair of a given SKU $m \in M$ increases in its additional regular lead time, i.e. $t^1_{m,0}$. Hence, given that repair resources are limited, we want to expedite the repair of parts with a greater additional regular repair lead time more often than parts with a smaller additional regular repair lead time.

If there were no restrictions on the aggregate mean fractions of failed parts that are expedited, then, irrespective of base stock levels matrix $\mathbf{S}$, the zero vector would be the optimal vector of expedite thresholds. Hence, in the first step of the greedy heuristic, we set all expedite thresholds T_m, $m \in M$, to zero and then start with greedy steps, in which we increase T_m leading to the largest decrease in distance to the set of feasible expedite vectors per acquisition cost and additional regular repair lead time.

The first step of the greedy heuristic is formally described as follows. We first partition the set of all expedite thresholds vectors T into a subset T^{feas} of expedite thresholds vectors that are feasible for Problem (P) and a subset $\mathbb{N}_0^{|M|} \setminus T^{feas}$ of expedite thresholds vectors that are infeasible. Next, for each expedite thresholds vector, we define the distance $d(T)$ to T^{feas} as

$$d(T) = \sum_{r \in R} \left(EXP_r(T) - \mathcal{E}_r^{max}\right)^+ .$$

In each greedy step, we have a current solution $T \in \mathbb{N}_0^{|M|} \setminus T^{feas}$, and we look at the ratio of the decrease in distance to T^{feas} if T_m, $m \in M$, is increased by one unit and the product of the acquisition cost and the additional regular repair lead time. To this end, let $-\Delta_m d(T)$ denote the decrease in distance to the set of feasible vectors of expedite thresholds. For a given SKU $m \in M$ that uses repair resource $r \in R$ in the repair of its failed parts, we obtain

$$\begin{aligned} \Delta_m d(T) &= d(T + \mathbf{e}_m) - d(T) \\ &= \left(EXP_r(T + \mathbf{e}_m) - \mathcal{E}_r^{max}\right)^+ - \left(EXP_r(T) - \mathcal{E}_r^{max}\right)^+ , \end{aligned}$$

where $\mathbf{e}_m$ is an $|M|$-dimensional vector with a 1 on position m and zero otherwise.

Since the Erlang loss formula, and thus $EXP(T_m)$, is convex and decreasing in T_m

(e.g., Messerli, 1972), it follows that $-\Delta_m d(T) \geq 0$ for all $m \in M$. The ratio

$$\Gamma_m^T = \frac{-\Delta_m d(T)}{t_{m,0}^1 c_a^m}$$

denotes the decrease in distance to the set of feasible vectors of expedite thresholds per both the acquisition cost and the additional regular repair lead time. During each greedy step, we increase the expedite threshold of SKU m with the highest Γ_m^T to $T_m + 1$. We continue with these steps until we arrive at a feasible solution T and we denote this solution by $\bar{T}$.

We now proceed with the second step of the greedy heuristic. If there were no restrictions on the aggregate mean numbers of backorders, then, irrespective of the vector of expedite thresholds, the zero matrix would be the optimal base stock levels matrix. Hence, in the second step of the greedy heuristic, we set all base stock levels $S_{m,n}$, $m \in M$, $n \in N$, to zero and then start with greedy steps, in which we increase $S_{m,n}$ leading to the largest decrease in distance to the set of feasible base stock levels matrices per acquisition cost.

The second step of the greedy heuristic is formally described as follows. We first partition the set of all base stock levels matrices $\mathscr{S}$ into a subset $\mathscr{S}^{feas}$ of base stock levels matrices that are feasible for Problem (P) and a subset $\mathscr{S} \setminus \mathscr{S}^{feas}$ of base stock levels matrices that are infeasible. Next, for each base stock levels matrix, we define the distance $d(\mathbf{S}, \bar{T})$ to $\mathscr{S}^{feas}$ as

$$d(\mathbf{S}, \bar{T}) = \sum_{c \in C} \left(EBO_c(\mathbf{S}, \bar{T}) - \mathcal{B}_c^{max} \right)^+ .$$

In each greedy step, we have a current solution $\mathbf{S} \in \mathscr{S} \setminus \mathscr{S}^{feas}$, and we look at the ratio of the decrease in distance to $\mathscr{S}^{feas}$ and the acquisition cost if $S_{m,n}$, $m \in M$, $n \in N$, is increased by one unit. To this end, let $-\Delta_{m,n} d(\mathbf{S}, \bar{T})$ denote the decrease in distance to the set of feasible base stock levels matrices. For each SKU $m \in M$ and warehouse $n \in N$, let $\mathbf{E}_{m,n}$ be an $|M| \times |N|$ matrix with positions (m', n'), $m' \in M$, $n' \in N$, with ones on positions m and n and zero otherwise. Then, for a given SKU $m \in M$ of capital good type $c \in C$, we obtain

$$\begin{aligned} \Delta_{m,n} d(\mathbf{S}, \bar{T}) &= d(\mathbf{S} + \mathbf{E}_{m,n}, \bar{T}) - d(\mathbf{S}, \bar{T}) \\ &= \left(EBO_c(\mathbf{S} + \mathbf{E}_{m,n}, \bar{T}) - \mathcal{B}_c^{max} \right)^+ - \left(EBO_c(\mathbf{S}, \bar{T}) - \mathcal{B}_c^{max} \right)^+ . \end{aligned}$$

Increasing the base stock level of a given SKU $m \in M$ at the central warehouse has a decreasing effect on the expected backorders at all local warehouses $n \in N_l$, and no effect on the expected backorders of all other SKUs. Moreover, increasing the base

stock level of a given SKU $m \in M$ at some local warehouse $n \in N_l$ has a decreasing effect on the expected backorders for that SKU at that local warehouse and no effect on all other expected backorders. These assertions are easily verified along similar lines as the proof of Lemma 2.2. It then immediately follows that $-\Delta_{m,n} d(\mathbf{S}, \bar{T}) \geq 0$ for all $m \in M$ and $n \in N$. The ratio

$$\Gamma^{\mathbf{S}}_{m,n} = \frac{-\Delta_{m,n} d(\mathbf{S}, \bar{T})}{c^m_a}$$

denotes the decrease in distance to the set of feasible base stock levels matrices per acquisition cost. During each greedy step, we increase the base stock level of SKU m at warehouse n with the highest $\Gamma^{\mathbf{S}}_{m,n}$ to $S_{m,n}+1$. We continue with these steps until we arrive at a feasible solution $\mathbf{S}$.

2.6. Expediting repairs at additional costs

We have so far considered a constraint on the aggregate mean fraction of failed parts that are expedited per repair resource so far. This constraint models the agreements between repair shop managers and inventory managers that determine how much of the total stream of failed parts can be expedited per repair resource. When there is an internal repair shop, these agreements may for example relate to the available expedited repair capacity per repair resource. There might be settings where it is relatively easy to obtain the exact costs associated with expediting a repair (e.g. in case of an external repair shop). In this section, we show that our DCG algorithm can be applied in this alternate setting almost immediately. In fact, as we will see shortly, there exists an equivalence relation between this setting and our original setting for a certain expediting cost structure.

When expediting has an additional cost, the objective is to minimize a total cost rate per time unit consisting of the total depreciation cost rate in spare parts and the total expediting cost rate, where the depreciation cost rate is obtained by depreciating the total initial investment costs in spare parts over their useful lifespan. We note that an alternative, but mathematically equivalent, formulation is to minimize a total initial cost consisting of both the total investment costs in spare parts (as in the original setting) and the total expected discounted expediting costs over an infinite horizon (see Appendix 2.C).

We first introduce some additional notation. Let the cost of expediting the repair of one part of SKU $m \in M$ be denoted by c^m_e. We assume that this cost is linearly proportional to the acquisition cost of that SKU, so that expediting the repair of more expensive SKUs is also more expensive than expediting the repair of cheaper SKUs.

That is, $c_e^m = \kappa \cdot c_a^m$, where $\kappa > 0$ denotes what we refer to as the expediting cost multiplier. The depreciation cost rate of each part of SKU $m \in M$ is denoted by c_d^m, and can be obtained by linearly depreciating its acquisition cost c_a^m over its useful lifespan τ, i.e. $c_d^m = c_a^m/\tau$. For a given control policy $(\mathbf{S}, T)$, the total depreciation cost rate in spare parts is defined as

$$C_d(\mathbf{S}) = \sum_{m \in M} \sum_{n \in N} c_d^m S_{m,n}$$

and the total repair expediting cost rate as

$$C_e(T) = \sum_{m \in M} c_e^m \lambda_{m,0} EXP_m(T_m).$$

The mathematical formulation of the alternate decision problem, which we call problem $(\widehat{P})$, is given as follows:

$$\begin{aligned} (\widehat{P}) \qquad & \min && C_d(\mathbf{S}) + C_e(T) \\ & \text{subject to} && EBO_c(\mathbf{S}, T) \leq \mathcal{B}_c^{max}, && \forall c \in C \\ & && \mathbf{S} \in \mathscr{S}, \quad T \in \mathbb{N}_0^{|M|}. \end{aligned}$$

We note that problem $(\widehat{P})$ constitutes the Lagrangian relaxation of our original decision problem (P) when either (i) each SKU has a dedicated repair resource or more generally when (ii) expediting costs are identical for all SKUs that use the same repair resource. The cost of expediting the repair of SKU m (c_e^m) is then equivalent to the Lagrange multiplier (ρ_r) for the resource needed for the repair of SKU m divided by the constant $\lambda_{m,0}$.

We now proceed to show how the DCG algorithm can be altered so that it applies to problem $(\widehat{P})$. We note that the main structure of the DCG algorithm remains unchanged; only its master problem and column generation sub-problem should be modified. The master problem of the DCG algorithm is now given by problem $(\widehat{MP})$:

$$\begin{aligned} (\widehat{MP}) \quad \min \quad & \sum_{m \in M} \sum_{n \in N} \sum_{k \in K_m} c_d^m S_{m,n}^k x_m^k + \sum_{m \in M} \sum_{k \in K_m} c_e^m \lambda_{m,0} EXP_m(T_m^k) x_m^k \\ \text{subject to} \quad & \sum_{m \in M_c^C} \sum_{n \in N_l} \sum_{k \in K_m} EBO_{m,n}(\mathbf{S}_m^k, T_m^k) x_m^k \leq \mathcal{B}_c^{max}, && \forall c \in C \\ & \sum_{k \in K_m} x_m^k = 1, && \forall m \in M \\ & x_m^k \geq 0, && \forall m \in M, \\ & && \forall k \in K_m \end{aligned}$$

The corresponding column generation sub-problem for SKU $m \in M_c^C$ is now given by:

$$(\widehat{SP}(m)) \quad \min \sum_{n \in N} c_d^m S_{m,n} + c_e^m \lambda_{m,0} EXP_m(T_m) - p_c \sum_{n \in N_l} EBO_{m,n}(\mathbf{S}_m, T_m) - \upsilon_m$$
$$\text{subject to} \quad \mathbf{S}_m \in \mathbb{N}_0^{|N|}, \quad T_m \in \mathbb{N}_0.$$

For a fixed expedite threshold T_m, $(\widehat{SP}(m))$ has exactly the same structure as the column generation sub-problem for our original setting. Hence, all results on $(SP(m))$ presented in Section 2.5.2 remain to hold for $(\widehat{SP}(m))$ exactly as stated. In particular the solution method described at the end of that section immediately applies to $(\widehat{SP}(m))$.

2.7. Computational study

The computational study in this section consists of two parts. In Section 2.7.1, we report on a case study at NS and present managerial insights. In Section 2.7.2 and Appendix 2.E, we evaluate and benchmark the performance of our solution approaches in an extensive numerical study based on a large test bed of randomly generated instances. We do so both for constrained expediting, i.e. problem (P), and expediting at additional costs, i.e. problem $(\widehat{P})$. We programmed our solution approaches as single threaded applications in C with GLPK as the solver of both linear and integer linear programs. All computations were carried out on a Windows PC (32 bit) with an Intel Quad Core 2.20 GHz processor and 8 GB RAM.

2.7.1 Case study at NS

NS is the principal passenger railway operator in the Netherlands. Its fleet consists of 900 rolling stock units, divided over twelve different train series. The spare parts inventory system of NS consists of one central warehouse and twelve local warehouses. There is a large repair center incident to the central warehouse. This repair center consists of multiple repair shops, each responsible for a different repair resource.

2.7.1.1 Setup and objective

This case study is focused on the VIRM train series; our case study therefore involves one capital good type, i.e. $|C| = 1$. The VIRM series consist of 176 rolling stock units, all of which are being operated as intercity trains that connect most cities in

the Netherlands. We consider the six most important warehouses where the VIRM train series is maintained and leave a handful of locations with only incidental demand out of scope; hence, $|N_l| = 6$.

We select 74 critical SKUs that occur in the configuration of the VIRM series. Of these SKUs, 30 require a mechanical resource for their repair and 44 require an electronic resource for their repair; hence, $|M_r^R| = 2$. The regular and expedited repair lead time for both repair resources is three weeks and one week, respectively. The transportation time is one week and includes administration time and shipment time from the central warehouses to all local warehouses. The acquisition costs of all SKUs range between 150.52 and 23,399.64 euros, and are 2,282.54 euros on average. We applied Maximum Likelihood Estimation to 5 years of historical failure data to estimate the demand intensities for each SKU. This estimation procedure leads to demand intensities that vary between 1 and 174 per year. In Figure 2.4, we plot and classify each SKU based on its normalized demand intensity and normalized acquisition cost. This classification will be important when we discuss the results of our case study.

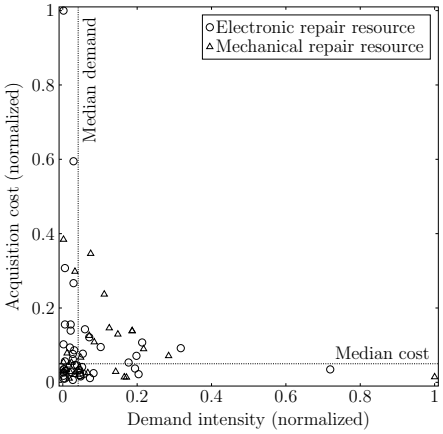

Figure 2.4 Scatterplot of SKUs in case study.

The current practice at NS can be described as follows. On a strategic level, inventory managers decide upon stocking levels using a single-item single-echelon model of the commercial package Servigistics (formerly Xelus Parts Planning). This model does not take into account that NS has the possibility to expedite the repair of parts in short supply. Expediting decisions are then made by inventory managers and repair shop managers together operationally on a weekly basis. For the electronic and mechanical repair shop, we observe from historical data that 30 percent of the total stream of failed parts is expedited.

The main objective of this case study is twofold. First, we want to determine the

reductions in investment costs that can be achieved when our solution approaches are used to achieve the same performance as the current approach of NS achieves. Second, and more importantly, we want to understand how a dynamic repair policy can be leveraged to reduce the total investment costs in spare parts while meeting availability targets.

Our benchmark for the case study is the current solution that NS uses. In this solution, the investment in each spare part is determined by the stocking model of Servigistics. For this investment decision, we determine the best achievable availability performance by optimizing expediting decisions and stock placement within our modeling framework. We then use the DCG algorithm and the greedy heuristic to find alternative investment decisions with at least the same availability performance. Further details regarding this benchmark, including the single-item model of Servigistics, are provided in Appendix 2.F.

2.7.1.2 Results and managerial implications

Table 2.2 shows the normalized investment costs of the three approaches to make investment decisions and the corresponding availability performance and expediting fractions. We observe that the greedy heuristic leads to an investment costs reduction of 52.43% compared to the current approach. As expected, the DCG algorithm has an even higher cost benefit with 53.55%. We remark that these savings are in the same order of magnitude as in other real-life applications that have compared a system approach with a single-item approach in a multi-item spare parts context. See, for example, Sherbrooke (2004) and Van Houtum and Kranenburg (2015) for various applications at the US Air Force and the Royal Dutch Navy, respectively. (We see later in Section 2.7.2.2 that our solution approaches also lead to high savings in the setting of NS when we benchmark against a state-of-the-art multi-item model).

Apart from the substantial investment costs reductions that can be reaped, it is interesting to note that the gap between the DCG algorithm and the greedy heuristic is small. Later when we discuss the results of our numerical experiments in Section 2.7.2, we will see that this holds across a large variety of industrial size problem instances.

Table 2.2 Main results case study at NS.

Solution approach	Investment costs	EBO_{VIRM}	$EXP_{\text{electronic}}$ (%)	$EXP_{\text{mechanical}}$ (%)
Approach NS	100	13.70	29.82	29.98
Greedy heuristic	47.57	13.65	29.99	29.97
DCG algorithm	46.45	13.65	29.99	29.94

Recall that we classified all SKUs into four distinct SKU groups based on their acquisition costs and demand intensities. To illustrate how the decisions of our new approaches realize the substantial cost reductions reported in Table 2.2, we will investigate the performance of each of these SKU groups. To facilitate presentation, we first introduce some additional notation. Let G denote the set of different SKU groups, hence $G = \{\text{high demand, low demand}\} \times \{\text{high cost, low cost}\}$ and thus $|G| = 4$. The SKUs that belong to group $g \in G$ are contained in the set M_g^G. For each of the three investment decisions consisting of acquired stock $\mathbf{S}$ and expedite thresholds T, we now calculate the following performance measures:

$$\begin{aligned}
\overline{EXP}_g(T) &= \left(\textstyle\sum_{m\in M_g^G} EXP_m(T_m)\right)/|M_g^G|,\\
EXP_g(T) &= \textstyle\sum_{r\in R}\sum_{m\in M_g^G} \delta_m^r EXP_m(T_m),\\
STOCK_g(\mathbf{S}) &= \left(\textstyle\sum_{m\in M_g^G}\sum_{n\in N} S_{m,n}\right) / \left(\textstyle\sum_{m\in M_g^G} \lambda_{m,0}\right),\\
EBO_g(\mathbf{S}, T) &= \textstyle\sum_{m\in M_g^G}\sum_{n\in N_l} EBO_{m,n}(\mathbf{S}_m, T_m), \text{ and}\\
COST_g(\mathbf{S}) &= 100\cdot\left(\textstyle\sum_{m\in M_g^G}\sum_{n\in N} c_a^m S_{m,n}\right)/C(\mathbf{S}),
\end{aligned}$$

which all provide meaningful information about an SKU group $g \in G$. The mean expedited repair utilization and the total aggregate mean expedited repair utilization of g are given by $\overline{EXP}_g(T)$ and $EXP_g(T)$, respectively. Note that $\overline{EXP}_g(T) \in [0, 1]$ and $EXP_g(T) \in [0, 0.6]$. $STOCK_g(\mathbf{S})$ provides a normalized measure of how much stock of all SKUs in g is acquired. The total mean number of backorders for g is given by $EBO_g(\mathbf{S}, T)$. Finally, $COST_g(\mathbf{S})$ measures the relative difference between the investment costs in g and the overall total costs under the investment decision.

Table 2.3 provides the performance measures for each SKU group $g \in G$ under each of the three investment decisions. For now, we only consider the performance measures of our solution approaches, and we turn our attention to the left upper quadrant: SKUs with low demand intensities and high acquisition costs. As the table indicates, the unavailability due to this group of SKUs is kept relatively low by providing full repair priority to failed parts rather than by investing in spare parts. Similar to Sherbrooke (2004, p.12) in an application at the US Air Force, we observe that for "high cost items [...] the model has allocated a zero or low stock level". Although failed parts always receive expedited repair, this SKU group utilizes only 9.22% of the total available expediting capacity. Conversely, if we look at the right lower quadrant, SKUs with low acquisition costs and high demand intensities receive almost no expedited repair. Instead, the unavailability due to this group of SKUs is kept relatively low by acquiring large amounts of spare parts.

For the other two SKU groups, our solution approaches neither solely invest in spare

Table 2.3 Performance measures per SKU group under each solution approach.

		Solution approach						
		DCG	Greedy	Current	DCG	Greedy	Current	Measure
Acquisition cost	High	100	100	9.34	67.08	65.56	34.63	$\overline{EXP}$ (%)
		9.22	9.22	1.25	46.73	41.96	27.25	EXP (%)
		0.00	0.00	92.57	15.62	16.17	23.98	$STOCK$
		0.00	0.00	37.75	52.51	54.37	46.18	$COST$
		1.94	1.94	0.92	8.39	8.54	5.76	EBO
	Low	55.57	61.39	13.43	5.11	10.83	36.60	$\overline{EXP}$ (%)
		2.04	4.18	1.26	1.94	4.60	30.04	EXP (%)
		66.00	66.00	104.5	44.24	44.24	23.78	$STOCK$
		5.08	4.22	4.92	42.41	41.41	11.14	$COST$
		1.09	1.05	0.96	2.22	2.12	6.06	EBO
			Low			High		
				Demand intensity				

parts nor solely expedite the repair of failed parts. If we look at the SKUs with low demand intensities and low acquisition costs, we indeed observe that this group has a large amount of normalized acquired stock as well as a high average expedited repair utilization. As a result, this group has the smallest mean number of backorders of all groups. The impact on the total investment costs and the total available expediting capacity is however small as both demand intensities and acquisition costs are low.

From the right upper quadrant, we observe that a large part of the available expediting capacity is utilized by the group of SKUs with both high demand intensities and high acquisition costs. Although the investment costs in this group are more than half of the total costs of the investment decision, the normalized acquired stock is relatively small. Finally, with more than 8 expected backorders, the unavailability due to this group of SKUs is significantly larger than all other groups.

Our integrated solution approaches thus lead to well-balanced investment decisions in which we acquire large amounts of spare parts of SKUs with low acquisition costs. In doing so, we maximize the availability of these SKUs at relatively low investment costs. Almost all available expediting capacity is then leveraged to dynamically prioritize the repair of failed parts with high acquisition costs, which allows us to refrain from excessively acquiring spare parts with such high costs. The current approach leads to a less balanced investment decision. As Table 2.3 indicates, the current approach invests heavily in spare parts with high acquisition costs and mainly prioritizes the repair of failed parts of SKUs with high demand intensities.

We emphasize that the dynamic repair policy requires real-time information about the repair pipeline in deciding upon expedited repairs. This should be taken into account when implementing our solution approaches in practice. In case of NS, this

real-time information is obtained through modern tracking technology that utilizes radio-frequency identification.

2.7.2 Numerical experiments

In the previous section, we have described how our solution approaches leverage a dynamic repair policy to reduce the total investment costs in spare parts while satisfying availability and expediting constraints. In this section and Appendix 2.E, we assess the value of having such an advanced dynamic repair policy in the first place. We also investigate whether our solution approaches find solutions that are close to optimal and whether they find such solutions within reasonable time. We answer these questions for both the original decision problem that was also faced by NS, i.e. problem (P), and the alternate decision problem in which expediting a repair comes at an additional cost, i.e. problem $(\widehat{P})$.

2.7.2.1 Test beds

Our test bed for decision problem (P) consists of 2592 randomly generated problem instances obtained by permuting all input parameters over multiple levels that are based on representative data for the capital goods industry. This test bed consists of both symmetric instances, in which the demand intensities across all local warehouses are identical but varied for different SKUs, and asymmetric instances, in which the demand intensities are varied across all local warehouses and SKUs. Our test bed for decision problem $(\widehat{P})$ also consists of 2592 instances, and is identical to the test bed for decision problem (P) except for the two input parameters that are specific to problem (P), i.e. $\mathcal{E}_r^{max}$ and $|R|$. We replace those two input parameters with the two input parameters that are specific to problem $(\widehat{P})$, i.e. the expediting cost multiplier κ and the useful life span τ. Further details regarding both test beds are relegated to Appendix 2.E.

2.7.2.2 Results when expediting is constrained

We first consider decision problem (P). To evaluate the effectiveness of our solution approaches in solving this problem, we compute a feasible solution for each generated instance using both solution approaches and we measure the relative difference between the total cost obtained by the solution approach and the corresponding lower bound. That is,

$$\%GAP = 100 \cdot \frac{(C_P^{UB} - C_P^{LB})}{C_P^{LB}},$$

where C_P^{LB} is obtained using the method described in Section 2.5.1 and where C_P^{UB} is obtained using the method described in Section 2.5.3 in case of the DCG algorithm, or using the method described in Section 2.5.4 in case of the greedy heuristic.

To quantify the value of our dynamic repair policy, we create a state-of-the-art benchmark instance for each original instance of problem (P) that we generate. This benchmark instance is identical to the original instance except that it is not possible to differentiate repair lead times through expediting. The mean repair lead time of this benchmark instance is then kept below the mean repair lead time of the original instance. This is achieved as follows: We set $\mathcal{E}_r^{max}$ to 1.0 for each repair resource $r \in R$ in the original instance such that it is feasible (and optimal) to expedite all repairs. We then change the expedited lead time $t_{m,0}^2$ of each SKU $m \in M$ to the shortest mean repair lead time possible in the feasible solution to the original instance. For a given SKU $m \in M$ that requires resource $r \in R$ for its repair, this shortest mean repair lead time is $(1-\mathcal{E}_r^{max})\cdot(t_{m,0}^1+t_{m,0}^2)+\mathcal{E}_r^{max}\cdot t_{m,0}^2$. For this benchmark instance, we compute a lower bound on the optimal cost using the method described in Section 2.5.1. We denote this lower bound by C_{BM}^{LB} and we compare it with C_P^{UB} of the original instance, obtained by the DCG algorithm. That is,

$$RED = 100 \cdot \frac{(C_{BM}^{LB} - C_P^{UB})}{C_{BM}^{LB}},$$

where $\%RED$ will indicate how much stock investment reductions can be achieved because of the possibility to expedite the repair of parts in short supply.

Aggregated results of the numerical experiments involving problem (P) are presented in Table 2.4. Detailed results are provided in Appendix 2.E. We note that the solutions to the problem instances generally exhibit the same behavior as extensively described in the case study.

Table 2.4 Aggregated results numerical experiments involving problem (P).

	DCG algorithm				Greedy heuristic				Benchmark	
	$\%GAP$		CPU time (s)		$\%GAP$		CPU time (s)		$\%RED$	
Instances	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max
Asymmetric	0.26	0.75	90.08	939.03	1.07	3.15	1.55	11.47	7.95	18.81
Symmetric	0.28	0.77	111.34	1271.85	3.69	8.43	1.75	12.85	7.92	19.61
Total	0.27	0.77	100.71	1271.85	2.38	8.43	1.66	12.85	7.94	19.61

The numerical experiments indicate that both solution approaches perform very well. The average optimality gaps of the DCG algorithm over the asymmetric and the symmetric problem instances are only 0.26 and 0.28, respectively. The optimality gaps of the greedy heuristic are slightly larger with 1.07 and 3.69 over the asymmetric

and the symmetric problem instances, respectively. The greedy heuristic is the most efficient heuristic in terms of computation time. Although the computation time of the DCG algorithm is considerably higher, it is still acceptable given the size and strategic nature of the decision problem.

The stock investment reductions that can be achieved because of the possibility to expedite the repair of parts in short supply are quite high with an average stock investment reduction of around 7.9 percent and even reductions of up to 19.61 percent. Before we continue with analyzing the alternate setting, we briefly return to the case study at NS. The value of a dynamic repair policy in their setting is substantial with a stock investment reduction of 36.40 percent. This is not surprising because our numerical experiments indicate that the value of a dynamic repair policy increases in the additional regular repair lead time or in the fraction of total demand that may be expedited. Both input parameters are slightly larger in the case study than in the problem instances of our test bed.

2.7.2.3 Results when expediting comes at additional costs

We now proceed with decision problem $(\widehat{P})$. The optimality gap of the DCG algorithm in this setting is defined as $\%GAP = 100 \cdot {}^{(C^{UB}_{\widehat{P}} - C^{LB}_{\widehat{P}})}/_{C^{LB}_{\widehat{P}}}$, where $C^{UB}_{\widehat{P}}$ and $C^{LB}_{\widehat{P}}$ are defined in a similar way as in the previous section, and can they be obtained using the methods described in Section 2.6.

Recall that in this alternate setting, there is an external (or internal) repair shop that charges an additional cost whenever we decide to expedite the repair of a failed part. Even though we thus have to pay a cost premium for a shorter repair lead time, we can leverage this flexibility when inventory is critically low. Our model anticipates precisely these future expediting decisions when deciding upon stocking levels. To quantify the value of our dynamic repair policy in this setting, we create a state-of-the-art benchmark instance that is identical to the original instance except that we do not have the possibility to shorten the repair lead time at the expense of a cost premium when backorders are imminent. Hence, similar to the benchmark for problem (P), we can only decide upon stocking levels in meeting availability constraints. This benchmark is created as follows: In the original instance we set $c^m_e = \infty$ for each SKU $m \in M$ such that it is optimal to not expedite any repairs. For this benchmark instance, we compute a lower bound on the optimal cost using the method described in Section 2.6. We denote this lower bound by $C^{LB}_{\widehat{BM}}$ and we compare it with $C^{UB}_{\widehat{P}}$ of the original instance, obtained by the DCG algorithm. That is, $\%RED = 100 \cdot {}^{(C^{LB}_{\widehat{BM}} - C^{UB}_{\widehat{P}})}/_{C^{LB}_{\widehat{BM}}}$, where $\%RED$ will now indicate by how much the total cost rate can be reduced because of the possibility to dynamically expedite repairs at the expense of a cost premium when inventory is critically low.

The aggregated results of the numerical experiments involving problem $(\widehat{P})$ are presented in Table 2.5. Detailed results are again relegated to Appendix 2.E. We can

Table 2.5 Aggregated results numerical experiments involving problem $(\widehat{P})$.

DCG algorithm				Benchmark		Expediting	
$\%GAP$		CPU time (s)		$\%RED$		$EXP(\%)$	
Avg	Max	Avg	Max	Avg	Max	Avg	Max
0.18	0.54	176.21	2757.51	4.86	29.29	17.37	94.02

draw three main conclusions from these results. First, the DCG algorithm performs even better when it is applied to problem $(\widehat{P})$: The average and maximum $\%GAP$ are only 0.18 and 0.54 percent, respectively. Second, although the computation time of the DCG algorithm is higher for problem $(\widehat{P})$ than for problem (P), it is still well within acceptable bounds given the size and strategic nature of the decision problem. Finally and most importantly, we find that anticipating expediting decisions that will be made later with investment decisions in repairable spare parts leads to substantial savings, even when those expedited repairs come at a cost premium. Indeed, the possibility to expedite the repair of failed parts at additional costs is effective in reducing the total cost rate with average reductions of around 4.9 percent and even reductions up to 29.29 percent.

2.8. Concluding remarks

We have considered a multi-item two-echelon spare parts inventory system, where each warehouse keeps multiple repairable types to maintain several types of capital goods, and where the repair shop at the central warehouse has two options for the repair of each defective part: a regular repair option and an expedited repair option. Irrespective of the repair option, each defective part uses a certain resource for its repair. Assuming a dual-index policy at the central warehouse and base stock control at the local warehouses, we have proposed an exact evaluation procedure for a given control policy.

To find an optimal control policy, we have formulated an optimization problem aimed at minimizing the total investment costs under constraints on both the aggregate mean number of backorders per capital good type and the aggregate mean fraction of repairs that are expedited per repair resource. We have shown how this non-linear non-convex integer programming problem can be decomposed into independent Newsvendor type sub-problems per repairable type, which subsequently allows us to use column generation algorithms. As an alternative solution approach, we have

presented an efficient greedy heuristic. Both solution approaches perform very well across a large test bed of industrial size.

We have shown that a dynamic repair policy is effective in reducing the stock investment needed to meet availability requirements for multiple types of capital goods while also satisfying expedited repair constraints for multiple repair resources. Our numerical experiments further show that such reductions remain attainable when expediting repairs comes at additional costs rather than being constrained. Based on a case study at NS, we have shown that our solution approaches lead to well-balanced investment decisions in which large amounts of spare parts of SKUs with low acquisition costs are acquired. In doing so, the availability of these SKUs can be maximized at relatively low investment costs. Almost all available expediting capacity can then be leveraged to dynamically prioritize the repair of failed parts with high acquisition costs, which allows us to refrain from excessively acquiring spare parts with such high costs.

The research in this chapter can be extended in two important ways. The first class of possible extensions would consider additional transportation modes, either from the central warehouse to the local warehouses, or in between the local warehouses themselves. The former relates to expedited transportation while the latter relates to so-called lateral transhipments. Both can serve as an emergency mode in case a local warehouse is out of stock, or as an expedited transport mode in case of imminent downtime. The second class of possible extensions would allow for demand to be non-stationary. Such processes can capture demand fluctuations for repairable spare parts over time, which might occur in practice due to for instance periodic inspections or revisions of equipment. Similar to Arts et al. (2016), assuming that demand for each repairable type is a Markov modulated Poisson process would then be a promising approach.

2.A. Proofs

Proof of Theorem 2.1. Let $S_{m,0}$ and T_m be fixed. Let $f : \mathbb{N}^{|N_l|} \to \mathbb{R}$ be the part of objective function (2.7) that depends on $S_{m,n}$, $n \in N_l$. Then, by omitting constants, objective function (2.7) reduces to

$$f(S_{m,1}, S_{m,2}, \ldots, S_{m,|N_l|}) = \sum_{n \in N_l} \left[c_a^m S_{m,n} - p_c EBO_{m,n}(S_{m,n}) \right],$$

where $EBO_{m,n}$ now depends only on $S_{m,n}$ because $S_{m,0}$ and T_m are fixed. By observing that each term in f is precisely the cost of an independent Newsvendor type problem, one for each local warehouse $n \in N_l$, the desired result directly follows. □

Proof of Lemma 2.1. Let T_m be fixed and $S_{m,n} = 0$ for all $n \in N_l$. Let $f : \mathbb{N} \to \mathbb{R}$ be the part of objective function (2.7) that depends on $S_{m,0}$. Then, by omitting constants, objective function (2.7) reduces to

$$f(S_{m,0}) = c_a^m S_{m,0} - p_c \sum_{n \in N_l} EBO_{m,n}(S_{m,0}),$$

where $EBO_{m,n}(S_{m,0})$ now depends only on $S_{m,0}$ because T_m is fixed and $S_{m,n} = 0$ for all $n \in N_l$.

Recall that the number of parts outstanding at local warehouse $n \in N_l$ is the sum of the demand during transport and shipping time $t_{m,n}$ from the central warehouse to local warehouse n, $D_{m,n}$, and the number of backorders at the central warehouse that belong to local warehouse n. Hence, since $S_{m,n} = 0$ for all $n \in N_l$, $EBO_{m,n}(S_{m,0})$ is equal to the sum of the expected backorders at the central warehouse that are from local warehouse n and the expectation of $D_{m,n}$ (see Equation (2.2)).

Then, since the number of backorders at the central warehouse that belong to local warehouse n is binomially distributed for a fixed total number of backorders (see Equation (2.1)), we have

$$f(S_{m,0}) = c_a^m S_{m,0} - p_c \sum_{n \in N_l} \frac{\lambda_{m,n}}{\lambda_{m,0}} EBO_{m,0}(S_{m,0}) = c_a^m S_{m,0} - p_c EBO_{m,0}(S_{m,0}),$$

where we have used the definition of $\lambda_{m,0}$ and the fact that $\mathbb{E}[D_{m,n}]$ is constant and can thus be omitted. By observing that f is precisely the cost of a Newsvendor type problem, the desired result directly follows. □

Proof of Lemma 2.2. Let T_m and $S_{m,0}$ be fixed and take some local warehouse $n \in N_l$. Let $Y \geq_{st} \tilde{Y}$ denote that a random variable Y is stochastically larger

than another random variable $\tilde{Y}$ in the usual stochastic order. Then, observe that $BO_{m,0}(S_{m,0}, T_m) \geq_{st} BO_{m,0}(S_{m,0}+1, T_m)$. This implies that $BO^n_{m,0}(S_{m,0}, T_m) \geq_{st} BO^n_{m,0}(S_{m,0}+1, T_m)$, and thus $BO^n_{m,0}(S_{m,0}) + D_{m,n} \geq_{st} BO^n_{m,0}(S_{m,0}+1) + D_{m,n}$, which is equivalent to $X_{m,n}(S_{m,0}, T_m) \geq_{st} X_{m,n}(S_{m,0}+1, T_m)$. Hence, in particular it holds that $P\{X_{m,n}(S_{m,0}, T_m) \leq x\} \leq P\{X_{m,n}(S_{m,0}+1, T_m) \leq x\}$ for any $x \in \mathbb{N}$. Hence, as $S^*_{m,n}(S_{m,0}, T_m)$ is the smallest $S_{m,n}(S_{m,0}, T_m)$ that satisfies Equation (2.9), we must have that $S^*_{m,n}(S_{m,0}+1, T_m) \leq S^*_{m,n}(S_{m,0}, T_m)$. □

Proof of Theorem 2.2. This follows directly from Lemma 2.1 and Lemma 2.2. □

2.B. Pseudo-code of solution approaches

This section provides pseudo-code of the DCG algorithm, including the exact solution method for $(SP(m))$, as well as the two-step greedy heuristic. Note that in the pseudo-code of the exact solution method for $(SP(m))$ (i.e. Algorithm 2), we continue the search over T_m by examining 4 values beyond the last observed local minimum, that is $N^{max} = 4$.

Algorithm 1 DCG algorithm for problem (P).

Step 1: Initialization

Determine an initial set of trivial policies $K^{res}_m \subseteq K_m$ for each SKU $m \in M$;

Step 2: Master Problem

Solve the restricted master problem (RMP) (2.3) – (2.6) with K_m replaced by K^{res}_m;

Obtain primal and dual solution;

Step 3: Column generation sub-problem

For the dual variables obtained in Step 2, execute Algorithm 2 for each SKU $m \in M$;

Step 4: Termination test

If Step 3 results in any policies with negative costs, add these to K^{res}_m and go to Step 2;

Else solve final version of (RMP) as an integer linear program and obtain a solution for problem (P);

Algorithm 2 Solution method for $(SP(m))$

Step 1: Initialization
Set T_m and N to 0, and N^{max} to 4;
Step 2: Initialization per T_m
Determine upper bound $\bar{S}_{m,0}(T_m)$ using Lemma 2.1 and set $S_{m,0}$ to 0;
Step 3: While $S_{m,0} \leq \bar{S}_{m,0}(T_m)$
Determine $S^*_{m,n}(S_{m,0}, T_m)$ using Theorem 2.1;
Determine corresponding reduced costs using Equation (2.7);
If lowest reduced costs per T_m so far then store policy $(S^*_{m,0}, T_m)$ and corresponding reduced costs;
Increase $S_{m,n}$ by 1;
Step 4: Termination test
If reduced costs of $(S^*_{m,0}, T_m) > (S^*_{m,0}, T_m - 1)$ then $N = N + 1$ else $N = 0$;
If $N \geq N^{max}$ then stop else increase T_m by 1 and go to Step 2;

Algorithm 3 Greedy heuristic for problem (P)

Step 1: Determine vector of expedite thresholds $\bar{T}$
For each repair resource $r \in R$
Set T_m to 0 $\forall\ m \in M^R_r$;
Calculate $\Gamma^T_m\ \forall\ m \in M^R_r$;
While $d(T) > 0$:
Determine m' with $\Gamma^T_{m'} \geq \Gamma^T_m\ \forall\ m \in M^R_r$;
Increase $T_{m'}$ with 1;
Calculate $\Delta_m d(T)$ and update $\Gamma^T_m\ \forall\ m \in M^R_r$;
Set $\bar{T}$ to T;
Step 2: Determine matrix of base stock levels $\mathbf{S}$
For each capital good type $c \in C$
Set $S_{m,n}$ to 0 $\forall\ m \in M^C_c$, $n \in N$;
Calculate $\Gamma^{\mathbf{S}}_{m,n}\ \forall\ m \in M^C_c$, $n \in N$;
While $d(\mathbf{S}, \bar{T}) > 0$:
Determine (m', n') with $\Gamma^{\mathbf{S}}_{m',n'} \geq \Gamma^{\mathbf{S}}_{m,n}\ \forall\ m \in M^C_c$, $n \in N$;
Increase $S_{m',n'}$ with 1;
Calculate $\Delta_{m,n} d(\mathbf{S})$ and update $\Gamma^{\mathbf{S}}_{m,n}\ \forall\ m \in M^C_c$, $n \in N$;

2.C. Alternative formulation expediting repairs at additional costs

Rather than minimizing a total cost rate, one might also be interested in minimizing a total initial cost consisting of both the total investment costs in spare parts (as in the original model) and the total expected discounted expediting costs. To this end, let $\beta > 0$ denote the discounting factor. For a given vector of expedite thresholds T, the total expected discounted expediting costs over an infinite horizon is then given by:

$$\begin{aligned}\tilde{C}_e(T) &= \sum_{m\in M} \int_0^\infty e^{-\beta t} c_e^m \lambda_{m,0} EXP_m(T_m)\,\mathrm{d}t \\ &= \sum_{m\in M} \frac{1}{\beta} c_e^m \lambda_{m,0} EXP_m(T_m),\end{aligned}$$

where the second equality follows from assuming that the system starts in steady state. The remaining analysis is now identical to the case with a total cost rate, with $C_e(T)$ changed to $\tilde{C}_e(T)$.

2.D. Allowing for commonality

In this section, we relax the assumption that M_c^C and M_r^R partition M, that is, the assumption that each SKU $m \in M$ occurs in the configuration of only one capital good type $c \in C$ and uses only one resource $r \in R$ for its repair. We first introduce additional notation to differentiate between demands for the same SKU that stem from different capital good types. We then briefly describe how problem (P) and its solution approaches change when commonality between SKUs is allowed.

Let $\lambda_{m,n,c}$ denote the demand intensity for SKU $m \in M$ at warehouse $n \in N$ originating from capital good type $c \in C$. If SKU m does not occur in the configuration of capital good type c, then $\lambda_{m,n,c} = 0$ by definition. Let the fraction of demands for SKU m at warehouse n that originate from capital good type c over all demands for that SKU at that warehouse be denoted by $\delta_{m,n}^c = \frac{\lambda_{m,n,c}}{\sum_{k\in C}\lambda_{m,n,k}}$.

The aggregate mean number of backorders for each capital good type $c \in C$ is now given by a weighted sum of the mean number of backorders for all SKUs occurring in the configuration of capital good type c, with the fractions $\delta_{m,n}^c$ as weights. That is,

$$EBO_c(\mathbf{S}, T) = \sum_{m\in M_c^C} \sum_{n\in N_l} \delta_{m,n}^c EBO_{m,n}(\mathbf{S}_m, T_m). \tag{2.11}$$

The definition of the aggregate mean fraction of failed parts that are expedited per repair resource $r \in R$, i.e. $EXP_r(T)$, remains however the same: SKUs now simply contribute to multiple aggregate mean fractions of expedited repairs whenever they require multiple resources for their repair. Hence, with $EBO_c(\mathbf{S}, T)$ now being defined as in Equation (2.11), we readily generalize our decision problem to the setting where commonality between SKUs is allowed.

The rest of the analysis goes along similar lines as for the setting without commonality. In particular, constraint (2.4) in the master problem of the DCG algorithm should be reformulated to

$$\sum_{m \in M_c^C} \sum_{n \in N_l} \sum_{k \in K_m} \delta_{m,n}^c EBO_{m,n}(\mathbf{S}_m^k, T_m^k) x_m^k \leq \mathcal{B}_c^{max}, \quad \forall c \in C,$$

which now incorporates our new definition for the aggregate mean number of backorders.

As SKUs may now belong to multiple capital good types and may now use multiple resources for their repair, the column generation sub-problem for SKU $m \in M$, refererred to as $(\widetilde{SP}(m))$, is now formulated as follows:

$$\begin{aligned} \min \quad & \sum_{n \in N} c_a^m S_{m,n} - \sum_{c \in C} \sum_{n \in N_l} p_c \delta_{m,n}^c EBO_{m,n}(\mathbf{S}_m, T_m) - \sum_{r \in R} \rho_r \delta_m^r EXP_m(T_m) - \upsilon_m \\ \text{subject to} \quad & \mathbf{S}_m \in \mathbb{N}_0^{|N|}, \qquad T_m \in \mathbb{N}_0. \end{aligned}$$

Note that the structure of $(\widetilde{SP}(m))$ is identical to $(SP(m))$. It is therefore readily verified that all properties as well as the exact solution method presented in Section 2.5.2 also hold for $(\widetilde{SP}(m))$, with the critical fraction $\frac{p_c + c_a^m}{p_c}$ in Theorem 2.1 and Lemma 2.1 changed to $\frac{\sum_{c \in C} \delta_{m,n}^c p_c + c_a^m}{\sum_{c \in C} \delta_{m,n}^c p_c}$.

In addition to the DCG algorithm, the two-step greedy approach can also be applied almost immediately to the setting where commonality between SKUs is allowed. The only difference is that the decreases in distances $\Delta_m d(T)$ and $\Delta_{m,n} d(\mathbf{S}, \bar{T})$ should now be calculated over multiple repair resources and multiple capital good types, respectively, with the additional note that in calculating $\Delta_{m,n} d(\mathbf{S}, \bar{T})$, we use Equation (2.11) for the aggregate mean number of backorders for each capital good type $c \in C$.

2.E. Numerical experiments

In this section, we report on our numerical experiments. The main objective of these experiments is to examine how the performance of our solution approaches, in terms of $\%GAP$ and $\%RED$, is affected by the input parameters of the decision problem. To this end, we consider a large test bed of randomly generated instances based on data representative for the capital goods industry. We first do so for the original decision problem, i.e. problem (P), and we subsequently treat the alternate setting in which additional costs are associated with expediting as well, i.e. problem $(\widehat{P})$.

2.E.1 Numerical experiments involving the constrained model

The test bed for the numerical experiments of this section consists of 2592 instances obtained through all combinations of the parameter values in Table 2.6. For each instance, we first use an uniform distribution $U[0.005, 0.25]$ to generate the demand intensity for each SKU $m \in M$ at all local warehouses $n \in N_l$, and subsequently multiply this generated demand intensity at each local warehouse $n \in N_l$ with a demand intensity multiplier, denoted ℓ. We consider symmetric instances in which demand intensities are identical across all local warehouses but varied for different SKUs, i.e. $\ell = 1$, as well as asymmetric cases in which demand intensities are varied across both local warehouses and different SKUs, i.e. $\ell = U[0.5, 1.5]$. Note that in each instance, we assign all SKUs uniformly at random to a repair resource set M_r^R for $r = 1, \ldots, |R|$.

The results of our numerical experiments involving problem (P) are summarized in Table 2.7. In this table, we present the average and maximum $\%GAP$ and computation times (in seconds) of both solution approaches as well as the average and maximum $\%RED$. We first distinguish between subsets of instances with the same value for a specific input parameter of Table 2.6 and then present the results for all instances. That is, the bottom row of Table 2.7 contains the results computed over all instances, and the other rows present the results computed over subsets of instances that have the same value for a specific input parameter.

The main observations drawn from Table 2.7 can be summarized as follows:

- The DCG algorithm performs very well. The average and maximum $\%GAP$ are 0.27 and 0.77 percent, respectively.
- The greedy heuristic performs very well when demand intensities are asymmetric. The average and maximum $\%GAP$ in this case are only 1.07 and 3.15, respectively. The greedy heuristic performs slightly worse with symmetric demand intensities: The average $\%GAP$ is 3.69 but instances with 8 or more

Table 2.6 Input parameter values for test bed involving problem (P).

	Input parameter	No. of choices	Values
1	Number of local warehouses, $\|N_l\|$	3	2, 4, 6
2	Number of capital good types, $\|C\|$	2	2, 4
3	Number of repair resources, $\|R\|$	2	2, 4
4	Number of SKUs per capital good type, $\|M_c^C\|$	3	20, 50, 100
5	Lead time from the central warehouse to local warehouse $n \in N_l$ of SKU $m \in M$, $t_{m,n}$	1	1
6	Expedited repair lead time of SKU $m \in M$, $t^2_{m,0}$	2	1, 2
7	Additional regular repair lead time of SKU $m \in M$, $t^1_{m,0}$	2	3, 5
8	Acquisition cost of SKU $m \in M$, c_a^m	1	$U[100, 1000]$
9	Demand intensity for SKU $m \in M$ at each local warehouse $n \in N_l$, $\lambda_{m,n}$	1	$U[0.005, 0.25]$
10	Demand intensity multiplier, ℓ	2	1, U[0.5,1.5]
11	Maximally allowed mean number of backorders over all SKUs $m \in M_c^C$ for capital good type $c \in C$, $\mathcal{B}_c^{max}$	3	$\nu \sum_{m \in M_c^C} \sum_{n \in N_l} \lambda_{m,n}$ for $\nu = 0.04, 0.06, 0.08$
12	Maximally allowed mean fraction of expedited repairs over all SKUs $m \in M_r^R$ that use repair resource $r \in R$ during their repair, $\mathcal{E}_r^{max}$	3	0.05, 0.10, 0.20

do occur. This observation is in line with previous research which examined greedy heuristics in multi-item spare parts problems (e.g., Topan et al., 2017). A possible explanation for this slightly worse performance is due to how the second step of the greedy heuristic works. With symmetric demand intensities, we have the property that if in a given iteration the base stock level of a specific SKU is increased at one local warehouse, then also the base stock levels of the same SKU at all other local warehouses are most likely increased in the succeeding iterations. However, in most practical situations in which each local warehouse serves a distinct market with a different demand structure, one will most likely encounter asymmetric demand intensities and hardly ever symmetric demand intensities.

- The average $\%GAP$ of both solution approaches seem to decrease as the instance size (in terms of the number of local warehouses, capital good types and SKUs per capital good type) becomes larger. This is very convenient since we typically face large-sized instances in practice. The average $\%GAP$ percent of the DCG algorithm tends to increase with the number of repair resources. This is not surprising, because problem (MP) has $|M|+|C|+|R|$ constraints and the same number of basic variables in an optimal solution. Since constraint (2.6) assures that for each SKU $m \in M$ a convex combination of policies is chosen, there is

Table 2.7 Summary of numerical results involving problem (P).

		DCG algorithm				Greedy heuristic				Benchmark	
		$\%GAP$		CPU time (s)		$\%GAP$		CPU time (s)		$\%RED$	
Input	Value	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max
$\|N_l\|$	2	0.32	0.77	10.75	71.39	3.07	8.43	0.24	0.77	9.66	19.61
	4	0.26	0.63	61.38	292.86	2.23	5.75	1.23	4.42	7.54	14.70
	6	0.23	0.55	230.01	1271.85	1.85	4.77	3.49	12.85	6.62	13.27
$\|C\|$	2	0.29	0.77	62.42	611.31	2.36	8.43	1.10	6.49	7.86	19.61
	4	0.26	0.56	139.00	1271.85	2.41	8.04	2.20	12.85	8.01	18.63
$\|R\|$	2	0.25	0.57	100.54	1271.85	2.40	8.43	1.65	12.85	8.02	19.61
	4	0.29	0.77	100.89	1135.98	2.37	8.01	1.65	12.78	7.85	18.32
$\|M_c^C\|$	20	0.36	0.77	35.37	328.09	2.49	8.43	0.58	2.70	7.85	18.81
	50	0.24	0.51	88.36	731.60	2.34	7.64	1.45	6.43	8.01	19.61
	100	0.21	0.50	178.42	1271.85	2.32	6.52	2.92	12.85	7.96	18.49
$t_{m,0}^2$	1	0.28	0.77	79.87	926.74	2.57	8.43	1.43	9.99	8.98	19.61
	2	0.26	0.75	121.55	1271.85	2.20	7.27	1.88	12.85	6.89	14.91
$t_{m,0}^1$	3	0.28	0.77	67.21	682.63	2.48	8.43	1.38	9.39	6.38	14.24
	5	0.27	0.75	134.22	1271.85	2.28	6.70	1.92	12.85	9.50	19.61
ℓ	1	0.28	0.77	111.34	1271.85	3.69	8.43	1.75	12.85	7.92	19.61
	U[0.5,1.5]	0.26	0.75	90.08	939.03	1.07	3.15	1.55	11.47	7.95	18.81
ν	0.04	0.27	0.75	100.48	1104.62	2.43	8.43	1.74	12.85	7.98	18.37
	0.06	0.27	0.66	100.49	992.75	2.40	7.64	1.65	11.72	7.87	19.61
	0.08	0.28	0.77	101.17	1271.85	2.32	8.04	1.57	11.69	7.97	18.63
$\mathcal{E}_r^{max}$	0.05	0.27	0.62	104.42	1271.85	2.13	6.09	1.70	12.85	5.21	9.83
	0.1	0.28	0.69	98.99	1135.98	2.29	8.01	1.66	12.43	7.79	14.81
	0.2	0.27	0.77	98.72	1125.12	2.73	8.43	1.59	11.88	10.81	19.61
Total		0.27	0.77	100.71	1271.85	2.38	8.43	1.66	12.85	7.94	19.61

a basic variable for each SKU m. Hence, there are at most $|C| + |R|$ SKUs for which the optimal solution to problem (MP) is fractional. This explains why the GAP percent increases with the number of repair resources. Note that this does not hold for the number of capital good types because the number of basic variables that increase with the number of capital good types is clearly more than the corresponding increase in the maximum number of SKUs for which the optimal solution to problem (MP) is fractional.

- The average $\%GAP$ of the DCG algorithm tends to decrease as the fraction of total demand that may be expedited or backordered decreases. This also seems to hold for the greedy heuristic, except with symmetric demand intensities. In the latter case, the average $\%GAP$ of the greedy heuristic increases when the fraction of total demand that may be backordered decreases.

- The greedy heuristic is the most efficient heuristic in terms of computation time.

The computation time of the DCG algorithm is considerably higher. Over 98 percent of that computation time is spent on solving the sub-problems. This task can also be parallelized using a multi-threaded approach, which would reduce the computation time of the DCG algorithm even further. The computation time of both solution approaches increases as the problem size (in terms of the number of local warehouses, capital good types and SKUs per capital good type) gets larger and decreases when the means of the repair lead times get smaller.

- The stock investment reductions that can be achieved because of the possibility to expedite the repair of parts in short supply are quite high with an average stock investment reduction of around 7.9 percent and even reductions of up to 19.61 percent.
- The stock investment reductions due to our dynamic repair policy increase when the additional regular repair lead time or the fraction of total demand that may be expedited increase, and decrease when the expedited repair lead time increase.

2.E.2 Numerical experiments involving the cost model

In this section, we consider a test bed that is identical to the test bed used in the previous section except for the two input parameters that are specific to problem (P), i.e. $\mathcal{E}_r^{max}$ and $|R|$. We replace those two input parameters with the two input parameters that are specific to problem $(\widehat{P})$, i.e. the expediting cost multiplier κ and the useful life span τ. We vary the expediting cost multiplier over three levels: $\kappa \in \{0.002, 0.004, 0.006\}$. Assuming that a regular repairs costs 5 to 10 percent of the acquisition cost of a part, these expediting cost multipliers imply a relative cost premium for expedited repairs of 2 to 12 percent. Such relative cost premiums are often reported in the dual-sourcing literature, where a cost premium is paid to utilize an expedited supplier over a regular supplier (e.g., Klosterhalfen et al., 2011; Sun and Van Mieghem, 2019). The useful lifespan (in years) is varied over two levels, both of which are representative for the capital goods industry: $\tau \in \{8, 10\}$ (we now explicitly assume that demand intensities are per day). Note that the total number of instances remains unchanged and hence this test bed also consists of 2592 instances.

Table 2.8 summarizes the results of our numerical experiments involving problem $(\widehat{P})$. We present the average and maximum $\%GAP$ and computation times (in seconds) of the DCG algorithm, the average and maximum $\%RED$ as well as the average and maximum EXP (expressed as percentages), i.e. the aggregated mean fraction of failed parts that receive expedited repair. Similar to Table 2.7 in the previous section, we first distinguish between subsets of instances with the same value for a specific input parameter and then present the results for all instances.

Table 2.8 Summary of numerical results involving problem $(\widehat{P})$.

		DCG algorithm				Benchmark		Expediting	
		$\%GAP$		CPU time (s)		$\%RED$		$EXP(\%)$	
Input	Value	Avg	Max	Avg	Max	Avg	Max	Avg	Max
$\|N_l\|$	2	0.22	0.54	11.69	75.27	7.59	29.29	23.13	94.02
	4	0.18	0.51	95.15	676.21	4.10	18.44	16.18	84.33
	6	0.15	0.52	421.78	2757.51	2.88	14.05	12.81	72.81
$\|C\|$	2	0.15	0.54	111.81	1767.52	4.89	28.98	17.24	94.02
	4	0.21	0.52	240.61	2757.51	4.82	29.29	17.51	93.64
τ	8	0.19	0.54	161.15	2757.51	5.74	29.29	21.93	94.02
	10	0.17	0.52	191.26	2587.05	3.97	21.11	12.81	65.99
$\|M_c^C\|$	20	0.22	0.54	59.22	661.44	4.82	29.29	17.55	93.85
	50	0.16	0.52	151.53	1515.09	4.87	28.98	17.34	94.02
	100	0.16	0.50	317.87	2757.51	4.87	28.81	17.24	92.57
$t^2_{m,0}$	1	0.19	0.54	135.44	2499.62	5.70	29.29	18.45	94.02
	2	0.17	0.52	216.97	2757.51	4.01	22.08	16.29	93.24
$t^1_{m,0}$	3	0.17	0.54	129.29	1922.82	2.38	12.07	9.05	43.07
	5	0.19	0.52	223.13	2757.51	7.33	29.29	25.70	94.02
ℓ	1	0.19	0.54	191.42	2757.51	4.75	28.49	16.96	94.02
	U[0.5,1.5]	0.16	0.51	160.99	2300.56	4.96	29.29	17.78	93.64
ν	0.04	0.18	0.52	173.86	2757.51	5.10	29.29	17.77	93.24
	0.06	0.18	0.52	178.50	2458.69	4.84	28.59	17.41	94.02
	0.08	0.18	0.54	176.26	2587.05	4.63	29.23	16.94	93.85
κ	0.002	0.19	0.54	93.84	1081.79	9.68	29.29	40.13	94.02
	0.004	0.19	0.52	188.04	2304.80	3.26	12.83	8.45	26.59
	0.006	0.16	0.51	246.74	2757.51	1.62	7.64	3.54	12.59
Total		0.18	0.54	176.21	2757.51	4.86	29.29	17.37	94.02

The numerical results in Table 2.8 are mostly in line with the numerical results of the previous section. We will now briefly discuss the main differences and important observations:

- The DCG algorithm also yields excellent results when it is applied to problem $(\widehat{P})$: The average and maximum $\%GAP$ are 0.18 and 0.54 percent, respectively. These optimality gaps are smaller than the optimality gaps of the DCG algorithm applied to problem (P), which is not surprising as problem $(\widehat{P})$ has $|R|$ fewer constraints than problem (P) and therefore fewer SKUs for which the optimal solution to problem $(\widehat{MP})$ is fractional (see also our discussion related to the optimality gaps in the previous section).

- Although the computation time of the DCG algorithm is higher for problem $(\widehat{P})$ than for problem (P), it is still well within acceptable bounds given the size and strategic nature of the decision problem. We note that this increase

in computation time can in large part be explained by the fact that we use the two-step greedy heuristic to create initial policies for the DCG algorithm; this heuristic is not devised specifically for problem $(\widehat{P})$ and these initial policies are therefore not necessarily good. This in contrast with problem (P) where the two-step greedy heuristic leads to very good policies for initializing the DCG algorithm.

- The possibility to expedite the repair of failed parts at additional costs is effective in reducing the total cost rate with average reductions of around 4.9 percent and even reductions up to 29.29 percent. Hence, anticipating expediting decisions that will be made later with investment decisions in repairable spare parts leads to substantial savings, even when those expedited repairs come at an additional price.
- The reductions in the total cost rate due to our dynamic repair policy decrease when expediting repairs becomes more costly, and increase when the useful lifespan of SKUs becomes shorter.

2.F. Approach to determine benchmark for case study

Our benchmark for the case study is the current solution that NS uses. In this solution, the investment in each spare part is determined by the stocking model of Servigistics. This stocking model is a single-item model that essentially ensures that for each SKU $m \in M$ sufficiently many spare parts are acquired to cover the lead time demand plus a safety level to protect against variability in demand. Hence, the amount of stock of SKU $m \in M$ determined by this stocking model is given by:

$$S_m^{ns} = \mu + k \cdot \sigma,$$

where k is a safety factor, and where μ and σ are the average and standard deviation of the demand over the total lead time (regular repair lead time + transportation time). The safety factor k for NS is quite involved but is set such that the fill rate for each SKU is 98%. See also Section 7.2 of Drent (2017) for more details.

Given S_m^{ns}, we determine the best achievable availability performance by optimizing expediting decisions and stock placement within our modeling framework. That is, we want to determine an expedite threshold $\tilde{T}_m \in \mathbb{N}_0$ and a base stock levels vector $\tilde{\mathbf{S}}_m \in \left\{\mathbb{N}_0^{|N|} : \mathbf{1} \cdot \mathbb{N}_0^{|N|} = S_m^{ns}\right\}$ such that $EBO(\tilde{\mathbf{S}}, \tilde{T})$ is minimized while $EXP_{\text{mechanical}}(\tilde{T}) \leq 0.3$ and $EXP_{\text{electronic}}(\tilde{T}) \leq 0.3$. This results in the following mathematical formulation

of the optimization problem:

$$(NS) \qquad \min_{\{\tilde{\mathbf{S}},\tilde{T}\}} \quad EBO(\tilde{\mathbf{S}},\tilde{T}) \tag{2.12}$$

$$\text{subject to} \quad EXP_{\text{mechanical}}(\tilde{T}) \leq 0.3, \tag{2.13}$$

$$EXP_{\text{electronic}}(\tilde{T}) \leq 0.3, \tag{2.14}$$

$$\tilde{\mathbf{S}} \in \mathscr{C}, \quad \tilde{T} \in \mathbb{N}_0^{|M|}, \tag{2.15}$$

where $\mathscr{C} = \left\{\tilde{\mathbf{S}} : \mathbf{1} \cdot \tilde{\mathbf{S}}_m = S_m^{ns} \ \forall m \in M\right\}$. We solve problem (NS) using a decomposition and column generation approach similar to our approach to solve problem (P), described in Section 2.5. The resulting corresponding sub-problem, however, does not allow for an easy solution method other than enumeration over $\tilde{T}_m$, and for each $\tilde{T}_m$, enumerating over all possible allocations of $\tilde{\mathbf{S}}_m$ over all local warehouses; see Section 7.2 of Drent (2017) for further details.

Chapter 3

Efficient emission reduction through dynamic mode selection

3.1. Introduction

The transportation sector has consistently been one of the most polluting European sectors for more than a decade now, and it is projected to remain so for the foreseeable future (European Environment Agency, 2020). This, unfortunately, appears to be a trend that stretches beyond Europe. Recent analysis indicates that the G20 countries, currently responsible for 80% of the global GHG emissions, will see an increase of 60% in their transportation sector emissions by 2050 (Vieweg et al., 2018). Prominent global climate targets, such as the ones outlined in the Paris agreement, will soon become unattainable (European Environment Agency, 2020; United Nations Environment Programme, 2020).

In light of the above, the European Union (EU) recently announced the Green Deal, a framework containing climate targets and policy initiatives that sets the EU on a path to reach carbon neutrality by 2050. The Green Deal is legally enshrined in the European Climate Law, which states that member states are legally committed to meet the targets, and face penalties in case they do not meet these targets. Being among the most polluting sectors, a key part of the Green Deal relates to policy initiatives that impact the transportation sector. For instance, the EU plans to extend the European emissions trading scheme (EU ETS) to include both road and maritime transport (Abnett, 2020; European Commission, 2020). Under the ETS, which until now includes only air transport, the EU enforces a cap on the total amount of GHG emissions from sectors covered by the scheme. The EU also investigates whether to

increase fossil fuel taxation, thereby effectively raising the price of GHG emissions. More and more companies are also reducing their emissions voluntarily as part of their corporate social responsibility. If not penalized by governments, companies that excessively pollute might still lose revenues as environmentally conscious customers take their business elsewhere (Dong et al., 2019).

The developments described above highlight the urgency for companies to explicitly incorporate GHG emissions in their supply chain decision making. In this chapter, we study the inbound transport and inventory decision making of a company that sells an assortment of products which are sourced from outside suppliers. The company wishes to keep the total GHG emissions from transport below a certain target level in the most economically viable manner. As is often the case in practice, the company relies on a third party logistics (3PL) provider for the inbound transport of the products. 3PLs typically offer several transport modes for the transportation of products – these may differ in terms of transportation costs, transit times, and GHG emissions.

The company can utilize the heterogeneity in the fleet of the 3PL to its advantage. While some transport modes are low emitting but slow, others may be fast but result in more emissions. Fast transport modes also typically come at the expense of a cost premium, and yet they are often relied upon when responsiveness is required (e.g. in case of imminent stock outs). Thus the company should rely dynamically on both transport modes. Implementing this holistically across the entire assortment of products allows the company to reduce emissions significantly for products for which it is relatively cost-efficient to do so and less for products for which this is more expensive. It additionally enables the company to reduce the total inventory and transportation costs by shipping the majority of products with a relatively cheap but slow transport modes while simultaneously resorting to faster but more expensive and often more polluting transport modes whenever expedited shipments are needed. While the advantages of dynamically selecting different transport modes are evident, two important and interrelated questions remain:

1. When should the company ship how many units of which product with what transport mode?
2. What is the value of dynamically shipping products with different transport modes?

These questions are interesting but also intricate when one wishes to answer them for an entire assortment of products where the combined total of GHG emissions from transportation is not allowed to exceed a certain target level.

To tractably answer the questions above, we focus on the setting where the 3PL offers two distinct transport modes for the transport of each product (or, equivalently, the

setting where the company has already decided on the two transport modes for each product). These transport modes need not be the same for every product; they will depend on the characteristics of the suppliers as well as the 3PL (e.g., some products can be transported using aircraft or rolling stock while others can be shipped via inland waterways or ocean shipping). The company decides periodically how many units it wishes to transport with what transport mode and incurs mode specific unit transportation costs. Shipments arrive at the company after a deterministic transit time that depends on the transport mode that is used. Demand for each product in every period is stochastic and independent and identically distributed across periods. Any demand in excess of on-hand inventory is backlogged and satisfied in later periods. The company incurs per unit holding and backorder penalty costs, and the specific cost parameters may vary from product to product. The company seeks to minimize the long-run average holding, backorder, and transportation costs while keeping the total long-run average GHG emissions from transportation of the entire assortment below a certain target level.

It is well-known that the optimal policy for the inventory/distribution system described above is complex, even in the simplest case of a single product and absent of the emission constraint (Whittemore and Saunders, 1977; Feng et al., 2006). For the control of each product, we therefore use a heuristic policy that is originally due to Veeraraghavan and Scheller-Wolf (2008). They show that their so-called dual-index policy performs quite well compared to the optimal policy. The dual-index policy tracks two inventory positions for each product: The slow inventory position, which equals the on-hand inventory plus all in-transit products minus backlog, and the fast inventory position, which is defined similarly but includes only those in-transit products that are due to arrive within the transit time of the fastest transport mode. Under the dual-index policy, we place orders with both modes such that these inventory positions are kept at (or above) certain target levels certain target levels, also referred to as base-stock levels. As such, the dual-index policy dynamically prescribes shipment quantities for both transport modes based on the net inventory level and the number of products that are still in-transit. To find the optimal base-stock levels for the entire assortment of products, we formulate the decision problem as a non-linear non-convex integer programming problem. A partition reformulation of this problem allows us to use column generation techniques to solve the decision problem. These techniques enable us to decompose the complex multi-product decision problem into simpler sub-problems per product. Leveraging a separability result of Veeraraghavan and Scheller-Wolf (2008), we show that this sub-problem constitutes a special Newsvendor problem that can be solved efficiently through a simulation-based optimization procedure.

The main contributions of this chapter are:

1. We are the first to study dynamic mode selection for an assortment of products with stochastic demand where the total average GHG emissions from the inbound transport of those products must be kept below a certain target level.

2. We provide a tractable optimization model that finds a tight lower bound on the optimal solution as well as near-optimal feasible solutions within reasonable time. We show that our mathematical formulation of the decision problem allows us to decompose the non-linear non-convex integer programming problem into sub-problems per product. We show that these sub-problems can be solved efficiently through a one dimensional search procedure in which each instance constitutes a Newsvendor type problem that is readily solved through simulation.

3. We perform an extensive computational experiment based on data from different industries. Through these experiments, we establish the value of dynamic mode selection by comparing our model with a model in which only one transport mode per assortment item can be used. This value can go up to 15 percent in cost savings. We also show that decomposing an aggregate carbon emission reduction target into targets for each item in the assortment individually is financially detrimental. Our holistic approach can lead to cost savings of over 40 percent relative to the approach with reduction targets per individual item.

The remainder of this chapter is organized as follows. In Section 3.2, we review the existing literature and position our work within the literature. Section 3.3 contains the model description as well as the mathematical formulation of the decision problem. A column generation procedure to solve the decision problem is provided in Section 3.4. We subsequently report on an extensive computational experiment in Section 3.5, and we provide concluding remarks in Section 3.6.

3.2. Literature review

This chapter integrates carbon emissions from inbound transportation into an inventory/distribution system with two transport modes. As such, our work contributes to the large stream of literature that studies multi-mode or multi-supplier inventory/distribution systems. For an excellent overview of such systems, we refer the reader to the review papers of Thomas and Tyworth (2006), Engebrethsen and Dauzére-Pérés (2019) and Svoboda et al. (2021), and references therein. We also contribute to the extensive body of literature that revolves around the integration of environmental aspects into supply chain decision making; see Dekker et al. (2012), Brandenburg et al. (2014), Barbosa-Póvoa et al. (2018), and references therein, for

an overview of this field. In what follows, we focus on contributions that are most relevant to the present chapter.

The decision how many products to order from which supplier is considered a canonical problem in the inventory management literature. It has been studied extensively since the sixties, mostly under the assumption that lead times are deterministic, that unmet demand is backlogged, and that only two distinct suppliers are at the disposal of the decision maker; the fastest being more expensive than the slowest. Fukuda (1964) and Whittemore and Saunders (1977) were the first to study this system. Assuming periodic review, they show that its optimal policy is a simple base-stock rule only under the assumption that the difference between the lead times of both suppliers is one period. For general lead time differences, the optimal policy is complex and can only be computed through dynamic programming for small instances. Since then, most researchers have focused on developing well-performing heuristic polices for which the best control policy parameters can be tractably obtained.

In this chapter, we rely on the so-called dual-index policy to decide upon the shipment sizes for both transport modes for each product. Under this policy, which is originally due to Veeraraghavan and Scheller-Wolf (2008), two different inventory positions are kept track off: One that includes all outstanding shipments and one that includes only those outstanding shipments that are due to arrive within the lead time time of the fastest mode. Veeraraghavan and Scheller-Wolf (2008) show numerically that the dual-index policy performs well compared to the optimal policy. In fact, Drent and Arts (2021) show that the dual-index policy is asymptotically optimal as the cost of the fastest transport mode and the backorder penalty costs become large simultaneously. The policy has received quite some attention in recent years (see, e.g., Sheopuri et al., 2010; Arts et al., 2011; Sun and Van Mieghem, 2019). We employ the dual-index policy because it is intuitive, has good performance, and can be optimized efficiently. Unlike the present chapter, the dual-index policy has so far been studied exclusively in single product settings under conventional cost criteria absent of any emission considerations.

Within the transportation literature, inventory/distribution systems with multiple transport modes have received considerable attention too. To properly embed the present chapter in the existing literature, we group contributions to this field into two categories depending on the modelling assumptions regarding the usage of the available transport modes (c.f. Engebrethsen and Dauzére-Pérés, 2019). The first category, which we refer to as *dynamic* mode selection, is concerned with inventory/distribution systems in which multiple transport modes are used simultaneously over a given (possibly infinite) planning horizon. Since we study an infinite horizon periodic review inventory model in which products can be transported

with two distinct modes in each period, our work falls into this category – as do all the inventory papers with two suppliers described so far. Only few papers exist in this category that explicitly account for carbon emissions, and the few that do differ substantially from the present work in terms of modelling choices as well as analysis. They either assume deterministic demand and a finite horizon (Palak et al., 2014) or study the closely related yet different problem of splitting an order among several transport modes (Konur et al., 2017). While not explicitly modeling carbon emissions, Lemmens et al. (2019) also rely on a dual-index policy to dynamically switch between different transport modes in the context of inter-modal transport. They show that this can lead to more usage of less polluting transport modes without compromising on costs or responsiveness. Different from the present work, all papers mentioned above consider the inventory control and transport mode decisions for a single product only.

The second category concerns inventory/distribution systems in which a single transport mode is selected a priori at the start of a planning horizon; all replenishment orders until the end of that planning horizon are then shipped with this mode. We refer to this category as *static* mode selection. Two papers belonging to this category are particularly relevant to our work. Hoen et al. (2014b) study a periodic review inventory system under backlogging where inbound transport is outsourced to a 3PL that offers multiple transport modes. Assuming base-stock control for each mode, they are interested in selecting the transportation mode that leads to the lowest long-run average total cost consisting of holding, backlogging, ordering, and emission costs. For calculating transportation emissions, they rely on the well-known NTM methodology (we discuss this methodology in more detail later in Section 3.3.1 and Appendix 3.A). We extend Hoen et al. (2014b) in two important directions. First, we move from static to dynamic mode selection, thereby incorporating the flexibility to dynamically switch between different transport modes for each product. Second, we consider an assortment of products under a single constraint on the total average transportation emissions from those products. Hoen et al. (2014a) consider a similar constraint in a multi-product variant of the setting of Hoen et al. (2014b) under the assumption that demand is deterministic and inversely related to the price set by the decision maker. They show that because of the portfolio effect of such an assortment-wide emission constraint, carbon emissions from transportation can be reduced substantially at hardly any additional cost.

The dual-index policy studied in this chapter has the appealing feature that it can mimic static mode selection. This is useful in our computational experiment where we establish the added value of dynamic mode selection over static mode selection. A closely related paper in that respect is Berling and Martínez-de Albéniz (2016) who study dynamic speed optimization of a single transport mode in a single-product stochastic inventory problem. They show that the value of dynamically controlling the speed of outstanding shipments, as opposed to a static speed policy, can be significant,

both financially and from a carbon emission perspective.

Our review so far has almost exclusively revolved around papers on multi-period inventory/distribution systems. We note that there is also a stream of literature that integrates carbon emissions into single period multi-supplier models, see e.g., Rosič and Jammernegg (2013), Arıkan and Jammernegg (2014), and Chen and Wang (2016). Similar to the majority of the papers discussed so far, these papers focus on single-product settings.

3.3. Model description

In this section, we first provide a description of the inventory/distribution system under consideration and introduce the notation that we use throughout this paper. We then describe the policy we propose to dynamically ship products with two transport modes. We conclude with providing a mathematical formulation of the decision problem.

3.3.1 Description and notation

We consider a company that sells an assortment of products. The inventories for these products are replenished from external suppliers through a third party logistics provider (3PL). A 3PL often offers several transport modes. We focus on the setting where the company has already decided upon two distinct transport modes that it would like to use for the transport of each product. These two transport modes will differ in terms of costs, lead times, emissions, or a combination thereof. Given these two transport modes for each product, the operational question that remains is how many units of each product the company should transport using which transport mode at what time so that costs –holding, backlog, and ordering– are minimized and an overall emission constraint is met. Companies will increasingly impose such constraints, either voluntarily or due to government regulation.

The inventory/distribution system under consideration runs in discrete time with $t \in \mathbb{N}_0$ denoting the period index. Without loss of generality, we assume that the period is of unit length and coincides with the review epoch. Let $J = \{1, 2, \ldots, |J|\}$ denote the nonempty set of products that the company offers for sale. Demand for product $j \in J$ per period is a sequence of non-negative independent and identically distributed (i.i.d.) random variables $\{D_j^t\}$. Any demand in excess of on-hand inventory is backlogged. Let I_j^t denote the net inventory level (on-hand inventory minus backlog) of product j at the beginning of period t after any outstanding orders have arrived. Each unit of product j in on-hand inventory $(I_j^t - D_j^t)^+$ carried over to the next period incurs

a holding cost $h_j > 0$. Similarly, each unit of product j in backlog $(D_j^t - I_j^t)^+$ incurs a penalty cost $p_j > 0$. Here we use the standard notation $x^+ = \max(0, x)$.

Each product can be shipped using two distinct transport modes from one supplier (or, equivalently, using one or two distinct transport modes from two distinct suppliers). Let $M = \{f, s\}$ denote the set of available transport modes, where we use f and s to refer to the faster and slower transport mode, respectively. Associated with the transport of one unit of product $j \in J$ with mode $m \in M$ is a cost $c_{j,m} \geq 0$, a deterministic lead time $l_{j,m} \in \mathbb{N}_0$, a distance traveled from the supplier to the company $d_{j,m} > 0$, and a certain number of units CO2 emission $e_{j,m} \geq 0$. The weight of one unit of product j is denoted $w_j > 0$. Recall that the company outsources its transport to a 3PL provider and hence has no control on the actual shipping. We therefore consider variable emissions that depend only on product and transport mode specific characteristics as well as on distance traveled, and we refrain from incorporating a fixed emission factor per actual shipment. In line with previous literature that models transportation emissions in the context of mode selection (e.g., Hoen et al., 2014a,b), we endow $e_{j,m}$ with the following structure which is based on the NTM methodology:

$$e_{j,m} = w_j(a_j + d_{j,m}b_j), \tag{3.1}$$

where $a_j \geq 0$ and $b_j > 0$ are a fixed and a variable transport mode specific emission constant, respectively.

We define the lead-time difference between the fast and slow mode as $l_j = l_{j,s} - l_{j,f} \geq 0$ for each product $j \in J$. Observe that, contrary to the conventional literature on dual-mode problems, we do not impose any assumptions on the unit transportation costs or emission units. Thus our model allows for situations where, e.g., the expensive transport mode is either the fastest and most polluting or the fastest and least polluting. Finally, the amount of items of product $j \in J$ to be shipped with transport mode $m \in M$ in period t is denoted $Q_{j,m}^t$. With this notation, observe that shipments $Q_{j,f}^{t-l_{j,f}}$ and $Q_{j,s}^{t-l_{j,s}}$ arrive in period t so that we can write the following recursion for the inventory level I_j^t of each product j:

$$I_j^t = I_j^{t-1} - D_j^{t-1} + Q_{j,f}^{t-l_{j,f}} + Q_{j,s}^{t-l_{j,s}}.$$

All notation introduced so far as well as notation that we will introduce later is summarized in Table 3.1.

Table 3.1 Overview of notation.

Notation	Description
Sets	
J	Assortment; Set of all products.
M	Set of available transport modes, i.e. $M = \{f, s\}$.
Input	
D_j^t	Random demand for product $j \in J$ in period $t \in \mathbb{N}_0$.
p_j	Penalty cost for one unit of product $j \in J$ in backlog carried over to the next period.
h_j	Holding cost for one unit of product $j \in J$ in on-hand inventory carried over to the next period.
$c_{j,m}$	Cost of shipping one unit of product $j \in J$ with transport mode $m \in M$.
$l_{j,m}$	Transportation lead time for product $j \in J$ by transport mode $m \in M$.
l_j	Transportation lead time for product $j \in J$ by transport mode $m \in M$.
$d_{j,m}$	Distance for the transport of product $j \in J$ with transport mode $m \in M$.
a_j	Fixed emission constant corresponding with transport mode $m \in M$.
b_j	Variable emission constant corresponding with transport mode $m \in M$.
w_j	Unit weight of product $j \in J$.
$e_{j,m}$	Total units CO2 emission associated with shipping one unit of product $j \in J$ with transport mode $m \in M$.
$\mathcal{E}^{max}$	The maximally allowable carbon emissions for the transport of the entire assortment of products.
Decision variables	
$S_{j,m}$	Base-stock level for product $j \in J$ and transport mode $m \in M$.
Δ_j	Difference between the slow and fast base-stock level for product $j \in J$, i.e. $S_{j,s} - S_{j,f}$.
$\mathbf{S}_f$	The vector $(S_{1,f}, S_{2,f}, \ldots, S_{\lvert J\rvert,f})$.
$\mathbf{\Delta}$	The vector $(\Delta_1, \Delta_2, \ldots, \Delta_{\lvert J\rvert})$.
State variables	
I_j^t	Inventory level of product $j \in J$ at the beginning of period $t \in \mathbb{N}_0$ after orders have arrived.
$IP_{j,f}^t$	Fast inventory position of product $j \in J$ in period $t \in \mathbb{N}_0$ before shipping orders.
$IP_{j,s}^t$	Slow inventory position of product $j \in J$ in period $t \in \mathbb{N}_0$ after shipping orders with the fast transport mode.
$Q_{j,m}^t$	Amount of product $j \in J$ shipped with transport mode $m \in M$ in period $t \in \mathbb{N}_0$.
O_j^t	The overshoot of product $j \in J$ in period $t \in \mathbb{N}_0$, i.e. $(IP_{j,f}^t - S_{j,f})^+$.
Output of model	
$C(\mathbf{S}_f, \mathbf{\Delta})$	Total long-run average holding, backlog, and ordering costs under a given control policy $(\mathbf{S}_f, \mathbf{\Delta})$.
$E(\mathbf{S}_f, \mathbf{\Delta})$	Total emissions from transportation under a given control policy $(\mathbf{S}_f, \mathbf{\Delta})$.
$C_P^{UB}(C_P^{LB}$	Upper (lower) bound for the optimal solution to Problem (P).

3.3.2 Control policy

It is well-known that even for the simplest case where $|J| = 1$ and absent of the emission constraint, the policy that prescribes the optimal shipment quantities is complex and can only be computed through dynamic programming for very small instances that are arguably not representative for practice (Whittemore and Saunders, 1977; Feng et al., 2006). For the control of this inventory/distribution system, we therefore use a heuristic policy that is originally due to Veeraraghavan and Scheller-Wolf (2008). They show numerically that their so-called dual-index policy performs quite well compared to the optimal policy. The dual index policy tracks two indices: One that contains all orders that are still in-transit and one that contains only

those in-transit orders that are due to arrive within the lead time of the fast mode. Based on these outstanding orders, the policy dynamically ships orders with both modes to keep these indices at certain target levels. In line with standard inventory management nomenclature, we also refer to these target levels as base-stock levels. More specifically, the policy operates as follows. At the beginning of every period t after orders $Q_{j,f}^{t-l_{j,f}}$ and $Q_{j,s}^{t-l_{j,s}}$ have arrived, we review the *fast* inventory position

$$IP_{j,f}^t = I_j^t + \sum_{k=t-l_{j,f}+1}^{t-1} Q_{j,f}^k + \sum_{k=t-l_{j,s}+1}^{t-l_j} Q_{j,s}^k,$$

and, if necessary, ship an order $Q_{j,f}^t$ with the fast transport mode to raise the fast inventory position to its target level $S_{j,f}$. That is, the amount of product j shipped in period t with the fast transport mode equals:

$$Q_{j,f}^t = (S_{j,f} - IP_{j,f}^t)^+.$$

After placing the fast shipment order, we inspect the *slow* inventory position

$$IP_{j,s}^t = I_j^t + \sum_{k=t-l_{j,f}+1}^{t} Q_{j,f}^k + \sum_{k=t-l_{j,s}+1}^{t-1} Q_{j,s}^k,$$

and ship an order with the slow transport mode such that this inventory position is raised to its target level $S_{j,s}$. Thus the amount of product j shipped with the slow transport mode in period t equals:

$$Q_{j,s}^t = S_{j,s} - IP_{j,s}^t.$$

After shipping both orders, demand D_j^t is satisfied or backlogged, depending on whether there is sufficient inventory available or not. The period then concludes with charging holding or backlog costs.

The order of events in a period t for each product j is thus as follows:

1. Orders $Q_{j,f}^{t-l_{j,f}}$ and $Q_{j,s}^{t-l_{j,s}}$ arrive with the fast and slow transport mode, respectively, and are added to the on-hand inventory I_j^t.
2. Review the fast inventory position and ship order $Q_{j,f}^t$ with the fast transport mode at unit cost $c_{j,f}$.
3. Review the slow inventory position and ship order $Q_{j,s}^t$ with the slow transport mode at unit cost $c_{j,s}$.
4. Demand D_j^t occurs and is satisfied from on-hand inventory if possible, and

otherwise backlogged.

5. Incur a cost h_j for any unit in on-hand inventory $(I_j^t - D_j^t)^+$ and a cost p_j for any unit in backlog $(D_j^t - I_j^t)^+$.

Observe that under a dual-index policy, slow orders entering the information horizon of the fast transport mode may cause the fast inventory position to exceed its target level. The amount by which the fast inventory position exceeds its target level is referred to as as the overshoot. The fast inventory position in period t after placing orders with both modes thus equals $S_{j,f} + O_j^t$, where O_j^t denotes the overshoot for product j in period t:

$$O_j^t = IP_{j,f}^t + Q_{j,f}^t - S_{j,f} = (IP_{j,f}^t - S_{j,f})^+.$$

Later, in Section 3.4.2, we shall see that computing the steady state distribution of the overshoot is crucial for determining the performance of a given control policy for a single product.

We furthermore define $\Delta_j = S_{j,s} - S_{j,f}$, $j \in J$, so that the control policy for a product can be specified in terms of its base-stock levels $S_{j,s}$ and $S_{j,f}$ or in terms of its base-stock level for the fast transport mode $S_{j,f}$ and the difference Δ_j. We mostly use the latter specification in our subsequent analysis. A control policy $(\mathbf{S}_f, \Delta)$ for the entire assortment of products consists of the vectors $\mathbf{S}_f = (S_{1,f}, S_{2,f}, \ldots, S_{|J|,f})$ and $\mathbf{\Delta} = (\Delta_1, \Delta_2, \ldots, \Delta_{|J|})$.

In what follows, for all sequences of random variables X^t, we define their stationary expectation as $\mathbb{E}[X] = \lim_{T\to\infty}(1/T)\sum_{t=0}^{T} X^t$ and their distribution as $\mathbb{P}(X \le x) = \lim_{T\to\infty}(1/T)\sum_{t=0}^{T} \mathbb{1}\{X^t \le x\}$, where $\mathbb{1}\{A\}$ is the indicator function for the event A. Whenever we drop the period index t we refer to the generic stationary random variable X with expectation and distribution as defined above.

3.3.3 Decision problem

For a given control policy $(\mathbf{S}_f, \mathbf{\Delta})$, we define the total long-run average holding, backlog, and ordering costs per period for the entire assortment of products as

$$\begin{aligned} C(\mathbf{S}_f, \mathbf{\Delta}) &= \sum_{j\in J} C_j(S_{j,f}, \Delta_j) \\ &= \sum_{j\in J} \left(h_j \mathbb{E}[(I_j - D_j)^+] + p_j \mathbb{E}[(D_j - I_j)^+] + \sum_{m\in M} c_{j,m} \mathbb{E}[Q_{j,m}] \right), \end{aligned}$$

and the total emissions as

$$E(\mathbf{S}_f, \mathbf{\Delta}) = \sum_{j \in J} E_j(S_{j,f}, \Delta_j) = \sum_{j \in J} \sum_{m \in M} e_{j,m} \mathbb{E}[Q_{j,m}],$$

where it is understood that the expectation operators are conditional on the control policy $(\mathbf{S}_f, \mathbf{\Delta})$. Veeraraghavan and Scheller-Wolf (2008) show that $C(\mathbf{S}_f, \mathbf{\Delta})$ is well-defined for any control policy $(\mathbf{S}_f, \mathbf{\Delta})$ as long as $\mathbb{E}[D_j] < \infty$ for all products $j \in J$.

The objective of our decision problem is to minimize the total long-run average costs while keeping the total emissions below a target level $\mathcal{E}^{max}$. Combining the above-mentioned leads to the following mathematical formulation of our decision problem which we refer to as problem (P):

$$\begin{array}{lll} (P) & \min & C(\mathbf{S}_f, \mathbf{\Delta}) \\ & \text{subject to} & E(\mathbf{S}_f, \mathbf{\Delta}) \leq \mathcal{E}^{max}, \\ & & \mathbf{S}_f \in \mathbb{R}^{|J|}, \quad \mathbf{\Delta} \in \mathbb{R}_0^{|J|}. \end{array}$$

Let $(\mathbf{S}_f^*, \mathbf{\Delta}^*)$ denote an optimal solution to problem (P) and let C_P be the corresponding optimal cost. Note that Problem (P) can be interpreted as a non-linear non-convex knapsack problem where more than one copy of each item can be selected. It is well-known that even the simplest types of such knapsack problems are $\mathcal{NP}$-hard (e.g. Kellerer et al., 2004). Since our knapsack is more complex, we conclude that Problem (P) also falls in that same complexity class; it is hence likely that also for our problem no exact polynomial time solution algorithm exists.

3.4. Analysis

This section focuses on finding the optimal control policy for Problem (P). Our approach relies on the technique of column generation – also named Dantzig-Wolfe decomposition after its pioneers (Dantzig and Wolfe, 1960). This technique enables a natural decomposition of the original multi-product decision problem into smaller single-product problems that have more structure. Below, we first explain how we apply column generation to Problem (P), and we then describe a simulation-based optimization method for solving the sub-problem of this column generation procedure.

3.4.1 The column generation procedure

We first reformulate decision problem (P) as an integer linear program in which each binary decision variable corresponds to a certain combination of values for the decision

variables of our original decision problem. We subsequently relax the integrality constraint and we call this problem the master problem (MP). Formally, let K_j be the set of all possible dual index policies for product $j \in J$. Each policy $k \in K_j$ is determined by its policy parameters $S^k_{f,i}$ and Δ^k_j. Let $x^k_j \in \{0,1\}$ denote the decision variable that indicates whether policy $k \in K_j$ is selected ($x^k_j = 1$) for product $j \in J$ or not ($x^k_j = 0$). By relaxing the integrality constraint on this binary decision variable, we arrive at the mathematical formulation of the master problem (MP):

$$
\begin{aligned}
(MP) \qquad \min \quad & \sum_{j\in J}\sum_{k\in K_j} C_j(S^k_{j,f},\Delta^k_j)x^k_j && && (3.2)\\
\text{subject to} \quad & \sum_{j\in J}\sum_{k\in K_j} E_j(S^k_{j,f},\Delta^k_j)x^k_j \leq \mathcal{E}^{max}, && && (3.3)\\
& \sum_{k\in K_j} x^k_j = 1, && \forall j \in J && (3.4)\\
& x^k_j \geq 0, && \forall j \in J, \forall k \in K_j.
\end{aligned}
$$

Let C^{LB}_P denote the optimal cost for master problem (MP). Due to the linear relaxation of the integrality constraint on x^k_j, an optimal cost C^{LB}_P is also a lower bound on the optimal cost for Problem (P), C_P.

Due to its large number of decision variables, master problem (MP) is solved using column generation. To this end, we first restrict master problem (MP) to a small subset $\tilde{K}_j \subseteq K_j$ of policies per product $j \in J$ (i.e. columns). This problem is referred to as the restricted master problem (RMP). We then solve (RMP) to optimality, and we are interested in new policies $k \in K_j \setminus \tilde{K}_j$, $j \in J$, that will improve the objective value of (RMP) if they are added to $\tilde{K}_j$. Such policies $k \in K_j \setminus \tilde{K}_j$ are identified through solving a column generation sub-problem for each product $j \in J$. The column generation sub-problem for product $j \in J$ has the following form:

$$
\begin{aligned}
(SP(j)) \quad \min \; & h_j\mathbb{E}[(I_j - D_j)^+] + p_j\mathbb{E}[(D_j - I_j)^+] + \sum_{m\in M}(c_{j,m} - \eta e_{j,m})\mathbb{E}[Q_{j,m}] - \upsilon_j,\\
\text{subject to} \quad & S_{j,f} \in \mathbb{R}, \quad \Delta_j \in \mathbb{R}_0,
\end{aligned}
$$

where η denotes the dual variable of (RMP) that corresponds with the emission constraint (3.3) and υ_j denotes the dual variable of (RMP) that corresponds with the constraint (3.4) that assures that for each product $j \in J$ a convex combination of policies is chosen. If for product $j \in J$ a feasible solution for ($SP(j)$) exists with a negative objective value, then this policy is added to $\tilde{K}_j$ since the objective value of (RMP) can be improved when solved with the enlarged set $\tilde{K}_j$.

We continue with iterating between optimizing (RMP) and finding new policies through solving ($SP(j)$), $j \in J$, until no product for which there is a policy with a

negative reduced cost exists. An optimal solution for (RMP) is then also an optimal solution for (MP). If this optimal solution contains integer values only, then it is also an optimal solution for (P). If this is not the case, then we solve (RMP) one last time as an integer program to find an integer solution for (RMP), which is then also a feasible solution for (P). Recent inventory literature has shown that solving the restricted master program as an integer program to arrive at an integer solution leads to good performance in terms of optimality gaps (e.g., Drent and Arts, 2020; Haubitz and Thonemann, 2021), and often outperforms alternative approaches such as local searches or rounding procedures (Alvarez et al., 2013). The corresponding cost of the resulting feasible solution is also an upper bound, denoted C^P_{UB}, for C^P.

In the next section, we provide a simulation-based optimization procedure to solve the column generation sub-problem $(SP(j))$.

3.4.2 Solving the column generation sub-problem

The column generation sub-problem $(SP(j))$ has the same structure as the problem studied by Veeraraghavan and Scheller-Wolf (2008). We follow their simulation-based optimization procedure to solve $(SP(j))$. This procedure is grounded in the following separability result that allows us to find the optimal $S_{f,i}$ for a given Δ_j as the solution to a special Newsvendor problem.

Lemma 3.1. *(Veeraraghavan and Scheller-Wolf, 2008, Proposition 4.1) The distributions of the overshoot O_j, the fast transport mode shipment size $Q_{f,i}$, and the slow transport mode shipment size $Q_{s,i}$ are functions of Δ_j only, independent of $S_{f,i}$.*

Let $O_j^t(\Delta_j)$ denote the overshoot of product j in period t for a given Δ_j. Recall that the fast inventory position of product j in period t after shipping equals $S_{j,f}+O_j^t(\Delta_j)$. Consequently, for the net inventory level of product j in each period t, we can also write

$$I_j^t = S_{j,f} - \left(\sum\nolimits_{k=t-l_{j,f}}^{t-1} D_j^k - O_j^{t-l_{j,f}}(\Delta_j) \right). \tag{3.5}$$

By plugging (3.5) in the objective function of $(SP(j))$, and using Lemma 3.1 as well as the fact that in each period the overshoot is independent of the demand and $S_{j,f}$, we readily recognize that for given Δ_j the objective function is convex in $S_{j,f}$. This implies the following result.

Lemma 3.2. *(Veeraraghavan and Scheller-Wolf, 2008, Theorem 4.1) The optimal base-stock level $S^*_{j,f}$ for a given Δ_j, denoted $S^*_{j,f}(\Delta_j)$, equals*

$$S^*_{j,f}(\Delta_j) = \inf \left\{ S_{j,f} \in \mathbb{R} : \mathbb{P}\left(\sum\nolimits_{i=1}^{l_{j,f}+1} D_j - O_j(\Delta_j) \leq S_{f,i} \right) \geq \frac{p_j}{p_j + h_j} \right\}.$$

It now remains to calculate the objective value of $SUB(j)$ for given Δ_j and corresponding $S^*_{j,f}(\Delta_j)$. Observe that in each period immediately after shipping orders with both modes, the slow inventory position equals the fast inventory position plus the overshoot and all remaining outstanding slow orders. Since these inventory positions are equal to their respective base-stock levels following order placement, we have for each product j in each period t:

$$S_{j,s} = S_{j,f} + O^t_j + \sum_{k=0}^{l_j-1} Q^{t-k}_{j,s}. \tag{3.6}$$

From (3.6) it follows that $\mathbb{E}[Q_{j,s}] = (\Delta_j - \mathbb{E}[O_j])/l_j$. Since under backlogging the sum of both orders must on average be equal to the period demand, we finally find $\mathbb{E}[Q_{j,f}] = \mathbb{E}[D_j] - \mathbb{E}[Q_{j,s}]$.

To solve $(SP(j))$, $j \in J$, to optimality, it thus suffices to perform a one-dimensional search over Δ_j. For each Δ_j, we compute the stationary distribution of the overshoot. With this stationary distribution we readily find the optimal base-stock level $S^*_{j,f}(\Delta_j)$ through Lemma 3.2 and the total reduced cost through the identities following Equation (3.6). As there is in general no closed-form expression for the stationary distribution of the overshoot, we follow Veeraraghavan and Scheller-Wolf (2008) and rely on simulation to compute this distribution.

Note that our optimization model and analysis readily extends to settings where the slow transport modes of all (or some) products are operated according to any other rule that depends only on the current overshoot as well as all in-transit orders that are not yet included in the fast inventory position. That is, any other rule that depends only on the information state $(O^t_j, Q^{t-1}_{j,s}, Q^{t-2}_{j,s}, \ldots, Q^{t-l_j+1}_{j,s})$, $j \in J$, $t \in \mathbb{N}_0$. Most well-performing control policies satisfy this condition, e.g., the Capped Dual-Index policy (Sun and Van Mieghem, 2019), the Tailored Base-Surge policy (Allon and Van Mieghem, 2010), and the Projected Expedited Inventory Position policy (Drent and Arts, 2021). Sheopuri et al. (2010) show that for such control policies, the stationary distribution of the overshoot is a function of only the parameter(s) for operating the slow transport mode, and that consequently a Newsvendor result similar to Lemma 3.2 holds for all such policies.

3.5. Computational experiment

Our test bed has three different types of assortments of products, each representing a different type of industry. The first assortment type consists solely of products for which emissions from the slowest transport mode are less than the emissions from the

faster transport mode. The opposite holds for the second assortment type. The third assortment has products of both types.

We perform a parametric computational experiment. The base case is set up as follows. We consider 100 products for each assortment type, i.e. $|J| = 100$. The input parameters in the base case are identical for all three assortments, except for the carbon emissions from transportation. For each product $j \in J$, the period demand D_j follows a negative binomial distribution. To create heterogeneous assortments, the parameters of this negative binomial distribution are randomly drawn from two separate distributions for each product j. The mean $\mu_{D_j} := \mathbb{E}[D_j]$ is randomly drawn from a gamma distribution with mean 100 and coefficient of variation of 0.5. The coefficient of variation $CV_{D_j} := \sqrt{\mathbf{Var}[D_j]}/\mu_{D_j}$ is randomly drawn from a shifted beta distribution with mean 0.9, standard deviation of 0.25, and shifted to the right by 0.3. Since low demand products typically have higher holding cost, the holding cost h_j is negatively correlated with the mean demand μ_{D_j} of each product j through a Gaussian copula with a fitting covariance matrix. In particular, h_j is drawn from a gamma distribution with mean 1 and coefficient of variation equal to CV_{D_j}, with a Pearson correlation coefficient of -0.5. Details regarding our approach to generate correlated random numbers are relegated to Appendix 3.B.

We set $l_{j,f}$ and $c_{j,s}$ to 0 for all products, and focus on $l_{j,s}$ and $c_{j,f}$, which now coincide with the lead time difference and the cost premium of product $j \in J$, respectively. We set $l_{j,f} = 3$ for all products. The back-order penalty cost p_j for product j is a function of its holding cost h_j. The ratio between the holding and the penalty cost is an important determinant of the service level in an inventory system. Therefore we set $p_j = \psi_j \chi_j^p h_j$ where ψ_j is a parameter we use to control the ratio between p_j and h, and χ_j^p is a random perturbation. That is, χ_j^p has a shifted beta distribution with mean 0.98, standard deviation 0.1, and shifted 0.02 to the right.

The cost premium $c_{j,f}$ of the fast transport mode of product j equals $\chi_j^c p_j l_{j,f}$, where χ_j^c has a beta distribution with mean 0.25 and standard deviation 0.1. Table 3.2 provides a summary of how we randomly generated the products of the base case; we shortly explain how we randomly generated the emission units for these products for all the three assortment types. In Table 3.2, $NB(\mu, cv)$ denotes a negative binomial random variable with mean μ and coefficient of variation cv, $\Gamma(\mu, cv)$ denotes a gamma random variable with mean μ and coefficient of variation cv, and $B(\mu, \sigma, s)$ denotes a beta random variable with mean μ and standard deviation σ that is shifted to the right by s; if we drop s then this beta random variable is not shifted.

As explained in Section 3.3, we rely on the NTM framework (NTM, 2015) to set emissions based on the structure of equation (3.1). We apply the NTM framework to data from the UN Comtrade Database (United Nations Comtrade Database, 2020) to obtain sample emission units; details regarding this methodology are relegated

Table 3.2 Generating the base case input parameters.

Input parameter	Generation
$\|J\|$	100
D_j	$NB(\mu_{D_j}, \mathrm{CV}_{D_j})$, with $\mu_{D_j} \sim \Gamma(100, 0.5)$ and $\mathrm{CV}_{D_j} \sim B(0.9, 0.25, 0.3)$
$l_{j,f}$	0
$l_{j,s}$	3
h_j	$\Gamma(1, 0.5)$, with $\rho_{\mu_{D_j}, h_j} = -0.5$
p_j	$\psi_p \chi_j^p h_j$, with $\psi_p = 9$ and $\chi_j^p \sim B(0.98, 0.1, 0.02)$
$c_{j,s}$	0
$c_{j,f}$	$\chi_j^c p_j l_{j,f}$, with $\chi_j^c \sim B(0.25, 0.1)$

to appendix 3.A. We then apply maximum likelihood estimation on these sample unit emissions to obtain three distinct sets of two distribution functions; two for each assortment type of products. The number of carbon emission units $e_{j,m}$ for product $j \in J$ with transport mode $m \in M$ are then randomly drawn from these distributions. These distributions are presented in Table 3.3. In this table, $LN(\mu, \sigma)$ denotes a random variable whose logarithm is normally distributed with mean μ and standard deviation σ, and $WB(\lambda, k)$ denotes a Weibull random variable with scale λ and shape k.

Table 3.3 Generating the emission units for the base case.

Assortment type	Emission parameter	Generation
1	$e_{j,s}$	$\Gamma(0.35, 0.21)$
	$e_{j,f}$	$e_{j,s} + LN(1.52, 0.21)$
2	$e_{j,f}$	$\Gamma(0.19, 1.27)$
	$e_{j,s}$	$e_{j,f} + \Gamma(2.19, 1.27)$
3	$e_{j,f}$	$WB(0.87, 0.77)$
	$e_{j,s}$	$\Gamma(3.31, 1.34)$

The total allowable carbon emissions from transportation $\mathcal{E}^{\max}$ is set as a percentage of the total *reducible* carbon emissions. For each instance of the test bed, the reducible carbon emissions is defined as the difference between the total amount of carbon emissions of the control policy that is optimal for Problem (P) absent of the emission constraint and the total amount of carbon emissions of the control policy that leads to the lowest possible total carbon emissions. The latter implies that each product is only shipped with its least polluting transport mode. Under the dual-index policy, setting Δ_j to zero implies that all orders for product $j \in J$ are shipped with its fastest transport mode. Alternatively, letting $S_{j,f}$ go to $-\infty$ implies that all orders for product j are shipped with its slowest transport mode (Veeraraghavan and Scheller-Wolf, 2008).

To evaluate the effectiveness of the column generation procedure in solving Problem

(P), we compute for each instance of the test bed the relative difference between the total average cost under a feasible solution and the corresponding lower bound. That is,

$$\%GAP = 100 \cdot \frac{C_P^{UB} - C_P^{LB}}{C_P^{LB}},$$

where C_P^{LB} and C_P^{UB} are obtained using the methods described in Section 3.4.

As described above, the dual-index policy can mimic static mode selection. Hence, to find a feasible static mode selection solution to Problem (P), we apply our column generation procedure of Section 3.4 in which we restrict the solution space for each product $j \in J$ such that all orders are shipped with either the fastest or the slowest transport mode. (Note that in the Master Problem (MP) for this approach the set of possible policies K_j for each product $i \in J$ contains only two single transport mode policies.) The long run average cost of this solution is denoted C_P^{SMS}. To quantify the value of dynamic mode selection, we compare for each instance of the test bed the long run average cost of static mode selection with the long run average cost of dynamic mode selection. That is,

$$\%SMS = 100 \cdot \frac{C_P^{SMS} - C_P^{UB}}{C_P^{UB}},$$

where $\%SMS$ indicates the relative increase in the long run average cost when the company chooses to rely on only one transport mode for each product in meeting an assortment wide constraint on total emissions.

To quantify the portfolio effect, we define for each instance of the test bed an additional benchmark in which we enforce an emission constraint for each product $j \in J$. This emission constraint is set as a percentage of the total reducible emissions of the individual product rather than of the entire assortment of products. We compute a feasible solution for this benchmark instance with the column generation procedure of Section 3.4, which can readily be modified so that it can be applied to settings where we have additional emission constraints. (An alternate solution procedure is to use a simple search procedure per product.) We refer to this approach as blanket mode selection. The long run average cost of this solution is denoted C_P^{BMS}. To quantify the portfolio effect, we compare for each instance of the test bed the long run average cost of blanket mode selection with the long run average cost of dynamic mode selection. That is,

$$\%BMS = 100 \cdot \frac{C_P^{BMS} - C_P^{UB}}{C_P^{UB}},$$

where $\%BMS$ indicates the relative increase in the long run average cost when the

company enforces an emission constraint on each individual product rather than one single constraint for the emissions from the entire assortment.

In solving the column generation sub-problem for each product, we simulated 10 samples of 9500 time periods following a 5000 time periods warm-up . The width of the 95 percent confidence interval of the long run average cost per period for each product was no larger than 3 percent of its corresponding point estimate for each instance of the column generation sub problem that we solved. The average computational time of our column generation procedure was 23 minutes, 5 minutes for blanket dynamic mode selection, and less than 5 sec for static mode selection.

3.5.1 Results for the base case

Figure 3.1 presents the normalized optimal average costs of each approach for each assortment group under emission targets that range from from 0 to 100 percent of the total reducible carbon emissions. Observe that for each assortment group, all approaches have the same performance when the emission target is set at 100 percent of the total reducible emissions. In this case, all approaches solely utilize the least polluting transport mode. Alternatively, when we impose no target on the emissions from transportation, then both dual mode approaches perform equally well while the static mode selection approach seems to perform the poorest over all possible emission targets. Indeed, static mode selection is more than 15 percent more expensive than both dual dual mode approaches for all assortment types when transportation emissions are not constrained. Based on Figure 3.1, we conclude that dynamic mode selection, as opposed to static and blanket mode selection, has great potential to efficiently curb carbon emissions from transportation at relatively little additional costs across all assortment types.

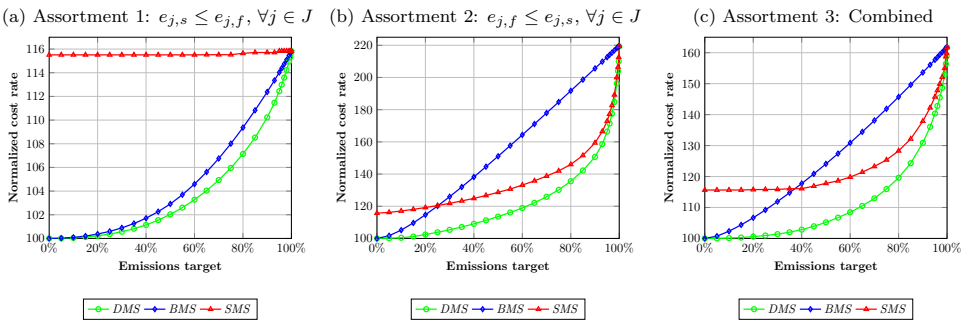

Figure 3.1 Optimal normalized absolute costs of each approach for different targets on the reducible emissions.

We explicitly compare our dynamic mode selection with the benchmark approaches in Figure 3.2, which presents the $\%SMS$ and $\%BMS$ percentages for each assortment group under emission targets that range from from 0 to 100 percent of the total reducible carbon emissions. The figure indicates that the performance of static mode selection over dynamic mode selection is consistent across all assortments. The relative increase in its total cost over dynamic mode selection is the largest when there is no emission target, and gradually decreases as the emission constraint tightens. At moderate carbon emission targets, around 40 to 60 percent of the total reducible emissions, static mode selection still leads to increases in the total average cost per period of around 10 to 15 percent for all assortment types.

The performance of the blanket mode selection approach depends on the specific assortment type. Figure 3.2(a) illustrates that when the unit emissions from the fast transport mode are more than those from the slow transport mode, the performance of the blanket mode selection approach seems to be quite reasonable. This can be explained as in this setting, the cost of the fastest transport mode is larger than the cost of the slowest transport mode. The most polluting transport mode is thus also the most expensive transport mode. The portfolio effect is limited for this assortment type.

For the other two assortment types, however, the cheapest transport mode is not necessarily also the least polluting transport mode, and the portfolio effect is more prevalent. Figures 3.2(b) and 3.2(c) show that $\%BMS$ can be more than 35 and 20 percent in assortment type 2 and 3, respectively, under carbon emission reduction targets of 50 percent. The static mode selection approach, which also takes advantage of the portfolio effect, even outperforms the blanket mode selection approach for quite some emission targets.

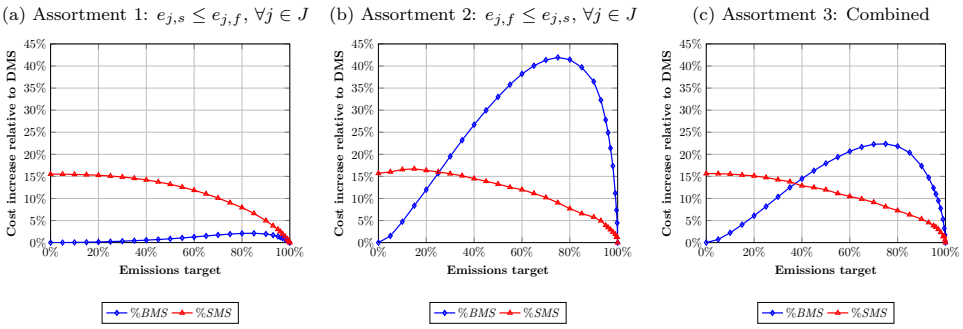

Figure 3.2 Relative surplus of the optimal cost of the alternative approaches (BMS and SMS) compared to the DMS approach for different targets on the reducible emissions.

To recapitulate, the value of dynamically shipping products with two transport modes

simultaneously rather than statically selecting one transport mode a priori is quite large. Regardless of the assortment type, $\%SMS$ is in between 5 and 15 percent for emission reduction targets up to 90 percent. The portfolio effect depends on the specific assortment type. If the least polluting transport mode of each product is also its cheapest transport mode, then the fastest and most polluting transport modes are typically only relied upon in case of imminent backorders. This behavior remains in case of an assortment-wide emission target, and the portfolio effect is consequently rather limited. If the least polluting transport modes are not necessarily the cheapest transport modes, then there is substantial value to be reaped in optimizing the assortment of products under a single emission constraint rather than under separate emission constraints for each individual product. Indeed, $\%BMS$ can go up to 40 and 20 percent for assortment type 2 and 3, respectively.

Table 3.4 below presents the average relative slack in the emission constraints for each assortment over the different emission targets considered in the base case analysis. The table shows that due to the binary nature of the static mode selection approach, the total average emissions under this approach are often substantially lower than the target level. This leads to particularly poor performance for assortment type 1. We observed in our computational experiments that for this assortment type, the static mode selection approach selects the cheapest and thus least polluting transport mode for almost all products under each emission target.

Table 3.4 Average slack in emission constraints for the base case analysis.

	Assortment type		
Approach	1	2	3
DMS	0.00%	0.00%	0.00%
BMS	0.08%	0.61%	0.12%
SMS	12.16%	0.17%	0.97%

The average $\%GAP$ of the base case over all emission targets is less than 0.01 percent, indicating that the column generation procedure finds feasible solutions that are close to optimal. Such a low average $\%GAP$ occurs because there can be at most 1 product for which the optimal solution to Problem (MP) is fractional. Indeed, Problem (MP) has $|J|+1$ constraints and an optimal solution for this problem has the same number of basic variables. Constraint (3.4) assures that for each product $j \in J$ a convex combination of policies is chosen. As such, there is at least one basic variable for each product j. This implies that there is at most 1 product for which the optimal solution to Problem (MP) is fractional.

3.5.2 Comparative statics

In this section, we study how changes in the input parameters with respect to the base case affect the performance of the blanket mode selection and the static mode selection approach. In what follows, we keep the emission target fixed at a 50 percent reduction of the reducible emissions, and we study the effects of changing a certain input parameter while generating the other input parameters as in the base case, i.e. as in Table 3.2. We also investigate the effect of scaling the emission differences between the most polluting and least polluting transport modes when the target on the total emissions is kept fixed. To achieve this, we first generate emission units as in the base case. We subsequently change the emission units of the most polluting transport mode through scaling $e_j :=\mid e_{j,s} - e_{j,f} \mid$ by a constant δ_e while keeping the emission units from the least polluting mode fixed at its base level. The changes in the parameters we investigate are summarized in Table 3.5.

Table 3.5 Changes in the base case input parameters.

Parameter	Generation	Base case	Changes
$\|J\|$		100	$\|J\| \in \{40, 60, 80\}$
μ_{D_j}	$\Gamma(100, CV_{\mu_{D_j}})$	0.5	$CV_{\mu_{D_j}} \in \{0.3, 0.4, 0.6, 0.7\}$
CV_{D_j}	$B(0.9, 0.25, s_{D_j})$	0.3	$s_{D_j} \in \{0.2, 0.25, 0.35, 0.4\}$
$\rho_{\mu_{D_j},h_j}$		-0.5	$\rho_{\mu_{D_j},h_j} \in \{-0.3, -0.4, -0.6, -0.7\}$
ψ_p		9	$\psi_p \in \{3, 4, 5, 19, 99\}$
χ_j^c	$B(0.25, \sigma_c)$	0.1	$\sigma_c \in \{0.15, 0.2, 0.3, 0.35\}$
$l_{j,s}$		3	$l_{j,s} \in \{2, 4\}$
δ_e		1	$\delta_e \in \{0.8, 0.9, 1.1, 1.2\}$

Figure 3.3 shows the effect of changing the coefficient of variation of the gamma distribution from which we sample the mean demands per period for each product. The figure indicates that this effect is relatively limited. With respect to the base case, both $\%BMS$ and $\%SMS$ only change up to 1 percent point for all three assortment types. We can draw a similar conclusion for the effect of changing the Pearson correlation coefficient between the holding cost and the mean demand per period of each product. Figure 3.4 shows that both $\%BMS$ and $\%SMS$ change at most 1 percent point with respect to the base case for all three assortment types.

Alterations in the shift parameter of the beta distribution from which we sample the coefficient of variation of the demand per period for each product has a relatively moderate effect on both $\%BMS$ and $\%SMS$. Figure 3.5 illustrates that for all assortment types, the $\%SMS$ tends to increase in the variability of the demand while the $\%BMS$ decreases. This indicates that the flexibility to dynamically ship products with two transport modes has particular merit in highly variable demand settings.

Figure 3.6 indicates that for all assortment types, an increase (decrease) in the lead

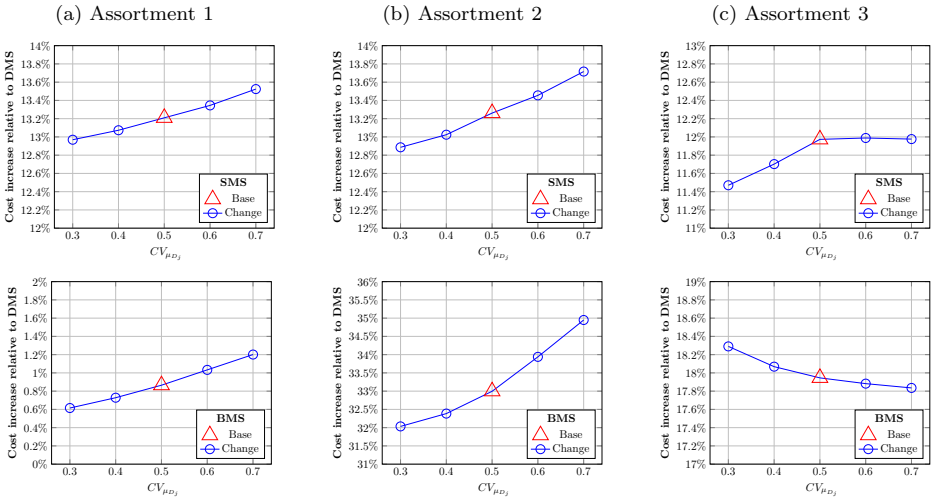

Figure 3.3 Effect of changing $CV_{\mu_{D_j}}$ while keeping the rest of the parameters as in the base case.

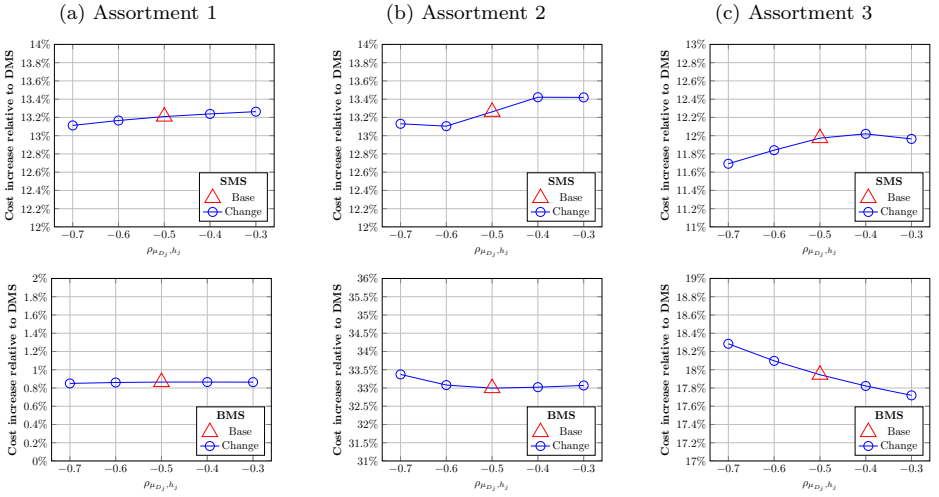

Figure 3.4 Effect of changing $\rho_{\mu_{D_j},h_j}$ while keeping the rest of the parameters as in the base case.

time difference between the fastest and the slowest transport modes of each product leads to an increase (decrease) in both $\%SMS$ and $\%BMS$. The blanket mode selection

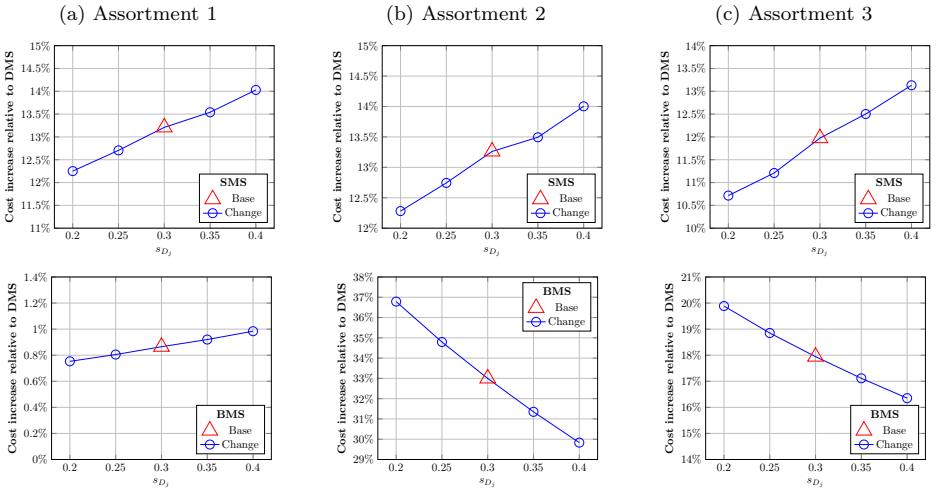

Figure 3.5 Effect of changing s_{D_j} while keeping the rest of the parameters as in the base case.

approach seems to be more susceptible to changes in the lead time difference than the static mode selection approach. For assortment type 2, for instance, an increase in the lead time difference to 4 leads to an increase in $\%BMS$ of 7 percent points with respect to the base case. By contrast, $\%SMS$ increases only slightly by 0.5 percent points. This can be attributed to the fact that the blanket mode selection approach imposes constraints on the emissions of each individual product while the static mode selection approach imposes a single constraint on the entire assortment of products.

Figure 3.7 illustrates the effect of changing the critical ratio for all products through varying ψ_p. We conclude that this effect is quite large. For assortment type 1, for instance, $\%SMS$ varies from 20 percent to 2 percent. While $\%SMS$ seems to decrease in the critical ratio for all products, $\%BMS$ tends to increase. For assortment type 2 and 3, for instance, $\%BMS$ increases from 15 percent to over 80 percent. These effects can be explained by the fact that as the critical ratios of all products approach 1, our dynamic mode selection approach will mimic the static mode selection approach in which the assortment wide emission constraint is met by relying on less polluting transport mode for product for which this is relatively cheap to do so. The blanket mode selection approach, however, must meet emission targets for each product individually which leads to poor performance if we increase the critical ratios for all products.

Figure 3.8 indicates that $\%SMS$ decreases in the cost of the fast transport mode. This can again be explained by the fact that our dynamic mode selection approach will also

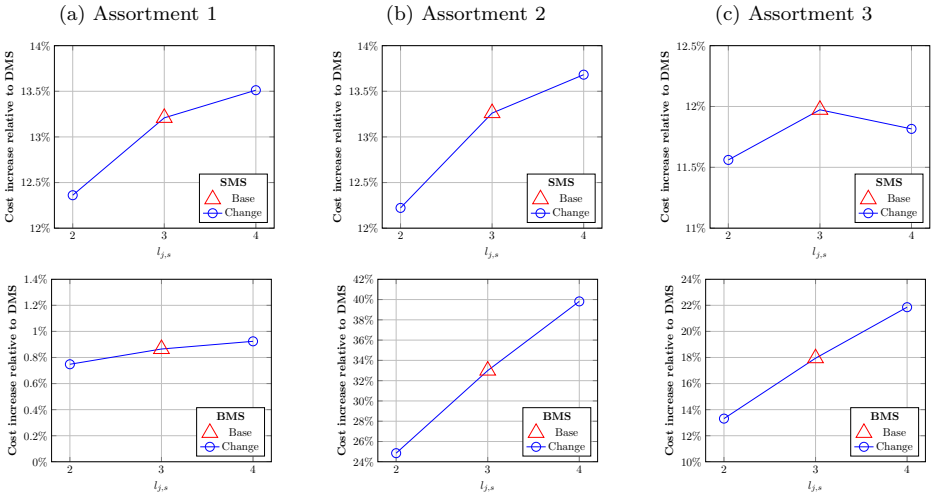

Figure 3.6 Effect of changing $l_{j,s}$ while keeping the rest of the parameters as in the base case.

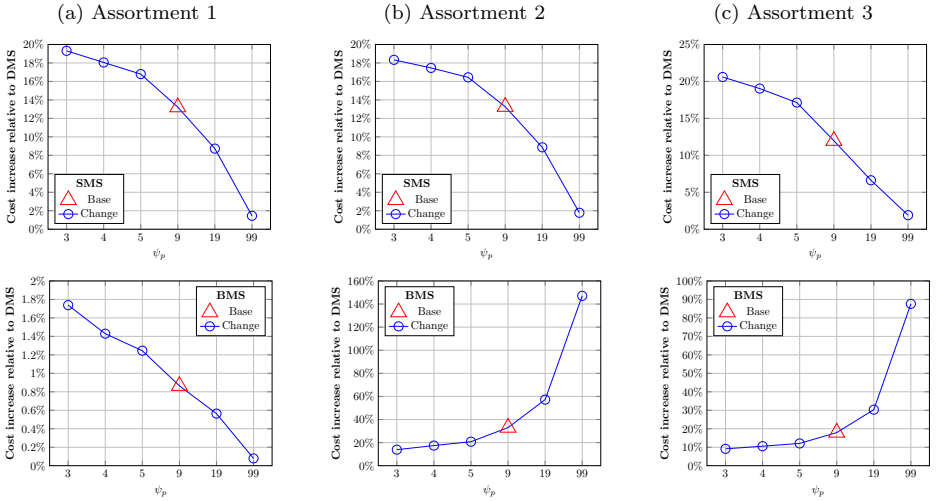

Figure 3.7 Effect of changing ψ_p while keeping the rest of the parameters as in the base case.

rely more on the cheaper transport mode as the cost premium for the fast transport mode increases, and that consequently the gap with the static mode selection approach

decreases.By contrast, $\%BMS$ increases in the cost of the fast transport mode. This can be attributed to the fact that the blanket mode selection approach, contrary to the other two approaches, imposes itemized emission constraints and relying on the most expensive but least polluting transport mode is inevitable. We can draw similar conclusions for the effects scaling the emission units of the most polluting transport mode, see Figure 3.9. Finally, Figure 3.10 indicates that the impact of the assortment size on both $\%SMS$ and $\%BMS$ is relatively limited. With respect to the base case, both $\%BMS$ and $\%SMS$ only change up to 2 percent point for all three assortment types.

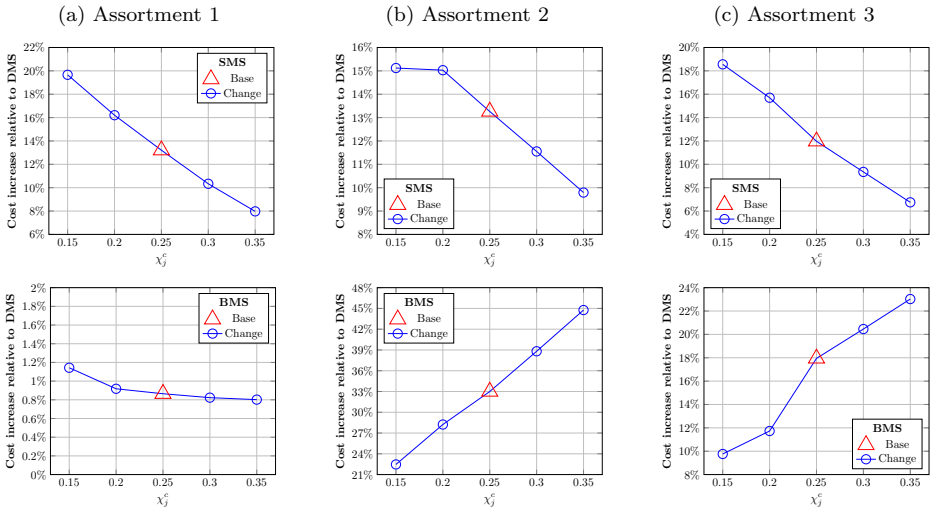

Figure 3.8 Effect of changing χ_j^c while keeping the rest of the parameters as in the base case.

3.6. Concluding remarks

As carbon emissions from the transportation sector are projected to increase over the next decades, it is important for companies to rethink their supply chain strategies and explicitly incorporate carbon emissions into their decision making. In this paper, we have studied the inbound transport and inventory management decision making for a company that sells an assortment of products. The company wishes to minimize inventory costs while keeping the total emissions from the inbound transport of the entire assortment below a certain target level. Each product can be shipped using two distinct transport modes. As each mode has its own merits, we have proposed a

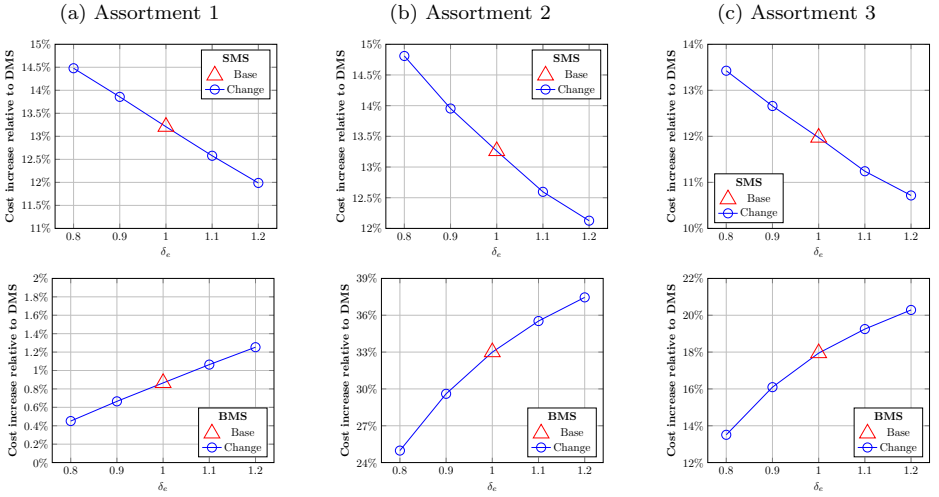

Figure 3.9 Effect of changing δ_e while keeping the rest of the parameters as in the base case.

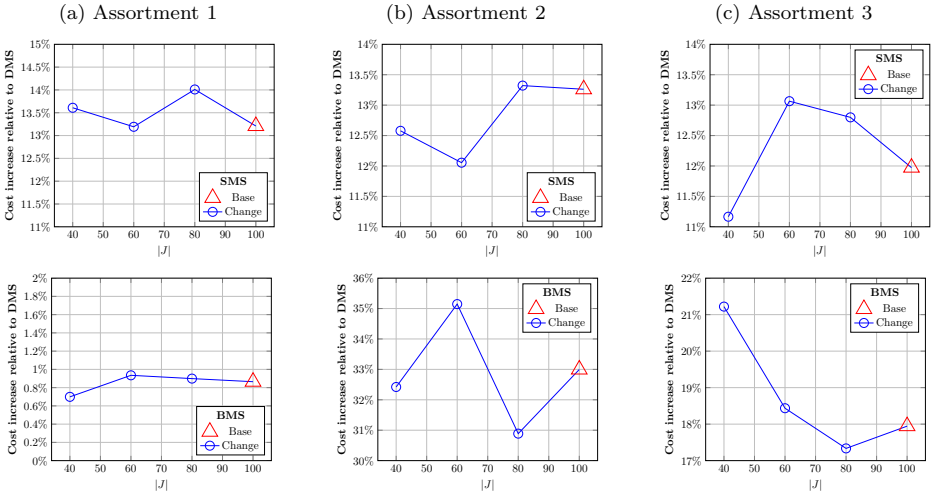

Figure 3.10 Effect of changing $|J|$ while keeping the rest of the parameters as in the base case.

dynamic mode selection model that allows the company to ship products with either mode depending on when one mode is more favorable than the other. Since the optimal policy for dual transport mode problems are known to be complex, we have

assumed that shipment shipment quantities for each product are governed by a dual-index policy. We have formulated the resulting decision problem as a mixed integer linear program that we have solved through a column generation solution procedure. This column generation procedure decomposes the complex multi-product problem into smaller sub-problems per product. These sub-problems are readily solved through a simple bisection search over Newsvendor type problems.

In an extensive computational experiment, we have compared the performance of our dynamic mode selection approach with two alternative approaches that are considered state-of-the-art. The first benchmark, static mode selection, lacks the flexibility to dynamically ship products with two transport modes; it rather selects one transport mode for each product a priori. The second benchmark, blanket mode selection, does have the flexibility to rely on two transport modes simultaneously but it makes transport decisions for each product individually rather than holistically for the entire assortment. Our computational experiments indicate that the value of our dynamic mode selection approach over the blanket mode selection approach is particularly high for assortments of products for which the fastest transport modes are not necessarily the most polluting transport modes. For such settings, our dynamic mode selection approach can reduce the long run average costs by 40 percent under moderate carbon emission targets. These huge savings can be attributed to the portfolio effect inherent to our approach. The computational experiments further indicate that dynamic mode selection can significantly outperform static mode selection. Under moderate emission targets, dynamically relying on two transport modes rather than a single transport mode can lead to cost savings of up to 15 percent.

3.A. Carbon accounting

In this section, we briefly explain how we determine the distribution functions that we use for pseudo-random generation of the unit emissions in our computational experiments. We utilize the United Nations Comtrade Database (2020) to calculate the average unit weights for 122 groups of products imported by The Netherlands in 2020. These product groups consist of two categories: (i) apparel goods that are imported from Vietnam and (ii) industrial goods that are imported from China and Germany. We consider air transport (from Tan Son Nhat international airport) and sea transport (from Haiphong port) as the fast and slow transport mode for the apparel category, respectively. For the industrial goods category, we assume sea transport from Shanghai, China, as the slow mode and road transportation from Stuttgart, Germany, as the fast mode.

We rely on the Network for Transport Measures methodology (NTM, 2015) to model and measure transportation emissions based on equation 3.1. This model has been widely used in literature (e.g., Hoen et al., 2014a,b). Following the NTM methodology, we first compute the overall carbon emissions generated by a single vehicle and then allocate a proportion of those emissions to each freight unit carried by the vehicle.

Sea transportation. We assume that all products are shipped via container. The average age of the container fleet worldwide is around 12 years and the average vessel size (dwt) of container ships with age 10-14 is 43,993 ton (United Nations Conference on Trade and Development, 2020). Based on section 7 of the NTM framework and resolutions of the Marine Environment Protection Committee (International Maritime Organization, 2011), we approximate sea transportation emissions in kilograms of CO_2 of one unit of a certain product with weight w (in kilogram) for a certain trip with distance d (in kilometers) using the following relation,

$$e_{sea} = w \cdot EI_{ship} \cdot 10^{-3} \cdot d$$

where EI_{ship} is kilograms of CO_2 emissions per kilogram weight per kilometer. Furthermore EI_{ship} is computed through,

$$EI_{ship} = \frac{(a \cdot dwt^{-c})/(PDR_{ship} \cdot LCU)}{1.852}$$

where a and c are constants, dwt is the deadweight tonnage of the ship, LCU is average load capacity utilization, $F(LCU)$ is fuel consumption as a function of load, and PDR_{ship} is the payload of the ship. 1.852 is the nautical mile to km conversion coefficient. For a container ship, NTM methodology states: $a = 0.17422, c = 0.201, LCU = 0.70, F(LCU) = 1$, and $PDR = 0.8$. Succinctly, we have for the total emissions in kilograms of CO_2 of one unit of a product with weight w (in kilogram)

for a sea trip with distance d (in kilometers)

$$e_{sea} = w(1.996 \cdot 10^{-5} \cdot d). \tag{3.7}$$

Air transportation. Our calculations for the emissions of air transportation are based on section 8 of the NTM Framework. We consider an Airbus A310-300 F as the aircraft. Based on the May 2021 Air Cargo Market Analysis of The International Air Transport Association, we assume an average international cargo load factor of 65%. Following the NTM Framework, we have the following relation for air transportation emissions

$$e_{air} = \frac{w}{c_{max}}(CEF + VEF \cdot d),$$

where c_{max} is the maximum freight load, CEF is the constant emissions factor, and VEF is the variable emissions factor. CEF and VEF are the outcomes of applying a linear regression on real data provided by the NTM. We obtain the CEF and VEF parameters via interpolation over the associated tables provided by the NTM. We furthermore assume $c_{max} = 39,000kg$ as per section 8.3.1 and perform the interpolation on table 4.1 of section 8.2.1. Succinctly, we have for the total emissions in kilograms of CO_2 of one unit of a product with weight w (in kilogram) for an air trip with distance d (in kilometers)

$$e_{air} = w(1.525 \cdot 10^{-1} + 4.938 \cdot 10^{-4} \cdot d). \tag{3.8}$$

Road transportation. We rely on Hoen et al. (2014a) to obtain the emission units of road transportation. They too rely on the NTM framework to estimate CO_2 emissions from road transportation in Europe. In particular, they approximate the total emissions in kilograms of CO_2 of one unit of a product with weight w (in kilogram) for a truck trip with distance d (in kilometers) as

$$e_{road} = w(3.214 \cdot 10^{-4} + 4.836 \cdot 10^{-5} \cdot d). \tag{3.9}$$

Distances. We use NTMCalc Basic 4.0 (NTM, n.d.), which is an online tool provided by the NTM for approximating emissions, to calculate travel distances between the origin destination pairs as described at the beginning of this section. Based on this tool we find that the sea distance between Haiphong and Rotterdam is 9,610 nautical miles (17,798 km) and the sea distance between Shanghai and Rotterdam is 10,525 nautical miles (19,492 km). The distance traveled by aircraft between Tan Son Nhat international airport and Rotterdam The Hague Airport is 10,073 km, and the road distance between Stuttgart and Rotterdam is 633 km. With these distances, we compute the total kilogram CO_2 emissions of one unit of product with weight w for each mode-trip category using Equations (3.7)-(3.9). We call these emission

coefficients and they are presented in Table 3.6 below.

Table 3.6 Emissions coefficients for each mode-trip category.

Industry	Slow Mode	Fast Mode
Apparel	$3.552 \cdot 10^{-1}$	5.127
Industrial	$3.891 \cdot 10^{-1}$	$3.093 \cdot 10^{-2}$

Fitting distribution functions. We use the emission coefficients from Table 3.6 to calculate the unit emissions for the 122 groups of products mentioned at the beginning of this section. We subsequently use maximum likelihood estimation on the resulting emission units to find distribution functions from which we can sample the emission units of the fast and slow transport modes for all three assortment types in our computational experiment. The emission unit distributions for assortment type 1 are based on the apparel category, the emission unit distributions for assortment type 2 are based on the industrial category, and the emission unit distributions for assortment type 3 are based on both categories. The final distribution functions are provided in Table 3.3.

3.B. Generating correlated random numbers

Suppose X and Y are two real random variables with marginal distribution functions F and G, respectively. Suppose their joint distribution is bi-variate standard normal $\mathcal{N}_\rho$ with Pearson's correlation coefficient $\rho = \mathbf{Cov}(X,Y)/(\sqrt{\mathbf{Var}[X]}\sqrt{\mathbf{Var}[Y]})$ (Nelsen, 2006). Let Z be a vector of size two with independent random elements that have standard normal distributions Φ, and let $W = AZ$ be a linear combination of Z with

$$A = \begin{bmatrix} 1 & 0 \\ \rho & \sqrt{1-\rho^2} \end{bmatrix}.$$

It can be shown that W has a bivariate normal distribution $\mathcal{N}_\rho$ with covariance matrix $\Sigma = AA^T$ (see, e.g., Gut, 2009). We use this result to sample from X and Y as follows:

1. Generate the vector $Z = \begin{bmatrix} Z_1 \\ Z_2 \end{bmatrix}$ by independently sampling from a standard normal distribution function,
2. Calculate the bivariate normal sample $W = \begin{bmatrix} W_1 \\ W_2 \end{bmatrix} = AZ$,
3. Generate the required samples by inversion $\begin{bmatrix} X \\ Y \end{bmatrix} = \begin{bmatrix} F^{-1}(\Phi(W_1)) \\ G^{-1}(\Phi(W_2)) \end{bmatrix}$.

Chapter 4

Effective dual-sourcing through inventory projection

4.1. Introduction

Dual-sourcing is a common and effective supply chain management practice employed by companies that operate both globally and domestically. It generally refers to situations where companies can replenish their inventories from a regular supplier as well as from an expedited supplier; the latter offering a shorter lead time than the former but at the expense of a cost premium. Even though the expedited supplier charges an additional cost, companies can leverage its responsiveness when inventories deplete precipitously.

Examples of companies that have utilized a dual-sourcing strategy are manifold. Allon and Van Mieghem (2010) describe a $10 billion high-tech U.S. company with suppliers in China and Mexico. The Chinese supplier charges lower costs than Mexico but shipping its products to the U.S. takes considerably longer. Rao et al. (2000) describe a similar situation at Caterpillar; they too use a more expensive expedited supplier whenever a shorter lead time is desired. While both examples concern two distinct suppliers, dual-sourcing may also refer to situations where companies can replenish their inventories from one supplier using two distinct transport modes. For example, Threatte and Graves (2002) describe a situation in which Polaroid Corporation relies on ocean transport to ship products from a supplier in Asia to the U.S., but occasionally uses expedited air transport when inventory levels are critically low.

In this chapter, we treat a setting similar to the ones described above. In particular, we consider a single-item, single-echelon dual-sourcing inventory system under periodic review facing stochastic demand, where excess demand is backlogged. We assume that demands across periods are independently and identically distributed (i.i.d.). The expedited supplier has a shorter lead time than the regular supplier but charges a higher unit price. There are no fixed costs associated with ordering at either supplier. At the end of each period, linear holding and backorder costs are charged for leftover inventory and backlogged demand, respectively. We seek to minimize the long run average total cost per period, which consists of ordering, holding, and backorder costs.

The dual-sourcing system introduced above has been studied for almost sixty year now, starting with the contributions of Barankin (1961) and Neuts (1964). They show that the optimal policy exhibits a simple order-up-to structure when the lead times of the regular and expedited supplier are one and zero periods, respectively. Fukuda (1964) later extends this result to the case where the difference between the lead times of both suppliers is one period, irrespective of their individual lead times. For lead time differences larger than one period, Whittemore and Saunders (1977) and more recently Feng et al. (2006) show that the structure of the optimal policy is complex, and depends in general on the entire vector of outstanding orders. Although the optimal policy for general lead time differences can be computed in principle by stochastic dynamic programming, this is not practical as the state space of the dynamic program grows exponentially in the lead time difference; that is, it suffers from the curse of dimensionality. Most researchers have therefore focused on devising relatively simple but non-optimal replenishment policies (we will provide a thorough review later in Section 4.2).

Even though some heuristic policies have been shown to perform numerically close to optimal in at least certain regimes, theoretical justifications for such useful asymptotic optimality performance guarantees are mostly lacking; the only exception being the Tailored Base-Surge (TBS) policy. This policy orders a fixed amount at the regular supplier every period and uses an order-up-to rule for placing expedited orders. The TBS policy was first proposed by Allon and Van Mieghem (2010), though Rosenshine and Obee (1976) and Janssen and de Kok (1999) have studied a closely related policy before. Numerical analysis by Klosterhalfen et al. (2011) and Janakiraman et al. (2015) indicates that the optimality gap of the TBS policy closes as the lead time difference grows large.

Xin and Goldberg (2018) have recently provided theoretical justification for these numerical observations by proving that a simple TBS policy is asymptotically optimal as the lead time of the regular supplier grows large, with the lead time of the expedited supplier held fixed. The intuition behind this result is that as the lead time difference grows, the amount of randomness between when a regular order is

placed and when this order arrives grows so large that a constant order, which ignores the entire vector of outstanding orders, performs nearly optimal. To our knowledge, this is the only heuristic policy for dual-sourcing inventory systems under periodic review with general lead time differences that has an asymptotic performance guarantee. Unfortunately, the TBS policy performs poorly outside this asymptotic regime, particularly when either the backorder cost becomes large or when the unit price charged by the expedited supplier becomes large. When both cost parameters become large simultaneously, the performance of the TBS drops even precipitously (Klosterhalfen et al., 2011). Yet, it is exactly this setting that is prevalent in practice, and it therefore remains an open research problem to determine control policies that are provably asymptotically optimal in this regime.

In this chapter, we propose such a new policy, which we call the Projected Expedited Inventory Position (PEIP) policy. The PEIP policy keeps track of the expedited inventory position (net inventory level plus outstanding orders arriving within the expedited lead time) and places expedited orders such that this inventory position reaches a certain target level. Regular orders entering the information horizon for placing expedited orders may cause the expedited inventory position to exceed its corresponding target level. This implies that the expedited inventory position in expectation exceeds the target level. Under the PEIP policy, regular orders are therefore placed such that the expected expedited inventory position at the time of the regular order entering the expedited inventory position is kept at a target level. In doing so, the PEIP policy leverages all the information contained in the pipeline – as does the optimal policy. This is in contrast with existing heuristic policies that either aggregate all outstanding orders or ignore them entirely and place a constant order every period regardless. We further provide a separability result that enables efficient optimization of the PEIP policy. Hence, the PEIP policy is vastly more advanced than existing heuristic policies, and yet can be optimized efficiently and retains an intuitive appeal for practitioners. It is also assuring to note that for the case of consecutive lead times, the PEIP policy reduces to two order-up-to levels, which is known to be the optimal policy for that specific case (Fukuda, 1964).

We show that, under some mild conditions on the demand distribution, the relative difference between the long average cost per period of the best PEIP policy and the optimal policy converges to zero when both the shortage cost and the price for expedited units become large, with their ratio held constant. Our approach relies on establishing lower and upper bounds on the average cost rates of the optimal policy and the best PEIP policy, respectively, in terms of the average cost rates of canonical inventory systems with only one supplier. We then show that the ratio of these bounds converge to 1 in the non-trivial asymptotic regime where both the shortage cost and the price for expedited units grow large. Our research is useful for practice as it directly leads to an easily implementable heuristic that has asymptotic performance

guarantees in high service regimes where utilizing an expedited source is expensive. In the capital goods industry, for example, it is very common that the shortage costs for critical parts as well as the costs associated with performing an expedited shipment are relatively high (see, e.g., Westerweel et al., 2021).

The main contribution of this chapter is threefold:

1. We propose the PEIP policy for periodic review dual-sourcing inventory systems with general lead time differences, and we present an efficient optimization scheme to find its optimal policy parameters. This optimization scheme hinges on a separability result that allows us to decompose the optimization into two simple one-dimensional optimization problems.

2. We prove that the PEIP policy is asymptotically optimal when both the shortage cost and the cost premium for expedited units become large, with their ratio held constant. Our proof technique can serve as a template that is readily applicable to other heuristic polices. We illustrate this by showing that several existing heuristic policies are also asymptotically optimal in this non-trivial regime of broad interest.

3. Through an extensive numerical investigation, we conclude that the PEIP policy outperforms the best performing heuristics in literature. Our numerical investigation further shows that, as the lead time difference between the regular and expedited supplier grows, the PEIP policy converges to a TBS policy which is asymptotically optimal in this regime (Xin and Goldberg, 2018).

The remainder of this chapter is organized as follows. In the next section, we position our research with respect to other contributions that have been made. After that, in Section 4.3, we formally define the decision problem and introduce the notation that we use throughout this chapter. Section 4.4 presents the PEIP policy, including analytical results that allow for an efficient optimization procedure to find the optimal policy parameters. We proceed with establishing our asymptotic optimality result in Section 4.5. After that, we report on a large numerical study in Section 4.6, and we provide concluding remarks in Section 4.7.

4.2. Literature review

Dual-sourcing systems operating in discrete time as well as their variants (e.g., continuous review models, multi-echelon systems, etc.) have been widely studied. For a broad overview of this rich field, we refer the interested reader to the surveys of Minner (2003) and Svoboda et al. (2021), as well as to the recent studies of, e.g.,

Gong et al. (2014), Boute and Van Mieghem (2015), Arts et al. (2016), Sapra (2017), Song et al. (2017), Drent and Arts (2020), and the references therein. Here we confine ourselves to contributions most relevant to our research, a large part of which revolves around heuristic policies.

Whittemore and Saunders show already in 1977 that optimal policies for the periodic review dual-sourcing inventory systems we typically face in practice (i.e. with lead time differences larger than one) are complex and difficult to compute. Yet it lasted until the beginning of this century before researchers began to devise various effective heuristic policies. Scheller-Wolf et al. (2007) propose the Single-Index (SI) policy that uses order-up-to rules for order placement at the regular and expedited supplier based on only the regular inventory position (inventory level plus all outstanding orders). By additionally tracking the expedited inventory position that includes only orders that will arrive within the expedited lead-time, Veeraraghavan and Scheller-Wolf (2008) introduce the Dual-Index (DI) policy. Regular and expedited orders are then placed according to order-up-to rules based on their respective inventory positions. They propose a simulation-based optimization and show that this is easily implementable and often results in near-optimal solutions in numerical examples.

Since then, the DI policy has been extended in various useful ways. Sheopuri et al. (2010) analyze the more general class of policies that use an order-up-to rule for placing expedited orders. Arts et al. (2011) incorporate stochastic lead times and provide an approximate evaluation method so that optimization can be done without simulation. Sun and Van Mieghem (2019) introduce the Capped Dual-Index (CDI) policy in which the DI policy is accompanied with a cap that smooths regular orders. Even though the CDI policy has been developed in a robust optimization setting under the worst-case cost performance criterion, numerical analysis indicates that it also performs well in terms of minimizing the long run average cost per period in a stochastic setting.

Sheopuri et al. (2010) show that the inventory system with lost sales is a special case of the dual-sourcing inventory system treated here given that the expedited supplier follows an order-up-to rule. The intuition behind this is that the orders placed at the expedited supplier can be thought of as demand lost from the regular supplier, with the cost premium of the expedited supplier being the lost sales penalty cost. Inspired by this relation, Sheopuri et al. (2010) propose the Vector Base-Stock (VBS) policy that uses an order-up-to rule for the expedited supplier, and a vector base-stock policy for the regular supplier. Hua et al. (2015) later generalize the VBS policy to the Best Weighted Bounds (BWB) policy. In addition, by using properties of $L^\natural$-convexity, they show that the optimal expedited orders are more sensitive to the older (i.e. soon-to-arrive) outstanding orders whereas the optimal regular orders are more sensitive to the younger outstanding orders. Extensive numerical investigations indicate that

the BWB policy as well as the CDI policy described earlier are the best performing heuristic policies in the current literature, and can even outperform advanced deep reinforcement learning algorithms (Sun and Van Mieghem, 2019; Gijsbrechts et al., 2019).

While the PEIP policy is similar to the aforementioned heuristic policies in that it also assumes an order-up-to rule for the expedited supplier, it differs significantly in how it operates the regular supplier. First, the PEIP policy projects the expedited inventory position in deciding upon regular orders. As such, it controls the expected expedited inventory position directly taking into account explicitly the excess by which the order-up-to level of the expedited supplier overshoots; existing heuristics do so only indirectly. Second, similar to the optimal policy, the PEIP policy leverages all available state information in deciding upon order placement at the regular supplier. Existing heuristics either aggregate all available information, or ignore it entirely and place a fixed amount every period regardless.

Our research also contributes to the growing body of literature that establishes asymptotic performance guarantees of heuristic policies for inventory systems whose optimal policies are difficult to analyze if not intractable. The relatively few studies that belong to this literature stream mostly focus on single-echelon inventory systems with lost sales. For such systems, Huh et al. (2009) and Bijvank et al. (2014) show that certain order-up-to policies are asymptotically optimal as the lost sales penalty cost approaches infinity while Goldberg et al. (2016) and Xin and Goldberg (2016) show that a simple constant order policy is asymptotically optimal as the lead time approaches infinity. Both heuristic policies perform however poorly outside their respective asymptotic regimes. This has recently led to the development of two heuristic policies that have asymptotic performance guarantees and additionally perform well in both regimes: the Capped Base-Stock (CBS) policy proposed by Xin (2021a,b), and the Projected Inventory Level (PIL) policy proposed by Van Jaarsveld and Arts (2021).

The CBS policy can be interpreted as a hybrid between a constant order policy and an order-up-to policy, and it thus enjoys the superior performance of both policies in their respective regimes. The PIL policy places orders such that the expected inventory level at the time of arrival of an order is raised to a target level. Van Jaarsveld and Arts (2021) show that the PIL policy is asymptotically optimal for large lost sales penalty costs as well as for large lead times, though the latter only under the assumption that demand is exponentially distributed. The PEIP policy is related to the PIL policy: While a PIL policy places orders by projecting the expected inventory level, a PEIP policy places regular orders by projecting the expected expedited inventory position (while retaining an order-up-to rule for the expedited supplier).

Unlike inventory systems with lost sales, there currently exists only one heuristic

policy for periodic review dual-sourcing inventory systems that has asymptotic performance guarantees. Indeed, a simple TBS policy is asymptotically optimal as the lead time difference between the regular supplier and expedited supplier grows large (Xin and Goldberg, 2018). In this chapter, we consider a different asymptotic regime and show that the best PEIP policy is asymptotically optimal as the unit cost of the expedited supplier and the shortage cost approach infinity simultaneously. Our main proof exploits the earlier described relation between lost sales inventory systems and dual-sourcing inventory systems; this allows us to draw on the asymptotic results for lost sales inventory systems of Huh et al. (2009). We further argue that our proof technique in itself has merit too since it can serve as a template that is readily applicable to existing as well as future heuristic policies. We illustrate this by showing that the CDI, DI, and SI policies are also asymptotically optimal in the non-trivial asymptotic regime under consideration. As such, we extend the literature on asymptotic optimality results for dual-sourcing inventory systems considerably.

We end this review by remarking that recent advances on asymptotic optimality results for inventory systems that have complex optimal policies stretch beyond the lost sales and dual-sourcing inventory systems described above. See, for instance, Rong et al. (2017) for such results in the context of distribution systems consisting of one central warehouse and multiple local warehouses, Janakiraman et al. (2018) on inventory systems with multiple products that share production capacities, and Fu et al. (2019) on inventory systems for perishable products. We refer the interested reader to Goldberg et al. (2021) for a recent and comprehensive review of this growing field.

4.3. Notation and problem formulation

This section provides a formal problem description that mostly follows the notational conventions of Sheopuri et al. (2010). We consider a single-item, single-echelon inventory system under periodic review facing stochastic demand, where excess demand is backlogged. Products can be ordered each period both at a regular supplier and at an expedited supplier. Let c^r (l^r) and c^e (l^e) denote the unit prices (lead times) from the regular and expedited supplier, respectively. We assume that $c^e > c^r$ and $l^e < l^r$, that is, the expedited supplier offers a higher unit price, but its lead time is shorter. Notice that the present inventory problem reduces to a standard inventory problem with only one supplier if these assumptions do not hold. Both lead times are assumed to be deterministic as well as a non-negative integer multiple of the review period. For notational convenience, we also define the lead time difference $l = l^r - l^e \geq 1$.

Let D_t denote the random demand in period t, with t indexed forward in time starting from 0, i.e. $t \in \{0, 1, \ldots\}$. Demand across periods is a sequence of non-negative i.i.d. random variables. Whenever we drop the period index of D_t, we refer to the generic random variable with the same distribution as D_t. The cumulative demand over n periods is denoted as $\mathbf{D}^n$. Let I_t denote the net inventory level (on-hand inventory minus backlog) at the beginning of period t after orders have arrived, but before demand has occurred. Let q_t^e (q_t^r) denote the order size placed with the expedited (regular) supplier in period t. The orders that arrive in period t are thus $q_{t-l^e}^e$ and $q_{t-l^r}^r$, and we can consequently write the following recursion for the inventory level:

$$I_t = I_{t-1} - D_{t-1} + q_{t-l^e}^e + q_{t-l^r}^r.$$

At the end of each period t, costs are levied as follows. Per unit in on-hand inventory $(I_t - D_t)^+$ carried over to the next period, a cost $h > 0$ is charged. Similarly, a cost $p > 0$ is charged for each unit in backlog $(D_t - I_t)^+$ carried over to the next period. Here we use the convention $x^+ = \max(x, 0)$.

Observe that if the unit price charged by the expedited supplier c^e is higher than the total backlog cost incurred on a unit over l periods, i.e. $l \cdot p$, then there is no cost benefit to be derived from ordering any units from the expedited supplier over the regular supplier. In fact, as shown in Sheopuri et al. (2010), the present inventory problem reduces to a standard single supplier inventory problem if $c^e \geq p \cdot l$. We therefore assume throughout the remainder that $c^e < p \cdot l$. We revisit this assumption in Section 4.5 when we expound on the particular asymptotic regime we are interested in.

Let the state of the system at the beginning of period t after orders have arrived be given by $\mathbf{x}_t = (I_t, q_{t-l^r+1}^r, \ldots, q_{t-1}^r, q_{t-l^e+1}^e, \ldots, q_{t-1}^e)$. The sequence of events in any period t is as follows:

1. Products that were ordered $l^e (l^r)$ periods ago at the expedited (regular) supplier are received and added to the on-hand inventory, i.e. order $q_{t-l^e}^e (q_{t-l^r}^r)$.
2. The size of the order placed with the expedited supplier q_t^e is determined based on the current state of the system $\mathbf{x}_t$.
3. The size of the order placed with the regular supplier q_t^r is determined based on the current state of the system $\mathbf{x}_t$ and the order just placed with the expedited supplier q_t^e.
4. The demand D_t is realized, and holding or shortage costs are charged, in case that $(I_t - D_t)^+ > 0$ or $(D_t - I_t)^+ > 0$, respectively.

The total cost incurred in period t is denoted by C_t, i.e. $C_t = c^e \cdot q_t^e + c^r \cdot q_t^r + h \cdot$

$(I_t - D_t)^+ + p \cdot (D_t - I_t)^+$. An admissible policy π is as a rule that determines the size of the orders placed with the regular and expedited supplier based only on the historical information available in that period as outlined in the sequence of events above. All such admissible policies π are contained in the set Π. Since the total cost incurred in a period t depends on the policy π, we henceforth write C_t^π to denote this dependence explicitly. Our performance measure of interest is the long run average cost per period

$$C^\pi = \limsup_{T\to\infty} \frac{1}{T} \sum\nolimits_{t=1}^{T} C_t^\pi.$$

We assume throughout the remainder that $\mathbf{x}_0 = \mathbf{0}$. That is, the initial inventory level equals zero, and there are initially no outstanding expedited or regular orders due to arrive within their corresponding lead times. This assumption simplifies our analysis later; it is however not strictly necessary as we are interested in the long run average cost per period, which is not affected by any specific choice for the initial state.

We conclude with two remarks on the optimal policy $\pi^\star = \operatorname{argmin}_{\pi\in\Pi} C^\pi$. First, as shown by Sheopuri et al. (2010), $\pi^\star$ depends on c_e and c_r only through the price premium of the expedited supplier, i.e. $c_e - c_r$. This result follows intuitively from the observation that any policy that has a finite long run average cost per period will incur at least a cost of $c^r \cdot \mathbb{E}[\mathrm{D}]$. We will therefore assume without loss of generality that $c_r = 0$. Second, even for small lead time differences, the structure of $\pi^\star$ remains poorly understood and depends in general on the entire vector of outstanding orders. The policy we propose in the next section also takes into account this entire vector when deciding on the size of the orders placed with both suppliers. However, different from $\pi^\star$, our policy does allow for tractable analysis and efficient optimization, even for large-sized problems often encountered in practice.

4.4. The projected expedited inventory position policy

In this section, we present the Projected Expedited Inventory Position (PEIP) policy for operating dual-sourcing inventory systems. We first formally define the rules under which both suppliers operate, and we subsequently present analytical results that allow for an efficient solution procedure to find the optimal policy parameters.

4.4.1 Order placement rules

Under a PEIP policy, the ordering rule for the expedited supplier is assumed to be an order-up-to rule that operates as follows. At the beginning of each period t after

orders $q^r_{t-l^r}$ and $q^e_{t-l^e}$ have arrived, we review the expedited inventory position

$$IP^e_t = I_t + \sum\nolimits_{i=t-l^e+1}^{t-1} q^e_i + \sum\nolimits_{i=t-l^r+1}^{t-l} q^r_i, \tag{4.1}$$

and, if necessary, place an expedited order q^e_t to raise IP^e_t to its target level S^e, i.e. $q^e_t = (S^e - IP^e_t)^+$. Note that regular orders entering the information horizon for placing expedited orders may cause the expedited inventory position to exceed this target level. That is, immediately after expedited orders have been placed, there is an overshoot defined as $O_t = IP^e_t + q^e_t - S^e = (IP^e_t - S^e)^+$. This implies that the expedited inventory position following expedited ordering is in expectation greater than or equal to the target level. In fact, strictly greater unless products are sourced exclusively from the expedited supplier, which we have argued is sub-optimal under the assumptions we make.

The PEIP policy aims to keep the expedited inventory position (and consequently the overshoot) at a fixed level. Although the expedited inventory position cannot actually be kept at a fixed level, it can be kept at a fixed level in expectation. More specifically, the PEIP places regular orders such that the expected expedited inventory position at the time those orders enter the information horizon for expedited order placement is raised to a target level. We henceforth refer to this target level as the projected expedited inventory position level, denoted U.

Before we provide a formal description of this ordering rule, we first state the following result related to the minimal required information for optimal regular order placement given that the expedited supplier follows an order-up-to rule. This information is sufficient to determine the expected expedited inventory position in the period that the regular order placed now enters the expedited inventory position.

Lemma 4.1. *(Sheopuri et al., 2010, Lemma 4.1) Under the assumption that the expedited supplier follows an order-up-to rule, the optimal regular order in period t, q^r_t, depends on $\mathbf{x}_t$ and q^e_t only through the l-dimensional state description $\mathbf{x}^r_t = (O_t, q^r_{t-1}, q^r_{t-2}, \ldots, q^r_{t-l+1})$.*

Observe from the definition of the expedited inventory position in (4.1) that this inventory position in period $t + l$, i.e. IP^e_{t+l}, includes the regular order q^r_t that we will place in the current period t. We write $\mathbb{E}[IP^e_{t+l}(q^r_t)|\mathbf{x}^r_t]$ to denote the conditional expected expedited inventory position l periods into the future in period t given the current state of the system $\mathbf{x}^r_t$ as a function of the regular order q^r_t that we will place in period t. The ordering rule for the regular supplier can now formally be described as follows. In each period t immediately after expedited ordering, we place a regular order q^r_t such that $\mathbb{E}[IP^e_{t+l}|\mathbf{x}^r_t]$ equals the projected expedited inventory position level

U. That is, the period t regular order q_t^r solves

$$\mathbb{E}[IP_{t+l}^e(q_t^r)|\mathbf{x}_t^r] = U. \tag{4.2}$$

The solution to (4.2) is unique and can be computed efficiently via a simple bisection search.

Proposition 4.1 below shows that under the PEIP policy, it is always possible to place such a non-negative order (possibly zero) to attain the projected expedited inventory position level U. We note that Van Jaarsveld and Arts (2021) have established a similar result in the context of lost sales inventory systems. The proof of the proposition can be found in Appendix 4.A.

Proposition 4.1. *Under the PEIP policy, the conditional expected expedited inventory position in period $t+l$ before regular ordering in period t is less than or equal to the projected expedited inventory position level U for all periods $t \in \mathbb{N}$.*

It is worth noting that the regular ordering rule under the PEIP policy is of the order-up-to type when the lead times of both suppliers are consecutive, i.e. when $l = 1$. To see this, recall that under such an order-up-to rule, regular orders are placed such that the usual inventory position is kept at a target level, say S^r. Now, observe that, following expedited ordering in some period t, we have $IP_{t+1}^e = S^e + O_t + q_t^r - D_t = S^r - D_t$, so that $\mathbb{E}[IP_{t+1}^e|\mathbf{x}_t^r] = S^r - \mathbb{E}[D]$. Thus, when $l = 1$, the ordering rule for the regular supplier under a PEIP policy with projected expedited inventory position level $U = S^r - \mathbb{E}[D]$ is equivalent to an order-up-to rule with target level S^r. This observation implies that for the special case of consecutive lead times, the PEIP policy reduces to the *optimal policy* (Fukuda, 1964). This feature is shared with the CDI, DI, SI, and BWB policy, but not with the TBS policy.

4.4.2 Optimization procedure

We now turn to the issue of finding the optimal PEIP policy parameters. Observe that the expedited inventory after ordering is bounded from below by the order-up-to level S^e. This implies that the projected expedited inventory position equals S^e plus the projected overshoot. A given projected expedited inventory position level U thus only scales IP^e for a given S^e, which suggests that order placement depends on S^e and U only through their difference. We state this observation formally in the lemma below, which is a special case of Lemma 4.2 of Sheopuri et al. (2010):

Lemma 4.2. *The distributions of O, q^r, and q^e depend on U and S^e only through their difference $V = U - S^e$.*

As our discussion preceding Lemma 4.2 already alludes to, the difference V has an explicit meaning that allows for an alternative but equivalent specification and operation of the PEIP policy. Rather than keeping the projected expedited inventory position at a fixed level, this alternative specification places regular orders such that the expected overshoot is kept at a fixed level (while retaining an order-up-to level S^e for the expedited supplier). To state this equivalence relation formally, observe that $\mathbb{E}[O_{t+l}|\mathbf{x}_t^r]$ denotes the expected overshoot l periods into the future at time t given the system state at time t following ordering, $\mathbf{x}_t^r$. In any period t after ordering, we have $\mathbb{E}[IP_{t+l}^e|\mathbf{x}_t^r] = S^e + \mathbb{E}[O_{t+l}|\mathbf{x}_t^r]$. Hence, placing regular orders such that $\mathbb{E}[IP_{t+l}^e|\mathbf{x}_t^r]$ equals the projected expedited inventory position level U is equivalent to placing regular orders such that $\mathbb{E}[O_{t+l}|\mathbf{x}_t^r]$ equals what we henceforth refer to as the projected overshoot level $V = U - S^e$.

While the original specification of the PEIP policy has intuitive appeal, the alternative specification outlined in the previous paragraph allows us to decompose the optimization of the policy parameters into two simple uni-variate optimization problems. Indeed, letting O^V denote the steady state overshoot random variable for a given V, the optimal order-up-to level S^{e*} for a given V has a simple form:

Lemma 4.3. *The optimal order-up-to level S^{e*} for a given V equals*

$$S^{e*} = \inf\left\{S^e \in \mathbb{R} : \mathbb{P}\left(\mathbf{D}^{l^e+1} - O^V \leq S^e\right) \geq \frac{p}{p+h}\right\}.$$

The result above implies that a simple search procedure over the projected overshoot level V suffices to find the globally optimal PEIP policy.

Lemma 4.3 is a special case of Lemma 4.3 of Sheopuri et al. (2010). Sheopuri et al. (2010) show that this result holds in general for all policies that have an order-up-to rule for the expedited supplier and operate some given regular ordering rule that uses only the information contained in $\mathbf{x}_t^r$. For all such policies, the optimal order-up-to level for the expedited supplier can be interpreted as a standard Newsvendor model with one additional feature. The ordering rule for the regular supplier essentially dictates a return stream in the form of the overshoot which must therefore be subtracted from the lead time demand to arrive at the net lead time demand observed by the said Newsvendor. The PEIP policy has particular intuitive appeal in the context of this Newsvendor with returns interpretation; it keeps the expected amount to be returned l periods into the future at a fixed level.

4.5. Asymptotic optimality

In this section, we prove that the best PEIP policy is asymptotically optimal when both the shortage cost and the cost premium for expedited units become large, with their ratio held constant. We first elaborate on why it is important and non-trivial to study this asymptotic regime. We then present the asymptotic optimality result itself, and we subsequently show that this result is robust in the sense that it holds true for a broad class of policies.

4.5.1 Regime of interest

We know from Xin and Goldberg (2018) that a simple TBS policy is asymptotically optimal as the lead time difference between the regular supplier and the expedited supplier approaches infinity. Unfortunately, the policy performs poorly outside this asymptotic regime, particularly when either the cost for a unit in backlog or the price charged by the expedited supplier is large. When these costs grow large simultaneously, the performance of the TBS drops precipitously (Klosterhalfen et al., 2011). It is however exactly this latter asymptotic regime that is prevalent in practice, and thus establishing asymptotic performance guarantees of the PEIP policy in this regime has important practical implications. In the capital goods industry, for instance, backordering a critical part comes at a high cost while expediting shipments from an emergency supplier are expensive too (see, e.g., Westerweel et al., 2021).

In this section, we thus focus on the asymptotic regime where the unit expediting cost c^e and the shortage cost p grow large simultaneously. We do, however, need the assumption that their ratio is held fixed at a constant strictly smaller than l (i.e. $c^e < p \cdot l$). Recall from Section 4.3 that if we would not impose this assumption, then ordering from the expedited supplier over the regular supplier would never be beneficial. That is, the dual-sourcing inventory problem would reduce to a trivial single supplier inventory problem if $c^e \geq p \cdot l$.

We remark that, following a reasoning similar to the above, the asymptotic regime where both l and c^e grow large with their ratio held fixed at a constant strictly smaller than p is, at least theoretically, interesting as well. Indeed, this too is a non-trivial regime in which the optimal policy does not reduce to a trivial single supplier inventory problem. Xin (2021a) has recently made advances on this latter regime by showing that for the continuous time variant of the problem, the CDI policy has a worst-case performance guarantee.

4.5.2 Main results

We first introduce extra notation. Let $C^{PEIP^*}(p,h,c^e,l^r,l^e)$ and $C^{\pi^*}(p,h,c^e,l^r,l^e)$ denote the long run average cost per period of the dual-sourcing inventory system under the best PEIP policy and the optimal policy, respectively, as functions of the problem specific input parameters. We may occasionally omit certain input parameters when there is no ambiguity. We define the functions $p(n) = p \cdot n$ and $c^e(n) = c^e \cdot n$, and we are, hence, interested in the asymptotic regime where $n \to \infty$ with $\frac{c^e(n)}{p(n)}$ fixed at a constant strictly smaller than l. For a given $n > 0$, we further define $C^{PEIP^*}(n) = C^{PEIP^*}(p(n), c^e(n))$ and $C^{\pi^*}(n) = C^{\pi^*}(p(n), c^e(n))$ to denote the cost as we scale p and c^e proportionally with n.

Let $\mathcal{L}(p,h,l)$ and $\mathcal{B}(p,h,l)$ denote the canonical single supplier inventory systems with lost sales and backlogged demand, respectively, as defined in Janakiraman et al. (2007), Huh et al. (2009), and Bijvank et al. (2014). Following their notation, let $C^{\mathcal{L}^*}(p,h,l)$ denote the long run average cost per period under an optimal policy in $\mathcal{L}(p,h,l)$ as a function of the lost sales penalty cost p, the holding cost h, and the lead time l. Similarly, let $C^{\mathcal{B}^*}(p,h,l)$ denote the long run average cost per period under the optimal policy in $\mathcal{B}(p,h,l)$, with p now denoting the penalty cost for backlogged demand. It is well-known that the optimal policy for $\mathcal{B}(p,h,l)$ is a simple order-up-to rule. It is equally well-known that the optimal policy for $\mathcal{L}(p,h,l)$ has a complex structure and depends in general on the entire vector of outstanding orders, very much akin to the present dual-sourcing inventory system.

We first establish a lower bound on the long run average cost per period of the optimal policy, i.e. $C^{\pi^*}(p,c^e,l^r,l^e)$, in terms of both the long run average cost per period of the optimal policy for $\mathcal{B}(p,h,l)$, i.e. $C^{\mathcal{B}^*}(p,h,l)$, and the optimal policy for $\mathcal{L}(p,h,l)$, i.e. $C^{\mathcal{L}^*}(p,h,l)$.

Lemma 4.4. $C^{\mathcal{B}^*}(c^e/(l^r+1),h,l^r) \leq C^{\mathcal{L}^*}(c^e,l^r) \leq C^{\pi^*}(p,h,c^e,l^r,l^e)$

Proof. Observe that a lower bound on the long run average cost per period of the optimal policy is given by the long run average cost per period cost of an optimal policy for a system identical to the present inventory system but in which any backlogged demand at the end of a period is satisfied instantly by the expedited supplier at unit price c^e, where we assume without loss of generality that $c^e < p$. According to Theorem 3.1 of Sheopuri et al. (2010), this system is exactly $\mathcal{L}$, where the lost sales penalty cost equals c^e, i.e. $C^{\mathcal{L}^*}(c^e,l^r)$. Hence, we have $C^{\pi^*}(p,c^e,l^r,l^e) \geq C^{\mathcal{L}^*}(c^e,l^r)$. We further know from Theorem 5 of Janakiraman et al. (2007) that $C^{\mathcal{L}^*}(c^e,l^r) \geq C^{\mathcal{B}^*}(c^e/(l^r+1),l^r)$, and thus $C^{\pi^*}(p,c^e,l^r,l^e) \geq C^{\mathcal{B}^*}(c^e/(l^r+1),l^r)$. □

We note in passing that easily computable lower bounds for the optimal cost of dual-

sourcing systems are relatively scarce. In fact, we are aware of only one such lower bound in terms of the optimal policy for an equivalent dual-sourcing system where the lead time difference is 1 (Sheopuri et al., 2010). For this equivalent system, the optimal policy is myopic in that both suppliers use an order-up-to level. In this respect, the lower bound in terms of $\mathcal{B}(c^e/(l^r+1), l^r)$ in Lemma 4.4 is of independent interest as it is easy to compute and leads to an upper bound on the optimality gap for any given policy.

Our main result requires the following technical assumption regarding the cumulative demand over $l^r + 1$ periods:

Assumption 4.1. The random variable $\mathbf{D}^{l_r+1}$ has a positive and finite mean, i.e. $0 < \mathbb{E}[\mathbf{D}^{l_r+1}] < \infty$, and satisfies one of the following conditions:

i. $\mathbf{D}^{l_r+1}$ is bounded, or

ii. $\mathbf{D}^{l_r+1}$ is unbounded and $\lim_{d\to\infty} \mathbb{E}[\mathbf{D}^{l_r+1} - d|\mathbf{D}^{l_r+1} > d]/d = 0$.

This assumption first appeared in Huh et al. (2009). They show that many commonly used distributions in inventory models satisfy this assumption, including Poisson distributions, geometric distributions, negative binomial distributions with parameters $r > 0$ and $0 < p < 1$, and (log-)normal distributions. A sufficient condition for this assumption to be satisfied is that the demand distribution has an increasing failure rate.

We are now in the position to state the main result of this section.

Theorem 4.1. *Let the backorder penalty cost p and the expedited cost premium c^e be such that $c^e/p < l$. Then, under Assumption 4.1, the best PEIP policy is asymptotically optimal as c^e and p approach infinity simultaneously, that is*

$$\lim_{n\to\infty} \frac{C^{PEIP^*}(n)}{C^{\pi^*}(n)} = 1.$$

Proof. Observe that single sourcing from the regular supplier is a feasible PEIP policy by setting S^e equal to ∞. This observation implies that the numerator is bounded from above by the long run average cost per period of an optimal policy in a single supplier inventory system with backlogged demand whose lead time is equal to the lead time of the regular supplier, i.e. $C^{\mathcal{B}^*}(p, h, l^r) \geq C^{PEIP^*}(p, h, c^e, l^r, l^e)$. Combining this observation with Lemma 4.4, we have

$$\frac{C^{PEIP^*}(n)}{C^{\pi^*}(n)} \leq \frac{C^{\mathcal{B}^*}(p(n), h, l^r)}{C^{\mathcal{B}^*}(c^e(n)/(l^r+1), h, l^r)}.$$

By Theorem 2b of Huh et al. (2009), we know that under Assumption 4.1

$$\lim_{n\to\infty} \frac{C^{\mathcal{B}^*}(p(n), h, l^r)}{C^{\mathcal{B}^*}(c^e(n)/(l^r+1), h, l^r)} = 1,$$

which gives the desired result. □

Define the subset $\tilde{\Pi} = \{\pi \in \Pi | C^{\pi^*}(p, h, c^e, l^r, l^e) \leq C^{\mathcal{B}^*}(p, h, l^r)\}$. Note that it readily follows from the definition of $\tilde{\Pi}$ and the proof of Theorem 4.1, that any policy $\pi \in \tilde{\Pi}$ is asymptotically optimal in the regime under our consideration. This is an important observation that generalizes our result to a wide class of policies. We therefore state it as a theorem.

Theorem 4.2. *Let the backorder penalty cost p and the expedited cost premium c^e be such that $c^e/p < l$. Then, under Assumption 4.1, any policy $\tilde{\pi} \in \tilde{\Pi}$ is asymptotically optimal as c^e and p approach infinity simultaneously, that is*

$$\lim_{n\to\infty} \frac{C^{\tilde{\pi}^*}(n)}{C^{\pi^*}(n)} = 1,$$

where $C^{\tilde{\pi}^*}(n)$ is defined in a way similar as $C^{\pi^*}(n)$.

We conclude this section by noting that the CDI, DI, and SI policy are all contained in the set $\tilde{\Pi}$. They are thus asymptotically optimal when both c^e and p grow large with their ratio held fixed at a constant strictly smaller than l.

4.6. Numerical investigation

The theoretical results of the previous section are important as they provide performance guarantees for the PEIP policy in high service regimes where expedited orders are expensive. While this is particularly useful for practitioners operating in those regimes, it does not imply that the PEIP policy consistently performs well, especially outside those asymptotic regimes. In this section, we therefore undertake an extensive numerical investigation to examine the overall effectiveness of the PEIP policy, thereby complementing the theory of the previous section. This investigation mainly revolves around the comparison of the performance of the PEIP policy with the performance of several existing heuristic polices.

We first describe the benchmark heuristic policies in Section 4.6.1. We then continue with the numerical investigation itself, which consists of two parts. The first part, Section 4.6.2, compares the performance of the PEIP policy with the performance of

the benchmark heuristic policies based on a full factorial test bed of industrial size. The second part, Section 4.6.3, analyzes the performance of the PEIP policy as well as the benchmark heuristic policies as the lead time difference between both suppliers grows large. It is exactly this regime where a simple TBS policy is asymptotically optimal (Xin and Goldberg, 2018).

4.6.1 Benchmark heuristic policies and computational aspects

The following five heuristic policies will serve as benchmarks: (i) The DI policy, (ii) the CDI policy, (iii) the VBS policy, (iv) the BWB policy, and (v) the TBS policy. Although we have discussed these five heuristic policies in our literature review, we revisit a few points here that are important to bear in mind while considering the results of the numerical investigation.

First, the CDI policy and the BWB policy are the current best performing heuristic policies. Their computational requirements are however significantly higher than the PEIP policy. Indeed, while the PEIP policy requires only a one-dimensional search to find the optimal policy parameters, both the CDI policy and the BWB policy require a search over two dimensions to arrive at their best policy parameters. By contrast, the DI policy, VBS policy, and the TBS policy have the same computational requirements as the PEIP policy. Second, the CDI policy is a generalization of the DI policy, and it is therefore to be expected that the former policy outperforms the latter. The same holds for the BWB policy and the VBS policy – the latter being a special case of the former. Third, all five benchmark heuristic policies use an order-up-to rule for the expedited supplier. As such, echoing our discussion after Lemma 4.3, the optimal order-up-to level under each policy for given regular ordering rule parameters follows a Newsvendor type equation. This is convenient numerically as it only requires us to simulate the steady state overshoot distribution for given regular ordering rule parameters to arrive at the corresponding optimal order-up-to level. We describe this simulation-based optimization scheme in detail next.

For all problem instances in this numerical investigation, we find the best policy parameters for the PEIP policy as well as for the benchmark heuristic policies via golden-section searches over their respective regular ordering rule parameters. In each iteration of this search procedure, we first compute the steady state overshoot distribution via simulation, we then determine the optimal order-up-to level via the Newsvendor equation, and we finally compute the corresponding long run average cost per period. We use common random numbers across all policies per problem instance, and we continue simulating the steady state overshoot distribution until the width of the 95 percent confidence intervals for the mean and for the standard deviation of the overshoot are less than 1 percent of their respective point estimates.

After we find the best control policy parameters for all heuristic policies for a given problem instance, we compute the long run average cost per period for each heuristic policy under its best policy parameters via simulation. As in the simulation-based optimization scheme, this simulation uses common random numbers across policies, and it continues until the width of the 95 percent confidence intervals for the long run average cost per period is less than 1 percent of its point estimate.

4.6.2 Full factorial experiment

We start our numerical investigation by comparing the performance of the PEIP policy with the performance of the heuristic policies described in the previous section. We do so based on a large test bed of industrial size. Note that the optimal policy cannot be computed for such large instances due to the complexity of the corresponding dynamic program. We therefore exclude the optimal policy in this investigation.

In all instances of the test bed, the one period demand has a negative binomial distribution with mean 50 and the holding cost h is fixed at 1. We further keep the price of the regular supplier c^r fixed at 0 since only the cost difference between the suppliers affects the relative cost rate. The expedited lead time l^e is kept fixed at 0 periods so that the focus is on the lead time difference $l = l^r - l^e = l^r$. The remaining input parameters are varied over multiple levels. We vary the backorder penalty cost over five levels ($p \in \{4, 9, 19, 49, 99\}$), the coefficient of variation of the one period demand over five levels ($\sigma_{\mathrm{D}}/\mathbb{E}[\mathrm{D}] \in \{0.15, 0.25, 0.5, 1, 1.5, 2\}$), and the regular lead time over four levels ($l^r \in \{2, 3, 4, 5\}$). We set the price of the expedited supplier c^e equal to a fraction of the backorder penalty cost multiplied by the lead time difference of the specific instance, that is, $c^e = \delta pl$ for $\delta \in \{0.1, 0.2, 0.4\}$. While this approach is different from most numerical studies on dual-sourcing inventory systems, it does ensure that there is no instance in the test bed for which exclusively sourcing from the regular supplier is optimal. Table 4.1 summarizes all levels for each input parameter in our test bed that is not kept fixed. Permuting over all these levels leads to a large test bed of 360 instances.

Table 4.1 Input parameter values for our test bed.

	Input parameter	No. of choices	Values
1	Backorder penalty cost, p	5	4, 9, 19, 49, 99
2	Regular lead time, l^r	4	2, 3, 4, 5
3	Coefficient of variation of demand, $\sigma_{\mathrm{D}}/\mathbb{E}[\mathrm{D}]$	6	0.15, 0.25, 0.5, 1, 1.5, 2
4	Price of the expedited supplier, c^e	3	δpl for δ = 0.1, 0.2, 0.4

We are interested in the performance of the PEIP policy in comparison with the

performance of the benchmark heuristics. To that end, we compute for each instance of the test bed the percentage gap in the long run average cost per period of each benchmark heuristic policy $\pi \in \{\text{DI,CDI,VBS,BWB,TBS}\}$ relative to the PEIP policy: $100 \cdot \frac{C^{\pi^\star} - C^{\text{PEIP}^\star}}{C^{\text{PEIP}^\star}}$.

Table 4.2 summarizes the numerical results of the full factorial experiment. The table presents the average, maximum, and minimum gap percentages for each benchmark heuristic policy computed over all instances as well as the percentage of test instances in which the PEIP policy outperforms the specific benchmark heuristic policy. (A positive gap percentage implies that the PEIP policy outperforms the specific benchmark heuristic policy.) Detailed results are relegated to Appendix 4.B. Based on this table, we conclude that the PEIP policy outperforms the TBS, DI, CDI, and VBS policy. The performance of the PEIP policy seems to be comparable with the BWB policy, though on average the PEIP policy outperforms this policy as well. It is also interesting to note that the PEIP policy can outperform the BWB and CDI policy by up to 2.67% and 4.67%, respectively. The opposite never holds. Indeed, when the PEIP policy is outperformed by the BWB or CDI policy, it only occurs by much smaller margins of up to 0.29% and 0.08%, respectively. Note that these gaps are smaller than the confidence interval around the estimated performance for any single instance. This is an interesting observation, especially in view of the fact that the BWB and CDI policy are the current best performing heuristic policies, and have considerably higher computational requirements than the PEIP policy. Furthermore, these policies have more parameters and the BWB policy is arguably difficult to communicate to practitioners.

Table 4.2 Summary of full factorial experiment.

Heuristic	Avg	Max	Min	%
VBS	0.15	3.15	-0.05	90
BWB	0.03	2.67	-0.29	42
DI	0.71	4.97	0.00	100
CDI	0.19	4.67	-0.08	94
TBS	55.11	373.88	1.68	100

4.6.3 Long lead time difference regime

We will now examine the performance of the PEIP policy as well as the benchmark heuristic policies for increasingly larger lead time differences l. To that end, we will study a test instance in which we only vary l. The cost parameters in this test instance are fixed as follows: $h = 1$, $p = 19$, and $c^e = 5$. As in the previous section, the one period demand has a negative binomial distribution with mean 50, with $\sigma_{\text{D}}/\mathbb{E}[\text{D}]$ now fixed at 0.25.

Since the TBS policy is asymptotically optimal as the lead time difference l grows large (Xin and Goldberg, 2018), we will now use the relative difference between the long run average cost per period of each heuristic policy with the long run average cost per period of the TBS policy as the main performance measure. To that end, we compute the percentage improvement/reduction in the long run average cost per period of each heuristic policy $\pi \in \{\text{PEIP,BWB,VBS,CDI,DI}\}$ relative to the TBS policy:

$$\%\Delta_\pi = 100 \cdot \frac{C^{\text{TBS}^\star} - C^{\pi^\star}}{C^{\text{TBS}^\star}}.$$

Note that the long run average cost of a TBS policy is insensitive to l (see, e.g., Klosterhalfen et al., 2011). The long run average cost per period of the best TBS policy, i.e. $C^{\text{TBS}^\star}$, thus remains constant as we vary l.

Table 4.3 provides the gap percentages for increasingly larger values of l, along with the variability of the orders placed with the regular supplier computed through simulation. Based on this table, we can draw three main conclusions. First, the PEIP policy has the best relative improvement over the TBS policy for each value of l. Second, even for values of l that would be considered large in practice, the PEIP, BWB, CDI, and VBS policy outperform the TBS policy. Third, the variability of the orders placed with the regular supplier under the PEIP policy decreases in l. This suggests that the PEIP policy will increasingly mimic the TBS policy as the lead time difference grows. (Note that the TBS policy has zero variability in its regular orders.) The same holds for the BWB, CDI, and VBS policy. Whether these four heuristic policies are asymptotically optimal when l grows large is an open question.

Table 4.3 Long lead time regime.

	PEIP		BWB		VBS		CDI		DI	
l	$\%\Delta_\pi$	$\mathbb{V}[\mathrm{Q}^r]$	$\%\Delta_\pi$	$\mathbb{V}[\mathrm{Q}^r]$	$\%\Delta_\pi$	$\mathbb{V}[\mathrm{Q}^r]$	$\%\Delta_\pi$	$\mathbb{V}[\mathrm{Q}^r]$	$\%\Delta_\pi$	$\mathbb{V}[\mathrm{Q}^r]$
2	21.81	104.02	21.80	100.03	21.75	113.03	21.76	110.57	21.53	132.25
4	12.03	59.90	11.98	48.97	11.78	73.87	11.95	46.98	10.30	127.67
6	7.50	33.64	7.37	34.61	6.86	51.23	7.32	24.62	4.06	125.51
8	5.18	21.27	4.97	24.97	4.23	36.49	5.09	15.72	0.17	125.45
10	3.81	14.42	3.52	18.24	3.06	19.86	3.66	9.03	-2.53	124.25
12	3.05	9.23	2.72	14.85	2.55	11.38	2.84	8.81	-4.60	124.05
14	2.43	6.97	2.04	12.08	2.00	9.89	2.09	2.55	-6.26	122.81
16	2.06	5.29	1.78	10.47	1.78	3.63	1.83	2.74	-7.63	123.37
18	1.68	3.84	1.61	4.26	1.56	2.95	1.57	2.51	-8.57	123.48
20	1.45	3.01	1.36	2.79	1.34	2.50	0.09	0.02	-9.25	122.61

4.7. Concluding remarks

In this chapter, we have proposed the PEIP policy for dual-sourcing inventory systems. This policy uses an order-up-to rule for the expedited supplier, and places regular orders such that the expected expedited inventory position at the time of the regular order entering the expedited inventory position is raised to a target level. In doing so, the PEIP policy leverages all the information contained in the pipeline – as does the optimal policy. This is in contrast with existing heuristic policies that either aggregate all outstanding orders or ignore them entirely and place a constant order every period regardless. We have further established a separability result that reduces the optimization of the PEIP policy to two simple one-dimensional optimization problems. An extensive numerical investigation has indicated that the PEIP policy outperforms the current best performing heuristic policies, even the ones that employ more control parameters than the PEIP policy. The PEIP policy is thus vastly more advanced than existing heuristic policies, and yet retains an intuitive appeal for practitioners, has superior performance, and can be optimized efficiently.

We have shown that the PEIP policy is asymptotically optimal when the unit shortage cost and the unit price charged by the expedited supplier approach infinity simultaneously with their ratio held fixed at a constant. This result is useful for practice as it implies that the PEIP policy has asymptotic performance guarantees in high service regimes where utilizing an expedited source is expensive. We have further illustrated that this asymptotic optimality result is robust in the sense that it holds true for a wide class of dual-sourcing heuristic policies, including the CDI, DI, and SI policy. Our proof technique in itself has merit too since it can serve as a template that is readily applicable to future heuristic policies. As such, we have extended the literature on asymptotic optimality results for dual-sourcing inventory systems.

Our numerical investigation also poses an interesting question: Is the PEIP policy asymptotically optimal as the lead time difference between both suppliers grows large? Our numerical results provide strong evidence that this may be the case. An answer to this question for the PEIP and other policies is left for future research.

4.A. Proof of Proposition 4.1

Proof. Following our discussion after Lemma 4.2, it suffices to prove that in any period $t \in \mathbb{N}$, we can place a non-negative regular order so that $\mathbb{E}[O_{t+l}|\mathbf{x}_t^r]$ equals the projected overshoot level V. For ease of exposition, we will now denote the regular order placed in period t with q_{t+l}^r such that the order entering the expedited inventory position in period t is denoted q_t^r.

Due to our assumption on the initial state, i.e. $\mathbf{x}_0 = \mathbf{0}$, we have that $\exists q_0^r \geq 0$ such that $\mathbb{E}[O_l|\mathbf{x}_0^r] = \mathbb{E}[(O_{l-1} + q_0^r - D_{l-1})^+|\mathbf{x}_0^r] = V$. We define $Y_t = (O_t - D_t)^+$. It suffices now to show that $\mathbb{E}[Y_l|\mathbf{x}_0^r, D_0 = 0] \leq \mathbb{E}[O_l|\mathbf{x}_0^r] = V$ because $\mathbb{E}[Y_l|\mathbf{x}_0^r, D_t = 0]$ is the minimum projected overshoot level as seen in period 1 before a regular order is placed. Observe that

$$Y_l|(\mathbf{x}_0^r, D_0 = 0) = (((\ldots((O_0 + q_{-l+1}^r)^+ + q_{-l+2}^r - D_1)^+ \ldots)^+ + q_0^r - D_{l-1})^+ - D_l)^+,$$

and

$$O_l|\mathbf{x}_0^r = ((\ldots((O_0 + q_{-l+1}^r - D_0)^+ + q_{-l+2}^r - D_1)^+ \ldots)^+ + q_0^r - D_{l-1})^+.$$

It is easy to verify through tedious algebra and the fact that $(x)^+$ is monotone increasing that $Y_l|(\mathbf{x}_0^r, D_0 = 0) \leq O_l|\mathbf{x}_0^r$ state-wise and thus $\mathbb{E}[Y_l|\mathbf{x}_0^r, D_0 = 0] \leq \mathbb{E}[O_l|\mathbf{x}_0^r] = V$. □

4.B. Detailed results numerical investigation

Table 4.4 presents the percentage gap in the long run average cost per period of each benchmark heuristic policy relative to the PEIP policy. Table 4.5 presents the percentage of cases in which the PEIP policy outperforms each benchmark heuristic policy. In both tables, we first distinguish between subsets of instances with the same value for a specific input parameter of Table 4.1 and then present the results for all instances.

Table 4.4 Detailed results full factorial experiment.

		VBS			BWB			DI			CDI			TBS		
Input	Value	Avg	Max	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max	Min
p	4	0.32	0.65	0.08	0.09	0.45	-0.11	1.76	4.97	0.23	0.36	4.19	0.06	10.73	30.99	1.68
	9	0.20	0.62	0.02	0.01	0.27	-0.21	1.01	3.27	0.11	0.28	4.32	0.02	21.74	58.51	4.48
	19	0.08	0.33	-0.05	-0.04	0.66	-0.29	0.50	1.73	0.04	0.15	3.10	0.01	38.35	90.34	9.34
	49	0.12	3.15	-0.03	-0.03	0.46	-0.18	0.17	0.61	0.00	0.04	0.43	-0.07	76.82	199.62	23.71
	99	0.02	0.09	-0.02	0.12	2.67	-0.13	0.11	2.81	0.00	0.12	4.67	-0.08	127.90	373.88	32.34
l^r	2	0.07	0.29	0.00	0.08	2.67	-0.05	0.27	1.36	0.00	0.05	0.29	-0.01	48.75	235.21	1.75
	3	0.15	2.64	-0.01	0.02	0.73	-0.10	0.58	2.80	0.01	0.17	4.32	-0.02	54.35	286.08	1.68
	4	0.15	0.54	-0.01	0.02	0.85	-0.18	0.85	3.86	0.00	0.26	4.67	-0.07	57.45	334.64	1.76
	5	0.21	3.15	-0.05	0.00	1.31	-0.29	1.14	4.97	0.01	0.27	4.19	-0.08	59.87	373.88	1.70
$\frac{\sigma_D}{\mathbb{E}[D]}$	0.15	0.24	3.15	-0.02	0.09	2.67	-0.14	0.81	4.97	0.00	0.28	4.67	-0.08	76.50	373.88	4.29
	0.25	0.15	0.61	0.00	0.09	2.47	-0.18	0.81	4.87	0.01	0.13	1.46	-0.02	63.18	240.11	3.85
	0.5	0.12	0.62	0.00	-0.03	0.85	-0.20	0.79	4.39	0.01	0.24	4.32	-0.01	60.73	292.31	3.38
	1	0.10	0.60	-0.01	-0.03	0.44	-0.18	0.72	3.37	0.02	0.13	0.79	-0.01	49.06	228.77	2.68
	1.5	0.11	0.44	-0.02	0.00	0.20	-0.21	0.61	2.62	0.01	0.15	0.52	0.00	43.74	214.23	2.11
	2	0.15	0.55	-0.05	0.05	0.45	-0.29	0.54	2.81	0.01	0.19	2.79	0.00	37.43	177.20	1.68
δ	0.1	0.20	0.65	-0.02	0.02	0.45	-0.18	1.15	4.97	0.03	0.27	4.67	-0.08	24.68	88.41	1.70
	0.2	0.15	0.62	-0.03	0.02	2.67	-0.20	0.78	3.39	0.02	0.21	4.19	-0.01	48.60	182.03	4.48
	0.4	0.11	3.15	-0.05	0.06	2.47	-0.29	0.43	2.81	0.00	0.15	2.79	-0.01	89.73	373.88	9.82
Total		0.15	3.15	-0.05	0.03	2.67	-0.29	0.71	4.97	0.00	0.19	4.67	-0.08	55.11	373.88	1.68

Table 4.5 Detailed results full factorial experiment.

Input	Value	VBS	BWB	DI	CDI	TBS
p	4	100	78	100	100	100
	9	100	47	100	100	100
	19	96	22	100	100	100
	49	83	26	100	92	100
	99	72	35	100	79	100
l^r	2	94	51	100	90	100
	3	92	42	100	96	100
	4	90	38	100	97	100
	5	84	36	100	94	100
$\frac{\sigma_D}{\mathbb{E}[D]}$	0.15	97	52	100	92	100
	0.25	100	42	100	92	100
	0.5	98	30	100	95	100
	1	83	25	100	92	100
	1.5	78	43	100	98	100
	2	85	58	100	97	100
δ	0.1	93	51	100	90	100
	0.2	89	37	100	94	100
	0.4	89	44	100	97	100
Total		90	42	100	94	100

Chapter 5

Real-time integrated learning and decision making for stochastically deteriorating systems

5.1. Introduction

Advanced technical systems are critical for the smooth operation of public services such as public transport (e.g., aircraft, rolling stocks), utilities (e.g., power plants), and health care (e.g., MRI-scanners, X-ray machines) as well as for the primary processes of companies (e.g., data centers, lithography machines). Unavailability and failure of these systems – especially when unplanned – have severe consequences and can even lead to immediate safety hazard. For example, the failure of a deteriorating propeller blade led to the crash of a KC-130T aircraft with 16 casualties in the summer of 2017 (Insinna and Ziezuliwicz, 2018). A year before that, the nuclear power plant in Doel, Belgium, had to be shut down unexpectedly in December of 2015 and then again in January of 2016. This caused concerns for citizens in Belgium, the Netherlands and Germany, as well as expensive power imports by the operator Electrabel (Van Soest, 2016).

When unplanned downtime does not lead to immediate safety hazard, the consequences can still be severe financially. For example, the average cost of an unplanned data server outage was around 0.5 million in 2010 and was estimated at 0.74 million

in 2016. More than 60% of such outage costs are not related to repairing the failure, but to costs of (unplanned) downtime (Ponemon, 2016). Another striking example comes from the semiconductor industry that heavily relies on lithography equipment. Unplanned downtime of such equipment costs in the order of tens of thousands of dollars per hour (Kranenburg, 2006). Recent studies state that "unplanned downtime costs industrial manufacturers an estimated 50 billion annually and equipment failure is the cause of 42% of this unplanned downtime" (Studios, 2017; Coleman et al., 2017). The implication of these studies is twofold. First, unplanned downtime costs is a general problem faced by many industrial companies, stretching far beyond the two examples put forward above. Second, and arguably more importantly, unplanned downtime costs can be reduced significantly through minimizing equipment failures. This latter implication also seems to be well-understood by executives in asset management. In a recent survey, they perceive unplanned failures as the most important risk to manage (Pacquin, 2014).

Asset managers are responsible to minimize the risk of failure and unplanned downtime. This can be achieved by performing preventive replacement or maintenance regularly. However, too early performance of preventive maintenance leads to high capital expenditures as the useful lifetime of capital assets is cut short. The challenge faced by asset managers is therefore to optimize two conflicting objectives:

1. Minimize the risk of failure and unplanned downtime with all its adverse consequences.

2. Maximize the useful life of an expensive critical asset and of all its components.

The traditional approach asset managers would take to address the challenge posed above comprises two sequential steps. The first step is the estimation of a statistical model of the time to failure of an asset based on either condition or failure data. The second step is to optimize the decision of when to perform maintenance based on the statistical model, thereby trading off costly premature interventions with costly tardy replacements. The separation of estimation and optimization was necessitated by the fact that degradation data of a critical asset could only be obtained by performing expensive measurements during (planned) asset downtime. This separation of estimation from optimization and scarcity of data also necessitated the assumption that one asset is statistically identical to any other asset of the same type. The implicit assumption was therefore always that the population of assets is homogeneous. That is, the degradation paths or lifetimes of different assets are assumed to be statistically indistinguishable. In Section 5.7, we report on a case study with our industrial partner, a major medical equipment manufacturer and service provider. This case focuses on a component that is critical for the operation of an interventional X-ray system. The case study shows that the degradation paths of

components of different interventional X-ray systems differ from one another, that is, these degradation paths are not statistically indistinguishable. This implies that the traditional approach that hinges on the assumption of a homogeneous asset population is no longer appropriate in practice, increases failure risk, and reduces the useful lifetime of assets.

Recently built systems, including the interventional X-ray systems of our case study, often have integrated sensor technology that allows degradation data of a critical component to be gathered at hardly any additional cost (Jardine et al., 2006). These sensors measure the degradation of components in real-time and are increasingly integrated into the so-called Internet-of-Things (IoT) (Coopers, 2014). This means that each system generates a real-time stream of degradation data that offers new opportunities for both estimation and optimization of decision making as well as their integration. Asset managers can monitor and interpret the real-time data generated by each asset to assess degradation and intervene with maintenance or replacement when appropriate. The current best-in-class asset management companies reported that they already use real-time data to optimize their decision making (Pacquin, 2014).

Data that is generated by sensors and relayed in real-time through the IoT further allows us to relax the assumption that assets are statistically identical. Indeed, real-time degradation data allows us to learn degradation behavior on the individual component level. This assumption has been partially relaxed before in settings where measurements are only possible at planned down-times but not in real-time; see Elwany et al. (2011), Kim and Makis (2013), Chen et al. (2015), and Van Oosterom et al. (2017). When the condition of a system can only be measured at planned downtimes, the amount of information that can be learned from the degradation level is limited compared to the situation with real-time data. Accordingly, attention in previous literature is restricted to population heterogeneity within a finite sets of possibilities or heterogeneity in degradation drift only, with drift defined as the expected degradation increment per unit of time. By contrast, this chapter uses the *entire* degradation path of each component. This allows us to infer higher-order properties of the degradation behavior of the individual component, in particular the volatility, defined as the variance of a degradation increment per unit time. Hence, this chapter integrates estimation and decision making to tailor maintenance intervention decisions to assets individually.

Degradation of a component (or asset) is generally the result of random amounts of damage (e.g., wear, fatigue) through shocks that occur randomly over time. For instance, certain metal and ceramic components in trains and aircraft only degrade at events at which they are subjected to shocks (e.g., propagation of cracks in a propeller due to a heavy wind gust) rather than in a continuous fashion. Another

example is that of sea waves acting on offshore structures during extreme weather conditions (one can consider the offshore structure as the component), which too can be regarded as a sequence of random shocks. For such degradation processes, it is natural to model them as so-called jump processes, i.e., stochastic processes that have discrete movements at random times (cf. Sobczyk, 1987; Singpurwalla, 1995).

In this chapter, we consider a very general jump process. This jump process is often used to model the degradation of the type discussed in the previous paragraph (cf. Van Noortwijk, 2009). Specifically, we assume that the sequence of random shocks arrives as a Poisson process with a randomly varying shock size, that is, degradation is modeled as a compound Poisson process. Thus the inter-arrival times between two consecutive shocks are exponentially distributed, and every shock causes a random amount of damage to the component. The compounding distribution is quite general; the only restriction we impose is that this distribution belongs to the natural exponential family with non-negative support. This family includes many well-known distributions such as the geometric distribution and the Poisson distribution.

Components are subject to compound Poisson degradation, but the parameters of the Poisson process as well as the compounding distribution vary from one component to the next. That is, the population of components is heterogeneous. These parameters cannot be observed directly and they therefore need to be learned by observing the degradation signal that is relayed in real-time through sensors on the component. Although we observe the degradation level of a component continuously through condition monitoring, we can only interfere with the system at equally spaced discrete decision epochs. The costs to replace a component after failure is much higher than before because they include the costs of unplanned downtime. This decision problem can be modeled by a partially observable Markov decision process (POMDP). The entire past degradation path of a component is relevant state information in this setting. Dealing with the entire degradation path can lead to tractability issues. We circumvent these issues by using conjugate prior pairs to model the heterogeneity of the component population. We further collapse the state space by identifying structure in the prior to posterior updating procedure. This enables us to tractably compute optimal policies as well as prove structural results about optimal policies.

This chapter makes the following contributions:

1. We tractably model the situation where components are heterogeneous in their degradation process by using conjugate prior pairs. We collapse the high dimensional state space to only 3 dimensions while retaining all relevant information. This collapse gives insight into how all relevant information in a real-time degradation signal can be parsimoniously represented. Furthermore, this collapse makes the model both tractable numerically and amenable to

structural analysis.

2. We characterize the optimal replacement policy as a threshold replacement policy where the threshold is increasing in the age of a component and furthermore depends on the volatility of the observed degradation signal.

3. In our first simulation study, we study (i) the benefits of explicitly modeling heterogeneity, (ii) the benefits of integrating learning with decision making, and (iii) the impact of the amount of available historical degradation data for estimation of the population heterogeneity on their performance. The results of this simulation study indicate that the integration of learning and decision making leads to excellent results with gaps of only 0.60% on average relative to an oracle that knows the true population heterogeneity. By contrast, ignoring heterogeneity altogether leads to average gaps of 15.02% relative to an oracle that knows the true population heterogeneity. Failing to integrate learning with decision making leads to average gaps of 7.08% relative to that same oracle. Furthermore, we show that models that ignore population heterogeneity do *not* perform appreciably better when the amount of historical degradation data for model calibration increases.

4. In our second simulation study, we assess the performance of the optimal policy (under real-time, perfect data) when applied in a setting where (i) the degradation signal is not perfect, and (ii) the degradation signal is not relayed in real-time. The results indicate that having access to data in real-time is valuable, while at the same time, this data need not be perfect to achieve excellent performance.

5. We demonstrate the efficacy of integrated learning and decision making on a real data set of X-ray tube degradation in an interventional X-ray machine. This large normalized data set is made available together with this chapter to benchmark future approaches to perform maintenance on degrading systems. This is the first openly available data set containing real degradation data of components from a heterogeneous population. We find that integrated learning can save around 10.50% compared to approaches without learning and around 4.28% compared to an approach where learning is separated from decision making.

The remainder of this chapter is organized as follows. We start by providing a brief literature review in Section 5.2, and subsequently present the model formulation in Section 5.3. We characterize the optimal replacement policy in Section 5.4 and we report on the results of a comprehensive simulation study 5.5. We discuss the application of our approach to alternate settings in which the degradation signal is

imperfect or not relayed in real-time in Section 5.6. We establish the practical value of our model in Section 5.7, where we discuss a real life case study on the X-ray tube degradation in an interventional X-ray machine. Finally, we provide concluding remarks in Section 5.8.

5.2. Literature review

Maintenance and replacement models have been studied extensively in both management and engineering literature and have been reviewed over the decades (Pierskalla and Voelker, 1976; Valdez-Flores and Feldman, 1989; Scarf, 1997; Wang, 2002). The question of when to intervene to maintain or replace a component based on its condition or degradation has been studied already by Derman (1963) and Kolesar (1966). Both establish that, under specific conditions on the degradation process, there exists a threshold such that it is optimal to replace a component if and only if the degradation level of the component is found to have exceeded that threshold. Since these early results, many results about the optimality of threshold policies for the replacement of deteriorating components under many model variations have appeared; see e.g., Ross (1969), Kao (1973), Rosenfield (1976), Benyamini and Yechiali (1999), Makis and Jiang (2003), Maillart (2006) and Kurt and Kharoufeh (2010). Within this literature, the degradation process is assumed to be known to the decision maker and the focus is on proving structural properties of optimal replacement policies.

Elwany et al. (2011) started to study models in which the parameters of the degradation process are only partially known to the decision maker, and differ from component to component. They model the degradation process as a Brownian motion in which the drift parameter is initially unknown and comes from a known prior distribution upon replacement of a component. In the same spirit as our model, this prior distribution models the heterogeneity of the component population. During each planned downtime, the degradation is measured and the belief state regarding the drift parameter of the current component is updated. In this situation, the optimal policy is shown to be a threshold policy where the threshold is an increasing function of the age of the component. The intuition is that as a component ages without failing, it is more likely that the drift is not very high. Since then, this result has been extended to other degradation processes including the inverse Gaussian process (Chen et al., 2015), the gamma process (Zhang et al., 2016), and Markovian processes with a prior distribution on a finite number of possible transition matrices (Van Oosterom et al., 2017). Our work studies the situation where all parameters of the degradation process are initially unknown and come from a prior distribution. Furthermore, we consider the case where the degradation is not only measured during planned downtime but in real-time. This allows us to infer not only the drift of the degradation process but

also higher order properties, in particular the volatility, all of which are unknown to decision makers in real life. Thus, we add to the existing literature described above, by considering a setting where more relevant information can be used to learn the parameters, but also more parameters need to be learned.

All the papers in the previous paragraph use, as do we, a partially observable Markov decision process (POMDP) to study their models. This approach has also been frequently used to learn unknown demand distributions in inventory management; see e.g., Azoury (1985), Chen and Plambeck (2008), and Chen (2010). Furthermore, this approach has been used to study when expensive condition measurements should be made (e.g., Kim and Makis, 2013) and when spare part inventories should be replenished in anticipation of failures (e.g., Li and Ryan, 2011). A more related and recent paper that also involves a POMDP is Kim (2016). He studies a condition based maintenance setting akin to ours but in which the initial priors of the POMDP are mis-specified. The decision maker then seeks a policy that is robust against such mis-specifications.

The usage of condition monitoring to improve decision making is not limited to the condition based maintenance settings that we study. In fact, there exists a broad stream of literature that studies how condition monitoring can improve general decision making; this chapter contributes to this literature stream as well. For example, Uit het Broek et al. (2020) study how the production rate of a production system, which affects the degradation rate of such systems, should be dynamically adjusted based on condition monitoring such that production profits are maximized. Other examples of applications of condition monitoring include spare parts provision based on degradation of components (Olde Keizer et al., 2017), managing rental cars based on their condition (Slaugh et al., 2016), and simultaneously optimizing maintenance and production schedules for multiple products based on machine condition data (Batun and Maillart, 2012).

5.3. Model formulation

This section describes the degradation model and the integrated learning problem of learning the degradation behavior of a component and deciding when to replace it.

5.3.1 Compound Poisson degradation

We consider a component that degrades as random shocks arrive. Shocks arrive as a Poisson process and the damage that accumulates during a shock is random variable, i.e., degradation is a compound Poisson process. The Poisson intensity of

shock arrivals is denoted by λ. The random amount of damage incurred by a shock follows the law of a member of the one-parameter (denoted by ϕ) exponential family, supported on $\mathbb{R}_+$, where $\mathbb{R}_+$ denotes the non-negative real line. Hence, the probability density or mass function of this random amount can be expressed in the form

$$f(x|\phi) = h(x)e^{\phi T(x) - A(\phi)}, \tag{5.1}$$

where $T(x)$ is the sufficient statistic, and $h(x)$ and $A(\phi)$ are known functions. We assume that $T(x) \coloneqq x$, which enables a sate space collapse in our optimization problem (see Section 5.3.3). In the literature, this family of distributions is often referred to as the linear (due to the linear sufficient statistic) exponential family or natural exponential family and was first introduced by Morris (1982). This class encompasses many well-known distributions used in maintenance such as the geometric distribution, the Poisson distribution, the gamma distribution with known shape parameter, and the binomial distribution with known number of trials (see Morris (1982) for a complete overview). For simplicity and due to its practical appeal (see Section 5.7), throughout this chapter, we use the geometric distribution with support $\mathbb{N}_0 \coloneqq \{0, 1, \ldots\}$ to illustrate results. The following example illustrates how the geometric distribution with unknown success probability $p \in (0, 1)$ can be expressed in the canonical form of the natural exponential family.

Example 5.1 (Geometric distribution)**.** Let the damages be geometrically distributed with (unknown) success probability $p \in (0, 1)$. The probability mass function of the random amount of damage, denoted by $f(x|p)$, then takes the form

$$f(x|p) = (1-p)^x p = e^{x \ln(1-p) - \ln(1/p)}. \tag{5.2}$$

Comparing Equation (5.2) with Equation (5.1), we find that $h(x) = 1$, $T(x) = x$, $\phi = \ln(1-p)$, and $A(\phi) = \ln(1/p) = \ln(1/(1-e^{\phi}))$ for the geometric distribution. ♦

The degradation level is observed continuously, but it is only possible to interfere with the system at equally spaced decision epochs. These decision epochs correspond to planned downtimes. For convenience, we rescale time such that the time between two decision epochs equals 1. Furthermore, there exists a threshold $\xi \in \mathbb{N}_+$, where $\mathbb{N}_+ \coloneqq \{1, 2, \ldots\}$, such that a component has failed if its degradation is equal to or exceeds ξ. We let degradation level ξ represent all degradation states that are greater than or equal to some specific physical failure threshold.

Let $N_{(0,t]} \equiv N_t$ denote the total number of shocks received by a component from the start of its life (i.e., from the installation of the component) up to its age t. The number of shocks that arrive between age $t-1 \in \mathbb{N}_0$ and $t \in \mathbb{N}_0$ is denotes by $K_{(t-1,t]} \coloneqq N_t - N_{t-1}$. Observe that integer ages of components coincide with decision epochs. Moreover, let Y_i denote the damage incurred at the $i-$th shock since

the installation of the component. The compound Poisson process at component age $t \in \mathbb{N}_0$ satisfies

$$X_t = \sum_{i=1}^{N_t} Y_i = X_{t-1} + \sum_{i=N_{t-1}+1}^{N_t} Y_i, \quad t \in \mathbb{N}_0,$$

where $X_0 = N_0 = 0$ and $\sum_{i=1}^{0} \cdot \equiv 0$. We also let $\boldsymbol{Y}_t \coloneqq \left(Y_{N_{t-1}+1}, Y_{N_{t-1}+2}, \ldots, Y_{N_t}\right)$ and $Z_{(t-1,t]} \coloneqq \sum_{i=N_{t-1}+1}^{N_t} Y_i$.

5.3.2 Learning the degradation model

We assume that each component stems from a heterogeneous population of components in which each component has different degradation parameters λ and ϕ, which are unknown to the decision maker. Hence, the degradation parameters differ from one component to the next. This reflects the fact that the degradation process may be affected by the individual component's endogenous conditions. We treat the parameters λ and ϕ as random variables, denoted by Λ and Φ, which can be inferred with increasing accuracy by observing the degradation signal of the component in a Bayesian manner.

Λ has a Gamma distribution with shape α_0 and scale β_0 (i.e. $\Lambda \sim \text{Gamma}(\alpha_0, \beta_0)$) and Φ is distributed according to the general prior for a member of the exponential family (parameterized by a_0 and b_0), with $\alpha_0, \beta_0, a_0, b_0 > 0$, with prior density distribution denoted by $f_\Lambda(\lambda|\alpha_0, \beta_0)$ and $f_\Phi(\phi|a_0, b_0)$, respectively. We refer to α_0, β_0, a_0, and b_0 as the hyperparameters. Upon the installation of a new component, the parameters of the compound Poisson degradation process, λ and ϕ, are drawn from these distributions. Let k_t denote the observed number of shocks a components has sustained between ages $t-1$ and t, i.e., k_t is the realization of $K_{(t-1,t]}$. Furthermore let $\boldsymbol{y}_t \coloneqq \begin{pmatrix} y_t^1 & y_t^2 & \ldots & y_t^{k_t} \end{pmatrix}$ be the array of the observed amounts of damage of the shocks sustained, i.e., $\boldsymbol{y}_t$ is the realization of $\boldsymbol{Y}_t$. Finally, let $z_t \coloneqq \sum_{i=1}^{k_t} y_t^i$, be the sustained damage between ages $t-1$ and t, i.e., z_t is the realization of $Z_{(t-1,t]}$. The tuple $\boldsymbol{\theta}_t \coloneqq (k_t, \boldsymbol{y}_t)$ is then the observed degradation signal of a component between ages $t-1$ and t.

The sequential Bayesian updating procedure (also referred to as prior-to-posterior updating, (Ghosh et al., 2007)) works as follows. When a new component is installed, there is no observed degradation signal accumulated yet, and hence Λ and Φ follow independent prior distributions, respectively. This joint prior density distribution, denoted by $f^0_{\Lambda,\Phi}(\lambda, \phi) \coloneqq f_\Lambda(\lambda|\alpha_0, \beta_0) \cdot f_\Phi(\phi|a_0, b_0)$, may be obtained from historical or testing data. (In Section 5.C of the Appendix we discuss an appropriate estimation procedure.) At component age t, we use the observed degradation signal $\boldsymbol{\theta}_t$ and the

joint posterior density distribution of Λ and Φ updated at component age $t-1$, say $f^{t-1}_{\Lambda,\Phi}(\lambda,\phi) \coloneqq f_{\Lambda,\Phi}(\lambda,\phi|\boldsymbol{\theta}_0,\ldots,\boldsymbol{\theta}_{t-1})$, to derive the newly updated joint posterior distribution of Λ and Φ, denoted by $f^{t}_{\Lambda,\Phi}(\lambda,\phi) \coloneqq f_{\Lambda,\Phi}(\lambda,\phi|\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_t)$.

For tractability purposes, so-called conjugate pairs, which have the appealing computational property that the posterior is in the same family as the prior, are often of interest in prior-to-posterior updating. It is well-known that the gamma distribution is a conjugate prior distribution for the Poisson distribution and that a member of the exponential family has a conjugate prior whose density can be expressed in the form (cf. Ghosh et al., 2007)

$$f_\Phi(\phi|a_t,b_t) = H(a_t,b_t)e^{a_t\phi - b_t A(\phi)}.$$

However, since we infer the joint distribution of Λ and Φ using the same observed degradation signal, it is not evident which form the joint posterior distribution of Λ and Φ takes. Proposition 5.1 shows that this joint posterior distribution at component age t can be decomposed into two independent distributions of the same form with updated parameters that only depend on the information obtained in the last period ($\boldsymbol{\theta}_t$). All the proofs are given in Appendix 5.A.

Proposition 5.1. *Given the last observed degradation signal at component age t, $\boldsymbol{\theta}_t = (k_t, \boldsymbol{y}_t)$, and the joint prior distribution $f^{t-1}_{\Lambda,\Phi}(\lambda,\phi) = f_\Lambda(\lambda|\alpha_{t-1},\beta_{t-1}) \cdot f_\Phi(\phi|a_{t-1},b_{t-1})$, the joint posterior distribution, $f^{t}_{\Lambda,\Phi}(\lambda,\phi)$, is equal to $f_\Lambda(\lambda|\alpha_{t-1}+k_t,\beta_{t-1}+1)\cdot f_\Phi(\phi|a_{t-1}+z_t,b_{t-1}+k_t)$.*

Proposition 5.1 induces a simple scheme to infer the true parameters of the degradation process of a component with increasing accuracy. The following example illustrates how the parameter of the compounding distribution can be inferred in the case of geometrically distributed damages.

Example 5.1 (Geometric distribution, continued)**.** We endow a prior on the canonical parameter ϕ with density

$$f_\Phi(\phi|a_t,b_t) = H(a_t,b_t)e^{b_t\phi - a_t A(\phi)} = H(a_t,b_t)e^{\phi b_t}(1-e^\phi)^{a_t},$$

or equivalently, parameterized in terms of p using $p = 1-e^\phi$,

$$f_P(p|a_t,b_t) = H(a_t,b_t)p^{a_t}(1-p)^{b_t},$$

in which we recognize, after normalization, the beta distribution with shape parameter $a_t - 1$ and scale parameter $b_t - 1$. Note also that $a_t = a_0 + \sum_{i=1}^t z_i$ and $b_t = b_0 + \sum_{i=1}^t k_i$. ♦

We now determine the posterior predictive distribution at component age t of the

random variable $Z_{(t,t+1]}$ given the learned information contained in α_t, β_t, a_t, and b_t.

Lemma 5.1. *The posterior predictive distribution at component age t of the random variable $Z_{(t,t+1]}$ given the joint posterior distribution of Λ and Φ, $f_\Lambda(\lambda|\alpha_t,\beta_t)\cdot f_\Phi(\phi|a_t,b_t)$, is equal to:*

$$\mathbb{P}\Big[Z_{(t,t+1]}=z|\alpha_t,\beta_t,a_t,b_t\Big] = \sum_{k=0}^{\infty}\int_{-\infty}^{+\infty} f^{(k)}(z|\,\Phi=\phi)f_\Phi(\phi|a_t,b_t)d\phi\binom{k+\alpha_t-1}{k}\left(\frac{1}{\beta_t+1}\right)^k\left(\frac{\beta_t}{\beta_t+1}\right)^{\alpha_t},$$

where $f^{(k)}(z|\,\Phi=\phi)$ denotes the k-fold convolution of the probability density (or mass) function of the random variable $\{Z_{(t,t+1]}|\,\Phi=\phi\}$.

Lemma 5.1 can be used to construct an updated posterior predictive distribution at each component's age t of the next observed damage increment in real-time based on the observed degradation signal. Hence, the posterior distribution of the degradation parameters of the system is a Markov process whose evolution is induced by the degradation trajectory of the current component.

At first sight, the posterior predictive distribution in Lemma 5.1 seems rather intractable due to the convolution term involved. Fortunately, members of the natural exponential family with a linear sufficient statistic are closed under convolution with itself and hence possess a tractable form (Morris, 1982). Upon insertion of the expression for this convolution term and the corresponding conjugate prior, the posterior predictive distribution reduces to a closed-form expression that can be used for computational purposes. This is illustrated in the example below.

Example 5.1 (Geometric distribution, continued)**.** In this example we use the parameterization in terms of unknown success probability p, which we treat as a random variable denoted by P. Due to the discrete nature of the geometric distribution and as $p\in(0,1)$, we have by Lemma 5.1 that

$$\begin{aligned}\mathbb{P}\Big[Z_{(t,t+1]}=z\Big|K_{(t,t+1]}=&k,\alpha_t,\beta_t,a_t,b_t\Big]\\ &=\int_0^1\mathbb{P}\Big[\sum\nolimits_{i=1}^{k}Y_i=z\Big|P=p\Big]f_P(p|a_t,b_t)dp\\ &=\frac{1}{\mathrm{B}(a_t,b_t)}\binom{z+k-1}{z}\int_0^1 p^k(1-p)^z p^{a_t-1}(1-p)^{b_t-1}dp\\ &=\frac{\mathrm{B}(k+a_t,z+b_t)}{\mathrm{B}(a_t,b_t)}\binom{z+k-1}{z},\end{aligned}\tag{5.3}$$

where $\mathrm{B}(x,y) = \int_0^1 t^{x-1}(1-t)^{y-1}dt$ is the beta function. Note that the distribution of $K_{(t,t+1]}$ is a continuous mixture of Poisson distributions where the mixing distribution of the Poisson rate follows a Gamma(α_t, β_t) distribution, which is known to be the negative binomial distribution with success probability $q = \frac{1}{\beta_t+1}$ and $r = \alpha_t$ number of required successes. Hence, we have

$$\mathbb{P}\Big[K_{(t,t+1]} = k \big| \alpha_t, \beta_t\Big] = \binom{k+\alpha_t-1}{k} \left(\frac{1}{\beta_t+1}\right)^k \left(\frac{\beta_t}{\beta_t+1}\right)^{\alpha_t}. \tag{5.4}$$

Unconditioning Equation (5.3) using Equation (5.4) yields the closed form expression of the posterior predictive distribution:

$$\mathbb{P}\Big[Z_{(t,t+1]} = z | \alpha_t, \beta_t, a_t, b_t\Big] = \sum_{k=0}^{\infty} \frac{\mathrm{B}(k+a_t, z+b_t)}{\mathrm{B}(a_t,b_t)} \binom{z+k-1}{z} \binom{k+\alpha_t-1}{k} \left(\frac{1}{\beta_t+1}\right)^k \left(\frac{\beta_t}{\beta_t+1}\right)^{\alpha_t}. \quad \blacklozenge$$

5.3.3 Markov decision process formulation

Each component will incur a cost due to either corrective or preventive replacement. If the degradation level at a decision epoch is equal to the failure threshold ξ, then the failed component is replaced correctively at cost c_u. (Recall that degradation level ξ represents all degradation states that are greater than or equal to the failure threshold.) If the degradation level at a decision epoch does not exceed ξ, then we can either perform preventive replacement at cost c_p, or continue to the next decision epoch at no cost. We assume that replacements take negligible time and that $0 < c_p < c_u < \infty$ to avoid trivial cases.

Recall that each component stems from a heterogeneous population that includes components with different degradation parameters λ and ϕ. Note that these parameters cannot be observed; only the degradation signal $\boldsymbol{\theta}_i$ for $i = 1, \cdots, t$ is observable at component age t. We will therefore use a POMDP to model the integrated problem of learning the degradation parameters of a component and deciding when to replace. First we observe that due to the results in the previous section, the information state of a component at age t can be represented by $(\alpha_t, \beta_t, a_t, b_t)$. Furthermore, the decision maker knows the current degradation level x. The state at decision epoch τ is therefore given by $(x, \alpha, \beta, a, b, \tau)$ where x denotes the degradation level of the component that is in service and α, β, a, and b encode the most current degradation information of the component that is in service. This six dimensional state representation can be collapsed into an equivalent four dimensional state representation (x, n, t, τ) where n denotes the number of shocks that the current

component has sustained and t denotes its age. Indeed observe that by Proposition 5.1 we have $\alpha = \alpha_0 + n$, $\beta = \beta_0 + t$, $a = a_0 + x$, and $b = b_0 + n$. This representation is insightful: All the information in the degradation signal is encoded in the total degradation level, the number of shocks sustained, and the age of the component. The crucial assumption for this collapse is that the sufficient statistic $T(x)$ for the damage distribution is equal to x, see Equation (5.1). If the sufficient statistic would not be linear in the damage, then the state space would need to include all individual damage arrivals.

We are interested in finding the optimal replacement policy $\pi^* = \{\pi_\tau\}_{\tau \in \mathbb{N}_0}$ that minimizes the total expected discounted cost of corrective and preventive replacements over an infinite horizon, where the costs are discounted by a factor $\gamma \in (0, 1)$. This policy is a sequence of decision rules that prescribe whether or not to perform preventive maintenance if $x < \xi$. By Proposition 1.2.2. of Bertsekas (2007), there exists an optimal Markov (deterministic) policy depending only on state (x, n, t) independent of the decision epoch $\tau \in \mathbb{N}_0$. Let

$$V(\boldsymbol{s}) = \inf_{\pi \in \Pi} \lim_{T \to \infty} \mathbb{E}_{\boldsymbol{s}} \left[\sum_{\tau=1}^{T} \gamma^\tau C\left(\boldsymbol{S}_\tau, \pi(\boldsymbol{S}_\tau)\right) \right]$$

denote the expected total discounted cost given that the process starts in state $\boldsymbol{s} = (x, n, t)$ where $\boldsymbol{S}_\tau$ denotes the state of the component operating at decision epoch τ, Π denotes the set of Markov (deterministic) policies, $\mathbb{E}_{\boldsymbol{s}}$ is the conditional expectation given that the process starts in state $\boldsymbol{s} = (x, n, t)$, and $C(\boldsymbol{s}, \pi(\boldsymbol{s}))$ denotes the cost function defined as

$$C(\boldsymbol{s}, \pi(\boldsymbol{s})) = \begin{cases} c_p & \text{if } x < \xi \text{ and } \pi(\boldsymbol{s}) = \text{replace}, \\ 0 & \text{if } \ x < \xi \text{ and } \pi(\boldsymbol{s}) = \text{continue}, \\ c_u & \text{if } x = \xi. \end{cases}$$

Due to this state space collapse, the posterior distribution of $Z_{(t,t+1]}$ is function of the state $\boldsymbol{s} = (x, n, t)$. Therefore we will use the shorthand notation $Z(\boldsymbol{s})$ to indicate the random variable $Z_{(t,t+1]}$ given the state $\boldsymbol{s} = (x, n, t)$. Similarly, we use the shorthand notation $K(\boldsymbol{s})$ and $Y(\boldsymbol{s})_i$ for the random variable $K_{(t,t+1]}$ and the random variables $\{Y_i\}_{i \in \mathbb{N}}$, given state $\boldsymbol{s} = (x, n, t)$. Note that $Z(\boldsymbol{s})$ and $K(\boldsymbol{s})$ are dependent random variables. This is intuitively clear as damage can only accumulate when the component sustains shocks. It is convenient to define the random variable $A(\boldsymbol{s}) = (Z(\boldsymbol{s}), K(\boldsymbol{s}), 1)$ and note that its distributions can be determined with Proposition 5.1 and Lemma 5.1. The optimal replacement policy π^* is satisfies the Bellman optimality

equations

$$V(\boldsymbol{s}) = \begin{cases} c_u + \gamma \mathbb{E}\left[V(\boldsymbol{s_0} + A(\boldsymbol{s}_0))\right], & \text{if } x = \xi, \\ \min\Big\{c_p + \gamma \mathbb{E}\left[V(\boldsymbol{s}_0 + A(\boldsymbol{s}_0))\right]; \gamma \mathbb{E}_{\boldsymbol{s}}\left[V(\boldsymbol{s} + A(\boldsymbol{s}))\right]\Big\}, & \text{if } x < \xi, \end{cases} \tag{5.5}$$

where $\boldsymbol{s}_0 := (0, 0, 0)$.

The first case in equation (5.5) follows because failed components must be replaced correctively at cost c_u. If the component's degradation level is less than ξ, we can either perform a preventive replacement, which costs c_p, or leave the component in operation until the next decision epoch at no cost. Upon preventive or corrective replacement, the parameters of the new component are unknown and need to be learned as the replacement component ages.

5.4. Optimal replacement policy

This Section presents structural results on the optimal replacement policy. First we present two examples of optimal replacement policies that illustrate several structural properties. Motivated by practice (see Section 5.7), we use the geometric compounding distribution (with support $\mathbb{N}_0$) in these illustrations.

5.4.1 Examples of optimal replacement policies

Figure 5.1 illustrates the optimal replacement policy for two different degradation sample paths. In these examples, the solid black line depicts the observed degradation path as the component ages, where a star denotes a shock arrival. The dashed line depicts the optimal control limit, where at each decision epoch, the optimal action is to carry out a preventive replacement if the degradation level is at or above the optimal control limit. The failure level ξ depicted with the dot-dashed line, so that a corrective replacement has to be performed if the solid black line is at or above this dot-dashed line.

We can state three properties that we consistently observe, which are best explained by the illustrations in Figure 5.1.

1. The optimal replacement policy is a control limit policy where the control limit depends on the past degradation signal.

2. The optimal control limit is non-decreasing in the age of the component, at least for a fixed number of shock arrivals n.

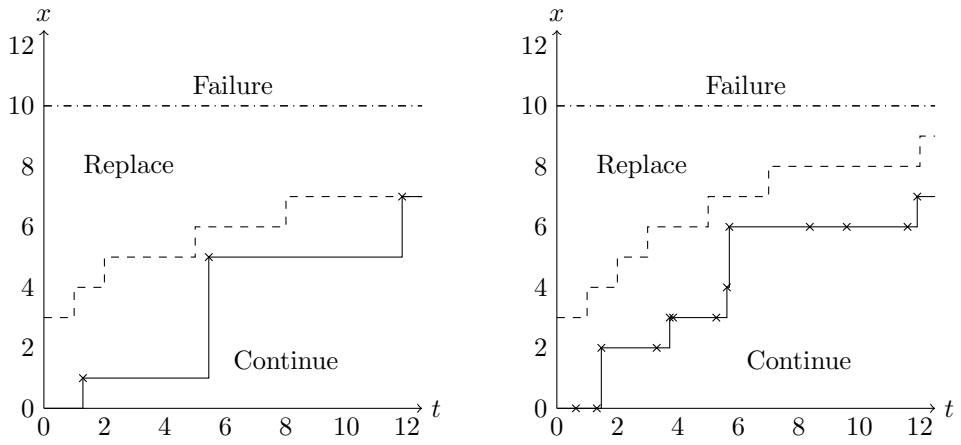

Figure 5.1 Two sample degradation paths and the optimal replacement policy, with $\alpha_0 = \beta_0 = a_0 = b_0 = 1$, $c_u = 4$, $c_p = 1$ and $\xi = 10$. For the compound Poisson processes that generate the sample paths, we have $\lambda = 0.25$ and $\phi = 0.2$ (left), and $\lambda = 1.5$ and $\phi = 0.8$ (right).

3. If we compare two components whose degradation is identical at age t, then the component with fewer shock arrivals has a lower control limit. This is illustrated in Figure 5.1, where both components have accrued the same cumulative degradation at age $t = 12$, yet the optimal policy prescribes to preventively replace the component on the left, but not the component on the right. Observe that the degradation path on the left may be described as more volatile.

   This monotonicity property with respect to the number of arrived shocks is usually but not always observed. An example where this property is violated is shown in Figure 5.2.

5.4.2 Structural properties

In this section, we establish both the optimality of a control limit policy, as well as the monotonic structure of this control limit, the observations above.

Lemma 5.1 shows how the decision maker can update the posterior predictive distribution of the upcoming degradation increment by utilizing the observed degradation signal $\boldsymbol{s}$. Proposition 5.2 presents two properties that provide insight into what the decision maker knows about the distribution of future degradation increments.

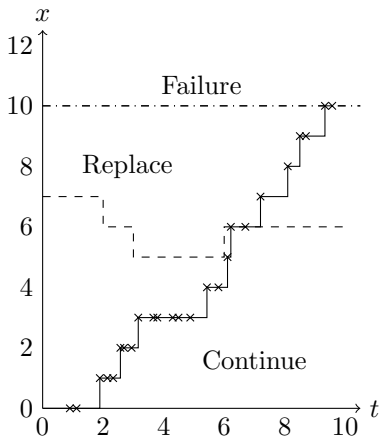

Figure 5.2 Sample degradation path and the optimal replacement policy, with $\alpha_0 = \beta_0 = 1$ and $a_0 = 60$ and $b_0 = 10$, $c_u = 4$, $c_p = 1$ and $\xi = 10$. For the compound Poisson process that generates the sample path, we have $\lambda = 1.75$ and $\phi = 0.9$.

Proposition 5.2. *The random variable $Z(\boldsymbol{s})$ satisfy the following stochastic orders:*

(i) $Z(\boldsymbol{s})$ is stochastically decreasing in t in the usual stochastic order;

(ii) $Z(\boldsymbol{s})$ is stochastically increasing in x in the usual stochastic order.

Part (i) of Proposition 5.2 shows that older components will accumulate – in expectation – less damage than younger components. As a component ages without failing, the decision maker infers that large increments are unlikely to happen for this component. The intuition behind Part (ii) of Proposition 5.2 is that when more damage has accumulated already, then we should expect more damage to accumulate in the future.

There is no monotone stochastic ordering in of $Z(\boldsymbol{s})$ in n in general as shown in the following example.

Example 5.1 (Geometric distribution, continued)**.** Consider again the geometric compounding distribution with a beta prior. Let $b_0 - \alpha_0 - 1 > 0$, then the expectation of $Z(\boldsymbol{s})$.

$$\mathbb{E}[Z(\boldsymbol{s})] = \frac{a(\alpha_0 + n)}{\beta(b_0 + n - 1)}$$

increases in n. However, the second moment

$$\mathbb{E}[Z(\boldsymbol{s})^2] = \frac{a(\alpha_0+n)}{\beta^2(b_0+n-1)^2} \times \left(\beta^3\left(\frac{2(a+1)}{b_0+n-2} + \beta(b_0+n) + 2b_0 + a - \beta + 2n\right) + a(\alpha_0+n)\right)$$

is not monotonically increasing. (The derivation of these two moments is provided in Appendix 5.B). For example, choosing $b_0 = 2.62$, $\alpha_0 = \beta = 1$ and $a = 10$ yields that the second moment is increasing for all $n \leq 15$ and decreasing for all $n \geq 16$. We can therefore conclude that there is no monotone stochastic ordering of the random variables $Z(\boldsymbol{s})$ in n (cf. Shaked and Shanthikumar, 2007, Theorem 1.A.3). ♦

The following two properties of the value function are essential to establish the structure of the optimal replacement policy. Proposition 5.2 is pivotal to establish these two properties. As such here too there is no monotonic behavior of the value function in n in general.

Lemma 5.2. *The value function $V(\boldsymbol{s})$ is*

(i) non-increasing in t;

(ii) non-decreasing in x.

With these result we can establish the main result of this section.

Theorem 5.1. *At each component age $t \in \mathbb{N}_0$, for a given number of shock arrivals $n \in \mathbb{N}_0$, there exists a control limit $\delta^{(n,t)} \leq \xi$, such that the optimal action is to carry out a preventive replacement if and only if $x \geq \delta^{(n,t)}$. The control limit $\delta^{(n,t)}$ is monotonically non-decreasing in t, for all n.*

The optimal control limit is non-decreasing in t because as the component ages without failing, the decision makes is increasingly assured that the current component degrades slowly relative to the general population of components.

The structural results are not only intuitive and convenient for the implementation of an optimal policy in practice. They can also be exploited to decrease the computational burden of finding the optimal policy by employing existing algorithms that rely on these monotonicity properties such as the monotone policy iteration algorithm (see Puterman, 2005, Section 6.11.2).

5.5. Simulation study

This section reports the results of a comprehensive simulation study. Although the established structural results hold for any one-parameter member of the exponential family with positive support, we assume in this section, motivated by practice, that the damages are geometrically distributed. This simulation study starts from the premise that the true hyperparameters of the degradation behavior are unknown to the decision maker; they only have access to historical degradation data for model calibration. This premise differs from previous contributions that use Bayesian techniques to model real-time learning, where hyperparameters for prior distributions are generally assumed to be given (e.g., Chen, 2010). This latter approach is, however, arguably not the case in practice. Indeed, decision makers only have access to historical degradation data that they should leverage in order to estimate hyperparameters. The main objective of this simulation study is, therefore, twofold:

1. To examine the value of integrating learning and decision making, which takes into account explicitly that degradation of components is heterogeneous (value of integration).

2. To assess how the amount of available historical degradation data that is used for model calibration affects the performance (value of data).

To assess the value of integration and data, we define three heuristic approaches, all of which start from the same historical degradation data, but differ in how they calibrate their models, learn from the degradation of components, and integrate learning with decision making. We shall compare the performance of each approach with the performance of an oracle who does know the true hyperparameters of the degradation behavior of different components. The oracle thus follows the replacement policy that solves the Bellman optimality equations in (5.5), calibrated with the true hyperparameters. We denote the oracle policy by $\pi_{\mathcal{O}}$.

For each instance of our simulation study, which we describe in detail later, the true hyperparameters of the gamma and beta prior distribution of Λ and Φ that model the population heterogeneity are denoted by $\tilde{\alpha}$, $\tilde{\beta}$, $\tilde{a}$, and $\tilde{b}$. These are known only to the oracle. The historical degradation data that serves as starting point for the heuristic approaches of that same instance is obtained by simulating degradation paths of components whose degradation parameters λ and ϕ are drawn from a gamma distribution with the true hyperparameters $\tilde{\alpha}$ and $\tilde{\beta}$, and a beta distribution with $\tilde{a}$ and $\tilde{b}$, respectively. We now proceed with defining the three heuristic approaches, after which we describe the simulation set-up and discuss the results.

5.5.1 Offline approach

The first heuristic approach ignores both the population heterogeneity and the real-time degradation signal, which is the current state-of-the-art. This approach assumes that the degradation of each component upon installation follows a compound Poisson process with the same parameters λ and ϕ. Under this assumption, the decision maker faces the classical replacement problem for which we know that the optimal replacement policy is given by a stationary control limit (e.g., Kolesar, 1966). The optimal control limit can be readily found by solving the following one dimensional Bellman optimality equations:

$$V(x) = \begin{cases} c_u + \gamma \mathbb{E}\left[V(\tilde{Z}(\lambda,\phi))\right], & \text{if } x = \xi, \\ \min\left\{c_p + \gamma \mathbb{E}\left[V(\tilde{Z}(\lambda,\phi))\right]; \gamma \mathbb{E}\left[V(x+\tilde{Z}(\lambda,\phi))\right]\right\}, & \text{if } x < \xi, \end{cases} \tag{5.6}$$

where $\tilde{Z}(\lambda,\phi) := \sum_{i=1}^{K(\lambda)} Y_i(\phi)$ denotes the degradation increment in between two consecutive decision epochs, $K(\lambda)$ is a Poisson distributed random variable with parameter λ and $Y_i(\phi)$ are geometrically distributed with parameter ϕ.

Hence, under the first approach, the decision maker approximates the degradation parameters by the corresponding point estimates $\bar{\lambda}$ and $\bar{\phi}$, which are obtained using Maximum Likelihood Estimation (MLE) based on the available historical degradation data, and then solves the optimality equations (5.6) with those estimates to obtain a single control limit, that is used for all components. In the remainder of this section, we refer to this approach as the offline approach because it ignores both the population heterogeneity and the real-time degradation signal. We denote this approach by $\pi_{\mathcal{N}}$ since it is the most naive heuristic approach of all three approaches.

5.5.2 Myopic online approach

The second heuristic approach does utilize the real-time degradation signal to update the point estimates of the degradation parameters of an individual component. As such, this approach takes into account that the population of components is heterogeneous. This approach is calibrated as follows. Given the historical degradation paths, we estimate the initial hyperparameters α_0, β_0, a_0, and b_0, by maximizing the likelihood of those degradation paths being induced by components stemming from a population whose heterogeneity is modeled through a gamma and beta distribution with these hyperparameters. Thus, based on the available historical degradation data, we estimate the initial hyperparameters α_0, β_0, a_0, and b_0, using MLE (further details regarding this MLE procedure are relegated to Appendix 5.C).

After calibration, this heuristic approach works as follows: At component age t, the

decision maker updates the information state encoded in $f_\Lambda(\lambda|\alpha_t, \beta_t)$ and $f_\Phi(\phi|a_t, b_t)$ in the same way as in the original Bayes model (cf. Proposition 5.1). The decision maker then updates the point estimates of the degradation increment based on the minimum mean square error (MMSE) estimator. In a Bayesian setting, MMSE estimates correspond to posterior means. Hence, at component age t, the MMSE estimates for the degradation parameters are given by $\bar{\lambda}_t := \int_0^\infty u f_\Lambda(u|\alpha_t, \beta_t)du = \alpha_t/\beta_t$ and $\bar{\phi}_t := \int_0^1 u f_\Phi(u|a_t, b_t)du = a_t/(a_t + b_t)$. The decision maker then computes a control limit by solving the optimality equations (5.6), where the parameters of the Poisson and of the geometric distribution of the degradation increment $\tilde{Z}(\lambda, \phi)$ are now given by $\bar{\lambda}_t$ and $\bar{\phi}_t$, respectively. Although the second approach partly captures the Bayesian learning benefits, it does not integrate learning with optimization but rather solves a myopic optimization problem repeatedly with the latest point estimates of the degradation parameters. In the remainder of this section, we therefore refer to this approach as the myopic online approach, denoted by $\pi_\mathcal{M}$.

5.5.3 Integrated Bayes approach

The third heuristic approach is the most sophisticated approach. It is similar to $\pi_\mathcal{O}$ in that it follows the replacement policy that solves the Bellman optimality equations in (5.5). However, as we have argued before, the true hyperparameters that model the population heterogeneity are unknown to the decision maker in practice. As such, this approach differs from $\pi_\mathcal{O}$ in that we calibrate this approach using MLE, in the same way as how we calibrate $\pi_\mathcal{M}$. That is, we estimate the initial hyperparameters α_0, β_0, a_0, and b_0, using MLE based on on the available historical degradation data. These estimated hyperparameters are then used as input for finding the optimal replacement policy through solving the Bellman optimality equations in (5.5). We henceforth refer to this approach as the integrated Bayes approach, denoted by $\pi_\mathcal{I}$. It is important to note that this approach is precisely how our model should be applied in practice, where the hyperparameters describing the heterogeneity of the component population should be estimated from historical degradation data.

5.5.4 Results

The main performance metric in this simulation study is the gap between the long run average cost rate induced by the oracle policy and the offline approach, the myopic online approach, and the integrated Bayes approach, where the latter three models are calibrated based on MLE. More formally, we are interested in $\%GAP_\pi = 100 \cdot (C_\pi - C_{\pi_\mathcal{O}})/C_{\pi_\mathcal{O}}$, where C_π is the long run average cost rate of approach $\pi \in \{\pi_\mathcal{I}, \pi_\mathcal{M}, \pi_\mathcal{N}\}$.

To achieve the two main objectives stated in the beginning of this section, we set

up a large test-bed consisting of instances obtained through all combinations of the parameter values in Table 5.1, with $\xi = 20$ and $c_p = 1$. To vary the heterogeneity in the population of components, we naturally vary the variance of the gamma and beta prior distributions of Λ and Φ, respectively, while keeping their respective means fixed at 1 and 0.5 for all instances.

Table 5.1 Input parameter values for simulation study.

	Input parameter	No. of choices	Values
1	Variance of prior gamma distribution, σ^2_Λ	2	0.5 1
2	Variance of prior beta distribution, σ^2_Θ	2	0.01, 0.02
3	Corrective maintenance cost, c_u	2	5, 10
4	Number of simulated degradation paths, ν	2	10, 50

For each instance of the test-bed, we determine the true hyperparameters $\tilde{\alpha}$, $\tilde{\beta}$, $\tilde{a}$, and $\tilde{b}$, based on the variance of the gamma and beta distribution of that instance (and their fixed means). We use these hyperparameters to compute $\pi_{\mathcal{O}}$ and we also simulate a number of degradation paths, denoted ν, corresponding to that instance. We subsequently use these simulated degradation paths to calibrate approaches $\pi_{\mathcal{I}}$, $\pi_{\mathcal{M}}$ and $\pi_{\mathcal{N}}$. That is, we use MLE to estimate α_0, β_0, a_0, and b_0, in case of $\pi_{\mathcal{I}}$ and $\pi_{\mathcal{M}}$, and point estimates $\bar{\lambda}$ and $\bar{\phi}$, in case of $\pi_{\mathcal{N}}$. We then simulate $15{\cdot}10^3$ components, where upon installation of a new component, its degradation parameters are drawn from a gamma and beta distribution with the true hyperparameters $\tilde{\alpha}$, $\tilde{\beta}$, $\tilde{a}$ and $\tilde{b}$. For each simulated component, we keep track of the relevant cost rates and subsequently calculate $\%GAP_\pi$ for each approach $\pi \in \{\pi_{\mathcal{I}}, \pi_{\mathcal{M}}, \pi_{\mathcal{N}}\}$. We repeat this procedure 30 times for each instance of the test-bed to ensure that the confidence intervals of the cost rates are sufficiently small. Throughout this simulation study, we use a discount factor of 0.99 in solving the corresponding optimality equations of each approach, and we truncate both state variables n and t at a sufficiently large value (i.e., 40) in computing the optimal policy under both $\pi_{\mathcal{I}}$ and $\pi_{\mathcal{O}}$.

The results of the simulation study are summarized in Table 5.2. In this table, we present the minimum, average, and maximum $\%GAP_\pi$ for each approach $\pi \in \{\pi_{\mathcal{I}}, \pi_{\mathcal{M}}, \pi_{\mathcal{N}}\}$. We first distinguish between subsets of instances with the same value for a specific input parameter of Table 5.1 and then present the results for all instances.

The following main observations can be drawn from Table 5.2: First, $\pi_{\mathcal{I}}$ yields excellent results with a gap of only 0.60% on average relative to the oracle. Both heuristic approaches $\pi_{\mathcal{M}}$ and $\pi_{\mathcal{N}}$ perform poorly, with gaps of 7.08% and 15.02% on average relative to the oracle. Ignoring both the degradation signal and the heterogeneity in the population can be quite detrimental, as gaps of up to 24.04% relative to the oracle do occur under $\pi_{\mathcal{N}}$. Although learning the degradation signal is beneficial, failing to integrate this with decision making directly can still lead to

Table 5.2 Results of simulation study.

		$\%GAP_\pi$								
		$\pi_{\mathcal{I}}$			$\pi_{\mathcal{M}}$			$\pi_{\mathcal{N}}$		
Input	Value	Min	Mean	Max	Min	Mean	Max	Min	Mean	Max
σ^2_Λ	0.3	0.05	0.58	1.26	3.52	6.14	10.05	7.91	13.97	22.93
	0.6	0.15	0.63	1.34	5.24	8.02	11.43	10.16	16.08	24.04
σ^2_Θ	0.01	0.05	0.46	1.02	3.52	5.54	8.11	7.91	11.31	15.25
	0.02	0.25	0.75	1.34	5.83	8.62	11.43	13.85	18.73	24.04
c_u	5	0.05	0.48	1.16	3.52	5.66	7.74	7.91	12.01	15.73
	10	0.24	0.72	1.34	5.21	8.50	11.43	10.83	18.03	24.04
ν	10	0.42	0.95	1.34	3.97	7.23	11.43	8.63	15.46	24.04
	50	0.05	0.26	0.49	3.52	6.93	11.36	7.91	14.58	23.00
Total		0.05	0.60	1.34	3.52	7.08	11.43	7.91	15.02	24.04

gaps with the oracle of up to 11.43%. All three approaches $\pi_{\mathcal{I}}$, $\pi_{\mathcal{M}}$, and $\pi_{\mathcal{N}}$ seem to perform worse when the heterogeneity in the population increases, and when the cost of performing corrective maintenance becomes higher.

Second, it is generally believed that increasing the amount of data available for model calibration leads to better decisions. However, when the underlying assumptions of the model are wrong, then this may not be true. Indeed, the performance of both heuristic approaches $\pi_{\mathcal{M}}$ and $\pi_{\mathcal{N}}$ does not increase considerably in the number of simulated degradation paths that serve as input to these models. The performance of $\pi_{\mathcal{I}}$ increases however significantly when this number increases. Furthermore, even when the amount of available data for estimating the population heterogeneity is limited, the integrated Bayes approach, $\pi_{\mathcal{I}}$, still yields excellent results.

5.6. Alternate settings

We have so far assumed that the degradation signal is relayed in real-time and that it provides a perfect observation of the actual degradation level. These assumptions are in line with what we have observed at our industrial partner who instigated this research (see also the case study in the next section). We note that there may be practical settings where such assumptions are not justified. It is however intractable to relax these assumptions and consequently compute optimal policies within our current modelling framework. In this section, we therefore study the performance of our integrated Bayes approach when it is applied to settings where the degradation signal is imperfect or relayed periodically. We do so in a simulation study that follows the same procedure as in the previous section.

5.6.1 Imperfect degradation signal

In line with most of the research on imperfect condition monitoring (e.g., Maillart, 2006; Kim and Makis, 2013), we model an imperfect degradation signal by constructing a state-observation matrix $\mathbf{Q}$ that captures the stochastic relationship between the actual degradation level X_t and the (imperfect) observation, denoted with $\bar{X}_t$, that the decision maker observes. More formally, let $\mathbf{Q} \coloneqq \big(q_{ij}\big)_{\xi\times\xi}$, whose entry $q_{ij} \coloneqq \mathbb{P}\left[\bar{X}_t = j \mid X_t = i\right]$ is equal to

$$\frac{\varphi_i(j|\sigma)}{\sum_{j=0}^{\xi-1}\varphi_i(j|\sigma)}, \tag{5.7}$$

where $\varphi_x(y|\sigma)$ is the density function of a normal random variable with mean x and standard deviation σ evaluated at y. Recall that in our simulation study, we assume that the damages are integer valued, so that we have ξ non-failed states, i.e. $0, 1, \ldots, \xi-1$, that are not observable, and one failed state that is observable. The matrix $\mathbf{Q}$ applies to the ξ non-failed states. Constructing this matrix using the parameterization in Equation (5.7) is an approach often used in literature (see, e.g., Maillart, 2006; Kim and Makis, 2013; Liu et al., 2021), where the standard deviation σ is a measure of the noise (or imperfectness) of the observations when the system has not failed yet. By varying the value of σ we can thus investigate the robustness of our approach to the level of imperfectness of the observations. Note that by setting σ equal to zero, we are in the situation that we have a perfect observation of the actual degradation level.

5.6.2 Intermittent degradation signal

When the degradation signal is relayed periodically at decision epochs, the decision maker has only access to the degradation level of the current component x_t (and its age t). However, to apply the integrated Bayes approach we require the number of sustained loading epochs n_t too. To resolve this, we rely on the following recursive formula to estimate n_t:

$$\bar{n}_t = \frac{x_t(a_0 + x_{t-1} - 1)}{b_0 + \bar{n}_{t-1}}, \qquad t > 0,$$

and $\bar{n}_0 = 0$. We round to the nearest integer in case $\bar{n}_t$ is not integral. This estimation procedure has intuitive appeal. To obtain an estimate for n_1 at $t = 1$, we divide the degradation level x_1 by the expected damage per loading epoch given the initial hyperparameters a_0 and b_0 that are obtained through the MLE procedure for model calibration, i.e. $\frac{b_0}{(a_0-1)}$. (Observe that this is the expectation of a geometric random

variable whose parameter is beta distributed with parameters a_0 and b_0, see Appendix 5.B for the derivation.) We then apply the updating rules of Proposition 5.1 using this estimate, that is $a_1 = a_0 + x_1$ and $b_1 = b_0 + \bar{n}_1$, and consequently apply the same logic at $t = 2$ to obtain an estimate for n_2. This procedure repeats until the component is replaced.

5.6.3 Results

The long run average cost rate of the integrated Bayes approach applied to the imperfect degradation signal setting and the intermittent degradation signal setting are denoted $C^{\text{imp}}_{\pi_{\mathcal{I}}}$ and $C^{\text{int}}_{\pi_{\mathcal{I}}}$, respectively. We are interested in how the integrated Bayes approach performs in these alternate settings, and we therefore compare said costs rates with the cost rate of the integrated Bayes approach that does have access to the true degradation signal in real time. That is, we compute $\%VAL_{\text{int}} = 100 \cdot (C^{\text{int}}_{\pi_{\mathcal{I}}} - C_{\pi_{\mathcal{I}}})/C_{\pi_{\mathcal{I}}}$ and $\%VAL_{\text{imp}} = 100 \cdot (C^{\text{imp}}_{\pi_{\mathcal{I}}} - C_{\pi_{\mathcal{I}}})/C_{\pi_{\mathcal{I}}}$ for each instance of the test bed. The instances of our test bed are identical to the test bed of the previous section, see Table 5.1. In the test bed of the imperfect degradation signal, we additionally vary σ over 5 different levels, i.e. $\sigma \in \{0.25, 0.5, 0.75, 1, 1.25\}$. Note that $C_{\pi_{\mathcal{I}}}$ is a lower bound on the cost rate of the *optimal* policy for both alternate settings.

The average $\%VAL_{\text{int}}$ and the average $\%VAL_{\text{imp}}$ are presented in Table 5.3 and Table 5.4, respectively. As before, we first distinguish between subsets of instances with the same value for a specific input parameter and then present the results for all instances.

Table 5.3 Results for intermittent degradation signals.

Input parameter	Value	$\%VAL_{\text{int}}$
Variance of prior gamma distribution, σ^2_Λ	0.3	6.51
	0.6	11.40
Variance of prior beta distribution, σ^2_Θ	0.01	10.37
	0.02	7.54
Number of simulated paths, ν	10	9.89
	50	8.02
Total		8.95

From Table 5.3 we see that when the degradation signal is not relayed in real-time, our integrated Bayes approach performs relatively poorly with an $\%VAL_{\text{int}}$ of almost 9 percent on average. This also implies that relaying a degradation signal in real-time rather than only periodically has considerable value. Indeed, it allows the decision

Table 5.4 Results for imperfect degradation signals.

Input parameter	Value	$\%VAL_{\text{imp}}$
Variance of prior gamma distribution, σ^2_Λ	0.3	3.29
	0.6	5.95
Variance of prior beta distribution, σ^2_Θ	0.01	4.59
	0.02	4.65
Number of simulated paths, ν	10	4.66
	50	4.58
Standard deviation of the noise, σ	0.25	0.01
	0.5	1.67
	0.75	3.50
	1	6.70
	1.25	11.23
Total		4.62

maker to not only learn the drift of a degradation signal (encoded in the degradation level and the age) but also the volatility of the degradation signal; the latter can only be inferred if one has access to the individual arrivals of loading epochs and their corresponding damages in between decision epochs.

By contrast, our integrated Bayes approach performs quite well when the degradation signal is imperfect. Indeed, from Table 5.4 we see that $\%VAL_{\text{imp}}$ is less than 5 percent on average, and even remains below 3.5 percent for moderate levels of noisiness.

5.7. Case study

Interventional X-ray (IXR) systems are used by physicians for minimally-invasive image-guided procedures to diagnose and treat diseases in nearly every organ system. X-ray tubes (denoted by the rectangle in Figure 5.3) are the most expensive replacement components of an IXR system and therefore of major concern. Philips Healthcare produces the IXR system and does maintenance and service for many hospitals that use the IXR system.

Unexpected downtime incidents have a major impact, especially for the patients whose medical procedure is cut short or postponed. The cost of premature maintenance is also substantial. Medical imaging equipment have list-prices on the order of one million US dollars and the annual maintenance expenses of such equipment are around 10% of the list-price (ECRI, 2013). Since these medical imaging systems generally last up to 10 years, roughly half of the total cost of ownership of such a system (excluding downtime costs) consists of maintenance costs.

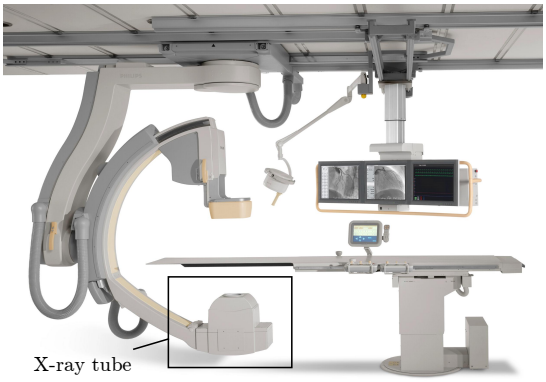

Figure 5.3 Example of IXR system with X-ray tube denoted by rectangle. (Philips, 2020)

Philips healthcare faces the challenge of replacing the expensive X-ray tube before failures occur but also to maximize their useful lifetime. Below we describe a case study on the most critical component of an IXR system: the X-ray tube. The full data set for this case-study is included in the online companion to this article.

This section is organized as follows: We describe the dominant failure mechanism of the X-ray tube (Section 5.7.1), give a description of the data set (Section 5.7.2), illustrate the operation of the optimal policy with real data (Section 5.7.3) and compare the three approaches outlined in Section 5.5 on real data with a bootstrapping study (Section 5.7.4).

5.7.1 IXR Filaments

A failure analysis performed by Philips Healthcare indicates that X-ray tube failures are predominantly caused by worn out filaments (Albano et al., 2019, Section 5.3.3.4). These tungsten filaments are heated to a high temperature by a voltage differential such that they emit electrons. These electrons are then accelerated by a high voltage potential differential towards the target so that they emit X-rays when they hit the target. The X-rays are then used to produce the desired image during image guided medical procedures. This process is depicted in Figure 5.4.

The tungsten evaporates slowly when the filament heats up. This filament usually develops a “hot-spot” at the thinnest location. The evaporation causes the hotspot to become thinner with every image taken. This continues until the tungsten melts at the hot-spot and the filament fails, see Covington (1973). The degradation state of a filament can be inferred from the resistance of the filament. Philips

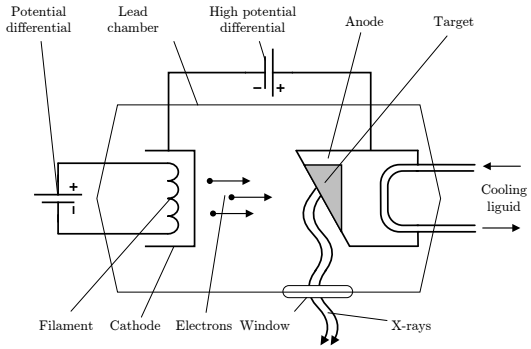

Figure 5.4 Simplified X-ray tube schematic.

Healthcare performed a statistical analysis to derive a single-dimensional health indicator that governs the degradation state over time of a filament (Albano et al., 2019, Section 7.9.7.12). This single-dimensional health indicator, the degradation, is recorded in a database each time that an IXR system is used for an image-guided procedure.

5.7.2 Degradation data of IXR X-ray tubes

The data set of X-ray tube degradation consist of 52 time-series of degradation levels. Let $\mathcal{I}$ be the set of all X-ray tubes for which there is available data; $|\mathcal{I}| = 52$. The time series of a single X-ray tube $i \in \mathcal{I}$ is denoted $\mathcal{J}_i$. Each datum $j \in \mathcal{J}_i$ in such a time-series is a tuple $(t_j, x_j)_i$ of the age of the X-ray tube t_j and the degradation level x_j at that age. Each tuple $(t_j, x_j)_i$ of X-ray tube i is generated when an IXR system is used and each time-series consists of 20.000-300.000 data points originating from a time period of 2-5 years. Due to confidentiality reasons we have left-truncated the data and normalized the data. That is, all time-series start with $x_0 = 0$ for $t_0 = 0$, and end at $x_{|\mathcal{J}_i|} = 50$ (i.e., $\xi = 50$) for all $i \in \mathcal{I}$. For each time-series, we computed the inter-arrival times (i.e., $t_j - t_{j-1}$) between succeeding data points and damage increments per data point (i.e., $x_j - x_{j-1}$). After removing outliers in the inter-arrival times due to either weekend or other prolonged non-operational periods and removing data points for which the image-guided procedure was considered too short to wear out the X-ray tube (i.e., shock arrival is regarded as non-critical), the assumption that shocks arrive as a Poisson process was not rejected based on the Kolmogorov-Smirnov test. Additionally, we normalized the time such that one unit of time corresponds to roughly the operational time that is considered as minimally achievable for performing maintenance practices from a practical perspective (e.g.,

sending a service engineer to the location of the hospital). Furthermore, based on the Akaike information criterion, the damage size distribution is within the class of the applicable probability distributions best represented by the geometric distribution. Finally, pair-wise Kolmogorov-Smirnov tests in our data set show that the parameters of the distributions for the inter-arrival time and damage size, differed from one component to another (i.e., there is heterogeneity).

5.7.3 Illustration of optimal replacement policy

Figure 5.5 shows two examples of the optimal replacement policy, applied a-posteriori to two time-series of the filament data set. In these examples, and also throughout the rest of this section, the ratio $\frac{c_u}{c_p}$ is set to 5 based on discussions with Philips Healthcare. Furthermore, the prior values of the hyperparameters are estimated using the MLE method described in Appendix 5.C applied on the remaining time-series, which resulted in $\alpha_0 = 44.88$, $\beta_0 = 32.43$, $a_0 = 4.29$, and $b_0 = 4.76$ for this example. Figure 5.5 visually confirms that there is heterogeneity in both the shock arrival

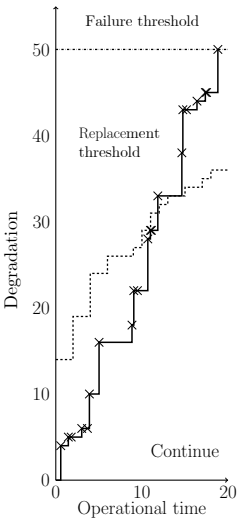

(a) Replacement when $(x, n, t) = (29, 14, 12)$ with $\delta^{(14,12)} = 29$.

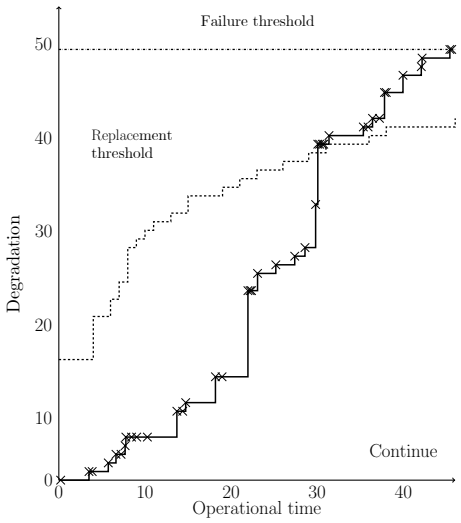

(b) Replacement when $(x, n, t) = (39, 28, 31)$ with $\delta^{(28,31)} = 39$.

Figure 5.5 Examples of optimal replacement policy applied to IXR filament degradation paths.

process and damage size distribution. In this example, the optimal replacement policy prescribes to preventively replace the X-ray tube at $x = 29$ and $n = 14$ at $t = 12$ (left)

and $x = 39$ and $n = 28$ at $t = 31$ (right), respectively, such that the useful lifetime of the X-ray tube is best utilized. The examples illustrate how integration of learning and decision making allows for the X-ray tube to be replaced early when necessary and late when possible.

5.7.4 Bootstrapping study

The goal of this section is to illustrate the optimal replacement policy, and to assess the performance of the integrated Bayes approach compared to both the offline and myopic online approach described in Section 5.5 on real-life data. The main performance metric is therefore the relative cost savings that can be attained by using the integrated Bayes approach instead of the two heuristic approaches commonly used in practice and described in Section 5.5. More formally, we are interested in

$$\%SAV_\pi = 100 \cdot (C_\pi - C_{\pi_\mathcal{I}})/C_\pi,$$

where C_π is the average cost rate of approach $\pi \in \{\pi_\mathcal{M}, \pi_\mathcal{N}\}$. In addition, we are interested in the impact of the amount of available historical degradation data on the $\%SAV_\pi$ of the two heuristic approaches.

We evaluate the performance metric $\%SAV_\pi$ retroactively on the data set described in Section 5.7.2. To evaluate the the impact of the amount of historical data available we bootstrap the amount of available data. That is, we first sample with replacement $s \in \{5, 10, 15\}$ time-series from the data set $\mathcal{I}$. We then use these time-series to estimate the required parameters for the integrated Bayes approach as well as for the heuristic approaches using MLE (see Appendix 5.C for further details regarding the MLE procedure). We then implement and evaluate both the integrated Bayes approach and the heuristic approaches on the remaining time-series that were not used for the estimation. This procedure is repeated 150 times per choice of s, such that the confidence intervals on the average cost rates are sufficiently small, resulting in 450 bootstrap instances. The resulting average values for $\%SAV_\pi$ are reported in Table 5.5. Observe that this bootstrapping study indicates the induced savings when the integrated Bayes approach is implemented instead of the offline approach or the myopic online approach in a real-life setting.

Table 5.5 shows that the integrated Bayes approach reduces the average cost rate with 10.50% on average compared to the state-of-the-art approach. Moreover, cost savings of 4.28% can be attained by integrating the learning with the decision making instead of using a data-driven approach that does not integrate the two. Moreover, the attainable savings are not significantly influenced by the amount of available historical degradation data. Hence, the integrated Bayes approach does not need a

Table 5.5 Results of bootstrapping study.

	$\%SAV_\pi$	
Number of time series	$\pi_{\mathcal{M}}$	$\pi_{\mathcal{N}}$
5	4.68	10.46
10	4.08	10.35
15	4.07	10.70
Total	4.28	10.50

large amount of historical data to perform well. These results highlight the value of the proposed method to integrate learning and decision making in a real-life setting.

5.8. Conclusion

In this chapter, we have considered the condition based maintenance of components that are subject to compound Poisson degradation, where the compounding distribution is a member of the one-parameter non-negative exponential family. The cost to replace a component after failure is much higher than before failure, because it includes the costs of unplanned downtime. We have assumed that the population of components is heterogeneous. That is, the degradation parameters of components may vary from one component to the other.

Since the degradation parameters of an individual component cannot be observed directly and need to be learned by observing its degradation signal, we have modeled this replacement problem as a partially observable Markov decision process. The entire past degradation path of a component is relevant state information in this setting, which can lead to tractability issues. We have shown that we can circumvent these issues by using conjugate prior pairs to model the heterogeneity of the component population. This allowed us to collapse a high dimensional state space to a 3 dimensional state space while retaining all relevant information. This collapse enabled us to tractably compute optimal policies as well as to characterize the optimal replacement policy.

We have characterized the optimal replacement policy as a threshold replacement policy, where the threshold is increasing in the age of a component. Furthermore, we have shown that the threshold also depends on the volatility of the observed degradation signal of a component. This volatility can only be captured by observing the *entire* degradation signal in real-time. Although we have assumed geometric compounding throughout this chapter, we show that these structural properties hold for a wide variety of compounding distributions.

We performed a comprehensive simulation study to assess (i) the benefits of explicitly modeling heterogeneity, (ii) the value of integrating learning with decision making, and (iii) the impact of the amount of available historical degradation data for model calibration on their performance. The results of this simulation study indicate that the integration of learning and decision making leads to excellent results with optimality gaps of only 0.60% on average. By contrast, ignoring heterogeneity leads to average optimality gaps of 15.02% while failing to integrate learning with decision making leads to average optimality gaps of 7.08%. Furthermore, we have shown that models that ignore population heterogeneity do *not* perform appreciably better when the amount of historical degradation data for model calibration increases.

Finally, we have established the practical value of integrated learning and decision making based on a real-life data set of filament wear in an interventional X-ray machine. We find that integrated learning can save up to 10.50% compared to approaches without learning and up to 4.28% compared to an approach where learning is separated from decision making. This normalized data set is made available to benchmark future approaches to perform maintenance on degrading systems.

5.A. Proofs

Proof of Proposition 5.1. The joint posterior distribution of Λ and Φ at component age t is proportional to the product of the joint likelihood function and the joint prior distributions on Λ and Φ at component age $t-1$. The joint likelihood of observing $\boldsymbol{\theta}_t$ given (λ, ϕ), denoted by $\mathcal{L}(\boldsymbol{\theta}_t|\lambda,\phi)$, is equal to

$$\begin{aligned}\mathcal{L}(\boldsymbol{\theta}_t|\lambda,\phi) &:= \mathbb{P}\Big[K_{(t-1,t]} = k_t, \boldsymbol{Y}_t = \boldsymbol{y}_t|\Lambda=\lambda,\Phi=\phi\Big] \\ &= \frac{\lambda^{k_t}e^{-\lambda}}{k_t!}\prod_{i=1}^{k_t}\Big[h(y_t^i)e^{\phi y_t^i - A(\phi)}\Big] \\ &= \frac{\lambda^{k_t}e^{-\lambda}}{k_t!}e^{\phi\sum_{i=1}^{k_t}y_t^i - k_t\cdot A(\phi)}\prod_{i=1}^{k_t}h(y_t^i).\end{aligned}$$

where $\prod_{i=1}^{0}\cdot \equiv 1$. This yields

$$\begin{aligned}&f_{\Lambda,\Phi}(\lambda,\phi|\boldsymbol{\theta}_0,\ldots,\boldsymbol{\theta}_t) \\ &\propto \mathcal{L}(\boldsymbol{\theta}_t|\lambda,\phi)\cdot f_\Lambda(\lambda|\alpha_{t-1},\beta_{t-1})\cdot f_\Phi(\phi|a_{t-1},b_{t-1}) \\ &= \frac{\lambda^{k_t}e^{-\lambda}}{k_t!}e^{\phi\sum_{i=1}^{k_t}y_t^i - k_t\cdot A(\phi)}\prod_{i=1}^{k_t}h(y_t^i)\cdot\frac{\beta_{t-1}^{\alpha_{t-1}}\lambda^{\alpha_{t-1}-1}e^{-\beta_{t-1}\lambda}}{\Gamma(\alpha_{t-1})}\times \\ &\qquad\qquad H(a_{t-1},b_{t-1})e^{a_{t-1}\phi - b_{t-1}A(\phi)} \\ &\propto \lambda^{\alpha_{t-1}+k_t-1}e^{-(\beta_{t-1}+1)\lambda}H(a_{t-1},b_{t-1})e^{\left(a_{t-1}+\sum_{i=1}^{k_t}y_t^i\right)\phi - \left(b_{t-1}+k_t\right)A(\phi)},\end{aligned}$$

which is after normalization (over hyperparameters) equal to

$$\begin{aligned}&\frac{(\beta_{t-1}+1)^{(\alpha_{t-1}+k_t)}\lambda^{\alpha_{t-1}+k_t-1}e^{-(\beta_{t-1}+1)\lambda}}{\Gamma(\alpha_{t-1}+k_t)}\times \\ &\qquad H(a_{t-1}+\sum_{i=1}^{k_t}y_t^i, b_{t-1}+k_t)e^{\left(a_{t-1}+\sum_{i=1}^{k_t}y_t^i\right)\phi - \left(b_{t-1}+k_t\right)A(\phi)}, \qquad (5.8)\end{aligned}$$

where $\Gamma(\cdot)$ denotes the gamma function and $H(a_{t-1}+\sum_{i=1}^{k_t}y_t^i, b_{t-1}+k_t)$ is the new normalization factor with the updated hyperparameters. Observe that the joint posterior distribution of Λ and Φ in Equation (5.8) is equal to the product of a Gamma$(\alpha_{t-1}+k_t, \beta_{t-1}+1)$ distribution and the general prior with updated hyperparameters $(a_{t-1}+\sum_{i=1}^{k_t}y_t^i, b_{t-1}+k_t)$ for a member of the exponential family, which completes the proof. □

Proof of Lemma 5.1. We first consider the posterior predictive distribution conditioned on $K_{(t,t+1]}$,

$$\mathbb{P}\Big[Z_{(t,t+1]} = z \Big| K_{(t,t+1]} = k, \alpha_t, \beta_t, a_t, b_t\Big] = \int_0^\infty f^{(k)}(z|\,\Phi = \phi) f_\Phi(\phi|a_t, b_t) d\phi, \quad (5.9)$$

where $f^{(k)}(z|\,\Phi = \phi)$ denotes the k-fold convolution of probability density (or mass) function of the random variable $\{Z_{(t,t+1]}|\,\Phi = \phi\}$. The distribution of $K_{(t,t+1]}$ is a continuous mixture of Poisson distributions where the mixing distribution of the Poisson rate follows a Gamma(α_t, β_t) distribution, which is known to be the negative binomial distribution with $p = \frac{1}{\beta_t+1}$ and $r = \alpha_t$. Hence, we have

$$\mathbb{P}\Big[K_{(t,t+1]} = k \Big| \alpha_t, \beta_t\Big] = \binom{k + \alpha_t - 1}{k} \left(\frac{1}{\beta_t + 1}\right)^k \left(\frac{\beta_t}{\beta_t + 1}\right)^{\alpha_t}. \quad (5.10)$$

Unconditioning Equation (5.9) using Equation (5.10) yields the desired result. □

Proof of Proposition 5.2. In this proof, we use the following two definitions quantifying the concept of one random variable being "bigger" than another random variable.

Definition 5.1 (1.A.1 Definition, Shaked and Shanthikumar (2007))**.** A random variable X is stochastically larger than a random variable Y in the usual stochastic order, denoted by $X \geq_{\text{st}} Y$, if and only if

$$\mathbb{P}[X \geq x] \geq \mathbb{P}[Y \geq x], \quad \text{for all } x \in \mathbb{R}. \quad (5.11)$$

Definition 5.1 implies that random variable X is less likely than random variable Y to take on large values, where large means any value greater than x, and that this holds for all $x \in \mathbb{R}$.

Definition 5.2 (1.C.1 Definition, Shaked and Shanthikumar (2007))**.** Let X and Y be continuous (discrete) random variables with probability densities f and g, respectively, such that

$$\frac{g(t)}{f(t)} \text{ increases in } t \text{ over the union of the supports of } X \text{ and } Y$$

(here $a/0$ is taken to be equal to ∞ whenever $a > 0$), or, equivalently,

$$f(x)g(y) \geq f(y)g(x), \quad \text{for all } x \leq y.$$

Then X is said to be smaller than Y in the likelihood ratio order (denoted by $X \leq_{\text{lr}} Y$). Note that the likelihood ratio order is stronger than the usual stochastic order, as such,

if $X \leq_{\text{lr}} Y$ then $X \leq_{\text{st}} Y$, cf. (Shaked and Shanthikumar, 2007, Theorem 1.C.1).

We now proceed with the proofs of parts (i) and (ii) of Proposition 5.2:

(i) We consider two different component ages, t^+ and t^- ($t^+ > t^-$), with the same degradation level x and the same total number of shocks, n, received by the component. By Proposition 5.1, we have $\beta_{t^+} > \beta_{t^-}$, whereas all other hyperparameters (i.e., α_t, a_t and b_t) remain fixed when only t changes. Note that

$$\mathbb{P}[K_{(t,t+1]} = u \,|\, \Lambda = \lambda] = e^{-\lambda}\frac{\lambda^u}{u!} \quad \text{and} \quad \mathbb{P}[\Lambda = \lambda \,|\, \alpha, \beta] = \frac{\beta^\alpha}{\Gamma(\alpha)}\lambda^{\alpha-1}e^{-\beta\lambda}.$$

We can then show that the random variables $\Lambda(\alpha, \beta) = \{\Lambda \,|\, \alpha, \beta\}$ are stochastically increasing in α and stochastically decreasing in β. This easily follows from the appropriate likelihood ratio.

We now show that the random variables $\mathcal{K}(\lambda) = \{K_{(t,t+1]} \,|\, \Lambda = \lambda\}$ are stochastically increasing in λ, i.e., $\lambda \leq \lambda'$ implies $\mathcal{K}(\lambda) \leq_{\text{st}} \mathcal{K}(\lambda')$. Consider the likelihood ratio

$$\frac{\mathbb{P}[\mathcal{K}(\lambda') = u]}{\mathbb{P}[\mathcal{K}(\lambda) = u]} = \frac{\mathbb{P}[K_{(t,t+1]} = u \,|\, \Lambda = \lambda']}{\mathbb{P}[K_{(t,t+1]} = u \,|\, \Lambda = \lambda]} = e^{-(\lambda'-\lambda)}\left(\frac{\lambda'}{\lambda}\right)^u.$$

It is immediately evident that for $\lambda \leq \lambda'$, the ratio is increasing in u. This yields the likelihood ratio order and as a consequence the usual stochastic order.

All in all, for $t^+ > t^-$, $\beta_{t^+} > \beta_{t^-}$, thus $\Lambda(\alpha, \beta_{t^+}) \leq_{\text{st}} \Lambda(\alpha, \beta_{t^-})$. From the stochastic monotonicity of $\mathcal{K}(\Lambda)$, the above yields $\mathcal{K}^{(t^+)} := \mathcal{K}(\Lambda(\alpha, \beta_{t^+})) \leq_{\text{st}} \mathcal{K}(\Lambda(\alpha, \beta_{t^-})) := \mathcal{K}^{(t^-)}$, cf. (Shaked and Shanthikumar, 2007, Theorem 1.A.2.). Finally, as the parameters associated with the random variables Y_i remain fixed, it is evident that

$$Z_{(t^+,t^++1]} = \sum_{i=1}^{\mathcal{K}^{(t^+)}} Y_i \leq_{\text{st}} \sum_{i=1}^{\mathcal{K}^{(t^-)}} Y_i = Z_{(t^-,t^-+1]},$$

cf. (Shaked and Shanthikumar, 2007, Theorem 1.A.4.). In conclusion, for $\boldsymbol{s} = (x, n, t)$, it follows that $Z(\boldsymbol{s})$ is stochastically decreasing in t.

(ii) We consider two degradation levels x^+ and x^- ($x^+ > x^-$) at the same component age t when we have observed the same total number of shocks n. By Proposition 5.1, we have $a_{x^+} > a_{x^-}$, whereas all other hyperparameters are equivalent as they remain fixed when only x changes.

It is well-known that a member of the one-parameter exponential family is likelihood-ratio ordered (and thus stochastically ordered in the usual stochastic order) according to its parameter (see e.g., Karlin and Rubin, 1956; Bapat and Kochar, 1994). Hence, we know that i) the random variables $\mathcal{Y}(\phi) = \{Y \mid \Phi = \phi\}$ are stochastically non-decreasing in ϕ, i.e., $\phi \leq \phi'$ implies $\mathcal{Y}(\phi) \leq_{\text{st}} \mathcal{Y}(\phi')$, and ii) the random variables $\Phi(a, b) = \{\Phi \mid a, b\}$ are stochastically non-decreasing in a, i.e., $a \leq a'$ implies $\Phi(a, b) \leq_{\text{st}} \Phi(a', b)$. Then, by Theorem 6 of Huang and Mi (2020), which relates the stochastic order of a posterior predictive distribution with the stochastic order of the posterior and the corresponding conditional distribution, we can conclude that $\mathcal{Y}^{(x^+)} := \mathcal{Y}(\Phi(a_{x^+}, b)) \geq_{\text{st}} \mathcal{Y}(\Phi(a_{x^-}, b)) =: \mathcal{Y}^{(x^-)}$. Finally, as the parameters associated with the random variables $K_{(t,t+1]}$ remain fixed, it is evident that

$$Z^{(x^+)}_{(t,t+1]} := \sum_{i=1}^{K_{(t,t+1]}} \mathcal{Y}_i^{(x^+)} \geq_{\text{st}} \sum_{i=1}^{K_{(t,t+1]}} \mathcal{Y}_i^{(x^-)} =: Z^{(x^+)}_{(t,t+1]}.$$

In conclusion, for $\boldsymbol{s} = (x, n, t)$, it follows that $Z(\boldsymbol{s})$ is stochastically increasing in x.

□

Proof of Lemma 5.2. We prove part (i) and omit the proof of part (ii) as its proof structure follows verbatim.

For $\boldsymbol{s} = (x, n, t) \in \Psi \times \mathbb{N}_0^2$, let $V^m(\boldsymbol{s})$ denote the value function at the m-th iteration of the value iteration algorithm, so that the value iteration algorithm produces the sequence $\{V^m(\boldsymbol{s})\}_{m \in \mathbb{N}_0}$. We use induction on the steps of the value iteration algorithm as a proof technique. Since our state space is countable, $\gamma \in (0, 1)$, and costs are bounded from above, the value function is guaranteed to converge point-wise to the optimal value function that satisfies Equation (5.5) (i.e., $V^m(\boldsymbol{s}) \to V(\boldsymbol{s})$ for all $\boldsymbol{s} \in \Psi \times \mathbb{N}_0^2$ as $m \to \infty$) from any arbitrary starting position through the value iteration algorithm (Bertsekas, 2007, Proposition 1.2.1).

For $\boldsymbol{s} = (x, n, t) \in \Psi \times \mathbb{N}_0^2$, we set $V^0(\boldsymbol{s}) = 0$. Note that $V^0(\boldsymbol{s})$ is non-increasing in t. We assume that the theorem holds for the m-th iteration, i.e., $V^m(\boldsymbol{s})$ is non-increasing in t. Then according to Equation (5.5), we have

$$V^{m+1}(\boldsymbol{s}) = \begin{cases} c_u + \gamma \mathbb{E}\left[V^m(\boldsymbol{s}_0 + A(\boldsymbol{s}_0))\right], & \text{if } x = \xi, \\ \min\Big\{c_p + \gamma \mathbb{E}\left[V^m(\boldsymbol{s}_0 + A(\boldsymbol{s}_0))\right]; \gamma \mathbb{E}_{\boldsymbol{s}}\left[V^m(\boldsymbol{s} + A(\boldsymbol{s}))\right]\Big\}, & \text{if } x < \xi. \end{cases} \tag{5.12}$$

Since $V^m(\boldsymbol{s})$ is non-increasing in t by the induction hypothesis and the random variable $A(\boldsymbol{s})$ is stochastically decreasing in t in the usual stochastic order (cf. Proposition 5.2 and the proof therein), the expectation $\mathbb{E}_{\boldsymbol{s}}\left[V^m(\boldsymbol{s}+A(\boldsymbol{s}))\right]$ holds this property as well (cf. Shaked and Shanthikumar, 2007, Theorem 1.A.3). Because the terms of the right-hand side of Equation (5.12) are non-increasing in t, $V^{m+1}(\boldsymbol{s})$ is also non-increasing in t. Due to the point-wise convergence in which the structure of $V^m(\boldsymbol{s})$ is preserved, we may conclude that $V(\boldsymbol{s})$ is also non-increasing in t (cf. Puterman, 2005, Theorem 6.3.1.). □

Proof of Theorem 5.1. Preventive replacement is optimal when the following equation holds

$$c_p + \gamma\mathbb{E}\left[V(A(\boldsymbol{s}_0))\right] \leq \gamma\mathbb{E}_{\boldsymbol{s}}\left[V(\boldsymbol{s}+A(\boldsymbol{s}))\right]. \tag{5.13}$$

Since the expectation $\mathbb{E}\left[V(A(\boldsymbol{s}_0))\right]$ is constant with respect to x, the left-hand side of Inequality (5.13) is constant with respect to x. Based on part (ii) of Lemma 5.2, the right-hand side of Inequality (5.13) is non-decreasing in x. Hence, if the optimal decision is to carry out preventive replacement in state $(\delta^{(n,t)}, n, t)$, then the same decision is optimal for any state (x, n, t) with $x \geq \delta^{(n,t)}$, which implies the control limit policy.

Similarly, the right-hand side of Inequality (5.13) is non-increasing in t by part (i) of Lemma 5.2. Hence, if it optimal to carry out a preventive replacement in state $(\delta^{(n,t)}, n, t)$, then it is optimal to carry out a preventive replacement in any state $(\delta^{(n,t)}, n, t')$ with $t' < t$, which implies that $\delta^{(n,t)}$ is monotonically non-decreasing in t. □

5.B. Deriving moments for the compound Poisson process

We first repeat some notation. Let $\boldsymbol{s} = (x, n, t) \in \Psi \times \mathbb{N}_0^2$ and let $Z(\boldsymbol{s}) = \sum_{i=1}^{K(\boldsymbol{s})} Y_i(\boldsymbol{s})$, where $K(\boldsymbol{s})$ is a Poisson random variable unknown rate λ, and $\{Y_i(\boldsymbol{s})\}_{i\in\mathbb{N}}$ is a sequence of independent and (identically) geometrically distributed random variables with unknown success probability p, independent from $K(\boldsymbol{s})$. We endow p and λ with prior beta and gamma distributions, denoted by P and Λ, respectively. We can then find the first and second moment of $Z(\boldsymbol{s})$ by repeatedly using the tower property of expectations and the law of total variance.

We have

$$\begin{aligned}\mathbb{E}[Y_i(\boldsymbol{s})] &= \mathbb{E}\big[\mathbb{E}[Y_i(\boldsymbol{s}) \mid P]\big] \\ &= \mathbb{E}\big[\frac{P}{1-P}\big] \\ &= \frac{\int_0^1 \frac{p}{1-p} p^{a_t-1}(1-p)^{b_t-1} dp}{\mathrm{B}(a_t, b_t)} \\ &= \frac{\mathrm{B}(a_t+1, b_t-1)}{\mathrm{B}(a_t, b_t)} \\ &= \frac{a_t}{b_t-1}, \end{aligned} \tag{5.14}$$

where the fourth equality follows from manipulating the term inside the integral to the density of another beta distribution.

Proceeding with $\mathbb{E}[K(\boldsymbol{s})]$, we have

$$\begin{aligned}\mathbb{E}[K(\boldsymbol{s})] &= \mathbb{E}\big[\mathbb{E}[K(\boldsymbol{s}) \mid \Lambda]\big] \\ &= \mathbb{E}[\Lambda] \\ &= \frac{\alpha_t}{\beta_t}. \end{aligned} \tag{5.15}$$

Combining (5.14) and (5.15) and using the updating rules of Proposition 5.1 for $\boldsymbol{s} = (x, n, t)$, we have

$$\begin{aligned}\mathbb{E}[Z(\boldsymbol{s})] &= \mathbb{E}\big[\sum\nolimits_{i=1}^{K(\boldsymbol{s})} Y_i(\boldsymbol{s})\big] \\ &= \mathbb{E}[K(\boldsymbol{s})]\mathbb{E}[Y_i(\boldsymbol{s})] \\ &= \frac{a_0+x}{b_0+n-1}\frac{\alpha_0+n}{\beta_0+t}. \end{aligned} \tag{5.16}$$

For the second moment, we first compute the variance. We have

$$\begin{aligned}\mathrm{Var}[Z(\boldsymbol{s})] &= \mathrm{Var}\big[\sum\nolimits_{i=1}^{K(\boldsymbol{s})} Y_i(\boldsymbol{s})\big] \\ &= \mathbb{E}\Big[\mathrm{Var}\Big[\sum\nolimits_{i=1}^{K(\boldsymbol{s})} Y_i(\boldsymbol{s}) \mid K(\boldsymbol{s})\Big]\Big] + \mathrm{Var}\Big[\mathbb{E}\sum\nolimits_{i=1}^{K(\boldsymbol{s})} Y_i(\boldsymbol{s}) \mid K(\boldsymbol{s})\Big]\Big] \\ &= \mathbb{E}\big[K(\boldsymbol{s})\big]\mathrm{Var}\big[Y_i(\boldsymbol{s})\big] + \mathrm{Var}\big[K(\boldsymbol{s})\big]\mathbb{E}\big[Y_i(\boldsymbol{s})\big] \\ &= \frac{\alpha_t}{\beta_t}\mathrm{Var}\big[Y_i(\boldsymbol{s})\big] + \frac{a_t}{b_t-1}\mathrm{Var}\big[K(\boldsymbol{s})\big], \end{aligned} \tag{5.17}$$

where $\mathbb{E}\big[Y_i(\boldsymbol{s})\big]$ and $\mathbb{E}\big[K(\boldsymbol{s})\big]$ are computed in (5.14) and (5.15), respectively. We now

proceed with the two variances in (5.17). For $\mathrm{Var}\big[Y_i(\boldsymbol{s})\big]$ we have

$$\begin{aligned}\mathrm{Var}\big[Y_i(\boldsymbol{s})\big] &= \mathbb{E}\Big[\mathrm{Var}\big[Y_i(\boldsymbol{s}) \mid P\big]\Big] + \mathrm{Var}\Big[\mathbb{E}\big[Y_i(\boldsymbol{s}) \mid P\big]\Big] \\ &= \mathbb{E}\Big[\frac{P}{(1-P)^2}\Big] + \mathrm{Var}\Big[\frac{P}{1-P}\Big]. \end{aligned} \tag{5.18}$$

For $\mathbb{E}\Big[\frac{P}{(1-P)^2}\Big]$, we have

$$\begin{aligned}\mathbb{E}\Big[\frac{P}{(1-P)^2}\Big] &= \frac{\int_0^1 \frac{p}{(1-p)^2} p^{a_t-1}(1-p)^{b_t-1}dp}{\mathrm{B}(a_t,b_t)} \\ &= \frac{\mathrm{B}(a_t+1,b_t-2)}{\mathrm{B}(a_t,b_t)} \\ &= \frac{a_t(a_t+b_t-1)}{(b_t-2)(b_t-1)}, \end{aligned} \tag{5.19}$$

where the second equality follows from manipulating the term inside the integral to the density of another beta distribution. Similarly, we can write

$$\begin{aligned}\mathrm{Var}\Big[\frac{P}{1-P}\Big] &= \mathbb{E}\Big[\frac{P^2}{(1-P)^2}\Big] - \mathbb{E}^2\Big[\frac{P}{(1-P)}\Big] \\ &= \frac{\mathrm{B}(a_t+2,b_t-2)}{\mathrm{B}(a_t,b_t)} - \Big(\frac{a_t}{b_t-1}\Big)^2 \\ &= \frac{a_t(a_t+1)}{(b_t-2)(b_t-1)} - \Big(\frac{a_t}{b_t-1}\Big)^2, \end{aligned} \tag{5.20}$$

which completes all the ingredients for $\mathrm{Var}\big[Y_i(\boldsymbol{s})\big]$. Using again the law of total variance, we can compute $\mathrm{Var}\big[K(\boldsymbol{s})\big]$ as follows

$$\begin{aligned}\mathrm{Var}\big[K(\boldsymbol{s})\big] &= \mathbb{E}\Big[\mathrm{Var}\big[K(\boldsymbol{s}) \mid \Lambda\big]\Big] + \mathrm{Var}\Big[\mathbb{E}\big[K(\boldsymbol{s}) \mid \Lambda\big]\Big] \\ &= \mathbb{E}\Big[\Lambda\Big] + \mathrm{Var}\Big[\Lambda\Big] \\ &= \frac{\alpha_t}{\beta_t} + \frac{\alpha_t}{\beta_t^2}. \end{aligned} \tag{5.21}$$

Plugging (5.19) and (5.20) in (5.18), and combining this with (5.21) in (5.17) gives us $\mathrm{Var}[Z(\boldsymbol{s})]$. We can then use the variance expansion $\mathrm{Var}[Z(\boldsymbol{s})] = \mathbb{E}[Z(\boldsymbol{s})^2] - \mathbb{E}^2[Z(\boldsymbol{s})]$ to obtain an explicit expression for the second moment in terms of the hyperparameters associated with state $\boldsymbol{s} = (x, n, t)$.

5.C. Maximum likelihood estimation

Let $\mathcal{I}$ denote the set containing all simulated sample paths (without maintenance). For each sample path $i \in \mathcal{I}$, we let x_i, n_i, and t_i, denote, respectively, the total damage accumulated by the component, the number of critical loading epochs sustained by the component, and the number of decision epochs that the component was operated until its failure. Thus the tuple $\boldsymbol{s}_i = (x_i, n_i, t_i)$ contains all relevant information for sample path i. Because the prior distributions of Φ and Λ are independent, we next look at their respective likelihoods given the set of simulated degradation paths $\mathcal{I}$ separately.

From Equation (5.9) we know that the likelihood of the degradation paths in $\mathcal{I}$ being induced by components stemming from a population whose degradation parameter Φ follows a beta prior distribution with parameters a_0 and b_0 is given by

$$\mathcal{L}_\Phi(\boldsymbol{s}_1, \ldots, \boldsymbol{s}_{|\mathcal{I}|} \mid a_0, b_0) = \prod_{i=1}^{|\mathcal{I}|} \left[\frac{B(n_i + a_0, x_i + b_0)}{B(a_0, b_0)} \binom{x_i + n_i - 1}{x_i} \right].$$

Observe that the number of critical loading epochs sustained by the i-th component of age t_i has a Poisson distribution with parameter $t_i\Lambda$, where $\Lambda \sim \Gamma(\alpha_0, \beta_0)$. By the scaling property of the gamma distribution, this number is therefore a continuous mixture of Poisson distributions where the mixing distribution of the Poisson rate follows a $\Gamma(\alpha_0, \beta_0/t_i)$ distribution, which is known to be a negative binomial (or Pascal) distribution with parameters $r = \alpha_0$ and $p = \frac{t_i}{\beta_0 + t_i}$. Hence, the likelihood of the degradation paths in $\mathcal{I}$ being induced by components stemming from a population whose degradation parameter Λ follows a gamma prior distribution with parameters α_0 and β_0 is given by

$$\mathcal{L}_\Lambda(\boldsymbol{s}_1, \ldots, \boldsymbol{s}_{|\mathcal{I}|} \mid \alpha_0, \beta_0) = \prod_{i=1}^{|\mathcal{I}|} \left[\binom{n_i + \alpha_0 - 1}{n_i} \left(\frac{t_i}{\beta_0 + t_i} \right)^{n_i} \left(\frac{\beta_0}{\beta_0 + t_i} \right)^{\alpha_0} \right].$$

We are interested in maximizing both likelihoods. For convenience, we maximize the log likelihoods. That is,

$$\underset{\alpha_0, \beta_0, a_0, b_0}{\arg\max} \ \ln \mathcal{L}_\Lambda(\boldsymbol{s}_1, \ldots, \boldsymbol{s}_{|\mathcal{I}|} \mid \alpha_0, \beta_0) + \ln \mathcal{L}_\phi(\boldsymbol{s}_1, \ldots, \boldsymbol{s}_{|\mathcal{I}|} \mid a_0, b_0),$$

which is a nonlinear multidimensional maximization problem that can be solved using standard numerical methods (e.g., the Nelder-Mead method).

Bibliography

K. Abnett. World's largest carbon market faces revamp under draft eu plan. *Reuters*, 2020.

I.J.B.F. Adan, A. Sleptchenko, and G.J.J.A.N. Van Houtum. Reducing costs of spare parts supply systems via static priorities. *Asia-Pacific Journal of Operational Research*, 26(04):559–585, 2009.

P.K. Aggarwal and K. Moinzadeh. Order expedition in multi-echelon production/distribution systems. *IIE transactions*, 26(2):86–96, 1994.

S. Alaswad and Y. Xiang. A review on condition-based maintenance optimization models for stochastically deteriorating system. *Reliability Engineering & System Safety*, 157:54–63, 2017.

M. Albano, E. Jantunen, G. Papa, and U. Zurutuza. *The MANTIS Book: Cyber Physical System Based Proactive Collaborative Maintenance*. River Publishers Series in Automation, Control and Robotics, 2019.

P. Alfredsson and J. Verrijdt. Modeling emergency supply flexibility in a two-echelon inventory system. *Management Science*, 45(10):1416–1431, 1999.

G. Allon and J.A. Van Mieghem. Global dual sourcing: Tailored base-surge allocation to near-and offshore production. *Management Science*, 56(1):110–124, 2010.

E.M. Alvarez and M.C. Van der Heijden. On two-echelon inventory systems with poisson demand and lost sales. *European Journal of Operational Research*, 235(1): 334–338, 2014.

E.M. Alvarez, M.C. Van der Heijden, and W.H.M. Zijm. The selective use of emergency shipments for service-contract differentiation. *International Journal of Production Economics*, 143(2):518–526, 2013.

E.M. Alvarez, M.C. Van der Heijden, and W.H.M. Zijm. Service differentiation in spare parts supply through dedicated stocks. *Annals of Operations Research*, 231 (1):283–303, 2015.

E. Arıkan and W. Jammernegg. The single period inventory model under dual sourcing and product carbon footprint constraint. *International Journal of Production Economics*, 157:15–23, 2014.

J.J. Arts. A multi-item approach to repairable stocking and expediting in a fluctuating demand environment. *European Journal of Operational Research*, 256(1):102–115, 2017.

J.J. Arts, M. Van Vuuren, and G. Kiesmüller. Efficient optimization of the dual-index policy using markov chains. *IIE Transactions*, 43(8):604–620, 2011.

J.J. Arts, R.J.I. Basten, and G.J.J.A.N. Van Houtum. Repairable stocking and expediting in a fluctuating demand environment: Optimal policy and heuristics. *Operations Research*, 64(6):1285–1301, 2016.

S. Axsäter. Modelling emergency lateral transshipments in inventory systems. *Management Science*, 36(11):1329–1338, 1990.

K.S. Azoury. Bayes solution to dynamic inventory models under unknown demand distribution. *Management Science*, 31(9):1150–1160, 1985.

R.B. Bapat and S.C. Kochar. On likelihood-ratio ordering of order statistics. *Linear Algebra and Its Applications*, 199:281–291, 1994.

E.W. Barankin. A delivery-lag inventory model with an emergency provision (the single-period case). *Naval Research Logistics*, 8(3):285–311, 1961.

A.P. Barbosa-Póvoa, C. da Silva, and A. Carvalho. Opportunities and challenges in sustainable supply chain: An operations research perspective. *European Journal of Operational Research*, 268(2):399–431, 2018.

R.J.I. Basten and G.J.J.A.N. Van Houtum. System-oriented inventory models for spare parts. *Surveys in Operations Research and Management Science*, 19(1):34–55, 2014.

S. Batun and L.M. Maillart. Reassessing tradeoffs inherent to simultaneous maintenance and production planning. *Production and Operations Management*, 21(2):396–403, 2012.

Z. Benyamini and U. Yechiali. Optimality of control limit maintenance policies under nonstationary deterioration. *Probability in the Engineering and Informational Sciences*, 13(1):55–70, 1999.

P. Berling and V. Martínez-de Albéniz. Dynamic speed optimization in supply chains with stochastic demand. *Transportation Science*, 50(3):1114–1127, 2016.

D.P. Bertsekas. *Dynamic Programming and Optimal Control, Vol. II.* Athena Scientific, 2007.

M. Bijvank, W.T. Huh, G Janakiraman, and W. Kang. Robustness of order-up-to policies in lost-sales inventory systems. *Operations Research*, 62(5):1040–1047, 2014.

S. Bitton, I. Cohen, and M.A. Cohen. Joint repair sourcing and stocking policies for repairables using erlang-a and erlang-b queueing models. *IISE Transactions*, 51 (10):1151–1166, 2019.

R. Boute and J.A. Van Mieghem. Global dual sourcing and order smoothing: The impact of capacity and lead times. *Management Science*, 61(9):2080–2099, 2015.

M. Brandenburg, K. Govindan, J. Sarkis, and S. Seuring. Quantitative models for sustainable supply chain management: Developments and directions. *European Journal of Operational Research*, 233(2):299–312, 2014.

R. Brooks and A. Geoffrion. Finding everett's lagrange multipliers by linear programming. *Operations Research*, 14(6):1149–1153, 1966.

K.E. Caggiano, J.A. Muckstadt, and J.A. Rappold. Integrated real-time capacity and inventory allocation for reparable service parts in a two-echelon supply system. *Manufacturing & Service Operations Management*, 8(3):292–319, 2006.

L. Chen. Bounds and heuristics for optimal Bayesian inventory control with unobserved lost sales. *Operations Research*, 58(2):396–413, 2010.

L. Chen and E.L. Plambeck. Dynamic inventory management with learning about the demand distribution and substitution probability. *Manufacturing & Service Operations Management*, 10(2):236–256, 2008.

N. Chen, Z.S. Ye, Y. Xiang, and L. Zhang. Condition-based maintenance using the inverse Gaussian degradation model. *European Journal of Operational Research*, 243(1):190–199, 2015.

X. Chen and X. Wang. Effects of carbon emission reduction policies on transportation mode selections with stochastic demand. *Transportation Research Part E: Logistics and Transportation Review*, 90:196–205, 2016.

M.A. Cohen, P.V. Kamesam, P. Kleindorfer, H. Lee, and A. Tekerian. Optimizer: Ibm's multi-echelon inventory system for managing service logistics. *Interfaces*, 20 (1):65–82, 1990.

M.A. Cohen, Y.S. Zheng, and V. Agrawal. Service parts logistics: a benchmark analysis. *IIE transactions*, 29(8):627–639, 1997.

C. Coleman, S. Damodaran, M. Chandramouli, and E. Deuel. Making maintenance smarter: Predictive maintenance and the digital supply network. Technical report, Deloitte, 2017.

Price Waterhouse Coopers. Sensing the future of the internet of things. Digital IQ snapshot, 2014.

E.J. Covington. Hot spot burnout of tungsten filaments. *Journal of the Illuminating Engineering Society*, 2(4):372–380, 1973.

M. Dada. A two-echelon inventory system with priority shipments. *Management Science*, 38(8):1140–1153, 1992.

G.B. Dantzig and P. Wolfe. Decomposition principle for linear programs. *Operations Research*, 8(1):101–111, 1960.

R. Dekker, J. Bloemhof, and I. Mallidis. Operations research for green logistics – an overview of aspects, issues, contributions and challenges. *European Journal of Operational Research*, 219(3):671–679, 2012.

C. Derman. Optimal replacement and maintenance under markovian deterioration with probability bounds on failure. *Management Science*, 9(3):478–481, 1963.

Q. Dong, C.and Li, B. Shen, and X. Tong. Sustainability in supply chains with behavioral concerns. *Sustainability*, 11(15), 2019.

C. Drent, M. Drent, J.J. Arts, and S. Kapodistria. Real-time integrated learning and decision making for deteriorating systems. Working paper, 2021a.

M. Drent. Stocking and expediting in two-echelon spare parts inventory systems under system availability constraints. Master's thesis, Eindhoven University of Technology, 2017.

M. Drent and J.J. Arts. Expediting in two-echelon spare parts inventory systems. *Manufacturing & Service Operations Management*, Forthcoming, 2020.

M. Drent and J.J. Arts. Effective dual-sourcing through inventory projection. Working paper, 2021.

M. Drent, P. Moradi, and J.J. Arts. Efficient emission reduction through dynamic mode selection. Working paper, 2021b.

ECRI. Healthcare product comparison system. Technical report, ECRI Institute, Plymouth Meeting, PA, 2013.

A. Elwany, N. Gebraeel, and L.M. Maillart. Structured replacement policies for components with complex degradation processes and dedicated sensors. *Operations Research*, 59(3):684–695, 2011.

E. Engebrethsen and S. Dauzére-Pérés. Transportation mode selection in inventory models: A literature review. *European Journal of Operational Research*, 279(1): 1–25, 2019.

European Commission. State of the Union: Q & A on the 2030 Climate Target Plan. 2020.

European Environment Agency. *Trends and projections in Europe 2020: Tracking progress towards Europe's climate and energy targets*. 2020.

Q. Feng, S.P. Sethi, H. Yan, and H. Zhang. Are base-stock policies optimal in inventory problems with multiple delivery modes? *Operations Research*, 54(4): 801–807, 2006.

M. Fischetti, F. Glover, and A. Lodi. The feasibility pump. *Mathematical Programming*, 104(1):91–104, 2005.

K. Fu, X. Gong, and G. Liang. Managing perishable inventory systems with product returns and remanufacturing. *Production and Operations Management*, 28(6):1366–1386, 2019.

Y. Fukuda. Optimal policies for the inventory problem with negotiable leadtime. *Management Science*, 10(4):690–708, 1964.

G. Gallego, Ö. Özer, and P. Zipkin. Bounds, heuristics, and approximations for distribution systems. *Operations Research*, 55(3):503–517, 2007.

J.K. Ghosh, M. Delampady, and T. Samanta. *An Introduction to Bayesian Analysis: Theory and Methods*. Springer Science & Business Media, 2007.

J. Gijsbrechts, R. Boute, J.A. Van Mieghem, and D.J. Zhang. Can deep reinforcement learning improve inventory management? performance on dual sourcing, lost sales and multi-echelon problems. Available at SSRN: https://papers.ssrn.com/abstract=3302881, 2019.

D.A. Goldberg, D. Katz-Rogozhnikov, Y. Lu, M. Sharma, and M. Squillante. Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. *Mathematics of Operations Research*, 41(3):898–913, 2016.

D.A. Goldberg, M.I. Reiman, and Q. Wang. A survey of recent progress in the asymptotic analysis of inventory systems. *Production and Operations Management*, 30(6):1718–1750, 2021.

X. Gong, X. Chao, and S. Zheng. Dynamic pricing and inventory management with dual suppliers of different lead times and disruption risks. *Production and Operations Management*, 23(12):2058–2074, 2014.

J. Grahovac and A. Chakravarty. Sharing and lateral transshipment of inventory in a supply chain with expensive low-demand items. *Management Science*, 47(4): 579–594, 2001.

S.C. Graves. A multi-echelon inventory model for a repairable item with one-for-one

replenishment. *Management Science*, 31(10):1247–1256, 1985.

J.F. Gross, J.M. Thompson, and C.M. Harris. *Fundamentals of Queueing Theory*. Wiley-Interscience, New York, NY, USA, 4th edition, 2008.

Allan Gut. *An Intermediate Course in Probability*. Springer, second edition, 2009.

C.B. Haubitz and U.W. Thonemann. How to change a running system—controlling the transition to optimized spare parts inventory policies. *Production and Operations Management*, 30(5):1386–1405, 2021.

K.M.R. Hoen, T. Tan, J.C. Fransoo, and G.J.J.A.N. Van Houtum. Effect of carbon emission regulations on transport mode selection under stochastic demand. *Flexible Service Manufacturing Journal*, 26:170–195, 2014a.

K.M.R. Hoen, T. Tan, J.C. Fransoo, and G.J.J.A.N. Van Houtum. Switching transport modes to meet voluntary carbon emission targets. *Transportation Science*, 48(4):592–608, 2014b.

C. Howard, J. Marklund, T. Tan, and I. Rijnen. Inventory control in a spare parts distribution system with emergency stocks and pipeline information. *Manufacturing & Service Operations Management*, 17(12):142–156, 2015.

Z. Hua, Y. Yu, W. Zhang, and X. Xu. Structural properties of the optimal policy for dual-sourcing systems with general lead times. *IIE Transactions*, 47(8):841–850, 2015.

K. Huang and J. Mi. Applications of likelihood ratio order in bayesian inferences. *Probability in the Engineering and Informational Sciences*, 34(1):1–13, 2020.

W.T. Huh, G. Janakiraman, J.A. Muckstadt, and P. Rusmevichientong. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Science*, 55(3):404–420, 2009.

V. Insinna and G. Ziezuliwicz. Investigation blames air force and navy for systemic failures in fatal marine corps c-130 crash that killed 16, December 2018.

International Maritime Organization. Resolution MEPC.203(62) amendments to the annex of the protocol of 1997 to amend the international convention for the prevention of pollution from ships, 1973, as modified by the protocol of 1978 relating thereto (inclusion of regulations on energy efficiency for ships in MARPOL annex vi). 2011.

G. Janakiraman, S. Seshadri, and G. Shanthikumar. A comparison of the optimal costs of two canonical inventory systems. *Operations Research*, 55(5):866–875, 2007.

G. Janakiraman, S. Seshadri, and A. Sheopuri. Analysis of tailored base-surge policies in dual sourcing inventory systems. *Management Science*, 61(7):1547–1561, 2015.

G. Janakiraman, M. Nagarajan, and S. Veeraraghavan. Simple policies for managing flexible capacity. *Manufacturing & Service Operations Management*, 20(2):333–346, 2018.

F. Janssen and T. de Kok. A two-supplier inventory model. *International Journal of Production Economics*, 59(1):395–403, 1999.

A. Jardine, D. Lin, and D. Banjevic. A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7):1483–1510, 2006.

E. Kao. Optimal replacement rules when changes of state are semi-Markovian. *Operations Research*, 21(6):1231–1249, 1973.

S. Karlin and H. Rubin. Distributions possessing a monotone likelihood ratio. *Journal of the American Statistical Association*, 51(276):637–643, 1956.

H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack problems.* Springer, Berlin, 2004.

M.J. Kim. Robust control of partially observable failing systems. *Operations Research*, 64(4):999–1014, 2016.

M.J. Kim and V. Makis. Joint optimization of sampling and control of partially observable failing systems. *Operations Research*, 61(3):777–790, 2013.

S. Klosterhalfen, G. Kiesmüller, and S. Minner. A comparison of the constant-order and dual-index policy for dual sourcing. *International Journal of Production Economics*, 133(1):302 – 311, 2011.

P. Kolesar. Minimum cost replacement under Markovian deterioration. *Management Science*, 12(9):694–706, 1966.

Dinçer Konur, James F. Campbell, and Sepideh A. Monfared. Economic and environmental considerations in a stochastic inventory control model with order splitting under different delivery schedules among suppliers. *Omega*, 71:46–65, 2017.

A.A. Kranenburg. *Spare Parts Inventory Control Under System Availability Constraints.* PhD thesis, Eindhoven University of Technology, 2006.

A.A. Kranenburg and G.J.J.A.N. Van Houtum. Effect of commonality on spare parts provisioning costs for capital goods. *International Journal of Production Economics*, 108(1):221–227, 2007.

A.A. Kranenburg and G.J.J.A.N. Van Houtum. A new partial pooling structure for spare parts networks. *European Journal of Operational Research*, 199(3):908–921, 2009.

M. Kurt and J.P. Kharoufeh. Optimally maintaining a Markovian deteriorating

system with limited imperfect repairs. *European Journal of Operational Research*, 205(2):368–380, 2010.

S. Lam. Queuing networks with population size constraints. *IBM J. Res. Devel.*, 21: 370–378, 1977.

H.L. Lee. A multi-echelon inventory model for repairable items with emergency lateral transshipments. *Management Science*, 33(10):1302–1316, 1987.

N. Lemmens, J. Gijsbrechts, and R. Boute. Synchromodality in the physical internet – dual sourcing and real-time switching between transport modes. *European Transport Research Review*, 1(10):518–526, 2019.

R. Li and J.K. Ryan. A Bayesian inventory model using real-time condition monitoring information. *Production and Operations Management*, 20(5):754–771, 2011.

X. Liu, Q. Sun, Z.S. Ye, and M. Yildirim. Optimal multi-type inspection policy for systems with imperfect online monitoring. *Reliability Engineering & System Safety*, 207:107335, 2021.

N. Loeffen. Repair shop and inventory control for spare parts: min-max versus a lead time interface agreement. Master's thesis, Eindhoven University of Technology, 2012.

M.E. Lübbecke and J. Desrosiers. Selected topics in column generation. *Operations Research*, 53(6):1007–1023, 2005.

L.M. Maillart. Maintenance policies for systems with condition monitoring and obvious failures. *IIE Transactions*, 38:463–475, 2006.

V. Makis and X. Jiang. Optimal replacement under partial observations. *Mathematics of Operations Research*, 28(2):382–394, 2003.

E.J. Messerli. Proof of a convexity property of the erlang b formula. *The Bell System Technical Journal*, 51(4):951–953, 1972.

S. Minner. Multiple-supplier inventory models in supply chain management: A review. *International Journal of Production Economics*, 81:265–279, 2003.

S. Minner, E.B. Diks, and A.G. De Kok. A two-echelon inventory system with supply lead time flexibility. *IIE Transactions*, 35(2):117–129, 2003.

K. Moinzadeh and P.K. Aggarwal. An information based multiechelon inventory system with emergency orders. *Operations Research*, 45(5):694–701, 1997.

K. Moinzadeh and C.P. Schmidt. An (s-1, s) inventory system with emergency orders. *Operations Research*, 39(2):308–321, 1991.

C.N. Morris. Natural exponential families with quadratic variance functions. *The Annals of Statistics*, pages 65–80, 1982.

J.A. Muckstadt. A model for a multi-item, multi-echelon, multi-indenture inventory system. *Management Science*, 20(4):472–481, 1973.

J.A. Muckstadt. *Analysis and algorithms for service parts supply chains*. Springer Science & Business Media, 2005.

R.B. Nelsen. *An Introduction to Copulas*. Springer, second edition, 2006.

M. Neuts. An inventory model with an optional time lag. *Journal of the Society for Industrial and Applied Mathematics*, 12(1):179–185, 1964.

NTM. Methods and Manuals, 2015. URL `https://www.transportmeasures.org/en/wiki/manuals/`.

NTM. NTMCalc Basic 4.0. `https://www.transportmeasures.org/ntmcalc/v4/basic/index.html#/`, n.d. Accessed: 2021-7-20.

M.C.A. Olde Keizer, R.H. Teunter, and J. Veldman. Joint condition-based maintenance and inventory optimization for systems with multiple components. *European Journal of Operational Research*, 257(1):209 – 222, 2017.

R. Pacquin. Asset management: The changing landscape of predictive maintenance. AberdeenGroup, March 2014.

G. Palak, S.D. Ekşioğlu, and J. Geunes. Analyzing the impacts of carbon regulatory mechanisms on supplier and mode selection decisions: An application to a biofuel supply chain. *International Journal of Production Economics*, 154:198–216, 2014.

C. Paterson, G. Kiesmüller, R. Teunter, and K. Glazebrook. Inventory models with lateral transshipments: A review. *European Journal of Operational Research*, 210 (2):125–136, 2011.

Philips. Allura Xper FD 10 - DS Interventional X-ray system, 2020. URL `https://www.usa.philips.com/healthcare/product/HC889021/diamondselectalluraxperfd10refurbishedxraysytem`.

W.P. Pierskalla and J.A. Voelker. A survey of maintenance models: the control and surveillance of deteriorating systems. *Naval Research Logistics*, 23(3):353–388, 1976.

Ponemon. Cost of data center outages. Technical report, Ponemon Institute, 2016.

M.L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2005.

D.F. Pyke. Priority repair and dispatch policies for reparable-item logistics systems.

Naval Research Logistics, 37(1):1–30, 1990.

U. Rao, A. Scheller-Wolf, and S. Tayur. Development of a rapid-response supply chain at caterpillar. *Operations Research*, 48(2):189–204, 2000.

Y. Rong, Z. Atan, and L.V. Snyder. Heuristics for base-stock levels in multi-echelon distribution networks. *Production and Operations Management*, 26(9):1760–1777, 2017.

D. Rosenfield. Markovian deterioration with uncertain information. *Operations Research*, 24(1):141–155, 1976.

M. Rosenshine and D. Obee. Analysis of a standing order inventory system with emergency orders. *Operations Research*, 24(6):1143–1155, 1976.

H. Rosič and W. Jammernegg. The economic and environmental performance of dual sourcing: A newsvendor approach. *International Journal of Production Economics*, 143(1):109–119, 2013.

S.M. Ross. A Markovian replacement model with a generalization to include stocking. *Management Science*, 15(11):573–739, 1969.

W.D. Rustenburg. *A system approach to budget-constrained spare parts management.* PhD thesis, Eindhoven University of Technology, 2000.

W.D. Rustenburg, G.J.J.A.N. Van Houtum, and W.H.M. Zijm. Systeemgericht spare parts management bij de nederlandse koninklijke marine. *Bedrijfskunde*, 73(2): 28–39, 2001.

A. Sapra. Dual sourcing in a serial system. *Production and Operations Management*, 26(12):2163–2174, 2017.

P.A. Scarf. On the application of mathematical models in maintenance. *European Journal of Operational Research*, 99(3):493–506, 1997.

A. Scheller-Wolf, S. Veeraraghavan, and G.J.J.A.N. Van Houtum. Effective dual sourcing with a single index policy. Working paper, 2007.

M. Shaked and G.J. Shanthikumar. *Stochastic Orders.* Springer Science & Business Media, 2007.

A. Sheopuri, G. Janakiraman, and S. Seshadri. New policies for the stochastic inventory control problem with two supply sources. *Operations Research*, 58(3): 734–745, 2010.

C.C. Sherbrooke. Metric: A multi-echelon technique for recoverable item control. *Operations Research*, 16(1):122–141, 1968.

C.C. Sherbrooke. *Optimal Inventory Modeling of Systems: Multi-Echelon Techniques.*

Kluwer Academic Publishers, Norwell, MA, USA, 2004. ISBN 1402078498.

R.M. Simon. Stationary properties of a two-echelon inventory model for low demand items. *Operations Research*, 19(3):761–773, 1971.

N.D. Singpurwalla. Survival in dynamic environments. *Statistical Science*, pages 86–103, 1995.

V.W. Slaugh, B. Biller, and S.R. Tayur. Managing rentals with usage-based loss. *Manufacturing & Service Operations Management*, 18(3):429–444, 2016.

A. Sleptchenko, M.C. Van der Heijden, and A. Van Harten. Using repair priorities to reduce stock investment in spare part networks. *European Journal of Operational Research*, 163(3):733–750, 2005.

K. Sobczyk. Stochastic models for fatigue damage of materials. *Advances in Applied Probability*, 19(3):652–673, 1987.

J. Song, L. Xiao, H. Zhang, and P. Zipkin. Optimal policies for a dual-sourcing inventory problem with endogenous stochastic lead times. *Operations Research*, 65 (2):379–395, 2017.

J.S. Song and P. Zipkin. Inventories with multiple supply sources and networks of queues with overflow bypasses. *Management Science*, 55(3):362–372, 2009.

Wall Street Journal Custom Studios. How manufacturers achieve top quartile performance, 2017.

J. Sun and J.A. Van Mieghem. Robust dual sourcing inventory management: Optimality of capped dual index policies and smoothing. *Manufacturing & Service Operations Management*, 21(4):912–931, 2019.

J. Svoboda, S. Minner, and M Yao. Typology and literature review on multiple supplier inventory control models. *European Journal of Operational Research*, 293 (1):1–23, 2021.

D.J. Thomas and J.E. Tyworth. Pooling lead-time risk by order splitting: A critical review. *Transportation Research Part E: Logistics and Transportation Review*, 42 (4):245–257, 2006.

K. Threatte and S.C. Graves. Tactical shipping and scheduling at polaroid with dual lead-times. Proceedings SMA Conference, 2002.

H.G.H. Tiemessen and G.J.J.A.N. Van Houtum. Reducing costs of repairable inventory supply systems via dynamic scheduling. *International Journal of Production Economics*, 143(2):478–488, 2013.

E. Topan, Z.P. Bayındır, and T. Tan. Heuristics for multi-item two-echelon spare parts

inventory control subject to aggregate and individual service measures. *European Journal of Operational Research*, 256(1):126–138, 2017.

M.A.J. Uit het Broek, R.H. Teunter, B. De Jonge, J. Veldman, and N.D. Van Foreest. Condition-based production planning: Adjusting production rates to balance output and failure risk. *Manufacturing & Service Operations Management*, 22(4): 792–811, 2020.

United Nations Comtrade Database, 2020. URL `https://comtrade.un.org`.

United Nations Conference on Trade and Development. Review of maritime transport. 2020.

United Nations Environment Programme. *Emissions Gap Report 2020*. UNEP DTU Partnership, 2020.

C. Valdez-Flores and R.M. Feldman. A survey of preventive maintenance models for stochastically deteriorating single-unit systems. *Naval Research Logistics*, 36(4): 419–446, 1989.

G.J.J.A.N. Van Houtum and A.A. Kranenburg. *Spare parts inventory control under system availability constraints*, volume 227. Springer, 2015.

W. Van Jaarsveld and J.J. Arts. Projected inventory level policies for lost sales inventory systems. Working paper, 2021.

J.M. Van Noortwijk. A survey of the application of Gamma processes in maintenance. *Reliability Engineering & System Safety*, 94(1):2–21, 2009.

C. Van Oosterom, H. Peng, and G.J.J.A.N. Van Houtum. Maintenance optimization for a Markovian deteriorating system with population heterogeneity. *IISE Transactions*, 49(1):96–109, 2017.

H. Van Soest. Wanneer gaat het echt fout in doel?, 2016. Algeneen Dagblad (Dutch Newspaper), January 4.

S. Veeraraghavan and A. Scheller-Wolf. Now or later: A simple policy for effective dual sourcing in capacitated systems. *Operations Research*, 56(4):850–864, 2008.

M. Vieweg, D. Bongardt, C. Hochfeld, A. Jung, E. Scherer, R. Adib, and F. Guerra. *Towards Decarbonising Transport – A 2018 Stocktake on Sectoral Ambition in the G20*. 2018.

H. Wang. A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research*, 139(3):469–489, 2002.

B. Westerweel, R.J.I. Basten, J. den Boer, and G.J.J.A.N. Van Houtum. Printing spare parts at remote locations: Fulfilling the promise of additive manufacturing.

Production and Operations Management, 30(6):1615–1632, 2021.

A.S. Whittemore and S.C. Saunders. Optimal inventory under stochastic demand with two supply options. *SIAM Journal on Applied Mathematics*, 32(2):293–305, 1977.

H. Wong, A.A. Kranenburg, G.J.J.A.N. Van Houtum, and D. Cattrysse. Efficient heuristics for two-echelon spare parts inventory systems with an aggregate mean waiting time constraint per local warehouse. *OR Spectrum*, 29(4):699, 2007.

L. Xin. 1.79-approximation algorithms for continuous review single-sourcing lost-sales and dual-sourcing inventory models. *Operations Research*, Forthcoming, 2021a.

L. Xin. Understanding the performance of capped base-stock policies in lost-sales inventory models. *Operations Research*, 69(1):61–70, 2021b.

L. Xin and D.A. Goldberg. Optimality gap of constant-order policies decays exponentially in the lead time for lost sales models. *Operations Research*, 64(6): 1556–1565, 2016.

L. Xin and D.A. Goldberg. Asymptotic optimality of tailored base-surge policies in dual-sourcing inventory systems. *Management Science*, 64(1):437–452, 2018.

L. Zhang, Y. Lei, and H. Shen. How heterogeneity influences condition-based maintenance for Gamma degradation process. *International Journal of Production Research*, 54(19):5829–5841, 2016.