

Additive Manufacturing in After-Sales Service Supply Chains



Nils Knofius

**ADDITIVE MANUFACTURING
IN AFTER-SALES SERVICE
SUPPLY CHAINS**

Nils Knofius

ADDITIVE MANUFACTURING IN AFTER-SALES SERVICE SUPPLY CHAINS

DISSERTATION

to obtain
the degree of doctor at the University of Twente,
on the authority of the rector magnificus,
Prof. dr. T.T.M. Palstra,
on account of the decision of the the Doctorate Board,
to be publicly defended
on Wednesday, 19th of December of 2018 at 10:45 hrs.

by

Nils Knofius

born on the 18th of January of 1988
in Henstedt-Ulzburg, Germany.

This dissertation has been approved by:

Promotor:

Prof. dr. W.H.M. Zijm

Co-promotor:

Dr. M.C. van der Heijden

Ph.D. thesis, University of Twente, Enschede, The Netherlands
Department of Industrial Engineering and Business Information Systems

This thesis is part of the PhD thesis series of the Beta Research School for Operations Management and Logistics (onderzoeksschool-beta.nl) in which the following universities cooperate: Eindhoven University of Technology, Maastricht University, University of Twente, VU Amsterdam, Wageningen University and Research, and KU Leuven. This research is part of the project “Sustainability Impact of New Technology on After sales Service supply chains (SINTAS)” and has been sponsored by the Netherlands Organization for Scientific Research under project number 438-13-207.

Cover photograph: Simon Buchou
Printed by: Ipskamp Printing, Enschede, The Netherlands
ISBN: 978-90-365-4697-3
DOI: 10.3990/1.9789036546973

© N. Knofius, 2018, Enschede, The Netherlands

All rights reserved. No parts of this thesis may be reproduced, stored in a retrieval system or transmitted in any form or by any means without permission of the author.

GRADUATION COMMITTEE:

Chairman and Secretary: Prof. dr. T.A.J. Toonen
University of Twente, The Netherlands

Promotor: Prof. dr. W.H.M. Zijm
University of Twente, The Netherlands

Co-promotor: Dr. M.C. van der Heijden
University of Twente, The Netherlands

Members: Prof. dr. K. Hoberg
Kühne Logistics University, Germany

Prof. dr. ir. G.-J. van Houtum
Eindhoven University of Technology, The Netherlands

Prof. dr. ir. T. Tinga
University of Twente, The Netherlands

Prof. dr. H. Schiele
University of Twente, The Netherlands

Dr. ir. M.R.K. Mes
University of Twente, The Netherlands

Contents

1	Introduction	1
1.1	Additive Manufacturing Technology	2
1.2	Spare Parts Supply under Additive Manufacturing	9
1.3	Research Design	15
1.4	Related Literature and Contribution	17
1.5	Research Techniques and Concepts	21
1.6	Thesis Outline	24
2	Identifying Spare Parts for Additive Manufacturing	25
2.1	Introduction	25
2.2	Ranking Method	26
2.3	Case Study	31
2.4	Sensitivity Analyses	34
2.5	Conclusion	36
	Appendices	38
3	Consolidating Spare Parts with Additive Manufacturing	41
3.1	Introduction	41
3.2	Literature Review on the METRIC Methodology	43
3.3	Effects of Consolidation on Life Cycle Costs	44
3.4	Model	48
3.5	Numerical Experiments	53
3.6	Conclusion	62
	Appendices	63
4	The Transition to Additive Manufacturing	65
4.1	Introduction	66
4.2	Model	66

4.3	Analysis	74
4.4	Conclusion	83
	Appendices	84
5	Additive Manufacturing as Dual Sourcing Option	89
5.1	Introduction	90
5.2	Dual Sourcing Literature	90
5.3	Model	92
5.4	Numerical Experiments	99
5.5	Case Study	104
5.6	Conclusion	108
	Appendices	109
6	Large-Scale Dual Sourcing Problems with AM Supply	113
6.1	Introduction	113
6.2	Problem Formulation and Notation	114
6.3	Markov Decision Problem Formulation	116
6.4	Iterative Procedure	120
6.5	Performance Analysis	122
6.6	Approximate Dynamic Programming Formulation	125
6.7	Conclusion	130
	Appendices	131
7	Conclusions	135
7.1	The Research Agenda Revisited	135
7.2	Roadmap for Implementation	140
7.3	Discussion of Further Research Areas	143
7.4	Concluding Remarks	144
	Bibliography	145
	Summary	154
	Acknowledgments	157
	About the Author	159
	Scientific Output	160

Chapter

1

Introduction

Additive Manufacturing (AM, also known as 3D printing) is rapidly gaining interest as a highly innovative manufacturing technology and is increasingly maturing into a powerful complement to more conventional manufacturing (CM) methods. In comparison to CM methods such as milling, drilling, casting and forging, AM technologies build complete parts by adding materials layer upon layer but without dedicated tooling. Most attention is given to the ability to produce complex structures that are readily customized to specific applications. For instance, the aerospace industry increasingly applies AM for the production of lightweight designs. Airbus already uses more than 2700 printed parts in the A350XWB airliner (Airbus, 2016).

As well to realizing designs that were previously infeasible, after-sales service supply chains are also often viewed as potential beneficiary of AM technology. After-sales service supply chains support the maintenance of advanced capital goods during their life cycle of typically several decades. This support consists of providing all resources needed for system upkeep, such as service engineers, tools, and spare parts. Spare parts management is usually demanding because of the combination of the large variety of parts, the presence of many expensive slow movers, a geographically dispersed installed base, and the often high costs of system downtime. This leads asset owners to request high service levels from maintenance providers, including the availability of sufficient spare parts. Examples of advanced capital goods can be found in manufacturing equipment for the high-tech industry, healthcare and communication systems, and defence equipment.

The potential of AM is explained best if we study common challenges in after-sales service supply chains first. For example, uncertain demand, long lead times and high downtime costs necessitate high spare parts stocks – resulting in large amounts of capital being tied up. Also, arranging spare parts supply is often a challenge once the regular production phase has ended. Suppliers demand high incentives for maintaining production capacities or may even decide to discontinue supply entirely.

Both problems may be overcome through the use of AM technology, at least in principle. First, short AM setup times and no requirement for dedicated tooling may support demand-driven spare parts provisioning and thus help to reduce the prevalence of large spare parts stocks. Second, utilizing generic AM processes may relax the dependency on suppliers and therefore decrease risks and costs associated with supply disruptions. Furthermore, AM may enable the implementation of a decentralized production concept which may increase supply chain responsiveness at low cost.

This thesis investigates how and to what extent after-sales service supply chains may profit from the application of AM technology. Building on experience gained from various organizations, we develop models and apply techniques from the field of Operations Research to categorize and quantify the effects of AM on after-sales service supply chains. In this introductory chapter, we discuss the basic concepts of AM technology and its potential effects on after-sales supply chains. Moreover, we position our work within the literature and present our research approach.

The chapter¹ is organized as follows. In Section 1.1, we introduce the reader to the key characteristics of AM technology. Next, in Section 1.2 we discuss the possible applications of AM technology in spare parts supply chains. Section 1.3 outlines the research design, before we discuss the relevant literature and the contribution of this thesis in Section 1.4. In Section 1.5, we discuss the research techniques and concepts applied before we close with Section 1.6 in which we give an overview of the thesis structure.

1.1 Additive Manufacturing Technology

Although only recently known to the general public, the first AM technology had already been commercialized in the late 1980s, when it was used as a technique for rapid prototyping, termed stereolithography. In that technology, a vat incorporating a vertically moving platform is filled with a photocurable liquid polymer. With the platform in its upper position, a laser focuses an ultraviolet beam on the upper surface layer, curing that part of the photopolymer to create a solid body. Next, the platform is lowered slightly and the cured polymer is covered with another layer of liquid polymer, after which the sequence is repeated (Kalpakjian and Schmid, 2014). By varying the shape of each new polymer layer, complex geometries can be built up through stereolithography.

Today, there exists a wide variety of industrial AM technologies, of which the most important ones are Selective Laser Sintering (SLS) and Selective Laser Melting (SLM), Electronic Beam Melting (EBM), Digital Light Processing (DLP) and Fused Deposition Modeling (FDM). This is not the place to give detailed descriptions of these technologies; it suffices to say that they differ widely in the amounts and types of materials used, in their speeds and accuracies, and in their domains of application. Here, we give a general overview of AM technologies to lay the foundation for a later discussion of supply chain matters.

¹This chapter is in part based on our contribution to Zijm et al. (2019).

For a more comprehensive study on the differences between the various AM technologies, we recommend Gibson et al. (2010), and Kalpakjian and Schmid (2014).

For most AM technologies, the AM process starts with a user-defined 3D CAD file of a component or product. Specific AM software then will be used to cut the 3D CAD into slices that are fed into an AM machine to “build-up” the component layer-upon-layer as shown in Figure 1.1. An increasing variety of raw materials have become available for AM applications, including ceramic powder, metal or even glass, as well as polymers. The CAD file may be generated from a design process, but may also result from a 3D scan of an existing object. As a result, design changes are easily incorporated. The thickness of the layers may be of the order of microns; naturally, the thinner the layer, the greater the detail and accuracy that can be achieved.




Figure 1.1: *The principle of Additive Manufacturing (EOS, 2015)*

1.1.1 Opportunities and Shortcomings

The unique production process of AM technologies has several implications for the future of manufacturing. Let us therefore first look at some potential benefits that are derived from the basic properties of AM technologies.

(+) *Design freedom to produce complex and tailored parts*

The design freedom afforded by AM technologies is certainly one of the main benefits. Design compromises to improve manufacturability are significantly less limiting when applying AM technologies rather than conventional manufacturing (CM) and thus facilitate designing parts for their intended use. Complex structures can be built that achieve a nearly optimal balance between strength and material usage which is not feasible using subtractive technologies. Benefits can be observed in the aerospace industry where light-weight designs, that are only producible using AM, lead to significant fuel savings. Figure 1.2 highlights this opportunity and shows the CM (a) and the AM design (b) of a hinge bracket that is used in aircraft. Overall, the AM-enabled topology optimization leads

to a weight reduction of 25%. Other common examples for design improvements concern heat exchangers or valves for which thermal control or flow resistance are improved.




Figure 1.2: CM (a) and AM hinge (b) design for a hinge bracket used in aircraft

However, not only for individual parts does AM's design freedom have significant implications. According to the Wohlers Report (2014), the most promising application of the new design freedom for operations is the integration of parts, i.e., the redesign of an assembled component with fewer, but inevitably more complex parts. This process is referred to as consolidation. Apart from reducing the number of assembly steps, and thereby reducing both production lead times and costs, consolidation may improve the reliability of assembled components; see Johnson and Kirchain (2009) and Wits et al. (2016). Couplings between parts, often the cause of various failure modes, can be removed. Furthermore, the performance of the consolidated part may be improved. In this context, performance refers to aspects such as reduced weight while fulfilling the same functionality, lower flow resistance or improved heat dissipation. Also, the supply chain might be simplified because the number of distinct parts that need to be sourced, tracked and inspected will decrease. Hence, operational complexities and often long parts supply lead times are reduced (Yang et al., 2015). We will discuss the implications of consolidation facilitated by AM technology in greater depth in Chapter 3.

(+) *Reduced material waste and operational energy consumption*

The reduction of materials usage is a clear result of applying additive processes. Note that the materials are fed to an AM machine in different modes to those used in subtractive processes. For example, Achillas et al. (2015) refer to cases in which a 40% reduction of material consumption was achieved. Combined with the more uniform requirement for raw materials, this characteristic may compensate for the often energy-intensive production process inherent to AM. Indirect effects caused by lower weight or optimized part properties may even further reduce the energy consumption. Hence, from a life cycle perspective, the energy balance of 3D-printed parts may well turn out to be positive.

(+) *High level of customization*

The opportunity to design products according to customer specifications and to manufacture them on-demand, using only basic materials, is entirely a result of the fact that both the design and the manufacturing process are highly digitalized. Tooling or product-dependent setup processes are usually not required, making AM a highly flexible technology. Thus, for instance, design changes or product changeovers are easily realized while the production process remains unaltered. As we will discuss later, especially for medical and dental applications where customer-specific solutions are paramount, AM has already transformed entire supply chains.

(+) *Faster time-to-market*

The fact that design and manufacturing process are so closely intertwined, together with the fact that the product is built up in one piece from only raw materials, simplifies the development phase and eliminates a number of steps in the assembly process. Hence, AM may significantly reduce the time-to-market which may well yield a competitive advantage. Risks associated with market failure decrease, given low setup and tooling costs. Accordingly, it is likely that AM may support an aggressive market strategy. Also, rapid design changes based on market feedback appear less demanding and thus give rise to more dynamic business models.

(+) *General purpose equipment*

Shifting between designs on conventional manufacturing equipment often requires both, a lengthy setup process and the change of dedicated tooling. By applying AM technology this process is likely to simplify significantly by quickly restarting the printing process from another digital file. In many printing processes, it is even possible to print several completely different parts in parallel. The resulting flexibility not only reduces the investment and storage costs of dedicated tooling but also increases productivity and asset utilization. Even more, business models that allow companies to rent out their excess production capacities may also become profitable or help to cover the fixed costs.

Despite the potential of AM technologies, it is unrealistic to assume that AM is about to replace conventional production methods. It appears more likely that AM technologies will complement rather than replace conventional production methods. Below, we list shortcomings of today's AM technologies, and the trade-offs involved when compared to conventional manufacturing methods.

(-) *Pre- and post-processing requirements*

Most prominent is the misperception that an AM process of itself produces industrial grade parts. The reality usually involves various process steps each of which may require software, equipment and high levels of expertise. Before starting the printing process, the printing design has to be generated and the AM equipment may require preparation. For instance, it may be necessary to change the feed stock which, in the case of metal printing equipment, involves extensive cleaning of the build chamber. Also, major post-processing steps are often required to meet quality standards. Support structures may need to be removed, or treatments may be required to improve material properties. Furthermore,

process variability, inherent in today's AM methods, often necessitates extensive quality controls that increase both production lead times and cost.

(–) *Limited range of printable parts*

While AM technology certainly offers a high degree of design freedom, it also has limitations. For instance, printing bulk structures remains challenging since both porosity and the risk of thermal stress may occur more frequently. Hence, most industrial printing processes have to adhere to printing size limitations. Even for moderate sizes, we may face the need to print a part in several pieces. Depending on the printing process, available materials and the inability to combine multiple material types also decrease the feasible range of printable parts. For instance, printing complex electronics is likely to remain infeasible in the foreseeable future due to the necessary composite structures.

(–) *Diminished part characteristics*

The characteristics of an AM part may not compare favorably to those of their conventionally manufactured counterpart. For example, based on the process characteristics of AM, the unit cost and reliability of a conventionally manufactured equivalent are often superior. Also, conventional parts that are assembled from components can often be repaired by replacing only a malfunctioning component after which the part can be re-assembled. However, a malfunctioning printed part may have to be discarded entirely, resulting in needless waste, and replacement by a complete new part is clearly more costly.

(–) *Design rights and liability*

Currently, a significant number of parts that are considered to be suitable for printing are already produced with different manufacturing methods. Therefore, it is not uncommon that the design rights are owed by another entity. To acquire the design rights may turn out cumbersome and potentially require high investments, in particular if it would impact future business opportunities of the existing design holder. In the future, when AM technology has become more common, design leasing concepts may reduce this problem.

Additional concerns are raised by the digital nature of AM methods. While increasing flexibility, businesses worry about both the protection of intellectual property rights and product liability. The latter is clarified by a simple example. Consider an innovative company that offers 3D-designs of its products for sale. If a customer printed this product (maybe with slight alterations to the design) and it subsequently failed, the question arises who is responsible for the failure: the company, the service provider, or the AM equipment manufacturer. At present, no standardized legal agreements are in place, and that creates uncertainty around otherwise promising new business models.

(–) *High marginal production costs*

Compared to production methods such as injection molding, it becomes apparent that AM technology is not applicable in every market segment. In fact, injection molding, which itself offers a high design freedom, is essentially the opposite of AM in terms of flexibility. High upfront investments in dedicated molds require a high degree of commitment, while changeovers to another product (which involves changing the molds) may be time-consuming. Despite this, unit costs are lower for high volumes, and since

the actual production time of a single product may be a matter of seconds, the technique is typically suitable for mass production, achieving economies of scale. On the other hand, AM is not at all suitable for mass production. The printing of large and complex product geometries may take several hours, and is often highly energy-intensive. In combination with low economies of scale, this leads to high marginal production costs.

(–) *Technological obsolescence and missing standards*

Another problem arises due to the novelty and short development cycles of AM technology. If a company invests in AM machinery, its equipment may well be outdated after only a short period of time. Although leasing or outsourcing concepts may lower these risks, rapid technological advancements also demand a high degree of organizational flexibility. In particular, the absences of standards often forces businesses to reorganize production processes on a per part basis. A roadmap published by a standards setting organization for the aerospace industry supports the conclusion that this problem is likely to persist for the next five to ten years (3ders.org, 2017).

1.1.2 Application Areas

Next, we discuss application areas of AM in more detail. We will look at the rapidly growing number of industries and sectors in which advantages such as high customization, light weight and short time to market count most. The strategic research agenda of the Additive Manufacturing Platform (AM Sub - Platform, 2014) mentions several of domains that have adopted 3D Printing as a key technology, including:

- *Medical/dental*

Titanium alloys have been extensively used as powder material for fabricating orthopedic/orthodontic implants. Other applications can be found in e.g. the hearing aid industry, which has made an almost 100% transition to AM. The key driver for these types of applications is typically the ability to provide customized solutions for an affordable price.

- *Aerospace*

Main business drivers are weight and attributed fuel savings. Projects like NASA's 'zero gravity' 3D printer meant to produce spare parts in the International Space Station (ISS) show interest to further expand AM's application area (TechCrunch, 2016). Future scenarios in which, for instance, downtime of airplanes are reduced with printed spare parts are likely.

- *Automotive*

In this sector, the design flexibility is one of the most important arguments to move to AM, next to the fast realization of prototyped or low-volume car parts. Experimentation with large scale prints may motivate further applications and indicate the interest to secure weight and thus fuel savings similar to the aerospace sector (Ford, 2017).

- *Consumer products*

The market for consumer products is to date the largest but also the most diverse sector that has embraced AM technologies. Mostly, the technology is still used for

prototyping but a quickly growing number of end-users has discovered AM to deliver highly personalized devices and products, ranging from toys and busts to home furnishings and fashion items (robes, shoes).

Other fields of application include industrial machinery, the military domain and architecture (prototyping). An overview of the most important application areas is found in Figure 1.3, based on the Wohlers Report (2014).

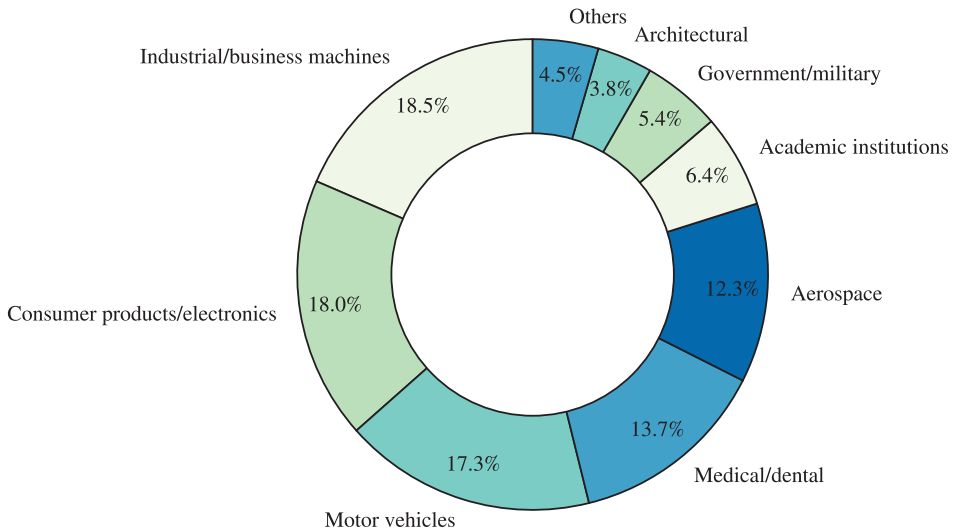


Figure 1.3: *Industrial and Public Sectors using Additive Manufacturing (Wohlers Report, 2014)*

In a Harvard Business Review publication, McCue (2015) reports that 30% of the Top 300 largest global enterprises are now using or evaluating the potential of AM. Some of the companies that are already exploiting AM technology are: General Electric (jet engines, medical devices), Lockheed Martin, Airbus and Boeing (aerospace and defense), and Aurora Flight Sciences (Unmanned Aerial Vehicles) (D’Aveni, 2015). In summary, AM has quickly gained a firm position in the manufacturing arena. Based upon the observed drop of the cost of AM systems with 50% in the last decade (Thomas, 2016) and a further expansion of the materials range, an increased penetration of the manufacturing arena is generally expected. Accordingly, forecasts of Siemens, a multi-national conglomerate, predict an AM market growth of 300% within the next 10 years (Siemens, 2014a).

1.2 Spare Parts Supply under Additive Manufacturing

As argued in the introduction, the possible benefits of AM for spare parts supply chains have attracted considerable interest. To a large extent, this interests relates to the unique characteristics of spare parts supply chains if compared to manufacturing supply chains, cf. Table 1.1. However, AM remains far from being established in spare parts supply chains.

Table 1.1: *Comparison manufacturing and spare parts supply chains. Based on Cohen et al. (2006)*

	<u>Manufacturing Supply Chain</u>	<u>Spare Parts Supply Chain</u>
Nature of demand	Predictable, can be forecasted	Unpredictable, sporadic
Required response	Standard, can be scheduled	ASAP (same day or next day)
Number of SKUs	Limited	15 to 20 times more
Product portfolio	Largely homogeneous	Always heterogeneous
Delivery network	Depends on nature of product; Multiple networks necessary	Delivering different service products; Single network
Purpose	Maximize velocity of resources	Pre-position resources
Reverse logistics	Does not handle	Handles return, repair, and disposal of failed components
Performance metric	Fill rate	Availability (uptime)
Inventory turns	Six to 50 per year	Up to four a year

In this section, we provide a more extensive background on envisaged applications of AM in this domain. To that end, we first describe spare parts supply chain characteristics in Section 1.2.1. Next, in Section 1.2.2, we discuss how the implementation of AM may affect spare part supply chains and outline resulting implications for sustainability in Section 1.2.3.

1.2.1 Characteristics of Spare Parts Supply Chains

To structure our discussion, we divide spare parts supply chain characteristics into four categories: demand, sourcing, service and life cycle characteristics. Later, in Section 1.2.2, we reuse this structure to classify the potential effects of AM for spare parts supply. Also, we rely on this framework to discuss gaps in the literature (Section 1.4).

(a) *Spare parts demand characteristics*

Spare parts are required for either preventive or corrective maintenance activities. In the case of preventive maintenance, spare parts demand may be known well in advance. Examples include the application of time- or age-based replacement policies under which spare parts may be ordered just-in-time and hence limit the need for spare parts inventories. However, difficulties may arise in the coordination process with other maintenance resources such as tooling or service engineers.

If preventive part replacement is based on inspections or condition monitoring, spare parts inventories become more important for achieving potentially short response times. For example, collected data may support reliable failure predictions only shortly before failure, or inspections may reveal damage that must be repaired at short notice. For corrective maintenance activities, spare parts demand is random and only observable once a system failure materializes. In such a case, if components are critical to system availability, spare parts stocks are essential. Low demand quantities combined with high demand variability typically lead to high inventory costs that may either emerge in terms of holding or backorder costs.

If the spare part is assembled, spare parts demand may occur at various hierarchical levels. For instance, repair options may become feasible in which a defective item can be repaired by just replacing a sub-component. In general, the service provider (i.e., an OEM or a third party) has to decide at which system hierarchy level to replace a defective part and, if possible, at which item hierarchy level to repair the part. Both decisions directly influence the spare part demand as the service provider has to design the inventory policy accordingly. As we will elaborate under Point (c), the replacement decision also has to be aligned with the service requirements.

(b) *Spare parts sourcing characteristics*

Gibson et al. (2010) note that low volume spare parts are mostly produced on generic (i.e. non-dedicated) equipment (e.g. CNC workstations). Hence, the production process typically requires set-up time in addition to special tooling and fixtures that lead to high sourcing costs. Depending on the level of integration, service providers may source either spare parts or their components from a supplier. Sourcing spare parts for asset maintenance typically involves highly specialized manufacturing supply chains. Only on rare occasions will the service provider have the option to choose from various supply sources. However, if multiple supply sources are available, it may be beneficial to rely on more than one. For instance, it may be possible to reduce service costs by combining an inexpensive but slow supply source with an expensive but fast supply source. Also, using multiple suppliers improves supply security and thus may protect against supply disruptions or improve the negotiation position. More commonly though, the service provider does not possess sufficient leverage nor control over its suppliers because of infrequent orders and low order quantities. As a result, service providers often have to accept long lead times and large minimum order quantities.

(c) *Spare parts service characteristics*

If the target response times are not met, service providers may incur high penalties and, probably even worse, see their reputation damaged. Hence, service providers usually aspire to run a highly responsive supply chain, but that is a complex task given the often wide range of parts. Moreover, spare parts in stock often represent a significant investment. In some industries, spare parts may easily cost tens of thousands of euros each. Under such circumstances, the service provider has to determine which spare parts and how many to store at which locations.

In a similar way, the service provider has to decide at which hierarchical level in the system structure to replace a defective part. While replacing broken parts at a high level in the system structure typically increases inventory costs, system downtime costs tend to be reduced. For various industries, the latter aspect is essential and one of the key reasons for stringent response times agreements. To give an example, in the semiconductor industry guaranteed asset repair times are often less than a few hours (Stein, 2012). To this end, service providers may also decide to keep high value spare parts close to (or even at) the installed bases, which may be spread across the world.

Under these conditions it is not surprising that service providers often store far more spare parts than actually required. This is exemplified in the findings of Cohen et al. (2006). They detected that more than 20% of spare parts are not used and become obsolete every year. Even though some leftover parts probably can be sold on the aftermarket, this has a significant negative economic impact. Cattani and Souza (2003) found that Hewlett-Packard's profits are reduced by about 1% of its annual revenue as a result of obsolescence.

(d) *Spare parts life cycle characteristics*

During the asset life cycle, the service provider will encounter various types of challenges. Usually these vary with the level of uncertainty involved, i.e., the better the service provider anticipates future demand, the easier it becomes to plan the appropriate parts supply, either through stock holding or through timely deliveries. If we consider various sources of uncertainty over the service horizon, we approximately find an uncertainty profile as shown in Figure 1.4.

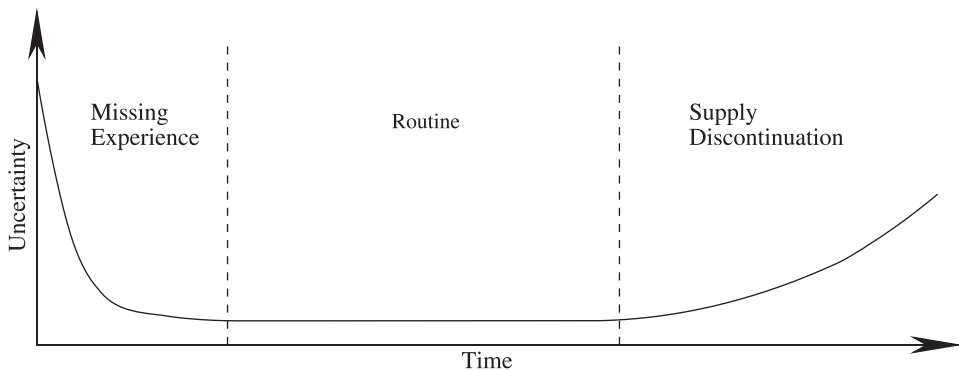


Figure 1.4: *Uncertainty over the service horizon of a spare part*

In the initial phase, the uncertainty is high since the experience with failure behavior and maintenance operations is limited or non-existent. Potential design changes and dynamic market developments may increase the pressure on the spare parts operations. Furthermore, the asset owner may experiment with the equipment utilization, which may lead to unstable and unpredictable asset deterioration. With the progression of

time, the conditions improve since more data and experience become available. Also, design changes occur less frequently as the market stabilizes.

However, towards the end of the service period, uncertainty typically rises again. In particular, the risk of supply disruption increases. For instance, a spare parts supplier may decide that the provision of legacy parts is no longer economic. At best, the service provider (or asset owner) is given the opportunity to purchase a final set of parts to meet possible demand during the expected remaining lifetime of its assets (Behford et al., 2015). Also, asset users may exploit their strong market position to demand an extension of the service period. The associated uncertainty about the service horizon length may further exacerbate the described problems encountered with supply disruptions.

Overall, complex trade-offs between setup, inventory, downtime and operational costs have motivated a wide range of research activities in the spare parts management domain. We refer to Muckstadt (2005), Sherbrooke (2004), Van Houtum and Kranenburg (2015) and Hu et al. (2018) for a more fundamental treatment of spare parts supply chains.

1.2.2 Effects on Spare Parts Supply

AM offers several opportunities for improvement in spare parts supply chains. Following the structure of Section 1.2.1, we elaborate on these opportunities.

(a) *Spare parts demand characteristics*

Spare parts demand that originates from maintenance activities may be positively affected by an *increased number of repair options*. For instance, worn-out parts that were previously discarded or too expensive to repair may well become repairable using AM which in turn may significantly reduce maintenance costs. As an example, Siemens (2014b) was able to reduce the repair lead time of burner tips in gas turbines by 90% and the associated repair cost by 30% after switching to AM. Figure 1.5 shows how a new burner tip is printed and afterwards attached to the burner.




Figure 1.5: *Repairing burner tip using Additive Manufacturing (Andersson et al., 2016)*

Additionally, AM technology may *improve maintenance performance*. For example, predictive maintenance concepts usually aim to optimize the trade-off between the risk of

failure and the risk of replacing a component that has a possibly long remaining lifetime. Short AM production lead times may allow the postponement of the maintenance decision without relying on spare parts inventories and thus may reduce both risk factors at low costs. In addition, an expanded monitoring period increases the statistical confidence in the component condition which further improves the decision quality.

On the other hand, the novelty of AM technology typically causes a *more uncertain failure behavior*. While for conventionally manufactured spare parts there may exist experience with operating the same or a comparable part, this is typically not the case for AM produced spare parts. In particular, such situations may arise if a company were to shift to AM technology during the service period.

(b) *Spare parts sourcing characteristics*

Compared to traditional spare parts supply chains, *sourcing and operational tasks become more intertwined*. For instance, at present we already observe that service providers expand their business models and offer on-demand printing capabilities to their customers (UPS, 2016). Although the implications of this business model are still not clear for capital goods, it holds various promises. For example, inventory costs may reduce across the supply chain because *work-in-progress inventory and safety stocks decrease*. Furthermore, shipping requirements reduce because printing hubs may offer a more local supply option. Given the long turnaround times for slow moving parts, it is important to note that low inventories simultaneously *reduce the risk of obsolescence*, i.e., storing parts that in the end will not be used.

Furthermore, sourcing tasks simplify greatly since spare parts become producible from raw material that can be shared among various products (Tsai, 2017). The alternative of printing spare parts may also *strengthen the negotiation position* of the service provider which, aside from decreasing sourcing costs and lead times, enables the service provider to have a greater influence on the spare part design. In future scenarios, service providers may even encourage a continuous improvement cycle of spare parts designs based on new information about materialized failures.

(c) *Spare parts service characteristics*

Bypassing specialized manufacturing supply chains increases the flexibility to serve customer demand. For example, Walter et al. (2004) discuss the concept of printing spare parts on location. They argue that this practice offers benefits if demand occurs at remote locations or if customer response times have to be short. So far, this could only be achieved by emergency shipments or by holding inventory close to the installed base as discussed in Section 1.2.1.

Also, it is conceivable that it may pay off to offer *backup supply solutions* using AM. Although AM sourced spare parts may be less reliable, customers are likely intrigued by short response times at lower prices which was previously not realizable through other expediting options. Under some conditions, AM-produced spare parts may even function as a *temporary fix*. That is, the printed part can bridge the interval until the intended replacement becomes available. Nowadays, first applications can be found in the military,

which often uses highly advanced equipment at remote locations, cf. McLearen (2015). However, for civil applications such as mining, humanitarian missions or space operations this type of service options may also become viable.

(d) *Spare parts life cycle characteristics*

Due to the flexibility of AM production methods, the varying degree of uncertainty over the service horizon becomes controllable. For example, missing experience of failure behavior and maintenance operations may be compensated for by a highly responsive supply chain. Also, possible spare parts design changes may be easily accommodated by AM technology since no investments in dedicated tooling would be incurred. Another application arises if spare part supply is discontinued. As elaborated in Section 1.2.1, discontinuation typically causes high costs and is more likely for low-volume parts. Applying AM technology, it may be possible to *reestablish the supply continuity* in an inexpensive way as mentioned by Sasson and Johnson (2016). Montero et al. (2018) discuss this option for military equipment operated at remote locations.

1.2.3 Implications for Sustainability

Apart from raw material waste reductions and effects originating from optimized part properties (cf. Section 1.1.1), the application of AM in spare parts supply chains may also improve sustainability. Forward flows of specific parts are replaced by the distribution of a limited diversity of raw materials that may be sourced at many locations. Stock keeping, but also the necessity to rely on unsustainable transportation modes such as air-cargo, may become less important to guarantee short response times since local production concepts may become feasible. Likewise, a demand-driven production approach mitigates the risk of spare parts obsolescence, which ultimately may reduce the disposal of unused spare parts.

Also, the design freedom of AM may increase the sustainability of spare parts supply chains. For instance, as we described in Section 1.2.2, it is possible that defective components (or their sub-components) become repairable by AM technology and therefore increase the usage period of parts. At present, it is not uncommon that a defective sub-component leads to the disposal of the entire component since there are no replacement parts available on the market. Based on the design freedom offered by AM, it may become feasible to also locate the repair process downstream in the supply chain since the requirement for dedicated repair equipment is likely to decrease. Overall, repair shops may offer a broader range of services, which may unlock economies of scale that further contributes to the sustainability of the repair process.

Finally, we observe increasing efforts that attempt the transformation of material waste into a feedstock for AM processes. For example, the Army in The Netherlands is experimenting with the shredding of PET bottles. The resulting PET granular then can be reused to print less demanding spare parts. In a mission context or at remote locations such a recycling option appears especially desirable. Not only does a local recycling option reduce the ecological footprint but it may also give organizations a monetary incentive to apply a more sustainable sourcing concept.

1.3 Research Design

This research is conducted as part of the SINTAS project, which is funded by the Netherlands Organisation of Scientific Research (NWO). In this project, a consortium of academic, governmental and industrial partners seeks to explore the impact of AM on the design and control of after-sales service supply chains. This thesis particularly addresses the project goal of examining the impact of AM on spare parts inventories at the various stages in the service period.

1.3.1 Research Objective

As the previous discussion revealed, the expectations arising from the advancement of AM technology are high. For spare parts supply chains, this enthusiasm may originate from the prospect of simplification. For instance, when confronted with “inspirational” talks, we are often encouraged to imagine a world without long-haul transportation, complex assembly processes or inventory. However, to leverage the potentials of AM technology, it is necessary to deconstruct these concepts and to separate the hype from reality. Through this research, we aim to contribute to this undertaking by offering a scientific perspective on how and to what extent after-sales service supply chains may benefit from AM technology. In particular, we formulate the following research objective:

To offer decision support for actors in after-sales service supply chains to identify and understand the value of AM technology for their organization, and to provide quantitative insights into both when and how AM technology may be used or combined with conventional manufacturing methods to improve the efficiency of service logistics.

1.3.2 Research Questions

In this section, we describe the conceptual framework that we have chosen to address the research objective. We do this on the basis of research questions that are presented below.

RQ 1. *How can organizations identify spare parts that are economically viable and technologically feasible for the application of AM technology?*

With a growing awareness of the potential of AM technology for spare parts supply, organizations attempt to identify actual use cases for their specific (business) environment. We propose a method that provides practitioners with a structural procedure to assess a large spare parts assortment on their potential impact on supply chain efficiency and responsiveness. The approach is based on the Analytic Hierarchy Process and relies on spare parts information that is easily retrievable from the company databases. This has two advantages: first, the approach can be customized towards specific company characteristics, and second, a very large number of spare parts may be assessed simultaneously. A field study is discussed in order to demonstrate and validate the approach in practice. Furthermore, sensitivity analyses are performed to evaluate the robustness of the method.

RQ 2. *How does the consolidation of spare parts with AM technology affect the total life cycle cost of capital goods and when is it valuable?*

As we discussed in Section 1.1.1, the consolidation of parts is perceived as one of the most promising applications of AM technology. Typically, consolidation with AM is chosen because of its functional benefits such as weight reductions. Consequences for asset maintenance, however, are not that well understood. We adopt a total life cycle cost perspective and investigate under which circumstances consolidation with AM technology is economically valuable. Therefore, we identify and study the root causes which are responsible for the economic value of consolidation using existing methods for spare parts optimization.

RQ 3. *When and how does a transition to AM technology become profitable for the low-volume spare parts business?*

After the identification of valuable spare parts for the application of AM technology, organizations may hesitate when and how to move to AM technology. Non-stationary effects such as decreasing AM production costs and a lack of experience with AM technology are compelling reasons to postpone the investment. Also, after the regular production phase, knowledge and product-specific tooling is often already available for sourcing spare parts with CM while AM still has to be prepared. We build a stochastic dynamic programming model to analyze the described situation. Based on a case study and numerical experiments we assess the value of different spare parts sourcing strategies and derive general guidelines for the transition to AM.

RQ 4. *Under what conditions does the sourcing of low-volume spare parts with a combination of AM and CM methods pay off if we acknowledge that part quality is largely influenced by the production method?*

One result of answering **RQ 3** was that a dual sourcing concept often appears favorable during a significant part of the total service horizon. To further study the underlying reasons and the benefits compared to single sourcing with either AM or CM, we construct a customized dual sourcing model. In particular, we respect the characteristic that the sourcing decision may influence future demand because of the different failure behavior between AM and CM parts. Using numerical experiments and a case study in the aviation industry, we explore under which conditions dual sourcing with AM performs best. For large problem instances though, the exact optimization of the proposed model encounters computational limitation. To this end, we formulate the following research question.

RQ 5. *How can we analyze large problem instances of the problem discussed in RQ 4, given the computational limitations of the exact analytical methods?*

We build an iterative procedure to extend (heuristic) dual sourcing methods to become applicable to cases with supply mode dependent failure behavior. The extension is easily implementable since it only relies on an estimate of a fraction of items ordered from either of the supply modes. To demonstrate its performance, we benchmark the procedures against exact results obtained for small problem instances and the case study discussed under RQ 4. Furthermore, we examine the option to model the problem with approximate dynamic programming by revealing the specific problem structure.

1.4 Related Literature and Contribution

After we have discussed the conceptual value of AM technology for after-sales service supply chains in Section 1.2.2, we review quantitative research that evaluates the value of AM technology for spare parts supply chains. The literature in this field is still limited and we only found sources that address spare parts sourcing and service characteristics, as introduced in Section 1.2.1. We review the corresponding literature in Section 1.4.1 and in Section 1.4.2. We close this section by stating our contribution in Section 1.4.3.

1.4.1 Sourcing Characteristics

Various authors assess the production costs when using AM technology. The first effort in this direction is made by Hopkinson and Dicknes (2003). They compare the production cost of AM with injection molding and show for various geometries that the total AM production costs are lower for small to medium production quantities. Lindemann et al. (2012) argue that it is essential to consider the life cycle costs when assessing the benefits of AM. They clarify this proposition with an example from the aerospace industry, in which weight reductions in the part design have an increasing impact over time due to reduced fuel consumption. They also identify a cost structure of AM in their research. Their work reveals that a large share of cost drivers are fixed, e.g., acquisition cost for AM machinery, support equipment and labor cost. Furthermore, Lindemann et al. (2013) present an approach to assess the life cycle costs for specific business environments and intended parts application. For a more extensive discussion about AM production costs we refer to Schröder et al. (2015) or Baumers et al. (2016).

In the literature, several methods are proposed for the identification of parts that are economically viable for production with AM technology. Common for these methods is that they use a bottom-up procedure, i.e., the method based on suggestions of employees. One example is the two-stage method suggested by Simkin and Wang (2014). In the first phase, it is examined whether the part suggested by an employee falls in at least one category of a defined list of potential benefits of AM technology. Examples from this list are improved functionality, lower sourcing costs, and lower import/export costs. If this is not the case, it is argued that printing the suggested part is almost certainly not worthwhile. In the second phase, it is examined which AM production methods can be used to manufacture the part. Unfortunately, the details of this assessment are not specified. Afterwards, cost-benefit analyses are performed with Monte Carlo simulation. For instance, Simkin and Wang (2014) compare the total life cycle costs of AM production methods with the costs of a conventional manufacturing process. Also, the impact of in-house manufacturing and outsourcing is compared. Again, it is not stated explicitly which factors are included in the life cycle costs, and how they are calculated.

Another method is proposed by Lindemann et al. (2015). They structure the entire bottom-up procedure with a workshop concept. During a first workshop, company representatives are informed about the advantages and limitations of AM technology. The purpose of this step is to qualify and inspire company representatives to independently identify parts for further analysis. During a second workshop, the resulting part candidates are evaluated by AM experts and the company representatives. To this end, Lindemann et al. (2015)

have developed a scoring method which assesses different part characteristics - primarily concerning technological constraints of AM such as part size and materials. Afterwards, economic aspects and possibilities for redesign of the best scoring parts are considered in more detail. This requires additional data collection and evaluation. The assessment is carried out by AM experts, though the details are not specified.

Liu et al. (2014) analyze the effect of using AM technologies instead of CM for a spare parts supply chain of aircraft. They compare central and decentral deployment of AM equipment with different demand characteristics and service level requirements. The spare parts design is considered to be identical for both production methods. In all experiments, the safety stock requirements are lower with AM technologies. Furthermore, they find that a central deployment of AM capacity is favorable for slow moving spare parts, with high demand variability and long AM production times. Otherwise, a distributed utilization of AM technologies appears favorable. The investment costs for AM equipment or personnel costs are not considered and therefore bias the analysis in favor of a decentralized deployment.

This critique is confirmed by the findings of Khajavi et al. (2014). They show that a decentralized layout only becomes attractive if the acquisition costs of AM equipment can be further reduced. Likewise, they identify a higher automation of AM equipment as crucial to reduce the required personnel costs in a decentralized AM supply chain. Finally, they demonstrate that a short production lead time of AM technologies is important - especially if short customer order lead times are demanded. Long production lead times enforce inventories and thus gradually reduce one of the key benefits of AM technologies. Later, Li et al. (2017) demonstrate that an AM supply chain typically outperforms conventional supply chains regarding carbon emission.

Barz et al. (2016) study the impact of a more efficient raw material utilization of AM technologies on the supply chain layout using mixed-integer programming to analyze a two-stage supply network. In the first stage, raw materials are delivered to production sites. In the second stage, the finished product is delivered to customer sites. Decision variables are the location of the production sites, the production site/customer site relations and the transportation quantities. They find that transportation costs decrease with the use of AM technologies. This result is explained by a lower requirement of raw materials and thus less transportation costs from the raw material source to the production site. Also, production sites tend to be located closer to customer sites due to this property. As a final observation they report that the number of opened production sites is rather independent of the raw material utilization. It needs to be mentioned, however, that some crucial assumptions are made: demand is deterministic and independent of the production technology, production capacity costs are independent of the production technology and no inventory is allowed. Also, it would be interesting to obtain insights in the consequences of a more uniform requirement of raw material with AM technologies, a topic that is not addressed in their paper.

1.4.2 Service Characteristics

Sirichakwal and Conner (2016) evaluate the influence of AM-imposed production lead times and holding costs reductions on the stock-out probability. To this end, they assume a single stockpoint and apply a continuous review base stock policy with emergency shipments. In general, they find that AM has positive effects on the total inventory costs. Furthermore, they argue that holding cost reductions decrease the stock-out probability because companies are incentivized to keep more stock. In particular, this finding holds for parts with low demand rates. Given that they do not adopt a costs perspective though, the magnitude of associated cost savings remains unclear.

Westerweel et al. (2018b) investigate which AM part reliability and AM production costs levels have to be achieved to reach a break-even point in the total life cycle costs compared to sourcing with CM methods. Therefore, they study a single stockpoint that follows a continuous review base stock policy with emergency shipments. Their model reveals that even if the reliability of the AM part is lower than that of the CM part, the AM version yields lower total costs under the assumption that the production costs of both design options are identical. Conversely, if the production costs are different but the mean times between failures (MTBF) are identical, the AM version is still preferable for cases with higher production costs. These findings are a consequence of the key assumption that AM always requires a shorter production lead time. Furthermore, they provide insights into the consequences of a large installed base size and the life cycle length. In case AM technologies require higher investment costs which cannot be offset by performance improvements in the short run, this may be mitigated by spreading the costs over a large installed base size and long life cycle.

Song and Zhang (2016) consider the parallel use of AM and CM methods. Therefore, they assume that AM technology functions as a capacitated emergency channel (modeled as an M/D/1 queue) but typically allows faster, though more expensive, resupply than the CM source. Also, they assume that AM parts show the same failure behavior as CM parts. Overall, they find that the production of parts on-demand with AM methods leads to cost savings and inventory reductions compared to the application of CM methods only. Especially for situations with large part variety the savings potential is significant.

Westerweel et al. (2018a) analyze the benefit of using AM spare parts supply as temporary expediting solution before a scheduled regular supply becomes available. To that end, printed parts are operated only until regular spare parts (with a higher reliability) are delivered via the regular supply mode. By means of an infinite horizon, discrete time Markov Decision Model, they analyze under which conditions it is advisable to use printed parts and which inventory policy should be applied. They show that the regular supply source should be operated according to a base stock policy, while the decision whether to print a replacement part follows a threshold policy. The value of using AM as temporary solution is established by means of two case studies in a military mission context and further supported via numerical experiments. They extend these results to show that printing parts remains beneficial even if an additional regular expediting option becomes available.

1.4.3 Contribution

The quantification of various opportunities of AM for spare parts supply chains demands substantial further research. First, we have identified few areas that have been explicitly modeled and analyzed compared to the wide range of concepts that we discussed in Section 1.2.2. Second, we did not find any study which evaluates potential challenges for spare parts supply chains that may originate from the increasing use of AM in the regular manufacturing phase. Third, the number of case studies on after-sales service supply chains that are essential to identify further problem characteristics is limited. Through this thesis, we aim to contribute to the scientific discourse by addressing several gaps. Below, we explain the five main contributions, one for each chapter.

1. We develop the first top-down approach to identifying spare parts suitable for the application of AM technology.

As discussed in Section 1.4.1, previously proposed methods rely on bottom-up procedures. That is, a practitioner realizes that AM technology might improve the characteristics of a specific part and proposes that part for further consideration. In contrast, a top-down approach initially considers the entire spare parts assortment and then systematically identifies the most promising parts. The top-down approach has proved both efficient and effective in several field studies.

2. We offer the first quantitative insights into how consolidation through the application of AM technology affects the total life cycle costs.

The quantification of indirect effects caused by AM-imposed design changes on spare parts supply chains have not yet received any attention in the literature. By evaluating the effects of consolidation on the total life cycle costs, we make the first contribution in this direction.

3. We develop a model to study non-stationary effects that may influence the decision to move to AM technology

To the best of our knowledge, this study presents the first results related to spare parts life cycle characteristics (cf. Section 1.2.1). In particular, we provide quantitative insights into how uncertain AM technology advancements may influence the decision to switch to an AM approach. Furthermore, we evaluate how available tooling to source spare parts by applying CM methods affects the benefit of transitioning to AM technology. A case study conducted at an OEM of radar systems extends the currently limited number of case studies in the field of after-sales service supply chains.

4. We propose a dual sourcing model under which sourcing decisions influence future demand.

Thus far, dual sourcing models do not consider the option that sourcing decisions may influence future demand. We develop an exact algorithm to analyze this situation. A case study conducted in the aerospace industry provides further insights into practical challenges that arise due to the implementation of AM technology.

5. We propose an extension for dual sourcing methods to evaluate large problem instances in which sourcing decisions influence future demand.

We propose and evaluate a new procedure to extend (heuristic) dual sourcing methods to analyze scenarios in which the failure behavior of parts sourced from both supply channels differs. Furthermore, we examine the specific problem structure to facilitate the development of more general solution frameworks with approximate dynamic programming.

1.5 Research Techniques and Concepts

In this thesis, we use several techniques and concepts from the field of Operations Research. Here, we briefly outline the applied techniques and concepts in the context of the conducted research in this thesis.

1.5.1 Multicriteria classification

As discussed above, AM technologies offer various opportunities for spare part supply chains while its successful implementation largely depends on the specific use case. Hence, in order to leverage the potentials of AM technology an essential task is the classification of possible use cases with respect to their (business specific) potential to profit from AM technology characteristics. Such classification problems are studied in the field of multicriteria analysis. Here, we refer to Zopounidis and Doumpos (2002) and Hu et al. (2018) for a more general review and subsequently direct our attention to the Analytic Hierarchy Process (AHP).

The AHP method was introduced by Saaty (1980) and supports decision makers to reveal the actual preference for a decision option relative to its alternatives. The underlying mechanism is best explained by a small example. Suppose we want to identify which characteristics of AM technology are most interesting for a specific company. For simplicity, let us assume that we have to decide between design freedom, supply responsiveness or reduced part usage cost. The goal of the company, say, is to improve customer satisfaction. This goal is decomposed in a number of attributes that influence the goal. Here, we focus on service quality, response time and service cost. Potentially we could further decompose each of these attributes, however, we restrict us to two hierarchy levels.

After the problem hierarchy is build, decision makers are asked (individually) to compare any two decision options with respect to their value for each attribute. For instance, one may ask *which decision option does increase the service quality more: a higher supply responsiveness or a reduced part usage cost? Rank your choice on a scale from 1 to 9, in which 1 means equally important and 9 that the supply responsiveness is clearly more important.*

After completing this activity for every decision option pair and attribute, we ascend one hierarchy level and ask the same type of question with respect to each attribute pair relative to the goal to improve customer satisfaction. The resulting scores then can be used to calculate which decision option has the highest preference among the decision makers. For details on the calculation we refer to Saaty (2008). However, we emphasize that the AHP method does

not eliminate subjective bias. The ability to decompose a decision problem in simple pairwise comparison tasks though, tremendously reduces decision complexity and makes decision consistency controllable. We will partially rely on the AHP method when addressing **RQ 1**.

1.5.2 Multi-Echelon Technique for Recoverable Item Control

The value of AM technology for after-sales service supply chains, largely depends on its effect on spare parts inventories, cf. Section 1.2.2. Yet, studying these effects in isolation may lead to false conclusions and therefore demands the assessment of fairly general spare parts networks. A pioneering work for the evaluation and optimization of such problems is the METRIC method by Sherbrooke (1968). In its original form, Sherbrooke considers a multi-item spare parts network consisting of one central repair location (depot) that supports various small repair shops (bases) that satisfy spare parts demand.

The evaluation is largely based on results from Queueing Theory. For instance, the estimation of the steady-state probabilities of the number of items in repair follows from Palm's Theorem which establishes a relation between the arrival process and the distribution of items in repair. Furthermore, using information about spare parts demand, lead times, and repair capabilities at each base, the METRIC method uses a marginal approach to optimize the inventory policy based on convexity properties.

Considering that real-life problem instances often concern thousands of spare parts and various repair locations, it is not surprising that the METRIC method adopts various assumptions, cf. (Sherbrooke, 2004). Furthermore, the performance evaluation is partially based on approximations. Nevertheless, as we will review in greater detail in Chapter 3, the METRIC method and its various extensions and improvements, allow the assessment and optimization of fairly general spare parts networks. In particular, we use the extension to consider hierarchical spare parts (multi-identure) to assess the effects of consolidating spare parts in Chapter 3.

1.5.3 Markov Decision Processes

The decision to use AM instead of conventional methods for spare parts supply may cause various changes to the service system. For example, we may have the option to order spare parts with a shorter replenishment lead time but therefore may need to accept a lower part quality. The merit of such changes usually only becomes clear if we study the service system over a longer time period. Furthermore, under certain conditions it may become necessary to take exogenous factors into account such as anticipated technological advancements or piece price reductions.

A common technique to model such systems is the use of Markov Decision Processes (MDP) which are widely applied in sequential decision making, i.e., which decision to take if that decision influences future decision making options. Each decision is associated with an expected cost (or reward). To identify the optimal decision at each decision moment, the system condition is represented by so-called state variables which contain all relevant system information. The decision and possible (stochastic) events cause a transition to the next state.

In principle, MDP's support the analysis of a wide range of problem types such as finite and infinite horizon problems and discrete and continuous time problems. Prerequisite for all MDPs, though, is the assumption that the Markov property holds, i.e., the transition to a new state depends probabilistically on the current state and on the action taken, but not on previous states. Based on this property it is possible to analyze fairly large systems and to determine optimal decision policies using mathematical programming or dynamic programming. For further details we refer to Tijms (2003), Bertsekas (2012) or Puterman (2014). We build a discrete time, finite horizon model to address **RQ 3**. Furthermore, we develop a continuous-time, infinite-horizon model to study **RQ 4**. For **RQ 5**, we model the problem as a discrete-time, infinite-horizon problem.

1.5.4 Approximate Dynamic Programming

For various real-life problems MDP formulations become computationally intractable. In particular, the number of system states or possible decision options grows rapidly with the problem size. Also, it is possible that the number of possible stochastic events leads to various transition options. In the literature these issues are coined as the three curses of dimensionality (Powell, 2011). We encounter this problem while addressing **RQ 4**.

One powerful solution framework to overcome these issues is approximate dynamic programming (ADP) which is also referred to as reinforced learning. By applying ADP it is typically the goal to approximate the value of each decision option in a state or to find a close-to-optimal policy. To that end, ADP combines techniques from various fields including dynamic programming, statistics, simulation and mathematical programming. For further details, we refer to Powell (2011) and Bertsekas (2012).

1.5.5 Discrete-Event Simulation

Discrete-event simulation models replicate the behavior of a real system. In this thesis, we apply discrete-event simulation for auxiliary purpose. In particular, we use it as modelling framework to test our assumptions under more general conditions and to validate our implementations. Each simulation run assesses a certain trajectory of random events which occur at discrete moments in time. In contrast to continuous time simulations, we only observe the system state if an event occurs. Between events, the system is supposed to remain unchanged from a logistical perspective. The accuracy largely depends on the selected transition mechanism and is commonly expressed by probability distributions or empirical data. At every event, a predefined control is applied. By performing a sufficiently large number of simulation runs, it is possible to obtain insights in how well the predefined control performs relative to a certain confidence level.

Compared to analytic methods, simulation usually facilitates the assessment of a system under more general conditions. Yet, to achieve a certain level of confidence with respect to the obtained results, simulation may require long computation times. Hence, it is often less attractive for extensive numerical experiments or optimization purposes. For an extensive treatment on discrete-event simulation we refer to Law (2007).

1.6 Thesis Outline

The structure of the thesis follows the research design presented in Section 1.3. Hence, each chapter addresses one research question.

In **Chapter 2**, we address **RQ 1** and elaborate how organizations may identify spare parts that appear promising for the application of AM technology. Furthermore, we report on our experience made with the application of the procedure in the aviation industry.

In **Chapter 3**, we examine the indirect effects of the advancement of AM technology on after-sales service supply chains. In particular, we study **RQ 2** and show that design changes caused by consolidation have various effects on the total life cycle cost.

In **Chapter 4**, we discuss when and how a organization should move to AM sourcing during the service period and thereby address **RQ 3**. A case study conducted in the defence industry exemplifies the situation and reveals, in combination with numerical experiments, how evolving conditions such as a decreasing AM piece price or a shrinking installed base size may effect the transition to AM.

In **Chapter 5**, we study the particularities of using AM as dual sourcing option for spare parts supply. Furthermore, we propose an exact model that we use to study **RQ 4**, i.e., we show under which conditions a dual sourcing approach pays off in comparison to single sourcing with CM or AM methods.

In **Chapter 6**, we address **RQ 5** and build and assess the performance of the iterative procedure that may be used to evaluate the benefit of using AM as dual sourcing option for large problem instance. Furthermore, we discuss the option to model the problem with approximate dynamic programming.

In **Chapter 7**, we draw our main conclusions and discuss options for future research.

Chapter 2

Identifying Spare Parts for Additive Manufacturing

In this chapter¹, a scoring method is developed for the identification of promising spare parts for the application of additive manufacturing (AM) technology. This corresponds to the first contribution mentioned in Section 1.4.3. The method is designed to rank several thousands of spare parts according to their possible benefit when produced by AM. We focus on economic and technological information that is available from standard databases in industry, such as ERP systems. The output of the method enables practitioners to prioritize and, therefore, to focus on the most promising parts first. As a result, this approach increases the effectiveness and efficiency of selecting promising business cases in after-sales service logistics. We demonstrate the approach by means of a case study conducted at a part supplier in the aviation industry. A validity and robustness study provides evidence that the method offers a suitable approach to prioritize a large spare part assortment.

2.1 Introduction

As we have discussed in Section 1.4.1, current methods for the identification of interesting parts for AM technology typically rely on a bottom-up procedure. That is, practitioners suspect that a part may be interesting for AM technology and propose the part for further consideration. However, for the spare part environment bottom-up procedures may entail disadvantages. First, they rely on the expertise of practitioners, which might be limited in after-sales service logistics and thus may lead to unsatisfactory results. Second, the evaluation tends to focus on technological feasibility rather than advantages from a supply chain perspective. As a consequence, practitioners typically consider only a relatively small

¹This chapter is based on Knofius et al. (2016).

part of the overall assortment and may overlook other parts for which the business case appears less intuitive. For example, consider a case where it is likely that the manufacturing costs increase, but the resupply lead time decreases. At first, such a case might be ignored. If the entire life cycle costs are considered though, the positive effects of a shorter lead time may outweigh the negative effects of higher costs (Van der Heijden et al., 2013).

We can avoid this type of problems if we rely on a top-down approach that can be initiated with a large part population. As such, it is possible to prioritize the analysis based upon potential economic benefit. This mitigates the risk of disregarding promising parts and additionally increases the efficiency. Furthermore, dependency on the expertise of practitioners can be decreased and thus the chance of underestimating logistical improvements is reduced.

The key contribution of the study reported in this chapter is to develop and validate a top-down approach to identify promising spare parts from a large assortment using information that is typically available in ERP systems. Note that the reference for the value to print a spare part is its current functionality. New functionalities that may be added using AM, are not considered since each part would have to be analyzed extensively, which would corrupt the efficiency of the top-down approach. Instead, the opportunity to add new functionalities should be addressed separately. To that end, it may be worthwhile to combine the top-down method with a procedure as proposed by Lindemann et al. (2015) (cf. Section 1.4). This has the advantage that a large part population is considered, while the attention of company representatives and AM experts is directed to the most promising parts.

The chapter is organized as follows. Section 2.2 describes the scoring method. In Section 2.3, we outline the results of a case study conducted in the aviation industry and validate the ranking method. In Section 2.4, we perform sensitivity analyses to examine the robustness and to demonstrate the flexibility of the approach. Finally, Section 2.5 states the conclusions to be drawn from this study.

2.2 Ranking Method

The objective of the top-down approach is to obtain a ranking which specifies the potential of AM for a spare part relative to the other analyzed spare parts from the perspective of supply chain management. In this section, we give an overview of the method. Later, we elaborate on the details in separate sections.

In the first step, we select the spare part assortment for the analysis. As will be clarified in Section 2.2.1, it is not recommended to always take the entire spare part population into account. Next, we score the resulting spare parts based on values of *spare part attributes*, which can be retrieved from the company databases. Table 2.1 gives an overview of relevant spare part attributes and summarizes which value level of a spare part attribute may indicate an improvement potential with AM. We explain the underlying logic of the assignment in further detail in Appendix 2.A.

Furthermore, we assess whether the spare part complies with technological constraints that are enforced by the current advancement of AM technology. For this purpose, we

Table 2.1: Value range of spare part attributes that indicate improvement potential with AM technology. Read: if spare part attribute ‘x’ belongs to value level ‘y’ then this indicates improvement potential ‘z’.

Attribute	Improvement potential						
	Reduce manufacturing/ordering costs	Reduce direct part usage costs	Reduce safety stock costs	Improve supply chain responsiveness	Post-ponement	Temporary fix	Reduce effect of supply disruptions
Demand rate	Low		Low		Low		
Resupply lead time			Long	Long	Long	Long	
Agreed response time			Short	Short		Short	
Remaining usage period		Long					
Manufacturing/order costs	High						
Safety stock costs			High		High		
Number of supply options	Few			Few			Few
Supply risk				High			High

introduce so-called *Go/No-Go attributes* to exclude parts which certainly do not qualify for AM from a technical perspective. Here, we consider the rather basic attributes *material type* and *part size*. Other constraints, for example associated with the geometric shape and tolerances for manufacturing, are usually more difficult to assess based solely on information that is easily obtainable from databases. In addition, such requirements may not be adequately represented by the conventionally manufactured part. For instance, it may be that the conventional manufacturing process may yield over-dimensional technical solutions.

However, the suggested Go/No-Go and spare part attributes are provided for orientation purposes only. Company-specific data availability may not allow utilizing exactly the same indicators. For example, in the case study we discuss in Section 2.3, we encountered difficulties to obtain data about the part size. Yet, the part identification number turned out to be a good proxy to obtain insights about the part size in this company. Furthermore, company-specific attributes may be available. For instance, we experienced a case in which a keyword indicated whether the company held the design rights for a spare part. Since the design rights indicate lower setup costs of the AM process, we used it as indicator for possible manufacturing cost reductions.

Next to considering Go/No-Go and spare part attributes, we also take *company goals* into account. This is motivated by different objectives of companies which may influence the preference to print specific spare parts. As such, some companies may focus on cost reductions, while others may prefer to improve their service despite higher costs. In the method, we use these preferences to derive a *weight* for each spare part attribute. These weights then determine how much influence a certain spare part attribute may have on the overall part score. In Section 2.2.2, we explain how these weights are derived from company goals.

Finally, based on the attribute weights and values, we compute weighted average scores for each spare part, and rank the analyzed spare part assortment accordingly. The ranking reveals which spare parts are more promising than others for the specific company. We discuss the scoring procedure in Section 2.2.3.

2.2.1 Determining the Spare Part Assortment

In order to allow the ranking of a large spare part assortment, the retrieval of data is limited to database queries or in-house analysis tools. Furthermore, to facilitate a proper comparison, one needs to have information on an attribute for a large portion of the analyzed spare parts. To achieve this may appear challenging if spare parts of different types of assets are considered. This holds even more if the company is operating in different supply chains. For example, an asset user or service provider may have more information about operational data than an OEM. As a consequence, in some companies only a subset of the entire spare part population may be selected. Also, separate analyses may become necessary if the data quality and availability varies over sub-populations. These decisions constitute the first step of the ranking method, and should be taken in close collaboration with company representatives.

The output of this step is an overview of the selected spare part assortment, the Go/No-Go attributes, the spare part attributes and the associated values. An example can be seen in Table 2.2. Note that only a subset of the relevant spare part attributes (cf. Table 2.1) is shown for illustrative purposes. This subset is used in the graphics throughout the remainder of this chapter.

Table 2.2: *Result first phase*

Part ID	1	2	3	...
Material type (<i>Electronic, Metal, Plastic</i>)	E	P	M	...
Part size (dm^3)	1	3	4	...
Supply risk (%)	21	50	35	...
Remaining usage period (month)	21	56	12	...
Supply options (#)	1	14	3	...
Manufacturing/ order costs (10,000 Euro)	5	15	1	...
...

2.2.2 Obtaining the Weight for the Spare Part Attributes

In this section, the company-specific attribute weights are derived from the company goals. To define suitable company goals, the classification scheme of Chopra and Meindl (2016) is used. They differentiate between responsive and efficient supply chains: one focusing on increasing flexibility and one focusing on reducing costs. To allow for more precision, the former is further distinguished into operational flexibility and strategic flexibility. While operational flexibility refers to the ability to match supply and demand, strategic flexibility here means the ability to handle potential supply disruptions in the future. Using a more practical terminology, this results in three company goals: *secure supply*, *reduce downtime* and *reduce costs*.

To evaluate the company-specific importance of each company goal, a pairwise comparison approach following the logic of the Analytic Hierarchy Process (AHP) is used. In the first stage, decision makers of the company give a score for each pair of company goals that

indicates which company goal has a higher priority. These scores allow for an approximation of an importance measure of each company goal relative to the other company goals. Due to the pairwise comparison, inconsistency becomes controllable and decision complexity is prevented. For a review of the AHP method, we refer to Section 1.5.1.

Next, the spare part attributes are assigned to the company goals. Given that the spare part attributes have already been allocated to improvement potentials (cf. Table 2.1), the relation between spare part attributes and company goals is established by assessing improvement potentials to company goals. This results in the allocation shown in Table 2.3. The motivation for this allocation is given in Appendix 2.B. Note that Go/No-Go attributes are not assigned to any company goal. They describe the technological feasibility of printing the spare part and are therefore independent of the company.

Table 2.3: *Spare part attributes assigned to company goals*

Attribute	Company goals		
	Secure supply	Reduce downtime	Reduce costs
Demand rate		X	X
Resupply lead time		X	X
Agreed response time		X	X
Remaining usage period			X
Manufacturing/order costs			X
Safety stock costs		X	X
Number of supply options	X	X	X
Supply risk	X	X	

In the second stage of the AHP method, pairwise-comparisons between the assigned spare part attributes for each company goal are performed. Accordingly, practitioners were asked the following type of question: *“If we improve both attribute values for the entire spare part assortment, which attribute does support the achievement of the company goal X better?”* This results in importance measures of the attributes.

Finally, to obtain the spare part attribute weights, importance measures of the attributes are multiplied by the importance measures of the associated company goal. Figure 2.1 provides an example of a typical result. Note that in case an attribute is allocated to more than one company goal, the weight equals the sum of all partial weights (cf. Figure 2.1: For example, the attribute supply risk obtains a weight of $0.22 + 0.105 = 0.325$). It has to be stressed that the resulting weights have to be understood as estimates. Even though the AHP method represents a well-established scientific approach, subjectivity in decision-making may lead to inaccuracies. In Section 2.4.1, sensitivity analyses will be used to quantify the consequences of these inaccuracies for the ranking.

2.2.3 Calculate the Overall Score of a Spare Part

After the spare part attribute weights have been computed, scores for each attribute value are calculated. For Go/No-Go attributes, a binary scoring is applied, i.e., if the attribute

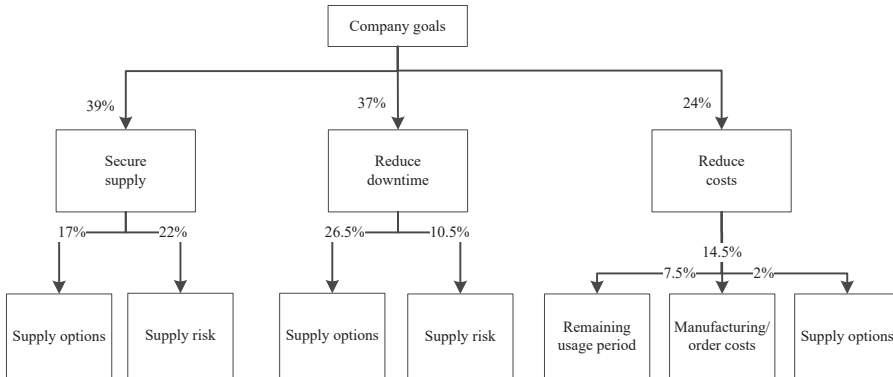


Figure 2.1: Example attribute weighting

value is located in the technologically feasible range, it is assigned a “1”, otherwise a “0”. A good estimation of the feasible range can be obtained through technical data sheets of recent AM machine releases.

For other spare part attributes, a linear scoring approach is used. That is, the value range of all spare parts is normalized: the best value receives a score of “1”, and the worst value a score of “0”. Values in-between receive a proportional score. Alternatively, one may use the 95% percentiles instead of the extreme values to protect against data pollution. In this case, all values exceeding the 95% percentiles obtain the score of the corresponding extreme value. Next, the weighted score for each spare part attribute is calculated by multiplying the score with the attribute weight. In a final step, the following procedure is applied in order to obtain the overall score for a spare part:

1. Multiply the scores of the Go/No-Go attributes.
2. Sum the scores of the spare part attributes.
3. Multiply the results of 1 and 2. Already one “No-Go” results in a score of “0”.

The final score for a spare part can range from “0” to “1”, where “1” represents the highest possible score. An example for one spare part can be seen in Table 2.4. Next to linear scoring of spare part attributes, other scoring methods might be employed. For example, consider a five or two point scale where the scale thresholds are determined by the value distribution of all spare parts. That is, for a two point scale, the worst 50% receives a score of “0”, and the other 50% a score of “1”.

It needs to be guaranteed, however, that the scoring approach provides a sufficient differentiation between the analyzed spare parts. To clarify, it would be undesirable if nearly all spare parts obtain the same score, as this would prohibit sufficient prioritization.

Table 2.4: Example final data of a spare part

Attribute	Value	Weight	Score	Weighted score
Material type	Metal	-	1	1
Part size	0.5	-	1	1
Supply risk	20	32.50%	0.21	0.06825
Remaining usage period	15	7.50%	0.31	0.02325
Supply options	5	45.50%	0.48	0.2184
Manufacturing/order costs	48	14.50%	0.24	0.0348
Total score:				0.3447

Section 2.4.2 provides an overview of the resulting differentiation for several scoring methods and study the impact on the final ranking.

2.3 Case Study

The ranking method outlined above was tested during a case study at a part supplier in the aviation industry, with more than 400,000 spare parts. Section 2.3.1 will elaborate on the application of the approach and highlight the key findings. In Section 2.3.2, the prioritization mechanism of the model is validated. For this purpose, a stratified sample of parts was selected from the ranking, and compared to the opinion of the implementation manager for AM technology at the company.

2.3.1 Application and Findings

After an evaluation of data availability and data cleaning, it was decided together with company representatives to base the analysis on 40,330 spare parts. As explained in Section 2.2.1, we selected this subset since it offered a rather consistent data availability across spare parts. The analysis was initiated with the eight spare part attributes as specified in Table 2.1 and the two Go/No-Go attributes.

As no suitable data source for the *agreed response time* was available, however, this attribute had to be dropped. Furthermore, the *part number* was used as a substitute for part size. This replacement was chosen because direct information about the *part size* was often not accessible. Fortunately, the company-specific numbering system relates part size to the part number and thus is a good proxy. In addition, the attribute *airplane type* was used instead of *number of supply options*. It was found that for spare parts which are exclusively used in specific airplane types, demand can be fulfilled by dismantling phased-out airplanes. Other information about the number of supply options was not easily retrievable.

Finally, the attribute *supply risk* was substituted by the attribute survival probability, where the survival probability defines the chance that a spare part supplier will be available within one year. This measure was available in this company for most of the analyzed spare

parts and was once calculated based on the model by Li et al. (2016). Table 2.5 gives an overview of all attributes associated with the weights derived from the AHP method.

Table 2.5: *Spare part attributes used in first case study*

Attribute	Weight	Explanation
Part number	-	The part number gives insights about spare part size
Material type	-	Indicates the material type e.g. electronic, composite or metal
Safety stock costs	18%	High safety stock may be reduced with AM
Manufacturing/order costs	17%	High sourcing costs may be reduced with AM technology
Demand rate	16%	For low volume production AM may reduce order costs
Survival probability	13%	Spare parts with high supply risk could be obtained with AM
Remaining usage period	13%	A longer usage period indicates higher operational saving potentials
Resupply lead time	13%	AM may reduce long resupply lead time which decrease safety stocks
Airplane type	10%	Specific airplanes obtain less spare parts from dismantling

Due to the Go/No-Go attributes, 34.140 of the analyzed spare parts were classified as not feasible to print from a technological perspective (in the near future). The remaining 6.190 spare parts were ranked, which resulted in a score distribution as shown in Figure 2.2.

Based on the ranking, the case company could already identify 1.141 technologically feasible and economically beneficial business cases. A typical example is a fitting stud used for the attachment of a safety belt as illustrated in Figure 2.3. For this case, it is estimated that it will be possible to reduce the resupply lead time by about 40% and the order costs by about 70% with AM. The prospect of this improvement potential stimulated a reengineering project for the fitting stud despite high costs for certification. This outcome demonstrates the benefit of the developed top-down approach: practitioners probably would have disregarded the fitting stud due to the high certification costs. In comparison, the




Figure 2.2: *Score distribution first case study*

ranking method typically exposes promising characteristics for high scoring items and thus justifies an assessment of the part in more detail.




Figure 2.3: *Safety belt with fitting stud (gray)*

2.3.2 Validation

The obtained ranking was compared with the opinion of the implementation manager for AM technology at the company in order to validate the prioritization mechanism. To that end, a stratified sample of 18 spare parts was selected from the ranking. The sample is divided in three subgroups: parts with scores larger than 0.8 (1), parts with scores between 0.4 and 0.6 (2) and parts between 0.01 and 0.4 (3), respectively. Note that a score of 0 indicates that the spare part is not printable. Therefore these cases are not considered in this analysis. Next, without knowledge about the rank, the AM manager was asked to assign a priority to each of the 18 spare parts — namely, most interesting (1), maybe interesting (2) and least interesting (3) for AM.

We show the results in Figure 2.4 (a). For clarification, we also show the results if AM manager and ranking method were to agree for each case in Figure 2.4 (b). As we observe in Figure 2.4 (a), the judgement of the AM manager appears more conservative than the ranking method, i.e. the AM manager assigned each of the 18 spare parts in the same or a lower priority class than the method. To that end, statistical tests (in so far feasible with the small sample size) would perform poor due to the overestimation with the ranking method. For instance, the Cohen’s Kappa, which measures the agreement between two classification mechanisms while accounting for chance agreement, is equal to 0.21. In the literature, this score level is typically regarded as indicative for a poor fit, cf. Landis and Koch (1977) and Fleiss et al. (2003). However, since the judgment of the AM manager can only offer an approximation for the true value of applying AM, and the sample size only allows an evaluation based on the score rather than the ranking itself, this result is less conclusive.

More detailed analysis showed that the different classification typically was a consequence of specific information which is difficult to include in a generic top-down approach. For example, consider the case which was assigned to Category (1) by the ranking method, yet the AM manager assigned it to the least interesting Category (3). In this case, the AM manager took into account that the supply of the spare part is about to be discontinued, but a cheap offer for a final order is available because the supplier of this spare part wants to sell the

		AM manager		
		1	2	3
Ranking method	1	2	4	1
	2	0	2	5
	3	0	0	4

(a)

		AM manager		
		1	2	3
Ranking method	1	2	0	0
	2	0	6	0
	3	0	0	10

(b)

Figure 2.4: Categorization of selected spare parts where (1) indicates most, (2) maybe and (3) least interesting for AM. Illustration (a) shows the result of the test. For clarification, Illustration (b) shows the results if ranking method and AM manager would agree for each case.

remaining inventory of this item. As a result, it does not appear interesting for the company to invest in an AM substitute, because the remaining demand can be covered economically.

For another case for which AM manager and ranking method disagree, the AM manager argues that the spare part has a demand rate of less than one part per year. Thus, the spare part would not be suitable for an “engaging” proof of concept to higher management. Instead, the AM manager prefers a part with a higher demand rate in order to demonstrate the benefits of AM on a more regular basis. From a political point of view this argument is reasonable, though it is rather questionable from an economical point of view. By adapting the scoring method, however, this political aspect could be taken into account. For instance, one may truncate the linear score for the demand rate and assign a score of “0” if the demand rate is below a certain threshold.

Both examples demonstrate that specific information or personal preferences may influence the identification of the right spare parts for the problem owner. However, some results also demonstrate the value of the ranking method. All four items categorized as least interesting by the ranking method were assigned to the least interesting category by the AM manager. Likewise, those items that were indicated as most interesting by the AM manager were scored highest by the model as well.

2.4 Sensitivity Analyses

As explained in Section 2.2.2, estimating the attribute weights may incur inaccuracies because of the subjectivity of decision-making. In Section 2.4.1, sensitivity analyses will be carried out to assess the consequences for the ranking. In Section 2.4.2, different scoring methods will be assessed with regard to their applicability. Also, the impact on the ranking will be evaluated.

2.4.1 Consequences of Inaccurate Weights

Inaccuracies are less worrisome if the ranking is rather insensitive to spare part attribute weight changes. Hence, the robustness of the ranking towards weight changes is assessed in this section. The following analysis is based on data from the case study carried out in the aerospace industry (cf. Section 2.3).

To evaluate the robustness, sensitivity analyses are performed on the spare part attribute weights where the change in the ranking is measured. This is achieved by computing the correlation between the actual ranking of the case study and the new ranking obtained by varying one weight. Note that varying one weight leads to change in all other weights because the relative importance of all weights changes. To measure the correlation, Spearman's rho is applied (Kornbrot, 2005). As a matter of course, Spearman's rho is equal to 1 for the weights used in the case study. Furthermore, a correlation of more than 0.5 is referred to as significant.

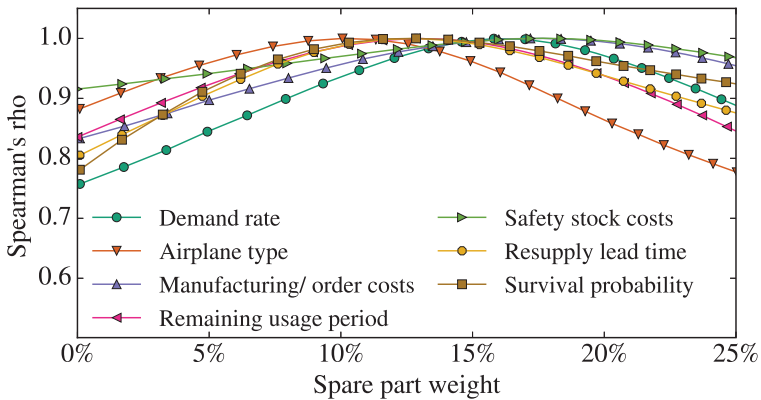


Figure 2.5: Sensitivity of attribute weights

As can be seen in Figure 2.5, the correlation between the rankings remains significant even if a spare part attribute weight is changed by more than 15%. Comparable behavior can be observed if the analysis is limited to the 10% top scoring spare parts. It can be concluded that the ranking appears robust against inaccuracies resulting from the AHP method.

2.4.2 Consequences of Different Scoring Methods

In Section 2.2.3, a linear scoring approach was proposed to evaluate the spare part attribute values. This has the advantage that most of the available information is considered, and therefore high differentiation among the spare parts can be achieved. In a practical setting however, it may appear useful to deviate from the linear scoring approach. Generally, this does not cause problems, as long as sufficient differentiation among the scores can be guaranteed. To clarify, if nearly all spare parts obtain the same score, the ranking is less useful.

Subsequently, it will be demonstrated that this requirement can be fulfilled with several scoring procedures. Also, it will be shown that the effect on the ranking is acceptable and does not yield considerable deviations. For both analyses, we again use data from the case study discussed in Section 2.3. Of course, it is impossible to conceive all possible scoring methods. Subsequently, the analyses will be limited to the comparison of linear scoring to scoring with a two point and a five point scale (cf. explanations in Section 2.2.3). Other scoring methods however, may be analyzed in the same manner.

First, the degree of differentiation of the overall scores is assessed by determining the number of unique scores for the 6.190 spare parts considered in the aerospace case study. As shown in Table 2.6, the two point scale approach yields the lowest differentiation with 143 unique scores. Even though this is substantially less than what can be achieved using the linear scoring approach (5753), this degree of differentiation should be sufficient to prioritize further analyses for most applications. Second, the effect of different scoring procedures on the ranking is evaluated. Using the same approach as in Section 2.4.1, Spearman's rho was computed between all three rankings using the linear scoring approach as a benchmark. Thus, the Spearman's rho is equal to 1 for this procedure.

As can be found in Table 2.6, all three scoring methods are significantly correlated. This indicates that it is acceptable to deviate from the proposed linear scoring approach and thus gives additional flexibility for the application in practice.

Table 2.6: *Unique scores depending on approach and ranking correlation*

Scoring type	Unique scores [#]	Unique scores [%]	Spearman's rho
Two point scale	143	5%	0.75
Five point scale	628	10%	0.81
Linear scoring	5753	94%	1

2.5 Conclusion

In this chapter, a method was developed to increase the transparency in the decision-making process of which spare parts may benefit from AM. An argument was made that such a method is required, as available concepts may underestimate the potential of AM, in particular for after-sales service supply chains.

A case study demonstrates the value of this method, as it facilitated the identification of more than 1000 technologically feasible and economically beneficial business cases. Simultaneously, this result shows the practical benefit of AM for after-sales service supply chains. A validation study gives evidence that the prioritization mechanism underlying the method supports distinguishing between interesting and less interesting cases. To that end, the method appears suitable to the task of prioritizing a large spare part assortment and thus makes the selection of spare parts more effective and efficient. The method was further shown to be robust against possible inaccuracies of spare part attribute weights that

may result from subjectivity of decision-making. Moreover, different scoring procedures were studied and found to be eligible for the ranking method, thus providing flexibility in terms of practical application. In conclusion, companies may be encouraged to use this approach in order to simplify the identification of promising spare parts for AM.

For future research efforts, it might be worthwhile to extend the proposed method by considering possible design improvements. For example, one might try to identify indicators that relate to the probability with which an assembled spare part can be printed as a single part and use these indicators as additional spare part attributes in the method. Furthermore, we found that the data collection step is usually the most challenging. In particular, organization-specific terminology or decentralized data storage often complicates this step. Statistical analyses or data mining techniques, such as clustering, may have the potential to simplify these steps.

Appendices

2.A Interpretation of Spare Part Attributes

A *low demand rate* often indicates high demand variability. This may lead to inefficiencies in the manufacturing process due to high tooling and setup costs (and thus incurring higher purchasing costs in case a supplier manufactures the spare part). Additionally, high demand variability results in relatively high safety stock costs to fulfill service level agreements. This may be reduced by decreasing the resupply lead time with AM. Furthermore, countermeasures against high uncertainties such as emergency shipments or supplier dedicated stock can be replaced, for example by printing on demand and therefore postponing the production decision. Note that the demand rate was chosen instead of the demand variability, because data about demand variability is typically more difficult to obtain – particularly for slow moving items.

A *long resupply lead time* may result in high safety stock costs or high system downtime, because the variability of the lead time demand is usually high. By reducing the resupply lead time with AM, one may therefore reduce the safety stock costs or downtime. Simultaneously, countermeasures against high safety stocks or long downtimes such as emergency shipments or supplier dedicated stock can be avoided because of an increased responsiveness. In a best case scenario, it is possible to print on demand and therefore transform the supply chain from a make-to-stock to a make-to-order setup. Finally, long resupply lead times may offer potential to use a temporary fix in order to reduce the safety stocks or downtime.

If the agreed *response time is short*, safety stocks are often located close to the customer site. This reduces pooling effects and therefore may lead to relatively high safety stock costs. AM technology may enable production on location or shorter resupply lead times, and thus decrease safety stock costs. Furthermore, printing on location or obtaining the spare part within a shorter resupply lead time yields a higher responsiveness of the supply chain. Accordingly, concepts like emergency shipments and supplier dedicated stock may become obsolete or may be replaced by a more efficient temporary fix.

If the *remaining usage period* of a spare part is *long*, the recurring direct usage costs of the spare part may be reduced more often and thus this situation offers the highest potential. For example, repair costs and assembly costs may be lower with an AM manufactured part. If the *manufacturing/order costs are high*, AM technology may offer a cheaper way to produce a spare part that can fulfill the same function. If the current *safety stock costs are high*, AM technology may reduce the resupply lead time and thus lead to lower safety stock costs. Also, service efforts like emergency shipments and supplier dedicated stock may be avoided. For example, it may be possible to print on demand (i.e. postpone the production decision) and therefore avoid safety stock costs.

If there are only a *few supply options* for a spare part, AM may offer a chance to reduce order costs because an additional supply option improves the negotiation position. Furthermore, AM might increase the flexibility, for example by employing a Dual Sourcing concept. Finally, the additional option to print the spare part may become important if

the regular supply is discontinued. If the *supply risk is high*, i.e. suppliers may permanently discontinue the production of the spare part soon, AM may be useful to obtain a more reliable supply source. Furthermore, a high supply risk implies less flexibility to deal with demand and supply variations. AM could increase this flexibility.

2.B Assignment of Spare Part Attributes

In order to obtain an importance measure, i.e. weight, for each spare part attribute, each spare part attribute needs to be assigned to a company goal. Given that the spare part attributes have already been allocated to improvement potentials (cf. Table 2.1), the relation between spare part attributes and company goals can be established by assessing which improvement potential is associated with which company goal.

The improvement potentials *reduce manufacturing/order costs*, *reduce direct part usage costs* and *reduce safety stock costs* describe the chance to improve the efficiency with AM. This aligns with the company goal to reduce costs. The remaining improvement potentials describe the ability to increase the flexibility with which a certain service function can be fulfilled. This affects both the company goal *reduce downtime* and the company goal to *secure supply*. As is explained in Section 2.2.2, however, the company goal *secure supply* is associated with the ability to handle potential supply disruptions in the future. This is represented by the improvement potential *reduce effect of supply disruptions*. The other improvement potentials describe the operational flexibility, i.e. the ability to match supply and demand. The assignment is visualized in Table 2.7.

Table 2.7: *Spare part attributes assigned to company goals*

Attribute	Reduce costs			Reduce downtime			Secure supply
	Reduce manufacturing/ordering costs	Reduce direct part usage costs	Reduce safety stock costs	Improve supply chain responsiveness	Post-ponement	Temporary fix	Reduce effect of supply disruptions
Demand rate	Low		Low		Low		
Resupply lead time			Long	Long	Long	Long	
Agreed response time			Short	Short		Short	
Remaining usage period		Long					
Manufacturing/order costs	High				High		
Safety stock costs			High				
Number of supply options	Few			Few			Few
Supply risk				High			High

After assigning the improvement potentials to company goals, the relation between spare part attributes and company goals can be established. That is, each spare part attribute which is assigned to a particular improvement potential is assigned to the respective company goal. This results in Table 2.3.

Chapter 3

Consolidating Spare Parts with Additive Manufacturing

While there are various arguments for considering additive manufacturing (AM) technology as an option to improve the performance of after-sales service supply chains, AM technology may also have indirect effects on after-sales service supply chains. Most evident indirect effects may originate from design changes which are the key driver for many applications with AM technology. As we have discussed in Section 1.1.1, companies may exploit the design freedom of AM technologies to shift to a more functional design. Common examples are a lower weight or a more efficient raw material utilization while achieving a comparable durability. Although design changes may hold the promise of operational costs savings, the consequence for the total life cycle costs are often less clear. In this chapter¹, we consider the most prominent application of the design freedom obtained with AM technology and study the effects of consolidation. With this study, we address the second contribution mentioned in Section 1.4.3.

3.1 Introduction

As we elaborated in Section 1.1.1, consolidation of parts is the redesign of an assembled component with fewer, but therefore more complex parts. While complex parts are often difficult to produce with conventional manufacturing (CM) technologies, the high degree of design freedom of AM facilitates consolidation. In Appendix 3.A, we illustrate a typical case of consolidation which demonstrates the technological feasibility of consolidating parts with AM and its associated potential for operational cost savings.

¹This chapter is based on the paper Knofius et al. (2018b).

Next to reducing the number of assembly steps, and thereby production lead time and costs, consolidation may improve the reliability of assembled components, see Johnson and Kirchain (2009) and Wits et al. (2016). Couplings between parts, often the cause of several failure modes, can be removed. Furthermore, the performance of the consolidated part may be improved. In this context, performance refers to aspects like less flow resistance or improved heat dissipation. Moreover, the supply chain might be simplified because the number of distinct parts that need to be sourced, tracked and inspected decreases. Hence, operational complexities diminish and long lead times are reduced (Yang et al., 2015). Finally, the general benefits of AM technologies apply, as elaborated earlier.

Despite the potentials of consolidation, it would not be realistic to assume that consolidation is always preferable. In fact, it is unclear under which conditions a consolidated part should be preferred over its CM and assembled counterpart. This is due to several potential disadvantages of consolidation. For instance, consolidation might remove the option to repair a defective system via the replacement of sub-components. In particular, the disassembly of the component would no longer be possible and thus the replacement of the entire component is inevitable. This may lead to additional acquisition costs and stocking of more complex parts compared to the assembly case.

Similarly, possible commonality effects, achieved by pooling stocks of common low level parts, are lost because of the higher level of customization of the consolidated part. Another shortcoming concerns the potentially higher purchasing costs because of more specific parts as well as the general novelty of industrial AM processes. Weller et al. (2015) point out that the latter may also increase the number of required production steps. For instance, current AM processes typically require support materials which have to be removed in separate production steps. In addition, post-processing steps are generally required to increase the surface quality or to reduce residual stress.

At present, process setup and calibration demand close cooperation between producer and user. These inherent dependencies limit supply options and may increase consolidation costs. In the future, however, it is likely that the number of suitable suppliers may rise significantly. Generic AM processes support the production of a wide product range with the same equipment while conventional suppliers are often limited to fewer products or dependent on suppliers themselves. Hence, costs associated with supplier dependencies are likely to decrease for consolidation but also for AM production in general. A final drawback of consolidation with AM concerns the spare parts characteristics which may be affected negatively. For example, metal based AM technologies often result in parts that suffer from porosity despite post-processing, cf. Cunningham et al. (2017). As a consequence, mechanical properties are compromised because pores act as a possible origin for cracks. Opportunities and drawbacks of consolidation are summarized in Table 3.1.

As discussed in Section 1.4, the literature does not quantify the consequences of consolidation. Instead, it is assumed that no design changes occur or that the effects of design changes are known a priori. Such assumptions are problematic for three reasons. First, design changes are imposed by AM technologies cf. Wits et al. (2016) and Lindemann et al. (2015). Second, as discussed by Thomas (2016), the merits of consolidation tend to be misinterpreted.

Table 3.1: *Characteristics of consolidation with AM technology*

Opportunities	Drawbacks
+ Less assembly steps	- No replacement of sub-components
+ Shorter lead times	- Lost commonality effects
+ Simplified supply chain	- Potentially higher purchasing costs
+ Potentially higher reliability	- Potentially lower reliability
+ Performance improvements	- Stocking of more complex parts
+ Reduced tooling and setup effort	- Post-processing
+ Lower raw material usage	- Currently high supplier dependencies
+ Future low supplier dependencies	

Third, consolidation may significantly influence the life cycle costs of capital goods as we will discuss later in this chapter. With this research, we aim to give a more complete perspective on the effects of consolidation with AM on the total life cycle costs. Therefore, we first identify cost drivers which are influenced by consolidation and then use existing methods for spare parts optimization to study the consequences of consolidation under different conditions. As our results will show, the value of consolidation often strongly depends on aspects other than just design improvements and its associated operational cost savings.

The remainder of the chapter is organized as follows. In Section 3.2, we discuss the literature related to the METRIC method in order to set the ground for the methodology applied in this chapter. We continue with a discussion about how consolidation may influence the life cycle costs of capital goods and identify essential cost drivers in Section 3.3. In Section 3.4, we explain the model and clarify our assumptions. Next, in Section 3.5 we conduct different experiments in order to quantify the effects of consolidation. We conclude with Section 3.6, in which we summarize the results and suggest directions for future research.

3.2 Literature Review on the METRIC Methodology

A seminal paper in the field of spare parts management has been written by Sherbrooke (1968) who introduces the METRIC methodology. In his paper, Sherbrooke considers a multi-item, two-echelon distribution network, which he evaluates with an approximation for the number of items in resupply. Based on convexity properties, Sherbrooke developed a marginal approach to optimize the inventory positions of each item. Later, this work was extended to more complex network structures and more accurate approximations for the number of items in resupply.

A first contribution on spare parts stocking decisions for assemblies was made by Sherbrooke (1971), who introduced the indenture level concept. The purpose of this concept is to categorize the material breakdown structure of an assembled spare part, more or less similar to a Bill of Material structure. A first indenture level part consists of second indenture level parts, etc. This categorization is used to organize the optimization of inventories

across different product or system hierarchy levels. A key assumption is that a failed first indenture level part can be repaired by detecting the failed second indenture level part(s), and either replacing them or continuing to search for the cause on the third level, etc. Sherbrooke's work was extended by Muckstadt (1973) to the MOD-METRIC model which considers a two-echelon, two-indenture system. Based on the approach of Graves (1985), Sherbrooke (1986) improved the MOD-METRIC model with a more accurate two-moment approximation for the number of items in resupply. The resulting approach is referred to as VARI-METRIC, and applicable to multi-echelon, multi-indenture systems with backordering. We base our analysis on the VARI-METRIC approach by Sherbrooke.

3.3 Effects of Consolidation on Life Cycle Costs

During the life cycle of expensive capital goods, different cost factors must be considered. Elmakis and Lisnianski (2006) distinguish the following cost categories: development costs, production costs, operation and service costs, and disposal costs. In this section, we describe how consolidation with AM may change these cost categories and explain how these changes are regarded in our analysis.

3.3.1 Development and Production Costs

Both the development and the production costs are largely influenced by the manufacturing process choice. Thus, the decision to use AM for the purpose of consolidation is likely to affect these cost categories. We justify this claim by the following observations: AM parts typically do not require tooling and less setup activities compared to CM. Additionally, assembly steps are fewer and in most cases we require less raw materials for the production of AM parts. The latter aspect decreases sourcing costs and supply chain complexity. On the other hand, industrial AM processes have not matured yet. Hence, the acquisition of AM machinery represents an expensive and risky investment which may lead to high piece prices of AM parts.

Moreover, quality constraints may demand extensive post-processing or rework of AM parts, which increases the production costs. However, developments in the field of process monitoring may reduce the need for post-processing in the future. An example is in-situ process monitoring for which features like temperature and size of the melt pool are measured with the goal to enable closed-loop control mechanisms; see Craeghs et al. (2012), Tapia and Elwany (2014) or Everton et al. (2016). These control mechanisms are able to reduce the high AM process variability and thus may decrease the number of post-processing steps required in the future. Also, various metal printing equipment producers incorporate automated post-processing functions in their machines. For instance, the MetalFab1 by Additive Industries supports automated heat treatment, part removal from the build plate and build plate storage Additive Industries (2018). While it seems probable that post-processing steps, in particular for metal prints, remain necessary in the foreseeable future, the time dedicated to post-processing activities is likely to decrease. Furthermore, such advancements may simplify standardization and qualification, which

have been identified as major challenges for the adoption of AM technologies in industry and significantly contribute to the typically high AM piece price (Frazier, 2014).

In our analysis, we encapsulate these considerations with the difference in unit costs between the AM part and the CM assembly. Depending on the supply chain layout, the unit costs may either be interpreted as purchasing or as production costs. In the following we focus our discussion on the situation of an external AM supplier. However, our results are similarly representative for the scenario in which AM production is carried out in-house. In this case, we would use the average AM throughput time instead of the contractually agreed AM replenishment lead times for our analysis, as is common in the METRIC-type of models. In fact, the assumption is that the inventory policy does not change the utilization of repair and resupply processes, so that it is sufficient to measure waiting times in these processes. Finite capacity models (e.g. Sleptchenko et al. (2003)) exist to facilitate the trade-off between investing in repair and production capacity or investing in spare part inventories. However, Song and Zhang (2016) find that in-house AM equipment utilization typically remains low in a multi-item, multi-source spare parts setting.

In either case, we do not regard fixed purchasing/production costs explicitly. AM parts typically do not require tooling and less setup since both the design and the manufacturing process are highly digitalized. To that end, low production quantities or even one-for-one production (no batching) becomes economically feasible and allows to incorporate setup costs in the unit costs. Of course, such a modeling decision does not appear appropriate in the presence of high fixed costs. For instance, consider the case in which tooling for the CM part was discarded and has to be rebuilt. Alternatively, a supplier may stop the support for a CM part. Today, these situations are often solved with expensive final orders, see e.g. Kleber et al. (2012). In both cases, the low AM preparation costs may favor the application of AM technologies. In this paper we do not regard these extreme cases, but note that future research is necessary to gain quantitative insights in the value of AM technologies under these circumstances.

The same holds for production technologies other than AM with which consolidation may be feasible. Examples are CNC processes and Flexible Manufacturing Systems (FMS). Although indeed the cutting tools exploited in CNC machinery and FMS are versatile (but constrained by their stand times), a part manufactured on a CNC machine does require special fixtures and clamping devices, often even build up or constructed for that particular part. In FMS, the combination of pallets and fixtures may be specific for particular products, and requires a setup performed outside the CNC machines (i.e. at the I/O station of the FMS). As a result, to extend our results for consolidation with other production technologies, it is perhaps more suitable to allow for batch production to account for fixed production costs explicitly. Since we focus on consolidation with AM exclusively, we defer these consideration to future research.

3.3.2 Operation and Service Costs

Often the operation and service cost category is the main contributor to the total life cycle costs of expensive capital goods (Öner et al., 2007). We separate this cost category into

operational costs, maintenance costs, and downtime costs.

Operational costs arise during, and as a result of, the direct operation of a system, such as electricity, personnel and fuel costs. Some of these cost factors may be affected by consolidation since the design change may cause different operational costs of the focal part. For example, consider the fuel nozzle discussed in Appendix 3.A. In this case, GE achieved a weight reduction through consolidation, and thus was able to reduce fuel consumption, which ultimately led to operational cost savings. However, in general these cost savings heavily depend on the specific case and are often difficult to measure exactly. To that end, we exclude the operational costs from the numerical experiments but note that, if for a specific case operational cost savings estimates were available, they can be easily incorporated as input to the model. As we show in Section 3.4.4 (Equation 3.4), the expected number of backorders is closely related to the expected number of operating parts. Hence by multiplying the expected number of operating parts with the average operational cost savings per part per period, we obtain an estimate for the associated operational cost savings per period.

Maintenance costs arise due to preventive and corrective maintenance activities and are required for the upkeep of capital goods. Maintenance activities and associated support functions, such as spare parts management, are likely to be affected by consolidation. As mentioned in Section 3.1, the failure behavior of consolidated and printed parts usually differs from those of CM parts. Failures of expensive capital goods caused by a dysfunctional part are typically solved by a repair-by-replacement policy where entire units are replaced to limit a system's downtime. Next, the removed units are sent to a specialized repair shop where the fault is diagnosed and solved by replacing a (typically far cheaper) sub-component. Even though this practice may lead to stocking of more expensive/complex parts, downtime and repair costs are typically reduced significantly. Furthermore, the stocking of lower-level items may reduce the repair lead time and thus decrease the necessity for a high base stock level of the higher-level items.

Consolidation of parts may obstruct the repair process, given that sub-components are not replaceable anymore. Instead, a new complex consolidated part has to be procured or manufactured. Hence, repair costs are likely to increase. On the other hand, if the consolidated part is replaced with a new version upon failure, part deterioration or repair process induced failures are avoided. For critical parts the consequences are limited, though, as detailed quality checks typically warrant a certain quality level for repaired items. Finally, shorter lead times may reduce stock levels and stock-out risks with consolidation even though the option to stock high level items only may increase the inventory costs again. As explained in Section 3.1, the lead time reductions are mainly achieved through a simplified supply chain and assembly process, and a (usually) shorter AM production lead time.

In order to capture the described differences in maintenance costs, we study the effects of different failure rates, replenishment lead times, holding costs and repair costs. Note that we focus on the corrective maintenance case in this chapter. For preventive maintenance, spare parts inventories are sometimes not required and thus the analysis is simplified. Alternatively, for example in case of condition-based monitoring, a more specific analysis becomes necessary that does not relate to the characteristics of consolidation specifically, but

rather to the short AM replenishment lead times in general. As we discussed in Section 1.2.2, it may become possible to postpone the repair decision without additional requirement for spare parts inventories and thus to collect further data about the condition of the monitored part. Future research should reveal if more evidence about the part condition justifies the application of AM technologies even despite possible drawbacks. Likewise, in our analysis we do not account for costs associated with the setup of possible repair shop infrastructure, but restrict ourselves to the case where repairs are carried out by external service providers.

Downtime costs arise if the capital good is non-operational; they are typically difficult to quantify because they relate to soft factors such as customer satisfaction and company image. We express the downtime costs in terms of an availability target for the system. This decision is justified by two observations. First, target availability is typically easier to assess intuitively than downtime costs and thus is more often specified in service contracts or functional specification documents. Second, there exists a relation between downtime costs and target availability, that is, higher downtime costs cause a higher target availability. In fact, it has been shown that in many cases a one to one relationship between downtime costs and service level constraint holds (Van Houtum and Zijm, 2000). Hence, if a company has insights in downtime costs, these are usable for determining a corresponding service level constraint. Such situations may for instance occur if the service provider reacts to a stock-out with an emergency supply that offers a suitable but more expensive substitute. Here we will compare the total life cycle costs of the AM part with the CM assembly, subject to a mutual availability constraint.

3.3.3 Disposal Costs

The last cost category are the disposal costs which arise during the phase out of the capital good. Consolidation may influence these costs due to differences in over- and underage costs. Overage costs refer to costs that arise if too many spare parts remain unused at the end of the life cycle. Often this situation leads to a depreciation of the value of the remaining stock. In case stock is remaining, consolidation may worsen this effect, as parts are more complex and specific, and hence more valuable. On the other hand, short AM replenishment lead times may permit smaller safety stocks which may compensate for this disadvantage.

Underage costs arise due to insufficient stock during the final phase of the product or system when regular supply is discontinued. This condition leads to downtime costs or additional charges to purchase spare parts during this final phase. Here, consolidation with AM technology may decrease the underage costs given that replenishment lead times are typically short and setup costs are low. As a consequence, the impact of supply discontinuations decreases. Unfortunately, it is difficult to predict during the design phase which of the two effects will eventually have a more significant impact, as aspects like decreasing demand and learning effects during the life cycle are difficult to foresee. Furthermore, potential over- or underage costs heavily depend on operational decisions, such as how to deplete the stock levels at the end of the life cycle of the capital good. For these reasons, we decided to exclude the disposal costs from the analysis. In an actual case where this insight is available however, one may include disposal cost differences between the AM part and the CM assembly as input into the analysis.

3.4 Model

In this section, we present the model that we will use to quantify the effects of consolidation. Therefore, we first present an outline of the model before we construct the mathematical model in the next subsection. Key assumptions are discussed in the text and are summarized in Section 3.4.3. We close this section with the description of the model evaluation and optimization.

3.4.1 Model Outline

Consider a single stock point which serves an installed base of systems with a critical component, i.e. once the component of a particular system fails, the entire system has to stop operation. The component is a multi-indenture item and may appear multiple times in each system. For ease of presentation, however, we assume that each component occurs only once in each system. Upon failure, the entire component is replaced by a stocked spare part to keep the downtime of the system short. Recall that for high-value assets, a repair-by-replacement is common practice in industry, see Section 3.3.

In case no stock is available, demand is backordered until a spare part becomes available. The failed component enters an uncapacitated repair process (if possible), which will return a ready-for-use spare part with a similar failure pattern as an entirely new item. We justify the modeling choice of an uncapacitated repair process with the common practice where throughput times are the subject of contractual agreements with the repair shop. This also means that the modeled repair times are in fact lead times in the repair shop, covering net processing times as well as waiting times for resources.

In our model, inventories are controlled according to a base stock policy. This modeling decision is justified by the fact that we focus on expensive, slow-moving items where batching independent of the replenishment process generally does not occur because the focus is on meeting tight service level requirements.

Depending on manufacturability constraints, different configurations of the component are feasible. Figure 3.1 illustrates a possible set of configurations of a component. Configuration A represents the design with the highest segmentation. That is, no consolidation of parts took place. In Configuration B, Parts 6 and 7 are consolidated with AM technology which results in a Part I with a specific failure rate, average replenishment lead time, holding and replenishment cost rate. Also, consolidation across indenture-levels is possible. This situation occurs for Configuration C where the functions of Parts 5, 6, 7 and 2 are replaced by Part II.

Finally, each component configuration may demand different replenishment processes. For instance, a consolidated part, like Part I and II in Figure 3.1, cannot be repaired by replacement of sub-components because these do not exist. Thus, a purchase takes place to obtain a substitute for Part I and II. In contrast, a failure of Part 2 can be solved by replacing the dysfunctional sub-component.

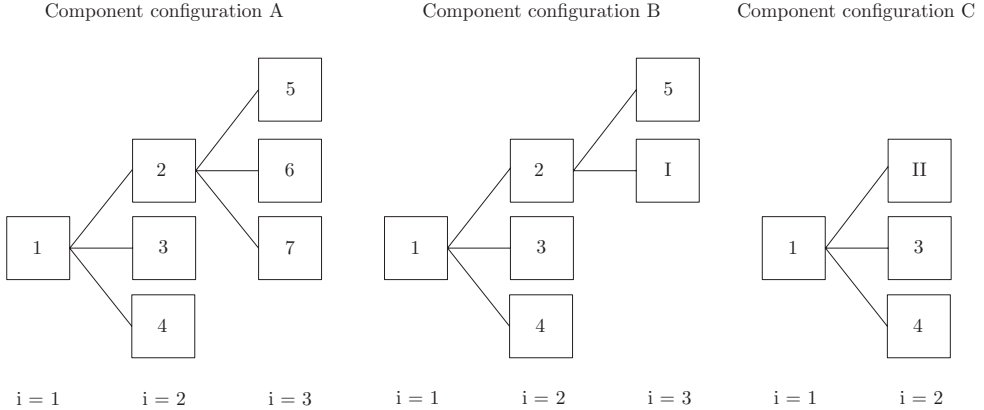


Figure 3.1: Possible component configurations

The goal of the model is to evaluate which component configuration minimizes the total costs subject to an availability constraint. Therefore, we jointly optimize the base stock level(s) and the component configuration.

3.4.2 Notation and Mathematical Model

In this section, we explain the mathematical model. Notations are introduced in the text and summarized in Table 3.2. The evaluation of the functions introduced in this sub-section, is presented in Section 3.4.4.

We consider an installed base of IB systems where we focus on one critical component. As we elaborate in Section 3.4.5, common methods allow an extension to a multi-item setting as well. The set of feasible configurations of that component is denoted by K . Each component configuration $k \in K$ is characterized by a multi-indenture structure and consists of a set of parts Z_k . For example, the set of parts of Configuration C in Figure 3.1 is represented by $Z_C = \{1, II, 3, 4\}$. Indenture levels are specified with i , $i = 1, 2, \dots, I$, where the set of parts at indenture level i is denoted by $Z_{ki} \subseteq Z_k$. Thus $Z_{C2} = \{II, 3, 4\}$. The set of children of part $z \in Z_k$ is denoted by $\Gamma_k(z)$. Accordingly, in Figure 3.1, the set of children of Part 2 in Configuration B is represented by $\Gamma_B(2) = \{5, I\}$. The set of parents of part $z \in Z_k$ is denoted by $\Psi_k(z)$ (note that since we assume that each component occurs only once in each system, $\Psi_k(z)$ consists of a single element, but we keep the set notation for the purpose of generalization).

The failure rate of the lowest indenture part $z \in Z_{kI}$ is described by λ_z . We assume that failures occur according to a Poisson process. For our purposes the Poisson demand assumption seems most appropriate, given that we consider a low-demand environment with mechanical parts that are dimensioned to outlast the intended maintenance interval or lifetime of the capital good (common practice for downtime critical mechanical components of capital goods). Accordingly, failures are random in nature and may, for instance, be

Table 3.2: Notation overview

Notation	Explanation
IB	Installed base size
K	Set of feasible configurations
k	Specific component configuration
Z_k	Set of parts in configuration k
i	Indenture level, where I denotes the lowest hierarchy level
Z_{ki}	Set of parts in configuration k at indenture level i
$\Gamma_k(z)$	Set of children of part z in configuration k
$\Psi_k(z)$	Parents of part z in configuration k
λ_z	Failure rate of a lowest indenture part $z \in Z_{kI}$
r_z	Replacement cost rate of part $z \in Z_k$
p_z	Purchasing cost rate of a lowest indenture part $z \in Z_{kI}$
a_z	Assembly cost rate of a higher indenture level part $z \in Z_k \setminus Z_{kI}$
$P_k(z)$	Total piece price of part z in configuration k
κ	Holding costs fraction
l_z	Average replenishment lead time/repair lead time of part z
x_k	Binary variable which indicates if configuration k is used
\mathbf{S}_k	Vector describing all base stock levels for configuration k
$s_k(z)$	Base stock level of part z in configuration k
$TC_k(\mathbf{S}_k)$	Average total costs of configuration k given \mathbf{S}_k
\bar{A}	Target availability of the installed base
$A_k(\mathbf{S}_k)$	Average availability for configuration k given \mathbf{S}_k
$m_k(z)$	Demand rate of part z
$EBO(\cdot)$	Expected number of backorders
$VBO(\cdot)$	Variance of backorders
$R_k(z)$	Parts in replenishment of part z in configuration k

caused by unintended stress levels, imprudence during maintenance/shipping activities or unobservable quality issues during the production. The failure rate of a higher indenture level part $z \in Z_k \setminus Z_{kI}$ is equal to the accumulated failure rate of its children.

If component configuration $k \in K$ is an assembly, i.e. $|Z_k| > 1$, the component can be repaired in the repair shop by replacing the failed part $z \in Z_{k2}$ with r_z being the cost rate of replacement. In case $\Gamma_k(z) \neq \emptyset$, the part $z \in Z_{k2}$ can be repaired by replacing the failed child $y \in \Gamma_k(z)$ with r_y being the cost rate of replacement, etc. Otherwise, if part $z \in Z_k$ has no children, i.e. $\Gamma_k(z) = \emptyset$, we discard the failed part and order a new part at a purchasing cost rate p_z . The holding costs per period (e.g. per year) of part $z \in Z_k$ are a fraction κ of the total piece price $P_k(z)$. In case we consider a part $z \in Z_{kI}$ on the lowest hierarchy level the total piece price is equal to the purchasing cost rate. For a part $u \in Z_k \setminus Z_{kI}$ on a higher hierarchy level, the total piece price is equal to the assembly cost rate a_z plus the total piece price of all its lower level parts included in $\Gamma_k(u)$. The average replenishment ($\Gamma_k(z) = \emptyset$) or repair lead time ($\Gamma_k(z) \neq \emptyset$) of $z \in Z_k$ is denoted by l_z .

We define the following decision variables:

$$x_k \forall k \in K, \text{ where } x_k = \begin{cases} 1, & \text{if configuration } k \text{ is used} \\ 0, & \text{otherwise.} \end{cases}$$

$\mathbf{S}_k = \{s_k(z) | \forall z \in Z_k\} \forall k \in K$, where $s_k(z)$ is the base stock level of part $z \in Z_k$.

The average total costs of configuration $k \in K$, given \mathbf{S}_k , are denoted by $TC_k(\mathbf{S}_k)$. The target availability for the installed base is given by \tilde{A} , while the average availability for a certain component configuration $k \in K$ and base stock policy \mathbf{S}_k is denoted by $A_k(\mathbf{S}_k)$. As a result, we have to solve the following non-linear integer optimization problem in order to find the minimum cost configuration. The evaluation of the functions in the optimization problem is discussed in Section 3.4.4.

$$\begin{aligned}
 & \underset{x_k, s_k(z)}{\text{minimize}} && \sum_{k \in K} x_k TC_k(\mathbf{S}_k) \\
 & \text{subject to} && \sum_{k \in K} x_k = 1 \\
 & && A_k(\mathbf{S}_k) \geq \tilde{A} \\
 & && s_k(z) \in \mathbb{N}_0 \\
 & && x_k \in \{0,1\}
 \end{aligned} \tag{3.1}$$

3.4.3 Assumptions

1. The critical component we study is used only once in each system. Also, each part occurs only once in the component.
2. For each part, a one-for-one replenishment policy (S-1,S) is applied.
3. Failures of parts occur according to a stationary Poisson process and are caused by a failure of a lower indenture part (if any).
4. The configuration of the component does not influence the failure rate of the lowest indenture parts.
5. Each part in the multi-indenture structure is critical, i.e. the entire component does not function if one part is defective and thus leads to a non-operational system.
6. If a failure occurs, it is always the entire component that is replaced.
7. Each component is repaired by replacement of a lower indenture part, except for parts without children; they are discarded upon failure, and a new part is purchased.
8. Lead times are independent and identically distributed for each part.
9. The repair lead time does not depend on lower-indenture parts. Thus, it captures the time to diagnose and replace a failed part if a spare part is in stock.
10. Repair capacity is not constrained; instead, we consider the contracted repair lead time.
11. Holding costs are encountered during the repair and ordering process.
12. No condemnation occurs, i.e. each repair is successful.

3.4.4 Model Evaluation

As clarified in Section 3.2, we use VARI-METRIC to evaluate our model (Sherbrooke, 1986). In this section, we will review the essential steps for our problem setting. To compute $A_k(\mathbf{S}_k)$, we must derive the demand rate for each part $z \in Z_k$ first. We begin

with computing the demand rate for parts without children ($\Gamma_k(z) = \emptyset$) with $m_k(z) = IB\lambda_z$. For an assembled part ($\Gamma_k(z) \neq \emptyset$), we use $m_k(z) = \sum_{y \in \Gamma_k(z)} m_k(y)$.

Next, we determine for each part the expected number of backorders $EBO(s_k(z), \cdot)$ and the variance of backorders $VBO(s_k(z), \cdot)$. For parts without children, we obtain these measures using a single site model, cf. Sherbrooke (2004). Afterwards, we compute the mean and variance of the number of parts in the replenishment process $R_k(z)$ for parts with children:

$$E[R_k(z)] = m_k(z)l_z + \sum_{y \in \Gamma_k(z)} EBO(s_k(y)) \quad (3.2)$$

$$Var[R_k(z)] = m_k(z)l_z + \sum_{y \in \Gamma_k(z)} VBO(s_k(y)) \quad (3.3)$$

To obtain the expected number of backorders $EBO(s_k(z), \cdot)$ and the variance of the backorders $VBO(s_k(z), \cdot)$ for these parts, we use a two moment approximation. That is, we fit a negative binominal distribution to the mean and the variance of R_z as given above. Analogously, we proceed with the higher indenture level parts until the expected number of backorders of the component, $EBO(\mathbf{S}_k)$, is obtained. Finally, to derive the average availability $A_k(\mathbf{S}_k)$, we use the following approximation:

$$A_k(\mathbf{S}_k) \approx \text{Max}\left\{1 - \frac{EBO(\mathbf{S}_k)}{IB}, 0\right\}. \quad (3.4)$$

In order to compute $TC_k(\mathbf{S}_k)$ we determine the total piece price $P_k(z)$ for every part $z \in Z_k$ first. For parts without children, we have $P_k(z) = p_z$. For assembled parts, we use $P_k(z) = a_z + \sum_{y \in \Gamma_k(z)} P_k(y)$. Next, we calculate $C_z(s_k(z))$, i.e. the average total costs for each part $z \in Z_{k1}$ and all its lower indenture parts. We have:

$$C_z(s_k(z)) = \begin{cases} m_k(z)(r_z + p_z) + s_k(z)\kappa p_z, & \text{if } \Gamma_k(z) = \emptyset \text{ or } |Z_k| = 1 \\ m_k(z)r_z + s_k(z)\kappa P_k(z) + \sum_{y \in \Gamma_k(z)} C_y(s_k(y)), & \text{otherwise} \end{cases} \quad (3.5)$$

As a result, we obtain the total costs with:

$$TC_k(\mathbf{S}_k) = \sum_{z \in Z_{k1}} C_z(s_k(z)). \quad (3.6)$$

3.4.5 Model Optimization

To optimize the formulated model, we follow the marginal analysis of Sherbrooke, cf. Sherbrooke (2004). The underlying idea of Sherbrooke's approach is to construct solutions by successively increasing the stock level of the most cost effective item, where cost effectiveness is captured by the backorder reduction per unit of capital (e.g per euro) invested. This procedure leads to a convex EBO-costs curve. Algorithm 3.1 formalizes this approach for component configuration $k \in K$.

Algorithm 3.1: Marginal analysis for component configuration $k \in K$

Set $s_k(z) := 0, z \in Z_k$ and define \mathbf{e}_z as a $|Z_k|$ -dimensional unit vector with the position of Part z equal to 1;

while $\tilde{A} > A_k(\mathbf{S}_k)$ **do**

Set $\mathbf{S}_{k_{\text{new}}} := \mathbf{S}_k + \mathbf{e}_z$;

Calculate $\Delta_z = \frac{EBO(\mathbf{S}_k) - EBO(\mathbf{S}_{k_{\text{new}}})}{TC_k(\mathbf{S}_{k_{\text{new}}}) - TC_k(\mathbf{S}_k)}$, $\forall z$;

Let $u := \text{argmax}_z \Delta_z$ and set $\mathbf{S}_k := \mathbf{S}_k + \mathbf{e}_u$;

Store efficient point $(EBO(\mathbf{S}_k), TC_k(\mathbf{S}_k))$;

Evaluate $A_k(\mathbf{S}_k)$;

Following Algorithm 3.1 for each configuration $k \in K$, we obtain a set of convex EBO-costs curves. The convexity property allows us to easily find the convex frontier of the EBO-costs curves, for example with the Graham scan (Graham, 1972). As a result, we can approximate which configuration leads to the most cost-efficient results given a desired availability \tilde{A} . Also, the resulting convex-frontier may lend itself to be utilized in more practical multi-item analyses, see e.g. Van Houtum and Kranenburg (2015). Here, however, we restrict ourselves to the single-item, single-location setup in order to evaluate the influence of consolidation with AM technology.

3.5 Numerical Experiments

In this section, we will conduct numerical experiments in order to gain insights in the effects of consolidation with AM. For this purpose, we will consider two setups: consolidation at the same indenture level ($I = 1$) in Section 3.5.1, and consolidation over two indenture levels ($I = 2$) in Section 3.5.2. For both setups, we will compare a CM non-integrated Configuration A with an AM integrated Configuration B. Throughout the experiments we will use three different performance indicators to measure the value of consolidation:

1. The percentage of instances where Configuration B has lower average total costs than Configuration A. This performance indicator is denoted by $\mathbf{B}\%$.
2. The average costs reduction with consolidation in case Configuration B has lower average total costs $\Delta TC[B] = \frac{TC_B(S_B) - TC_A(S_A)}{TC_A(S_A)}$, where the average of $\Delta TC[B]$ over several instances in percentage is denoted by $\Delta \mathbf{TC}[B]\%$.
3. The log difference between the total costs of both configurations, defined as $\ln \frac{TC_B(S_B)}{TC_A(S_A)}$, see Törnqvist et al. (1985). This measure has the advantage of being symmetric, i.e. $\ln(a/b) = -\ln(b/a)$. Thus, the average log difference is a suitable indicator of the relative costs difference over several instances. We denote the average by $\Delta \mathbf{ln}(\mathbf{TC})$.

3.5.1 Consolidation of Single-Indenture Components

In Figure 3.2, we illustrate possible component configuration structures in this section. Configuration A represents the CM manufactured component which consist of $|Z_A|$ parts. Configuration B is produced with AM and fulfills the same function with one part only.




Figure 3.2: Possible component configurations

To simplify the presentation of the model input and results, we relate the characteristics of parts in Configuration A to the characteristics of Part I. Thus, we derive the average replenishment lead time of Part I from the weighted average replenishment lead time incurred with Configuration A, i.e. $\tilde{l}_A = \sum_{z \in Z_A} \frac{m_A(z)}{m_A(A)} l_z$, with $m_A(A) = \sum_{z \in Z_A} m_A(z)$.

The measure \tilde{l}_A has the useful property of returning the same average replenishment lead time for both configurations if it holds that $l_I = \tilde{l}_A$. To experiment with differences in the average replenishment lead times as well, we introduce the factor α and define $l_I = \alpha \tilde{l}_A$. Following the same logic, we use the average failure rate experienced with Configuration A (i.e. $\lambda_A = \sum_{z \in Z_A} \lambda_z$) to compute the average failure rate of Part I and therefore have $\lambda_I = \beta \lambda_A$, with β describing the ratio between the average failure rates of both configurations. The purchasing cost rate of Part I is related to the purchasing cost rate of Configuration A (i.e. $p_A = \sum_{z \in Z_A} p_z$) with $p_I = \gamma p_A$, where γ models the ratio between the purchasing cost rate of both configurations. Therefore, if either of the factors (α, β, γ) is smaller than 1, consolidation offers a more favorable part characteristic.

The parameter ranges analyzed in the numerical experiments are based on the situation described in Section 3.3. Accordingly, we assume a component that is operated in expensive capital goods. For such components (and its subcomponents), repair and replenishment lead times often range in the order of magnitude of months, see Basten and Houtum (2014). Also, for components operating in expensive capital goods, the availability target is usually high, therefore we aim for 95% and 99.5% demand fulfillment from stock, respectively. For the holding cost fraction, we evaluate scenarios with low ($\kappa = 0.15$) and with high holding costs ($\kappa = 0.3$). Following results that suggest that AM technologies are valuable for low-demand rates in particular, see e.g. Khajavi et al. (2014) or Liu et al. (2014), we focus on low-demand environments. Also, we note that low-demand environments are more common for expensive capital goods.

For the lead time difference between both configurations, we typically find the assumption of a significantly shorter AM than CM lead time (i.e., $\alpha < 1$) in the literature. For example, Liu et al. (2014) report on data from the aerospace industry. For CM items they find lead times between 1.5 and 8 month, whereas, for AM parts, they typically assume lead times

of less than a month. Another example are the case studies discussed by Westerweel et al. (2018b). They observe AM resupply lead times of about two weeks while for the CM source, they find lead times of about 2 to 5 months. However, for the situation considered in this chapter, there are also arguments against such extreme lead time advantages. For example consolidated AM items typically exhibit a rather high geometric complexity which may cause additional processing times. In this context Löber et al. (2013) elaborate that complex structures often require more advanced post-processing procedures. To accommodate these findings in our numerical experiments, we primarily focus on lead time reduction ($\alpha < 1$) but also regard cases where the expected lead time increases ($\alpha > 1$).

The piece price for AM items is typically higher than for CM parts. Not only does this situation relate to the novelty of industrial AM processes, but also to the high raw material costs. Additionally, AM processes remain labor and skill intensive (Weller et al., 2015). Yet, as we discussed in Section 3.3, consolidation may also reduce the total piece price. For instance, costs associated with the assembly process no longer apply. Also, inefficiencies caused by sourcing sub-components from various suppliers can be reduced with consolidation. Accordingly, we analyze scenarios where the piece price decreases ($\gamma < 1$) but also scenarios where the piece price increases ($\gamma > 1$).

For the reliability, we adopt a similar approach since, following our discussion in Section 3.1, it is possible that consolidation influences the reliability positively ($\beta > 1$) or negatively ($\beta < 1$). In Table 3.3, we summarize the resulting parameter values, where $U[a,b]$ represents a continuous uniform distribution between the values a and b. Given that we can model arbitrary demand rates with λ_z , we set the installed base size equal to 1 (i.e. $IB = 1$). For each parameter combination, we sample 25 values from the uniform distributions, which results in 196,000 problem instances. The run time for each instance typically ranges in the order of milliseconds on a regular machine (Intel Core i5-4200U CPU, 8GB RAM, 64-bit Window 10).

Table 3.3: *Parameter values*

Parameters	Values
l_z	$U[0.2,1]$
λ_z	$U[0.2/ Z_k , 5/ Z_k]$
p_z	$U[100/ Z_k , 20000/ Z_k]$
κ	0.15, 0.3
\tilde{A}	0.95, 0.995
$ Z_A $	2, 3, 4, 5, 6
α	0.05, 0.2, 0.4, 0.6, 0.8, 1, 1.2, 1.4
β	0.4, 0.6, 0.8, 1, 1.2, 1.4, 1.6
γ	0.4, 0.6, 0.8, 1, 1.2, 1.4, 1.6

As a first step, we study the significance and the effect of changing input parameters with multiple linear regression. As the dependent variable we use the performance indicator $\Delta \ln(TC)$. The regression analysis revealed the following standardized relation with

log-transformed independent variables:

$$\begin{aligned} \Delta \ln(TC) = & -0.02 \ln(\tilde{l}_A) - 0.002 \ln(p_A) - 0.05 \ln(\tilde{A}) - 0.07 \ln(\kappa) \\ & + 0.02 \ln(\lambda_A) + 0.6 \ln(|Z_A|) + 0.13 \ln(\alpha) + 0.51 \ln(\beta) + 0.57 \ln(\gamma) \end{aligned} \quad (3.7)$$

In Appendix 3.B, we provide the full regression report and further diagnostic tests that qualify the model to be suitable for explanatory purposes.

We deduce several characteristics of consolidation from the regression model: the holding costs fraction (κ), the target availability (\tilde{A}), the weighted average replenishment lead time (\tilde{l}_A) and the purchasing cost rate of Configuration A (p_A) have a positive effect on the cost saving potential with consolidation (i.e. $\Delta \ln(TC)$ decreases). On the contrary, an increase of the other factor reduces the cost saving potential with consolidation. Also, it appears that β and γ have a considerably higher effect on $\Delta \ln(TC)$ than α . Finally, the number of integrated parts ($|Z_A|$) is the most important predictor for the value of consolidation. We will establish an interpretation of these findings in the remainder of this section.

Table 3.4: *Impact of changing parameters*

Parameter	Value	B%	$\Delta \text{TC}[\text{B}]\%$	$\Delta \ln(\text{TC})$
α	0.05	20%	-32%	0.46
	0.2	16%	-30%	0.54
	0.4	14%	-29%	0.60
	0.6	13%	-28%	0.65
	0.8	11%	-27%	0.70
	1	10%	-27%	0.74
	1.2	10%	-26%	0.77
	1.4	9%	-26%	0.80
β	0.4	41%	-34%	-0.8
	0.6	22%	-28%	0.28
	0.8	12%	-24%	0.53
	1	7%	-21%	0.74
	1.2	4%	-18%	0.91
	1.4	2%	-15%	1.05
	1.6	1%	-12%	1.18
γ	0.4	46%	-34%	-0.16
	0.6	22%	-26%	0.24
	0.8	11%	-22%	0.53
	1	5%	-19%	0.75
	1.2	3%	-15%	0.93
	1.4	1%	-13%	1.09
	1.6	1%	-11%	1.22

In Table 3.4 we present the effect of differences in replenishment lead time, failure rate and purchasing cost rate by varying α , β and γ . Overall, we find that the costs saving potential with consolidation is limited in the evaluated parameter range. Accordingly, the performance indicator $\Delta \ln(TC)$ is positive for most parameter values in Table 3.4. Nevertheless, the

results for $\Delta TC[B]\%$ clarify that, in case consolidation is superior, the cost saving potentials of consolidation are high. For example, consider the scenario where the purchasing cost rate is 1.6 higher for Configuration B than for Configuration A ($\gamma = 1.6$). On average, the higher purchasing cost rate leads to significantly higher total costs with consolidation ($\Delta \ln(TC) = 1.22$). However, for the few instances where Configuration B is preferable ($B\% = 1\%$), consolidation offers major cost savings on average ($\Delta TC[B]\% = -11\%$). This finding clarifies that assessing the value of consolidation by a single characteristic such as the purchasing cost rate is not advisable. Instead, a total costs perspective is required to evaluate the possible benefit of consolidation.

Furthermore, the results presented in Table 3.4 replicate the findings in the regression analysis: a short replenishment lead time (α) provides less incentive to make use of consolidation than a lower failure rate (β) or purchasing cost rate (γ). Accordingly, $B\%$ changes by only 11% points if we compare $\alpha = 0.05$ and $\alpha = 1.4$. This effect is significantly smaller than effects resulting from changes in β and γ .

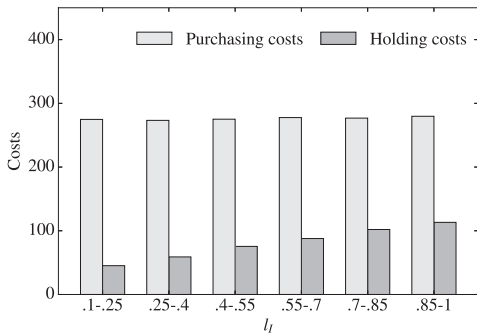


Figure 3.3: Effect of replenishment lead time changes on the average purchasing and the average holding costs




Figure 3.4: Effect of the number of functional integrated parts on the percentage of instances where Configuration B yields lower costs ($B\%$) and the average cost reduction if Configuration B yields lower cost in percentage ($\Delta TC[B]\%$)

The relation between costs and replenishment lead time allows us to deduce an explanation for this finding. The average purchasing costs defined by $P_B = m_B(I)p_I$ are independent of the replenishment lead time. Thus, changes in α affect the average holding costs ($H_B = s_B(I)\kappa p_I$) only, which limits the impact on the total costs. We illustrate this relation in Figure 3.3 for different values of l_I . On the contrary, β and γ affect both costs (P_B and H_B) and thus have a higher influence on the value of consolidation on average. Considering the fact that short lead times are a key benefit of AM, this finding may partially explain why consolidation does not appear recommendable for most test instances. Conversely, we may

reason that consolidation through AM becomes more valuable in cases where the average holding costs of a conventional configuration are high. This situation is more likely in cases with a high target availability (\tilde{A}), a high holding cost fraction (κ), a long weighted average lead time (\tilde{l}_A) or a high purchasing cost rate of Configuration A (p_A). Therefore, this result also offers an interpretation for findings from the regression analysis where the negative coefficients of \tilde{A} , κ , \tilde{l}_A and p_A indicate a positive correlation with the value of consolidation.

Next, we investigate the negative effect of the number of parts (i.e. $|Z_A|$) on the potential cost saving with consolidation, which we deduced from the regression analysis. In Figure 3.4, we illustrate this effect by plotting $B\%$ as a function of $|Z_A|$. As expected $B\%$ decreases with the number of consolidated parts. We hypothesize that this effect is related to the additional flexibility to fulfill the availability target (\tilde{A}). For instance, one may obtain the option to allocate stock unevenly among parts which leads to cost savings eventually. Given that this flexibility increases with $|Z_A|$, the result of the regression analysis appears reasonable.




Figure 3.5: Increase in failure rate differences favors segmented configuration illustrated by the increasing average log difference between the total costs of both configurations ($\Delta \ln(TC)$)

To further elaborate on this aspect, we conducted a small sub-experiment with $|Z_A|=2$ where we chose $\lambda_1=0.5$ while varying λ_2 . Other parameter values are chosen as they were in the previous experiment (cf. Table 3.3). In Figure 3.5, we show that $\Delta \ln(TC)$ is close to 0% if Part 1 and Part 2 have the same failure rate ($\lambda_1=\lambda_2$). The more the failure rates deviate, the larger $\Delta \ln(TC)$ becomes. The results exemplify the value of flexibility. In case both parts have a rather similar failure rate, the benefit of allocating stock unevenly between both parts is low. If, however, the difference is high, the benefit of allocating stock unevenly increases. Given that such differences are more likely to occur with a higher number of parts, $|Z_A|$ is a negative predictor for the benefit of consolidation. This hypothesis is further supported by comparable results in experiments where we varied the purchasing cost rate and the average replenishment lead time.

3.5.2 Consolidation of Two-Indenture Components

In this section, we focus on the effect of consolidation over two-indenture levels. The possible component configurations are illustrated in Figure 3.6. To simplify the presentation of the model input and results, we follow an approach comparable to the one in Section 3.5.1 and relate several input parameters to each other. For parameter relations that do not follow from the explanations in Section 3.5.2 immediately, we give a short outline subsequently. The remaining parameter relations are listed in Table 3.5.




Figure 3.6: Possible component configurations

We chose the weighted average replenishment lead time of parts at the lowest indenture level (\tilde{l}_{A2}) as orientation for the average repair lead time of Part 0, where $\tilde{l}_{A2} = \sum_{z \in Z_{A2}} \frac{m_A(z)}{m_A(A)} l_z$, with $m_A(A) = \sum_{z \in Z_{A2}} m_A(z)$. Accordingly, we obtain the average repair lead time for Part 0 by $l_0 = \xi \tilde{l}_{A2}$, where ξ models the relative deviation of the average repair lead time from \tilde{l}_{A2} . Likewise, we define the repair cost rate of Part 0 relative to the total piece price of Part 0 and therefore obtain $r_0 = \pi P_A(0)$, where π denotes the ratio between the repair costs rate and total piece price of Part 0.

In Table 3.5 we show the experimental settings. In case of the replenishment lead time factor (α), we restrict the parameter range to the case where the average replenishment lead time of the segmented configuration is at least as long as it is for the consolidated configuration. This restriction seems justified given that a shorter replenishment lead time for the AM produced component is more likely in practice, and we already observed that α has less impact on the value of consolidation (cf. Section 3.5.1). In case of β and γ , we reduced the parameter range as well, disregarding the most extreme values which allows us to put higher emphasis on the impact of the factors ξ and π . The experimental setup results in 270,000 problem instances because we sample from each uniform distribution 25 times.

As elaborated in Section 3.1, consolidation of assembly structures may eliminate a repair option. While an assembly structure can be repaired by replacing defective sub-components, this flexibility is lost with consolidation. To obtain further insights on this aspect we study the effect of the repair cost rate ratio π first. Next, we evaluate the impact of the lead time ratio ξ . In Figure 3.7, we depict $B\%$ and $\Delta TC[B]\%$ as a function of the repair cost ratio π . The results indicate that in case of low repair costs compared to the total piece

Table 3.5: Parameter values

Parameters	Values
l_z with $z \in Z_{k2}$	$U[0.5/ Z_{A2} , 1/ Z_{A2}]$
λ_z with $z \in Z_{k2}$	$U[0.2/ Z_{A2} , 5/ Z_{A2}]$
p_z with $z \in Z_{k2}$	$U[2000/ Z_{A2} , 20000/ Z_{A2}]$
κ	0.15, 0.3
\tilde{A}	0.95, 0.995
$ Z_{A2} $	1, 2, 3, 4, 5
α	0.2, 0.6, 1
β	0.8, 1, 1.2
γ	0.8, 1, 1.2
ξ	0.2, 0.4, 0.6, 1, 1.2
π	0.05, 0.1, 0.15, 0.2
p_0	$U[0, 2000]$
l_I	αl_{A2}
λ_I	$\beta \lambda_A$
p_I	$\gamma P_A(0)$

price of Part 0, consolidation is typically not advisable. Accordingly, $B\%$ and $\Delta TC[B]\%$ decrease with decreasing π . This observation relates to lower average purchasing costs with Configuration A in case π is small.




Figure 3.7: Effect of changing repair costs on the percentage of instances where Configuration B yields lower costs ($B\%$) and the average cost reduction if Configuration B yields lower cost in percentage ($\Delta TC[B]\%$)




Figure 3.8: Effect of changing repair lead time on the percentage of instances where Configuration B yields lower costs ($B\%$) and the average cost reduction if Configuration B yields lower cost in percentage ($\Delta TC[B]\%$)

While Configuration A allows for solving a failure by purchasing a new sub-component only, consolidation requires the purchase of the entire configuration. Next to the purchasing cost savings, the repair process may reduce the holding costs as well. This potential is illustrated

in Figure 3.8, where we show $B\%$ and $\Delta TC[B]\%$ as a function of the repair lead time ratio ξ . As we observe, the shorter the repair lead time, the less instances lead to consolidation as the preferred configuration ($B\%$). Also, the cost saving potential decreases slightly ($\Delta TC[B]\%$). These effects are a consequence of the possibility to resupply Part 0 quickly in case ξ is small. Accordingly, one may decrease the stock level of Part 0 while fulfilling the same availability target and thus decrease the average holding costs. If an assembly structure can be repaired by replacing sub-components only, the value of consolidation is limited. In particular, this conclusion holds if the repair cost rate is low compared to the purchasing costs of a new component and if the average repair lead time is short.




Figure 3.9: Effect of the number of functional integrated parts on the percentage of instances where Configuration B yields lower costs ($B\%$) and the average cost reduction if Configuration B yields lower cost in percentage ($\Delta TC[B]\%$)




Figure 3.10: Effect of changing stock allocation on the percentage of instances where Configuration B yields lower costs ($B\%$) and the average cost reduction if Configuration B yields lower cost in percentage ($\Delta TC[B]\%$)

Next, we address the question of how the number of consolidated parts at the second indenture level influences the value of consolidation. Therefore, we illustrate the percentage of instances where consolidation is preferable ($B\%$) depending on the number of sub-components ($|Z_{A2}|$) in Figure 3.9. We observe that $B\%$ decreases with increasing $|Z_{A2}|$. This finding appears reasonable, as nearly the same logic as in Section 3.5.1 applies. The higher the number of sub-components, the higher the flexibility to allocate stock unevenly among the parts to fulfill the required availability of the repair shop. One particularly noteworthy outcome can be observed when $|Z_{A2}|$ is equal to 1, because in this case more than 80% of instances favor consolidation. However, this finding further strengthens the flexibility argument given in Section 3.5.1. In case $|Z_{A2}|$ is equal to 1, Configuration A does not offer additional flexibility compared to Configuration B. Therefore, the benefit of keeping a segmented design needs to be justified by other aspects.

Finally, in Figure 3.10, we show that the allocation of stock among indenture levels

in Configuration A gives a good approximation whether consolidation is valuable. For instance, in case the stock level of Part 0 is high compared to the total stock of the sub-components (i.e. $s_0/\sum_{z \in Z_{A2}} s_z > 2$), consolidation appears more interesting. As a result, the stock allocation in a multi-indenture structure may give a first indicator for which parts consolidation might be worthwhile from a total costs perspective.

3.6 Conclusion

In this chapter, we have investigated the consequences of consolidation through AM technology. Therefore, we have quantified the total costs differences between a CM assembly and a consolidated AM part for different scenarios. Our results show that consolidation using AM often leads to higher total costs than the CM design. The key reason is that replacing a consolidated AM part is often more expensive than repairing a conventional part by replacing a subcomponent. This may easily wipe out the advantages of a shorter lead time or a lower failure rate of the consolidated AM part. This assessment is based on the following observations:

1. Typically, a combination of higher reliability, reduced replenishment lead time and lower price is required to achieve cost reductions with consolidation. The improvement of one characteristic alone was found to be insufficient in most cases (cf. Table 3.4).
2. The benefit of a shorter replenishment lead time through consolidation is less valuable than a lower price or higher reliability, whereas a shorter lead time is often pointed out as a major advantage of AM (cf. Table 3.4).
3. The more parts are consolidated into a single one, the less likely consolidation is to be beneficial due to restrictions on possible stock allocation (cf. Figure 3.4).
4. The higher the difference in part characteristics such as failure rate or purchasing costs, the less likely consolidation of these parts is to be beneficial (cf. Figure 3.5).
5. The option of repairing a CM assembly by replacing failed sub-components leads to a lower benefit of consolidation, especially when the repair lead time is short and/or the repair costs are low (cf. Figure 3.7 and Figure 3.8).

As a consequence, even though pointed out as a major opportunity for operations, consolidation with AM technology comes along with a number of hidden costs – in particular for the service business. We argue that, in the light of these findings, industry has to be convinced that the adoption of AM technologies is not only a task for the design department, but also for logistics and maintenance engineers. By disregarding this perspective in the early stages of AM adoption, unexpected costs may arise that in the long run may damage the perceived value of AM technologies. On the other hand, by including service costs in the analysis, the application of AM technologies may become justifiable also in less obvious cases. For instance, we found that even if the failure rate of a consolidated AM configuration is 1.6 times higher than that of a CM configuration, consolidation may still lead to significantly lowered total costs.

Appendices

3.A Example of Consolidation

One of the more popular business cases of consolidation is a fuel nozzle used in General Electric's CFM LEAP engines. Using AM technology it was possible to reduce the part count from 18 to 1 and, as a consequence, to decrease the weight by 25%, as well as increasing the estimated life time by a factor of 5. By 2020, GE expects to have printed 100,000 fuel nozzles. An aircraft equipped with the new LEAP engine unlocks savings of around US \$3 million per year compared to an engine that does not use the printed fuel nozzles (GE Aviation, 2015).



Figure 3.11: *Printed fuel nozzle used in CFM LEAP engines of General Electric*

3.B Regression Analysis

Table 3.6 shows the results of the multiple regression for the dependent variable $\Delta \ln(TC)$ where the independent variables are log-transformed. The analysis is based on the 196,000 instances specified in Section 3.5.1.

The R-squared for this model is 0.97. Furthermore, based on a t-test, each regressor is significant with a confidence level of more than 99% while White standard errors are consistent in the presence of heteroskedasticity. Finally, a Ramsey RESET test revealed that the model is sufficiently specified ($P=0.37$) and thus supports the choice of linear regression model for explanatory purposes.

Next to the main effects, we also considered the interactions. Preliminary results indicated that several interactions appear significant though usually with a low coefficient. To clarify the effects of interactions on our regression model, we subsequently also regard three interactions that appear most relevant. In particular, the number of integrated parts $|Z_A|$ with the holding cost rate κ , the number of integrated parts $|Z_A|$ with the total piece price

Table 3.6: *Results multiple regression*

	Unstandardized coefficient	White standard errors	Standardized coefficient	t	Sig.
Intercept	.3303	.0103		-32.1776	.0000
$\ln(\lambda_A)$.0532	.0016	.023	33.8593	.0000
$\ln(\tilde{l}_A)$	-.0526	.0021	-.018	-25.0332	.0000
$\ln(p_A)$	-.0031	.0011	-.002	-2.9321	.0034
$\ln(\alpha)$.1016	.0006	.132	170.6034	.0000
$\ln(\beta)$.9051	.0013	.511	698.4365	.0000
$\ln(\gamma)$	1.0013	.0013	.565	793.7871	.0000
$\ln(Z_A)$.8056	.0008	.603	984.5423	.0000
$\ln(\kappa)$	-.1699	.0017	-.073	-102.6065	.0000
$\ln(\tilde{A})$	-1.5557	.0248	-.045	-62.6956	.0000

of Configuration A p_A , and the holding cost rate κ with the expected lead time change α . Below, we show the regression equation including these three interaction effects.

$$\begin{aligned}
 \Delta \ln(TC) = & -0.02 \ln(\tilde{l}_A) - 0.05 \ln(\tilde{A}) - 0.14 \ln(p_A | Z_A) - 0.147 (\kappa | Z_A) \\
 & + 0.005 \ln(\kappa) + 0.004 \ln(p_A) + 0.02 \ln(\lambda_A) + 0.6 \ln(|Z_A|) \\
 & + 0.29 \ln(\alpha) + 0.51 \ln(\beta) + 0.57 \ln(\gamma) + 0.162 \ln(\kappa \alpha)
 \end{aligned} \tag{3.8}$$

The coefficient - 0.14 for the interaction $\ln(p_A | Z_A)$ indicates that a higher total piece price of Configuration A, slightly reduces the negative effect of the number of integrated parts. The same effect occurs if the holding cost fraction κ is large. Nevertheless, both interaction effects do not compensate for the main effect of $|Z_A|$ which clarifies that the number of integrated parts complicates consolidation with AM. Finally, a high holding costs fraction κ combined with a lead time increase through consolidation α further reduces the value of consolidation.

However, we decided not to study the interaction effects in further detail because of two reasons. First, the model is only used to offer a preliminary interpretation of the main effects before we detail our analysis in Section 3.5.1. Interaction effects obscure the interpretability since multicollinearity and overfitting effects become more significant. Second, our preliminary results showed that including interaction effects did not lead to a better model fit since we were not able to further improve the R-squared.

Chapter 4

The Transition to Additive Manufacturing

In Chapter 2, we discussed how companies may identify promising spare parts for the application of additive manufacturing (AM) technology. While the method has proven helpful in different environments, we noticed that companies were uncertain about how and when to perform the transition from the conventional manufacturing (CM) to the AM supply source. Typically companies possess the capability and knowledge to source spare parts with CM while experience still has to be gained on how to manage the supply of parts produced by AM.

Another aspect is the uncertainty associated with switching to a less mature manufacturing technology. For instance, a decrease in AM production costs is likely in the near future (cf. Gibson et al. (2010), Khajavi et al. (2014) and Thomas (2016)), though the order of magnitude and the timing is difficult to predict. In combination with a lack of in-house knowledge on AM technology such uncertainties may prompt a risk averse attitude at the management level. As a result, companies may rather trust in proven methods than to embrace and explore new opportunities presented by AM technology. In this chapter¹, we address this problem and evaluate under which conditions a transition to AM technology is advisable. By means of a stochastic dynamic programming model and numerical experiments, we examine which transition strategies are most viable under different conditions. As such, we address the third contribution mentioned in Section 1.4.3.

¹This chapter is based on the working paper Knofius et al. (2018c).

4.1 Introduction

The starting point for our analysis is a case study conducted at a world-wide operating OEM for radar systems. In particular, we focus on a protection cover that is used for shielding electronic components from moisture, heat, and accidental damage. As most of the case companies' spare parts that potentially appear producible with AM, the protection cover is a mechanical component that is designed to outlast the lifetime of the radar system. Hence, the protection cover exhibits (very) low failure rates that, in most cases, are caused by external factors such as unintended stress levels, imprudence during maintenance/shipping activities, or extreme weather conditions. With CM, the replenishment lead time of the protection cover exceeds more than half a year because of low demand quantities and the fact that the company is the sole customer. For the latter reason, the OEM and *not* the supplier owns product-specific tooling which induces additional inventory costs.

We demonstrate that moving to AM technology pays off under various conditions in the low-volume spare parts business. For example, in the case of the protection cover, a higher AM piece price and additional AM preparation costs are compensated by lower holding costs. Also, to the surprise of the OEM, an immediate investment in AM is the best strategy to minimize total service costs. Our contribution to the literature is as follows:

1. We investigate if the preparation of AM technology is recommendable and what potential gains are to be expected under various conditions. In contrast to other studies, we take into account that knowledge and product-specific tooling is generally available for sourcing CM spare parts (CM being the regular production method), while the transition to AM still needs some preparation.
2. We analyze how and when a transition to AM technology is advisable for spare parts supply during the remaining service horizon. Therefore, in comparison to other studies, we consider an evolving inventory system where AM piece price and demand rates may change over the course of the service horizon.
3. Apart from deriving general guidelines for decisions on how and when to switch to AM technology, we also conduct a case study in the defense industry. With that, we respond to the request for more case-based research in the application domain of AM technology, see, e.g., Weller et al. (2015) and Savastano et al. (2016).

The remainder of this chapter is organized as follows. In Section 4.2, we develop the mathematical model and discuss the underlying model assumptions. In Section 4.3, we analyze the situation encountered at the OEM and, next, conduct numerical experiments. In Section 4.4, we discuss possible extensions.

4.2 Model

To study the situation encountered at the OEM of radar systems, we first develop an analytical model that mimics the evolution and describes underlying trade-offs and decisions.

A general outline is presented in Section 4.2.1. In Section 4.2.2, we justify and motivate the model assumptions. Section 4.2.3 formalizes the situation encountered at the OEM and introduces the notation. In Section 4.2.4, we explain the costs evaluation, and we describe the transition logic of the stochastic dynamic program in Section 4.2.5.

4.2.1 Outline

We consider an OEM who is responsible for maintaining an installed base of (radar) systems. We focus on one specific component that is sourced initially with a CM method. The CM approach requires specific tooling which, because of low demand rates, is owned by the OEM.

Technological advancements, however, allow the OEM to source the component with a tool-less AM method as well. From the OEM's point of view, a tool-less AM approach associated with (significantly) shortened replenishment lead times may render a transition from CM to AM attractive. In fact, a shorter replenishment lead time may reduce the stock-level and may decrease risks associated with a stock-out. Yet the OEM is uncertain whether an investment in AM technology pays off in particular, given the possibly higher AM piece price. Moreover, if a transition to AM is deemed profitable, the OEM is uncertain when to invest in AM. For example, constantly decreasing AM piece prices and a lack of experience with AM technology are compelling reasons to postpone the investment.

Additionally, it is unclear how the ability to source items with AM could influence the sourcing strategy. For instance, one can envision a dual sourcing concept where AM and CM are used in parallel. Alternatively, one may argue that using two sourcing methods in parallel causes unnecessary costs. For example, by employing a single sourcing approach with AM, the OEM may realize a substantial reduction in holding cost as product-specific tooling required for the CM process can be discarded.

We develop a stochastic dynamic program to investigate the described trade-offs. Our decision variables reveal the preferred sourcing method and order quantity in each period, if (and when) to prepare AM and if (and when) to discard tooling necessary for CM. In particular, we study for each period which decisions minimize the service costs composed of setup, purchasing, holding, backorder, and discarding costs over the remaining service horizon.

4.2.2 Assumptions

Before outlining the model in detail, we describe the underlying key assumptions we used to construct the model.

1. All lead times are deterministic and shorter than one period.

Lead times are typically a matter of contractual agreements with suppliers. In case a supplier is unable to meet these agreements, delays are typically compensated. Hence it seems reasonable for our analysis to stick to the mutually agreed lead times. Also, given that we typically consider a period length of one year in our analysis, assuming lead times of less than one period appears justifiable.

2. Holding costs are encountered at the beginning of the period.

Equivalently, we may account for the holding costs at the end of each period. Here, however, we charge holding cost at the beginning of the period as it allows an evaluation independent of the demand realization.

3. Failures occur according to a Poisson process.

The Poisson demand assumption seems most appropriate, given that we consider a low-demand environment with mechanical parts that are dimensioned to outlast the intended maintenance interval or lifetime of the capital good (common practice for downtime critical mechanical components of capital goods). Accordingly, failures are random in nature and may, for instance, be caused by unintended stress levels, imprudence during system usage, maintenance, shipping activities, or unobservable quality issues during the production.

4. The expected number of failures are the same for an AM and CM part.

As discussed in Section 4.1, we consider parts with (very) low failure rates which are motivated by typical properties of printable parts. This characteristic has several consequences. First, possible failure rate differences between the AM and CM versions are small in absolute terms. Second, during the service horizon, the installed base composition only changes marginally as most parts outlast the life cycle of the (radar) system and thus leave the overall demand nearly unaffected. Finally, by assuming identical failure rates, it is not necessary to keep track of the installed base composition which reduces the state space substantially.

5. Failed parts are replaced in negligible time if stock is available.

The replacement time of a failed part is short in comparison to the order lead times. Furthermore, given that we assume identical failure rates (see Assumption 4), the associated replacement frequency is independent of the decisions.

6. The OEM takes any required measures to fulfill outstanding demand.

To analyze a long service horizon, we consider a long period length. Thus, in case of a stock-out occurring in some period, it seems realistic to allow for emergency shipments during that same period. In case tooling is no longer available or AM is not prepared, a stock-out may also lead to an immediate setup of the preferred method (which, if we stick to CM produced parts, means a purchase of tooling again).

7. Fixed order costs are incorporated in the piece price.

Given that we consider a low demand environment economies of scale are limited. Thus, we assume that ordering costs are incorporated in the piece price of each item. Note that this assumption appears not suitable for higher demand regimes since under these conditions CM setup costs would typically favor batch ordering. To that end, the CM ordering costs would have to be regarded explicitly similar to the model discussed in Chapter 6.

8. Emergency orders are carried out piece by piece.

See Assumption 6.

4.2.3 Notation and Mathematical Model

The remaining service horizon of the radar systems is composed of discrete time periods of equal length, denoted by $n=1,\dots,N+1$ (say, years). Depending on contractual agreements, the installed base size IB_n may change over the course of the service horizon. If stock is available, the component is supplied from a single stock point in negligible time. Otherwise, demand is backordered and the OEM incurs unit backorder cost b per period. To reduce the backlog time, an emergency order run may be initiated during the period. For each stored part, holding costs are encountered at the beginning of period n . Holding costs are expressed in terms of a holding costs fraction κ of the associated current piece price. Demand is modelled by the random variable D_n in period n and determined by the current installed base size IB_n and failure behavior of the installed components.

Originally, the OEM orders CM items from a supplier with a mean order lead time l_C and piece price c_C . The production of CM items requires product-specific tooling which, due to low demand rates, is owned by the OEM. If tooling is available, the OEM may decide to discard it at a cost fraction δ_T of the tooling piece price c_T . In the event the OEM wants to (re)purchase tooling (because the demand realization is higher than anticipated or tooling was not available in the first place), it is possible to order tooling at a cost price c_T associated with a mean order lead time of l_T . Tooling holding costs are expressed as a fraction κ_T of the piece price c_T and are charged at the beginning of the period if tooling is available and will not be discarded. At the end of the service horizon (i.e., $n=N+1$), tooling and any item stock are discarded. The costs for discarding a part are supposed to be a fraction δ of the associated current piece price.

The setup of the AM process takes l_P , costs c_P , and is required only once. The setup contains activities such as supplier and material selection, design modifications, qualification, and the determination of printing process parameters. The replenishment order lead time of an AM item is equal to l_A which is usually much shorter than l_C . The AM piece price is equal to $c_{A,n}$ in period n . We model the AM piece price as a function of n in order to describe anticipated piece price changes of the AM process. Typically, $c_{A,n}$ will be a decreasing function of n as AM production costs are expected to decrease rapidly. Later, we model the AM piece price by means of an experience curve defined as $c_{A,n} = c_{A,1}n^r$ with $r = \log(1-cf)/\log(2)$ and $0 \leq cf < 1$. For a more detailed discussion on the AM piece price development, refer to Appendix 4.A.

For each period n , we describe the model state with a four dimensional vector $\mathbf{i}_n = (i_{1,n}, i_{2,n}, i_{3,n}, i_{4,n}) \in \mathbf{I}_n$ where $i_{1,n} \in N_0$ refers to the current number of CM items in stock, and $i_{2,n} \in N_0$ to the current number of AM items in stock. Furthermore, $i_{3,n} = 0$ if no CM tooling is available and $i_{3,n} = 1$ otherwise, while $i_{4,n} = 0$ if the AM process has not been prepared yet and $i_{4,n} = 1$ otherwise. Note that in case AM has not been prepared yet ($i_{4,n} = 0$), we know that $i_{2,n} = 0$ (no AM parts in stock). Hence, it would be sufficient to use a three dimensional state space. For clarity of presentation, however, we stick to a four dimensional representation.

At the beginning of each period n , the OEM may face a number of decisions. Actions taken are recorded in the four dimensional vector $\mathbf{a}_n = (a_{1,n}, a_{2,n}, a_{3,n}, a_{4,n}) \in \mathbf{A}_n$ where

$a_{1,n} \in N_0$ describes the number of items to order and $a_{2,n}$ indicates the sourcing approach. Here, we use $a_{2,n}=0$ to describe a pure CM sourcing approach and $a_{2,n}=1$ to denote a pure AM sourcing approach. In case $a_{2,n}=2$, we use CM for regular supply and AM for emergency supply. Further, $a_{3,n}=1$ if tooling is discarded at the beginning of period n and $a_{3,n}=0$ otherwise. Similarly, $a_{4,n}=1$ if the AM process is prepared at the beginning of period n and $a_{4,n}=0$ otherwise.

Certain actions may be ruled out in advance and thus can be eliminated from the action space. As discussed in Section 4.2.2, CM items are always used first when both AM and CM items are available in stock. Furthermore, in case we do not carry out a regular order, i.e., $a_{1,n}=0$, we may limit $a_{2,n} \in \{0,1\}$. Also if we discard tooling $a_{3,n}=1$, either we do not make a regular order, i.e., $a_{1,n}=0$, or we use the AM supplier, i.e., $a_{2,n}=1$. Finally, if tooling is unavailable or AM has already been prepared we know that $a_{3,n}=0$ and $a_{4,n}=0$ respectively.

Table 4.1: *Notation overview*

Notation	Explanation
IB_n	Installed base size in period n
$n \in \{1, 2, \dots, N+1\}$	Period index
b	Backorder cost per time unit
κ	Holding cost fraction for items
D_n	Total demand in period n
$D_n(t)$	Demand during time interval t with $0 \leq t \leq 1$
l_C	CM replenishment lead time where $0 \leq l_C \leq 1$
c_T	Tooling order cost
κ_T	Holding cost fraction for the tool
δ_T	Tooling discarding cost fraction
l_P	Lead time for AM preparation where $0 \leq l_P \leq 1$
c_P	AM preparation costs
$c_{A,n}$	AM piece price in period n
c_C	CM piece price
l_A	AM replenishment lead time where $0 \leq l_A \leq 1$
δ	Item discarding cost fraction
$\mathbf{i}_n \in \mathbf{I}_n$	Four-dimensional state vector in period n , $\mathbf{i}_n = (i_{1,n}, i_{2,n}, i_{3,n}, i_{4,n})$
$i_{1,n}$	Number of CM items in stock at the start of period n
$i_{2,n}$	Number of AM items in stock at the start of period n
$i_{3,n}$	1 if tooling is available at the start of period n ; 0 otherwise
$i_{4,n}$	1 if AM has already been prepared at the start of period n ; 0 otherwise
$\mathbf{a}_n \in \mathbf{A}_n$	Four-dimensional action vector in period n , $\mathbf{a}_n = (a_{1,n}, a_{2,n}, a_{3,n}, a_{4,n})$
$a_{1,n}$	Order quantity in period n
$a_{2,n}$	Sourcing approach in period n ; 0 = CM; 1 = AM; 2 = Dual sourcing
$a_{3,n}$	1 if tooling is discarded at the start of period n ; 0 otherwise
$a_{4,n}$	1 if the AM process is prepared at the start of period n ; 0 otherwise
$C_n(\mathbf{i}_n, \mathbf{a}_n)$	Expected cost in period n
$p(\mathbf{i}_{n+1} \mathbf{i}_n, \mathbf{a}_n)$	Transition probability

Possible actions and cost assignments at the beginning of period n occur in the following order: 1) discard tooling, 2) incur holding costs, 3) begin preparation of AM, 4) order tooling, 5) order items, and 6) incur purchasing costs. The stochastic dynamic programming recursion for the total costs from period n until the end of the horizon reads as follows:

$$V_n(\mathbf{i}_n) = \min_{\mathbf{a}_n \in \mathbf{A}_n} \{C_n(\mathbf{i}_n, \mathbf{a}_n) + \sum_{\mathbf{i}_{n+1} \in \mathbf{I}_{n+1}} p(\mathbf{i}_{n+1} | \mathbf{i}_n, \mathbf{a}_n) V_{n+1}(\mathbf{i}_{n+1})\} \quad (4.1)$$

$$V_{N+1}(\mathbf{i}_n) = (i_{1,n}c_C + i_{2,n}c_{A,n})\delta + i_{3,n}c_T\delta_T \quad (4.2)$$

where $C_n(\mathbf{i}_n, \mathbf{a}_n)$ denotes the expected service cost in period n if we are in state $\mathbf{i}_n \in \mathbf{I}_n$ and take action $\mathbf{a}_n \in \mathbf{A}_n$ (details are discussed in the next section). The probability to jump to state $\mathbf{i}_{n+1} \in \mathbf{I}_{n+1}$ given that we are in state $\mathbf{i}_n \in \mathbf{I}_n$ and take action $\mathbf{a}_n \in \mathbf{A}_n$ is denoted by $p(\mathbf{i}_{n+1} | \mathbf{i}_n, \mathbf{a}_n)$. Equation 4.2 describes the costs encountered at the end of the service horizon, i.e., costs for discarding the remaining stock and tooling. In Table 4.1, we summarize the notation.

4.2.4 Cost Computations

The expected service costs $C_n(\mathbf{i}_n, \mathbf{a}_n)$ in period n are composed of holding costs $H_n(\mathbf{i}_n, \mathbf{a}_n)$, discarding costs for parts and tools $R_n(\mathbf{i}_n, \mathbf{a}_n)$, expected purchasing costs $P_n(\mathbf{i}_n, \mathbf{a}_n)$, setup costs $T_n(\mathbf{i}_n, \mathbf{a}_n)$, and expected backorder costs $B_n(\mathbf{i}_n, \mathbf{a}_n)$. Next, we show how we evaluate the cost factors for a given state vector \mathbf{i}_n and action vector \mathbf{a}_n .

The *holding costs* of parts and tools $H_n(\mathbf{i}_n, \mathbf{a}_n)$ are accounted for at the beginning of period n but after the possible discarding of the CM tool. Here, we value AM parts according to the present period cost price. Accordingly, we have:

$$H_n(\mathbf{i}_n, \mathbf{a}_n) = (i_{1,n}c_C + i_{2,n}c_{A,n})\kappa + i_{3,n}(1 - a_{3,n})c_T\kappa_T \quad (4.3)$$

In each period $n \leq N$, we encounter *discarding costs* $R_n(\mathbf{i}_n, \mathbf{a}_n)$ in case tooling is discarded. In the last period, we also discard unused parts. Thus the discarding costs are equal to:

$$R_n(\mathbf{i}_n, \mathbf{a}_n) = \begin{cases} a_{3,n}i_{3,n}c_T\delta_T & , \text{ if } n \leq N \\ (i_{1,n}c_C + i_{2,n}c_{A,n})\delta + i_{3,n}c_T\delta_T & , \text{ if } n = N + 1 \end{cases} \quad (4.4)$$

The expected *purchasing costs* $P_n(\mathbf{i}_n, \mathbf{a}_n)$ depend on the order quantity ($a_{1,n}$) and chosen supplier ($a_{2,n}$). Also, we have to consider the situation in which stock-outs require emergency order that lead to additional purchasing costs. The expected number of emergency orders in period n , $UP_n(i_{1,n}, i_{2,n}, a_{1,n})$, equals

$$UP_n(i_{1,n}, i_{2,n}, a_{1,n}) = E[(D_n - i_{1,n} - i_{2,n} - a_{1,n})^+] \quad (4.5)$$

Thus, the expected purchasing costs are represented by:

$$P_n(\mathbf{i}_n, \mathbf{a}_n) = \begin{cases} (a_{1,n} + UP_n(i_{1,n}, i_{2,n}, a_{1,n}))c_C & , \text{ if } a_{2,n} = 0 \\ (a_{1,n} + UP_n(i_{1,n}, i_{2,n}, a_{1,n}))c_{A,n} & , \text{ if } a_{2,n} = 1 \\ a_{1,n}c_C + UP_n(i_{1,n}, i_{2,n}, a_{1,n})c_{A,n} & , \text{ if } a_{2,n} = 2 \end{cases} \quad (4.6)$$

The *setup costs* $T_n(\mathbf{i}_n, \mathbf{a}_n)$ consist of AM preparation $T_{1,n}(\mathbf{i}_n, \mathbf{a}_n)$ and tooling procurement costs $T_{2,n}(\mathbf{i}_n, \mathbf{a}_n)$. Hence we have $T_n(\mathbf{i}_n, \mathbf{a}_n) = T_{1,n}(\mathbf{i}_n, \mathbf{a}_n) + T_{2,n}(\mathbf{i}_n, \mathbf{a}_n)$. AM preparation costs arise if the OEM decides to prepare the AM process at the beginning of period n , i.e., $a_{4,n} = 1$. Furthermore, in case we use AM for emergency supply, i.e., $a_{2,n} \in \{1, 2\}$, and AM has not been prepared ($i_{4,n} = 0$), preparation costs may arise to order items for backorder clearing. Accordingly, we have:

$$T_{1,n}(\mathbf{i}_n, \mathbf{a}_n) = (a_{4,n} + \min\{1, a_{2,n}\})(1 - a_{4,n})(1 - i_{4,n})PUP_n(i_{1,n}, i_{2,n}, a_{1,n})c_P \quad (4.7)$$

where $PUP_n(i_{1,n}, i_{2,n}, a_{1,n})$ represents the probability of an emergency shipment in period n , i.e., demand exceeds inventory ($i_{1,n} + i_{2,n}$) plus regular order ($a_{1,n}$):

$$PUP_n(i_{1,n}, i_{2,n}, a_{1,n}) = 1 - Pr\{D_n \leq i_{1,n} + i_{2,n} + a_{1,n}\} \quad (4.8)$$

The OEM encounters tooling procurement costs if we use CM supply, i.e., $a_{2,n} \in \{0, 2\}$ and $a_{1,n} > 0$, while tooling is not available ($i_{3,n} = 0$). Note that, as explained in Section 4.2.3, it is not possible to discard tooling ($a_{3,n} = 1$) and then order items from the CM supplier in the same period n . Furthermore, we may encounter tooling procurement costs for backorder clearing. This event may arise if no order is carried out ($a_{1,n} = 0$) and we use CM for emergency supply, i.e., $a_{2,n} = 0$ while tooling is not available, i.e., $i_{3,n} = 0$ or $a_{3,n} = 1$. Accordingly, we have:

$$T_{2,n}(\mathbf{i}_n, \mathbf{a}_n) = \begin{cases} I_{\{0,2\}}(a_{2,n})(1 - i_{3,n})c_T & , \text{ if } a_{1,n} > 0 \\ I_{\{0\}}(a_{2,n})(1 + a_{3,n} - i_{3,n})PUP_n(i_{1,n}, i_{2,n}, a_{1,n})c_T & , \text{ if } a_{1,n} = 0 \end{cases} \quad (4.9)$$

with $I_{\{\cdot\}}(\cdot)$ representing an indicator function that is equal to one if the condition is fulfilled and zero otherwise.

When we determine the *backorder costs*, we account for stock-outs in the period in which they arise. Thus, if a stock-out occurs towards the end of period n , the entire duration of a stock-out is charged in period n . We distinguish between four scenarios leading to downtime denoted by $\Omega_{1,n}(\mathbf{i}_n, \mathbf{a}_n)$, $\Omega_{2,n}(\mathbf{i}_n, \mathbf{a}_n)$, $\Omega_{3,n}(\mathbf{i}_n, \mathbf{a}_n)$ and $\Omega_{4,n}(\mathbf{i}_n, \mathbf{a}_n)$ respectively. Accordingly, we have:

$$B_n(\mathbf{i}_n, \mathbf{a}_n) = b(\Omega_{1,n}(\mathbf{i}_n, \mathbf{a}_n) + \Omega_{2,n}(\mathbf{i}_n, \mathbf{a}_n) + \Omega_{3,n}(\mathbf{i}_n, \mathbf{a}_n) + \Omega_{4,n}(\mathbf{i}_n, \mathbf{a}_n)) \quad (4.10)$$

The downtime calculations depend on the sourcing approach $a_{2,n}$. Here, we focus on the case $a_{2,n} = 1$ and refer to Appendix 4.B for the other cases. $\Omega_{1,n}(\mathbf{i}_n, \mathbf{a}_n)$ accounts for backorders that arise if demand cannot be filled from the initial inventory ($i_{1,n} + i_{2,n}$), but instead with the regular order ($a_{1,n}$). In this case, we incur downtime until order arrival. $\Omega_{2,n}(\mathbf{i}_n, \mathbf{a}_n)$ describes the scenario where the initial inventory and the regular order (if any) cannot cover demand. Subsequently, an emergency order becomes necessary which leads to downtime of at least the AM replenishment order lead time l_A . However, as we cover in $\Omega_{3,n}(\mathbf{i}_n, \mathbf{a}_n)$, it is also possible that AM has not been prepared yet. In this case, we have additional downtime l_P for the first emergency order while we may incur a part of

l_P for any additional demand arriving during AM preparation. Finally, as described by $\Omega_{4,n}(\mathbf{i}_n, \mathbf{a}_n)$, we may have started AM preparation at the beginning of the period which possibly has not finished before an emergency order becomes necessary. In this case, we encounter part of the AM preparation lead time l_P as well.

For $\Omega_{1,n}(\mathbf{i}_n, \mathbf{a}_n)$ we have:

$$\Omega_{1,n}(\mathbf{i}_n, \mathbf{a}_n) = \sum_{d=r+1}^{r+a_{1,n}} Pr\{D_n(t)=d\} \frac{0.5(d-r)(d-r+1)}{d+1} t, \text{ if } a_{2,n}=1 \quad (4.11)$$

with $r = i_{1,n} + i_{2,n}$ and $t = (l_A + l_P(1 - i_{4,n}))$ representing the relevant fraction of period n before order arrival. For the derivation of the formula we used that Poisson demand arrivals are uniform distributed over the relevant time interval, conditional on the number of arrivals in the interval. Accordingly, in case of d demand arrivals in t , the downtime is equal to $t \sum_{a=1}^d (d+1-a)/(d+1) = t \sum_{a=1}^d a/(d+1) = t(0.5d(d+1))/(d+1)$.

Next, $\Omega_{2,n}(\mathbf{i}_n, \mathbf{a}_n)$ is defined by:

$$\Omega_{2,n}(\mathbf{i}_n, \mathbf{a}_n) = UP_n\{i_{1,n}, i_{2,n}, a_{1,n}\} l_A, \text{ if } a_{2,n}=1 \quad (4.12)$$

assuming piece by piece ordering from the emergency source (cf. Section 4.2.2).

$\Omega_{3,n}(\mathbf{i}_n, \mathbf{a}_n)$ is only relevant if AM has not been prepared yet nor the AM preparation has started yet, i.e., $i_{4,n}, a_{4,n}, a_{1,n} = 0$. In this case, we have additional downtime l_P for the first emergency order while we may incur a part of l_P for any additional demand arrival during AM preparation. Given the Poisson assumption, the latter time fraction can be expressed by $\sum_{d=1}^{\infty} Pr\{D_n(l_P)=d\} d/2 = (\lambda_n l_P)/2$ where λ_n denotes the mean demand rate in period n . Intuitively, this relation follows because each demand arrival during the preparation lead time l_P , encounters $0.5l_P$ downtime on average. Accordingly, we have:

$$\Omega_{3,n}(\mathbf{i}_n, \mathbf{a}_n) = Pr\{D_n \geq i_{1,n} + i_{2,n} + 1\} l_P \left(1 + \frac{\lambda_n l_P}{2}\right), \text{ if } i_{4,n}, a_{4,n}, a_{1,n} = 0 \wedge a_{2,n} = 1 \quad (4.13)$$

Finally, $\Omega_{4,n}(\mathbf{i}_n, \mathbf{a}_n)$ accounts for the possibility of starting AM preparation at the beginning of the period (either because we ordered with the AM supplier or prepare AM proactively, i.e., $i_{4,n} = 0 \wedge (a_{1,n} > 0 \vee a_{4,n} = 1)$) which possibly has not finished before an emergency order becomes necessary. In this case, we encounter part of the AM preparation lead time l_P as well and thus have:

$$\Omega_{4,n}(\mathbf{i}_n, \mathbf{a}_n) = \sum_{d=r+1}^{\infty} Pr\{D_n(l_P)=d\} \frac{0.5(d-r)(d-r+1)}{d+1} l_P, \\ \text{if } i_{4,n} = 0 \wedge (a_{1,n} > 0 \vee a_{4,n} = 1) \wedge a_{2,n} = 1 \quad (4.14)$$

with $r = i_{1,n} + i_{2,n} + a_{1,n}$. The derivation follows the same line of reasoning as for $\Omega_{1,n}(\mathbf{i}_n, \mathbf{a}_n)$. To further clarify the backorder costs computations, we provide a numerical example in Appendix 4.B.

4.2.5 State Transitions

The state transitions depend on the current state, the actions, and the demand realization. To clarify the underlying logic, we discuss possible transitions for one state and two actions. Other transitions follow the same logic and are determined by the algorithm presented in Appendix 4.C. For both examples, we have the following state at the start of period n : one CM item is in stock (and no AM item), tooling is available and the AM process has not been prepared, i.e., $\mathbf{i}_n = (1,0,1,0)$. Figure 4.1 shows the possible transitions for the two selected actions.

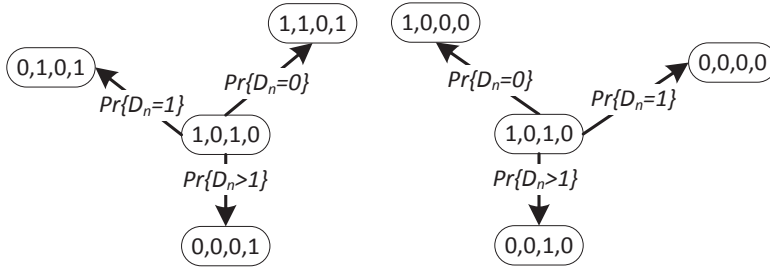


Figure 4.1: Transition with action $\mathbf{a}_n = (1,1,1,1)$ (left) and action $\mathbf{a}_n = (0,0,1,0)$ (right)

On the left side of Figure 4.1, we order one part from the AM supplier, decide to discard tooling and prepare AM, i.e., $\mathbf{a}_n = (1,1,1,1)$. For this example, it is important to realize that we always move to the state $\mathbf{i}_{n+1} = (0,1,0,1)$ but not to state $\mathbf{i}_{n+1} = (1,0,0,1)$ if $D_n = 1$. As explained in Section 4.2.3, by default we always use the CM part first if both part versions are in stock.

On the right side of the figure, we illustrate possible transitions for the action where we do not order any part, use CM for emergency supply, discard tooling and do not prepare AM, i.e., $\mathbf{a}_n = (0,0,1,0)$. As illustrated, despite discarding tooling at the beginning of period n , we still may end up in state $\mathbf{i}_{n+1} = (0,0,1,0)$, i.e., with tooling. This transition is a consequence of the assumption that the OEM solves a stock-out as soon as possible and thus, for this example, reproduces tooling if $D_n > 1$. Only in the next period ($n+1$) is it possible to discard tooling again.

4.3 Analysis

In Section 4.3.1, we evaluate the protection cover case. Next, in Section 4.3.2, we carry out a sensitivity analysis to evaluate the value of moving to AM technology for more general settings.

4.3.1 Case Study

The protection cover is used for shielding electronic components from moisture, heat and accidental damage in radar systems. Figure 4.2 shows the design of the protection cover.




Figure 4.2: *Design protection cover*

Originally, the protection cover is built with molding technology using carbon epoxy. In case of a failure, usually caused by external factors such as extreme weather conditions or imprudence during maintenance activities, the protection cover has to be replaced. Otherwise, the risk of damaging expensive components would be too high. In case of a stock-out, the exposed electronics is normally covered with a temporary solution that, at least, offers some protection. Overall, the company management translates the associated risks to backorder costs of approximately 43,800 euro per year (120 euro per day). Provided experience with AM technology is lacking, such low-risk cases are often preferred by companies or sometimes even enforced by regulations as, e.g., in the aerospace industry. Nevertheless, in Section 4.3.2, we also consider high backorder costs scenarios that are not uncommon in the capital goods spare parts business.




Figure 4.3: *AM piece price development compared to CM*

The CM replenishment lead time of the protection cover is relatively long and takes approximately seven months, which is caused by low demand quantities (< 1 per year), and the fact that the company is the only customer for the protection cover. For the latter reason, the company also owns the mold, which causes tool holding costs.

Evaluations showed that a Selective Laser Sintering (SLS) process, using glass-filled nylon, technologically qualifies for the production of the protection cover. Further, preliminary analyses by technical staff indicated that only minor design changes would be required (and desired) to provide functionality comparable to the CM version. Unfortunately, consultations with a supplier of glass-filled nylon prints revealed an estimated piece price of 13,300 euro, which is approximately 7,500 euro higher than the price of the CM version. As we discussed, such price differences are the norm rather than the exception today.

Table 4.2: *Input parameter case study*

$c_{A,1}$	13,300 euro	l_A	14 days
c_C	5,768 euro	l_C	196 days
c_P	10,000 euro	l_P	28 days
c_T	5,900 euro	l_T	65 days
b	43,800 euro/year	κ	0.1 euro/euro/year
λ	0.01 failures/system/year	κ_T	0.1 euro/euro/year
IB_n	60 systems	δ	-0.1 euro/euro
N	30 years	δ_T	-0.1 euro/euro

To model the AM piece price decrease over the service horizon, we follow the explanations in Appendix 4.A and assume an AM piece price reduction by about 30% (i.e., $cf=0.15$) within the next five years. The resulting AM piece price development is shown in Figure 4.3. Overall, the AM piece price remains higher than that for the CM version of the protection cover (5,768 euro) during the entire service horizon. Hence, at first sight, the value of moving to AM technology appears doubtful — in particular if we consider that about 10,000 euro are required for the preparation of AM. Currently, the case company possesses the mold, has one spare part in stock and has not prepared AM yet. Furthermore, based on contractual agreements, it is expected that the installed base size remains constant over the service horizon. All parameter values are presented in Table 4.2.

Using the model described in Section 4.2, we quantify the extent to which a transition to AM may influence the remaining service costs. Figure 4.4 displays the cost difference between servicing the installed base solely with CM (Do not move to AM) and the approach where a transition to AM is possible (Move to AM). Additionally, we show the cost composition where we omit cost factors that do not contribute more than 0.01% to the overall costs. As we observe in Figure 4.4 moving to AM turns out economically valuable with cost savings of more than 12% over the course of the remaining service horizon. This finding relates to holding cost savings that are obtained by discarding tooling and by reducing stock level and stock-out risk. The latter is a direct consequence of the relatively short AM replenishment lead time. Moreover, in Figure 4.4, we observe that holding costs reductions compensate for a higher AM piece price and additional AM preparation cost.




Figure 4.4: Cost-saving potential

To secure these cost savings, we find that the OEM should prepare AM in the first year and always use AM for emergency supply, i.e., $a_{2,n} \in \{1,2\}$. Figure 4.5 reveals more insights about the transition policy and, for each year n , shows the probabilities of CM tooling in stock, $Pr\{i_{3,n}=1\}$, and regular supply with AM, $Pr\{a_{2,n}=1|a_{1,n}>0\}$. Despite preparing AM in the first year, we find that it is not recommended to discard tooling before Year 4 in any case. Instead, depending on the demand realization, we may use CM for regular supply up to Year 9. Interestingly, according to our cost predictions, the piece price of AM will still be about 2,400 euros higher than with CM after the transition to AM is completed. Figure 4.6 shows the average inventory levels for both sourcing methods for each period. As expected, the average on-hand stock reduces with the transition to AM. The reason that we observe a lower average CM inventory level in Year 2 than in Year 3, is caused by the relative low demand. Since we ordered in Year 1, it is likely that we do not have to reorder in Year 2. This becomes more likely in Year 3 though.




Figure 4.5: Transition policy




Figure 4.6: Average inventory levels

For the company, the result clarifies that the economical valuation of AM should not be based on the AM piece price exclusively. Also, the awareness that an investment in AM technology provides immediate benefits, stimulated a more extensive search for AM applications in the company. In fact, the case study exemplified that the value of AM technology — in particular in the spare parts business — often appears to be underestimated. We conclude that conceptual insights alone appear insufficient to motivate the adoption of AM in the service business. Currently, business-specific use cases seem paramount to convincing management about the value of using AM technology for their spare parts operations.

4.3.2 Numerical Experiments

Some characteristics of the protection cover are rather specific and may deviate from other applications. In this section, we perform a sensitivity analysis to draw more general conclusions about the value of moving to AM technology. Therefore, we justify chosen parameter ranges first and then discuss the results.

Experimental Design

In the spare parts business for capital goods, we often encounter long service horizons. For example, Van Houtum and Kranenburg (2015) consider a service horizon between 10 and 40 years realistic. For our experiments, we analyze instances with a remaining service horizon (N) between 10 and 20 years. The choice to analyze a shorter time horizon is motivated by several aspects. First, today, it is likely that AM is considered in situations where the regular production phase ended several years ago. Second, in preliminary experiments we found that a transition to AM is typically considered in the first half of the service horizon, or not at all.

The installed base size in period n , IB_n , influences the demand rate and demand development over the service horizon. For the scenario encountered in the case company, we assumed a constant installed base size motivated by contractual agreements. After the regular production phase, a decreasing installed base size is also considered to be realistic in several cases, see e.g. Inderfurth and Mukherjee (2008) or Dekker et al. (2013). Thus, we will regard both options in our experiments, a constant installed base size and a decreasing installed base size. In order to model the installed base size reduction, we use a decreasing function of the form $IB_n = \lceil IB_1 - (n/x)^c \rceil$ with $c > 1$, yielding a concave decreasing curve with increasing n (see Figure 4.7). The value for x is set to obtain a desired final installed base size, i.e., $x = N / (IB_1 - IB(N))^{1/c}$. Figure 4.7 illustrates the different installed base development profiles we consider in the sensitivity analysis.

AM technology was identified to be the most valuable for low-volume production, i.e., for production where economies of scales do not play a significant role, see, e.g., Khajavi et al. (2014) and Liu et al. (2014). Accordingly, we chose the installed base size in combination with the failure rate such that the demand rate is low. Yet, compared to the case study where we encountered less than one demand arrival per year on average, we also investigate possible effects of higher demand rates.




Figure 4.7: *Installed base development*

The backorder costs (b) for the protection cover are minor if compared to other examples in the industry. For instance, in the semiconductor industry, failures leading to a standstill of the production line may cost tens of thousands of euros per hour (Kranenburg, 2006). In the numerical experiments, we will consider backorder cost between 36,500 and 1,825,000 euros per year. The reason behind the exclusion of even higher backorder costs in this study is twofold. First, companies might be inclined/forced to refrain from printing very critical parts, see Section 4.3.1. Second, in case of very high backorder costs, other solutions such as redundancy, design improvements, or predictive maintenance strategies become more attractive.

As we motivated for the protection cover, in case of low demand rates the replenishment lead time with CM methods is typically long, say several months. Liu et al. (2014) assume lead times between 1.5 and 8 months, whereas the AM lead time is typically assumed to be less than a month. In our experiments, we consider a comparable range. For the AM piece price in the first period, $c_{A,1}$, we use three values which, in combination with three values for the cost development factor (cf), lead to nine costs profiles. In Figure 4.8 we show the three cost profiles in case $c_{A,1}=10,000$ euro. The six other cost profiles exhibit the same pattern.

As is depicted, we study both: scenarios where AM remains more expensive over the entire service horizon and scenarios where, at a certain point in time, AM becomes less expensive than CM. Depending on the choice of the cost development factor (cf), the cost decline is faster or slower. Overall, we decided to investigate this wide range of possible cost developments, because various cost profiles are perceivable, depending on features such as material type, AM process, and geometric complexity. Finally, we vary the holding cost rate with $\kappa=0.1$ and $\kappa=0.15$. Remaining model parameters remain unchanged compared to the protection cover case.




Figure 4.8: Unit costs development with $c_{A,1}=10,000$ euro

Sensitivity Analysis

To evaluate the results, we use four performance indicators:

1. $\Delta \mathbf{C}_0$ describes, as an average over associated instances, the relative cost difference between the scenario where we limit ourselves to solely using CM, i.e., $a_{2,n} = 0$, for $n = 1, \dots, N+1$, and where transitioning to AM technology is possible.
2. $\Delta \mathbf{S}$ describes the difference between the maximum stock level in case we use CM and AM. For scenarios where it is optimal to never use AM, i.e., $a_{2,n} = 0$, for $n = 1, \dots, N+1$, we define $\Delta \mathbf{S} = 0$.
3. **Prep** denotes the time horizon that has passed before AM is prepared, averaged over all instances where preparation occurred. Instances, where AM is never prepared are excluded ($\sim 6\%$ of instances).
4. **Dis** describes the average time horizon that has passed before tooling is discarded. We consider the same instances as used for the calculation of the performance indicator **Prep** to allow for comparability between **Prep** and **Dis**.

Table 4.3 shows the factors and factor levels of the experimental design and the results for the performance indicators as average over the associated instances. Note that parameters not mentioned in Table 4.3, remain unchanged compared to the case study discussed in Section 4.3.1. Over all instances, we find average cost savings of about 35%. The results in Table 4.3 show that the cost-saving potential ($\Delta \mathbf{C}_0$) is most sensitive to (i) the lead time difference between the AM and CM method, (ii) the backorder costs and (iii) the AM and CM piece price difference (including the rate at which the AM piece price decreases). Subsequently, we interpret these findings.

The sensitivity of the cost-saving potential on the lead time difference and backorder costs are explained with holding cost savings. As we observe in Table 4.3, the lead time difference between AM and CM method significantly influences the stock level difference $\Delta \mathbf{S}$. That

Table 4.3: Results sensitivity analysis

Parameters	Parameter values	ΔC_0	ΔS	Prep	Dis
N	10 years	30%	2.5	17%	28%
	20 years	41%	2.4	4%	24%
IB_1	50 parts	35%	1.5	13%	14%
	150 parts	33%	2.4	10%	28%
	300 parts	38%	3.4	8%	36%
<i>Installedbase development</i>	$c=2; IB(N)=0$	35%	2.7	12%	22%
	$c=7; IB(N)=0$	35%	2.4	11%	27%
	<i>const.</i>	36%	2.3	9%	29%
b	36,500 euro/year	29%	2.4	10%	18%
	365,000 euro/year	32%	2.5	13%	28%
	1,825,000 euro/year	45%	2.6	9%	31%
l_C	60 days	23%	1.8	10%	29%
	180 days	48%	3.1	11%	23%
l_A	5 days	40%	2.8	6%	23%
	15 days	31%	2.1	15%	29%
κ	0.1 euro/euro/year	32%	2.6	14%	28%
	0.15 euro/euro/year	39%	2.3	7%	24%
$c_{A,1}$	6,000 euro	74%	2.3	0%	0%
	10,000 euro	23%	2.5	6%	26%
	14,000 euro	9%	2.6	28%	57%
cf	0.1	25%	2.5	17%	36%
	0.15	34%	2.5	11%	26%
	0.2	46%	2.4	5%	17%

is, the stock level difference increases, in case the AM lead time becomes shorter or the CM lead time becomes longer. On the other hand, the backorder costs affect ΔS only marginally (cf. Table 4.3). However, by increasing the backorder cost, the relative importance of saving holding cost increases. As explained during the case study, saving holding cost is a primary benefit of transitioning to AM technology. In fact, the relative holding cost fraction more than doubles if we compare $b=36,500$ euro/year and $b=1,825,000$ euro/year.

This result stresses the prospect of printing critical parts for the spare parts business. Currently, however, companies are reluctant (or not allowed) to consider high backorder costs cases. In particular, high AM process variability raises concerns about part reliability and, for the time being, renders the certification of critical parts cumbersome. In the near future, closed-loop control mechanisms which support online adjustments and corrections of the printing process may significantly decrease process variability, see, e.g. Craeghs et al. (2012) or Everton et al. (2016).

In Table 4.3, we also observe cost savings increasing if the holding cost fraction (κ) increases. Further analysis revealed that this result follows the same rational as for the

backorder costs: the higher the holding costs factor, the higher the relative importance of holding costs saving. We infer from these results that high inventory and backorder costs indicate the profitability of moving to AM technology in the low-volume spare parts business.

The initial AM piece price ($c_{A,1}$) is most influential on the cost savings (cf. Table 4.3). This result is explained by the fact that, in case the starting AM and CM piece price are comparable, CM becomes inferior quickly (equal or higher piece price and longer resupply lead time). Thus, a fast and complete transition to AM has few drawbacks and leads to immediate holding and purchasing cost savings. On the other hand, in case the AM piece price is high (which is more likely today), purchasing cost increase and solely the reduction of holding costs reduction may justify the transition to AM. Hence, in case the initial AM piece price is very high, cost savings are less substantial.

The trade-off between holding cost decrease and purchasing cost increase also influences the transition strategy. Figure 4.9 illustrates the situation by comparing the time horizon % before AM is prepared, **Prep**, and the time horizon % before CM tooling is discarded, **Dis**, for different values of the starting AM piece price ($c_{A,1}$).

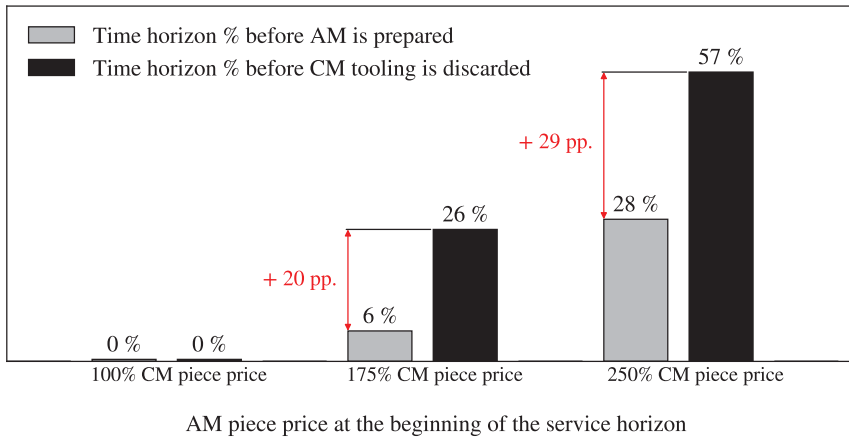


Figure 4.9: **Prep** and **Dis** for different values of the starting AM piece price $c_{A,1}$

As we observe in Figure 4.9, if the initial AM piece price is low, tooling is discarded at about the same time as AM is prepared. Hence, a dual sourcing approach as described in the case study is usually not profitable under these conditions. On the other hand, the higher the AM piece price becomes, the longer the time horizon where both sourcing methods are used in parallel (AM is prepared before tooling is discarded, cf. Figure 4.9). During the time horizon where dual sourcing is applied, we typically use CM for regular and AM for emergency supply ($a_{2,n} = 2$). Consequently, we are able to reduce stock while the purchasing cost increase is maintained within limits.

Similarly, as shown in Table 4.3, a large installed base size (IB_1) and high backorder costs (b) increase the time horizon where both methods are used in parallel. Remarkably, in case a transition to AM is profitable, AM preparation takes place relatively early during the service horizon ($\mathbf{Prep}=10\%$ on average, cf. Table 4.3). This observation motivates the conclusion that an investment in AM technology should not be postponed. Instead, before a complete transition to AM is economical, one could use both sourcing methods in parallel. Apart from the economical perspective, an early AM preparation also enables the accumulation of experience with AM which appears beneficial prior to a full transition to AM technology.

Finally, we observe that an increasing demand rate (determined by multiplying installed base size with failure rate) has a positive, though less predominant, effect on the value of AM. It remains an open research question when economies of scale reverse this trend. Here, we refrain from further analysis on this matter given that our model choices are tailored to a low-demand environment.

4.4 Conclusion

The findings can be summarized as follows:

1. Over all instances, we find average cost savings by moving to AM of about 35% despite a typically higher AM piece price and additional AM preparation costs. The cost-saving potential increases predominately with increasing backorder costs, increasing AM and CM lead time differences, and decreasing piece price differences between AM and CM.
2. The costs reduction is primarily achieved by holding cost savings with AM that are caused by decreasing stock levels, a reduced stock-out risk and the option to use a tool-less AM process. In case the backorder and inventory costs are high, the holding cost reduction becomes most beneficial.
3. If AM becomes competitive during the remaining service period, the preparation of AM should not be postponed. Instead, before a complete transition to AM is profitable, it is beneficial to use both sourcing methods in parallel for several reasons. A dual sourcing with AM enables holding cost savings while the purchasing cost increase (due to a typically high AM piece price) is kept within reasonable limits.
4. We find a long dual sourcing period if the AM piece price is high, the demand rate is high, or if the backorder costs are high. During this period, CM is used for regular production and AM is preferably used for emergency supply.
5. It appears that the value of AM for spare parts supply is underestimated. Company-specific business cases seem necessary to convince management about the value of moving to AM technology for spare parts supply. In particular, the opportunity to invest in AM technology at this early stage appears to be disregarded.

To further support our findings, it might be valuable to extend the proposed model by also considering stochastic AM piece price developments. Even though we base our analysis on predictions made in the literature, it is well-perceivable that uncertainty regarding piece

price development will influence the investment decision. Similar extensions are perceivable for the installed base development and the service horizon length. Finally, one may consider a scenario where supply disruptions occur during the service horizon. A straightforward extension in this regard might be to model an occasionally occurring loss of tooling. While conducting the case study, we were several times confronted with scenarios where the tooling was lost due to operational inadequacies.

Appendices

4.A AM Piece Price Predictions

Fast market growth, rapid technological advancements, patent expiration, and decreasing raw material prices are only some indicators that the production costs for AM parts will reduce significantly over the next years. A comparable effect for CM parts is unlikely. To account for this discrepancy, we model the AM piece price $c_{A,n}$ as a decreasing function over time.

A widely accepted approach to forecast technological costs, see e.g. Nagy et al. (2013) or Magee et al. (2016), is the use of an experience curve with the underlying logic that the costs reduce by a constant every time the experience doubles. Alberth (2008) explains that using experience curves qualify to obtain insights into potential prices during the growth phase in competitive markets which significantly resembles the AM market. For our model, we assume that experience is gained in every period and rely on estimated price reductions (see below) in order to specify the experience curve. More explicitly we have:

$$c_{A,n} = c_{A,1} n^r$$

where r describes the elasticity of the cost reduction to the experience, defined as

$$r = \log(1 - cf) / \log(2).$$

An analysis conducted by Roland Berger (2013) forecasts an AM production cost reduction potential of about 30% between 2018 and 2023. A study sponsored by the German government (Bechthold et al., 2015) and a report from Siemens (2014a) support these predictions. For the case study conducted in Section 4.3.1, we utilize the same cost development, i.e., $cf = 0.15$. Nevertheless, in order to understand the impact of more conservative or optimistic cost development predictions, we include the factor cf in the sensitivity analysis conducted in Section 4.3.1.

4.B Backorder Costs

Before we give a numerical example, we explain the calculation of $\Omega_{1,n}(\mathbf{i}_n, \mathbf{a}_n)$, $\Omega_{2,n}(\mathbf{i}_n, \mathbf{a}_n)$, $\Omega_{3,n}(\mathbf{i}_n, \mathbf{a}_n)$ and $\Omega_{4,n}(\mathbf{i}_n, \mathbf{a}_n)$ for sourcing with CM only ($a_{2,n}=0$) and where we apply a dual sourcing approach ($a_{2,n}=2$).

In case $a_{2,n}=0$, the calculations remain the same as for the case $a_{2,n}=1$ except that the associated lead times change. Therefore, we have to replace l_A by l_C and l_P by l_T in Equations 4.11–4.14. Furthermore, in case of $\Omega_{3,n}(\mathbf{i}_n, \mathbf{a}_n)$, we have to account for the possibility that we may decide to discard tooling in period n , i.e., $a_{3,n}=1$. Accordingly we have:

$$\Omega_{3,n}(\mathbf{i}_n, \mathbf{a}_n) = Pr\{D_n \geq i_{1,n} + i_{2,n} + 1\} l_T \left(1 + \frac{\lambda_n l_T}{2}\right), \text{ if } a_{1,n}, a_{2,n} = 0 \wedge (i_{3,n} = 0 \vee a_{3,n} = 1)$$

In case ($a_{2,n}=2$), $\Omega_{1,n}(\mathbf{i}_n, \mathbf{a}_n)$ is calculated similar to the case where $a_{2,n}=0$. The values $\Omega_{2,n}(\mathbf{i}_n, \mathbf{a}_n)$, $\Omega_{3,n}(\mathbf{i}_n, \mathbf{a}_n)$ and $\Omega_{4,n}(\mathbf{i}_n, \mathbf{a}_n)$ are calculated as if $a_{2,n}=1$.

Next, we give a numerical example to clarify the backorder costs calculations. Assume we are in state $\mathbf{i}_n = (2, 0, 1, 0)$ (two CM items, no AM items, tooling available, AM not prepared) and take action $\mathbf{a}_n = (2, 1, 1, 1)$ (order 2 items with AM, discard tooling, and prepare AM). We encounter backorders of type $\Omega_{1,n}(\mathbf{i}_n, \mathbf{a}_n)$ if $3 \leq D_n(t) \leq 4$ and have

$$\Omega_{1,n}(\mathbf{i}_n, \mathbf{a}_n) = Pr\{D_n(t) = 3\} \frac{1}{4}t + Pr\{D_n(t) = 4\} \left(\frac{2}{5} + \frac{1}{5}\right)t$$

with $t = l_A + l_P$. For the case that $D_n(t) = 3$, we know that the third demand arrives at $(3/4)t$ and is fulfilled after $(1/4)t$ on average using that demand arrivals are uniformly distributed. Following the same logic, in case $D_n(t) = 4$, the third and fourth demand arrival cause a backorder duration of $(2/5)t$ and $(1/5)t$, respectively. For the case that $D_n(t) > 4$, we encounter downtime of type $\Omega_{2,n}(\mathbf{i}_n, \mathbf{a}_n)$ and $\Omega_{4,n}(\mathbf{i}_n, \mathbf{a}_n)$ because AM preparation is initiated at the beginning of period n .

For the sake of explanation, let us assume $D_n(l_P) = 6$. Again using that Poisson arrival are uniform distributed, expected arrival times are equal to $(i/7)l_P$ ($i = 1 \dots 6$). Then, $\Omega_{4,n}(\mathbf{i}_n, \mathbf{a}_n) = (2/7 + 1/7)l_P$ because the first two demand arrivals are filled from stock and the second two with regular supply.

4.C State Transitions

Algorithm 4.1, shows the procedure to determine \mathbf{i}_{n+1} in case we use emergency supply in period n .

Algorithm 4.1: Determines \mathbf{i}_{n+1} based on \mathbf{i}_n , \mathbf{a}_n and D_n if $i_{1,n} + i_{2,n} + a_{1,n} - D_n < 0$

$i_{1,n+1}, i_{2,n+1} = 0, i_{3,n+1} = i_{3,n}, i_{4,n+1} = i_{4,n};$

if $a_{2,n} = 0$ **then**

$i_{3,n+1} = 1;$

if $a_{4,n} = 1$ **then**

$i_{4,n+1} = 1;$

if $a_{2,n} = 1$ **then**

$i_{4,n+1} = 1;$

if $a_{3,n} = 1$ **then**

$i_{3,n+1} = 0;$

if $a_{2,n} = 2$ **then**

$i_{4,n+1} = 1;$

if $a_{1,n} > 0$ **then**

$i_{3,n+1} = 1;$

else

if $a_{3,n} = 1$ **then**

$i_{3,n+1} = 0;$

Algorithm 4.2, shows the procedure to determine \mathbf{i}_{n+1} in case we do *not* require emergency supply in period n .

Algorithm 4.2: Determines \mathbf{i}_{n+1} based on \mathbf{i}_n , \mathbf{a}_n and D_n if $i_{1,n} + i_{2,n} + a_{1,n} - D_n \geq 0$

```

 $i_{3,n+1} = i_{3,n}$ ,  $i_{4,n+1} = i_{4,n}$ ;
if  $a_{3,n} = 1$  then
   $i_{3,n+1} = 0$ ;
if  $a_{2,n} \neq 1$  and  $a_{1,n} > 0$  then
   $i_{3,n+1} = 1$ ;
if  $a_{4,n} = 1$  or ( $a_{2,n} = 1$  and  $a_{1,n} > 0$ ) then
   $i_{4,n+1} = 1$ ;
if  $i_{4,n+1} = 0$  then
   $i_{1,n+1} = i_{1,n} + i_{2,n} + a_{1,n} - D_n$  and  $i_{2,n+1} = 0$ ;
else
  if  $a_{2,n} = 1$  then
    if  $i_{1,n} - D_n > 0$  then
       $i_{1,n+1} = i_{1,n} - D_n$  and  $i_{2,n+1} = i_{2,n} + a_{1,n}$ ;
    else
       $i_{1,n+1} = 0$  and  $i_{2,n+1} = i_{1,n} + i_{2,n} + a_{1,n} - D_n$ ;
  else
    if  $i_{1,n} - D_n > 0$  then
       $i_{1,n+1} = i_{1,n} + a_{1,n} - D_n$  and  $i_{2,n+1} = i_{2,n}$ ;
    else if  $i_{1,n} + i_{2,n} - D_n > 0$  then
       $i_{1,n+1} = a_{1,n}$  and  $i_{2,n+1} = i_{1,n} + i_{2,n} - D_n$ ;
    else
       $i_{1,n+1} = i_{1,n} + i_{2,n} + a_{1,n} - D_n$  and  $i_{2,n+1} = 0$ ;

```

Chapter 5

Additive Manufacturing as Dual Sourcing Option

In the previous chapter, we observed that dual sourcing often plays an important role in the transition from conventional manufacturing (CM) to additive manufacturing (AM). To that end, we study the implications of sourcing spare parts with a mix of AM and CM methods in more detail in this chapter¹ and develop an approach that is tailored towards the unique characteristics of dual sourcing with two production methods (which fills in Contribution 4, cf. Section 1.4.3).

As opposed to the previous chapter and in contrast with the literature, we consider different failure behavior of parts produced by AM and CM methods. This extension seems required since not only the production process but also design and used material are typically different for both part versions, cf. Wits et al. (2016). Predominantly for scenarios where the installed base size is relatively small compared to the demand rate, sourcing decisions will impact future demand, which makes the overall trade-off more complex. For example, fast resupply using AM helps reducing downtime in the short run, but in the long run a higher AM failure rate may offset this benefit. Using numerical experiments and a case study in the aviation industry, we explore under which conditions dual sourcing with two different production methods performs best.

¹This chapter is based on the working paper Knofius et al. (2017).

5.1 Introduction

The flexibility of selecting between AM and CM methods holds the potential of significant cost savings – even if AM part characteristics are (still) inferior to the CM version. For example, if the AM item fails three times more often than the CM item on average, there are still realistic cases where dual sourcing offers cost savings of up to 30%, cf. Section 5.4. Furthermore, we find that dual sourcing with AM is most promising for cases that are often observed in the spare parts business for capital goods: low demand rates, high holding and downtime costs. A case study conducted at an OEM in the aerospace industry extends our findings and exemplifies how practical considerations such as design costs and operational cost savings may influence the value of dual sourcing with AM. In particular, we clarify that the AM benefit of short lead times is often most interesting for parts that are not immediately natural candidates for design improvement with AM. This insight may demand a reconsideration of the current practice to select parts for AM based on the prospect of design improvements.

The remainder of this chapter is organized as follows. In Section 5.2, we position our work within the dual sourcing literature and clarify our contribution. Next, in Section 5.3, we develop a specific dual sourcing model and explain its evaluation and optimization. In Section 5.4, we conduct numerical experiments to study the value of dual sourcing in the spare parts business. Section 5.5 demonstrates the application of our model in practice. We close with Section 5.6.

5.2 Dual Sourcing Literature

Here, we review the dual sourcing literature where we focus on selected papers and refer to Minner (2003) and Zhou and Yang (2016) for a more extensive discussion. Dual sourcing models typically distinguish between two supply options: one that is cheap but with a long resupply lead time (regular supply), and one that is expensive but with a short resupply lead time (expedited order). The first contribution to the dual sourcing literature was made by Barankin (1961), who discusses a single-period model with emergency shipments. Whittemore and Saunders (1977) consider the periodic review case with deterministic lead times. They show that as soon as lead times between both supply options differ by more than one period, the optimal policy depends on delivery time and quantity of the in-transit parts. Although dynamic programming methods allow solving such problems, the large state space leads to computational intractability even for medium-sized problems. Hence, more recent contributions are devoted to approximations of the optimal policy.

For the periodic-review case, Veeraraghavan and Scheller-Wolf (2008) study a capacitated inventory model with deterministic lead times. They propose a dual-index policy that keeps track of one inventory position for each sourcing option. While the evaluation is partially based on simulation, the dual-index policy provides close to optimal results. Arts et al. (2011) extend the model and allow for stochastic lead times for the regular supply source. Also, their approach does not rely on simulation and thus decreases the computation time significantly. Scheller-Wolf et al. (2006) use a single-index policy where they monitor a single inventory

position for each sourcing option. In case the inventory position is below the expedited target level, an order is placed to raise the inventory position to this level. Next, a regular order is used to bring the inventory position up to the regular target level. Their results show that the single-index policy performs comparably well but can be computed 25-60 times faster than the dual-index policy by Veeraraghavan and Scheller-Wolf (2008). Interestingly Kiesmüller (2003), who proposed a dual-index policy for a remanufacturing system, argued that a single-index policy is expected to perform significantly worse under some conditions. Thus, further investigation of the single-index policy proposed by Scheller-Wolf et al. (2006) may be worthwhile to clarify under which conditions a single-index policy performs sufficiently well.

For the continuous-review case, Moinzadeh and Schmidt (1991) propose a dual-index policy with deterministic lead times to determine the order quantity for both supply options. Using information about arrival times of outstanding orders, they keep orders from the expensive sourcing option limited. Song and Zipkin (2009) extend the model of Moinzadeh and Schmidt (1991) for multiple supply options and stochastic lead times. Therefore, they construct a queueing network with overflow bypass. Zhou and Yang (2016) introduce a single-index (R,nQ) policy and study a more general setting regarding fixed order costs, batch ordering, and compound Poisson demand. Axsäter (2007) proposes a decision rule for placing expedited orders. Therefore, the regular source is operated according to a single-sourcing (R,Q) policy while, if the inventory level drops below zero, expedited orders are placed according to the decision rule. In follow up papers, the decision rule is improved (Axsäter, 2014) and the model is extended to incorporate the option of committed service times (Huang et al., 2011).

Allon and Van Mieghem (2010) propose a tailored base-surge (TBS) policy with stochastic lead times. Under the TBS policy goods are ordered at a constant rate from the cheap supply source while orders for the expensive supply source are issued according to a base stock policy. Song et al. (2017) propose an approach that finds the optimal policy for a continuous-review system for special cases. Therefore they transform the original problem into a simplified queueing system that shares the same optimal policy under specific conditions on the net inventory. For cases in which these conditions are violated, they exploit the results of the queueing system to construct a policy for the original system. Numerically they show that the resulting heuristic policy is close to optimal and typically outperforms the discussed methods of Song and Zipkin (2009) and Allon and Van Mieghem (2010).

The paper closest related to our work is the one by Song and Zhang (2016), who consider the use of AM methods as an emergency channel that may produce spare parts on-demand. In contrast to our model, stocking AM items is not allowed and thus a waiting time for each incoming demand is inevitable. In their model, AM equipment is capacitated (modeled as an $M/D/1$ queue) but typically allows faster, though more expensive, resupply than the CM source. Later, they find that AM equipment utilization typically remains low and therefore may support our assumption of an uncapacitated printing source, cf. Section 5.3.1. Also, they assume that AM parts have the same failure behavior as CM parts which is a fundamental simplification compared to our model. Overall, they find that the production of parts on-demand with AM methods lead to cost savings and inventory reductions compared to the application of CM methods only. Especially, for situations with large part variety, the

savings potential was significant. To the best of our knowledge, none of the many papers on dual sourcing addresses the impact of sourcing decisions on future demand which is essential for our analysis.

The contribution of the research presented in this chapter is as follows:

1. We develop a new dual sourcing model for a single-item, in which future demand depends upon the sourcing options being used via different failure rates for different spare part types. We develop an exact algorithm to analyze this situation.
2. In numerical experiments and a case study in the aviation industry, we explore the value of using a combination of AM and CM parts compared to single sourcing alternatives and study the structure of the optimal policy in case of dual sourcing.
3. We expose that printing on demand is typically not practical for downtime critical spare parts. Instead, spare parts stocks remain essential to realize high system availability.
4. We show that the current focus on design improvements with AM may cause an undervaluation of the logistical benefits with AM. That is, scenarios most interesting for design improvements are typically less promising for improving the effectiveness of spare parts supply with AM.

5.3 Model

We consider a company that is responsible for maintaining an installed base of a particular system type with a single stock point. Here, we focus on one specific component which can be sourced from either a CM or an AM supplier. Both options are characterized by differences in part reliability, unit costs and replenishment lead time. The company is uncertain on whether and how to combine both sourcing options. In particular, it is unclear whether to replace a failed part with a CM or AM spare part (if both spare parts versions are in stock) and when to reorder from the CM or the AM supplier instead.

Certainly, if the stock is (nearly) depleted, it may appear better to order the AM version to exploit the typically faster AM replenishment lead time. In other cases, though, sufficient stock might be available. Here, despite a long lead time, it may be better to order a usually more reliable and cheaper CM version of the item instead. These trade-offs are further complicated by the difference in part reliability: if increasingly more AM parts are operated, it is likely that future demand increases, assuming that AM items fail more often.

We investigate under which conditions it is advisable to use AM as dual sourcing option. Therefore, we develop a model which minimizes the long-run average cost by determining the optimal sourcing, maintenance, and inventory policy. In particular, we consider purchasing costs, maintenance costs, holding costs and backorder costs. In principle, we can solve this problem by different methods. Probably, the most common approach is the utilization of a Markov Decision Process (MDP). Here, we employ a continuous-time Markov Chain analysis in combination with linear programming methods to obtain a more efficient algorithm as we will elaborate in Section 5.3.4.

5.3.1 Key Assumptions

Before outlining the model in detail, we describe and motivate the key assumptions.

1. Failures follow a Poisson process.

This assumption appears justified given that we typically deal with mechanical parts that are designed to outlast the lifetime of the capital good or the preventive maintenance interval. For the latter case, preventive maintenance activities are typically scheduled well in advance and dedicated spare parts are usually ordered on demand. Hence, we only consider spare parts demand for corrective maintenance activities where failures are random in nature and may, for instance, be caused by unintended stress levels, imprudence during system usage, maintenance, shipping activities, or unobservable quality issues during the production, cf. Section 1.2.1. Nevertheless, the model can easily be extended to a wear-out process, for example by modeling the failure behavior by an Erlang-2 distribution.

2. One-for-one replenishment for both the CM and AM source.

In the dual sourcing literature, we often find replenishment policies other than one-for-one replenishment, cf. Section 5.2. In particular, many authors propose batch ordering which is motivated by high fixed order costs. For expensive, slow-moving items fixed order costs are typically insignificant, see, e.g. Van Houtum and Kranenburg (2015) and Hu et al. (2018), hence a one-for-one replenishment policy seems most reasonable. In case of AM, the absence of tooling and setup costs further reduces the relevance of fixed costs. Note that in Chapter 6, we propose a more general model that also is suitable for higher demand regimes since we explicitly regard the fixed ordering cost of the CM source.

3. Resupply rates are exponentially distributed.

We justify this assumption with two key reasons: first, this assumption facilitates the use of a continuous time Markov chain analysis. A drawback of exponential resupply lead times may be that ordering many items at once reduces the mean waiting time for parts (= lead time / number of parts ordered). Obviously, that is not realistic. However, we do not allow for ordering multiple parts at once, since we apply one-for-one replenishment policies only, cf. Assumption 2. Second, it is known that the performance of inventory systems for slow-moving spare parts is not very sensitive to the shape of the lead time distribution, cf. Alfredsson and Verrijdt (1999), Alvarez et al. (2013) and Sleptchenko and Van der Heijden (2016). As we show using a simulation study in Appendix 5.A, the same holds for our specific problem.

4. Uncapacitated AM equipment.

As justified in Section 5.3.2, we assume that spare parts are sourced from an external supplier whose capacity is shared between multiple customers. In such a situation, lead times are typically agreed upon in contracts, and waiting times for capacity are covered in these agreements. Furthermore, as shown by Song and Zhang (2016), even in case production is conducted on-demand, congestion at the AM equipment only marginally influences their results. Overall, they find a capacity utilization of less than 8% on average and 25% in the most extreme scenarios.

5.3.2 Notation and Model Outline

We model a single-item inventory system which serves an installed base of k systems, where each system requires one unit of the item to be operational. The CM and AM version fail according to a Poisson process with rates λ_C and λ_A respectively. Upon failure, we encounter maintenance costs m which include aspects such as spare parts transportation costs, repair costs, and service engineer costs. Maintenance costs are independent of the failed and replacement part. The unit costs of a new CM or AM part are denoted by c_C and c_A , and the resupply rates are exponentially distributed with rates μ_C and μ_A .

The installed base is supported by a single stockpoint carrying S non-repairable spare parts. Upon failure, a CM or AM spare part is taken from stock, and a new CM or AM part is ordered immediately. Hence, the total number of parts in the system (operational, in stock or in resupply) equals $N = k + S$. In case we run out of stock, demand is backlogged, and we incur backorder costs b per item per time unit. In fact, these backorder costs can be interpreted as penalty costs for system downtime. Otherwise, if a spare part is available, the replacement of the failed part takes place instantaneously (maintenance related downtime is covered in the maintenance costs). Holding costs are modeled as a fraction κ of the associated unit costs of the items in stock.

The state \mathbf{i} of the inventory system is described by the tuple $(n_{C,\mathbf{i}}, n_{A,\mathbf{i}}, r_{C,\mathbf{i}}, r_{A,\mathbf{i}}, s_{C,\mathbf{i}}, s_{A,\mathbf{i}})$ where $n_{C,\mathbf{i}}$ ($n_{A,\mathbf{i}}$) refers to the number of CM (AM) parts in operation, $r_{C,\mathbf{i}}$ ($r_{A,\mathbf{i}}$) refers to the number of CM (AM) parts in resupply, and $s_{C,\mathbf{i}}$ ($s_{A,\mathbf{i}}$) refers to the number of CM (AM) spare parts in stock. The set of feasible states is equal to

$$\begin{aligned} \Omega = \{ & (n_{C,\mathbf{i}}, n_{A,\mathbf{i}}, r_{C,\mathbf{i}}, r_{A,\mathbf{i}}, s_{C,\mathbf{i}}, s_{A,\mathbf{i}}) : \\ & n_{C,\mathbf{i}} + n_{A,\mathbf{i}} + r_{C,\mathbf{i}} + r_{A,\mathbf{i}} + s_{C,\mathbf{i}} + s_{A,\mathbf{i}} = N \\ & s_{C,\mathbf{i}} + s_{A,\mathbf{i}} = \max\{N - k - r_{C,\mathbf{i}} - r_{A,\mathbf{i}}; 0\} \\ & n_{C,\mathbf{i}} + n_{A,\mathbf{i}} \leq k \\ & n_{C,\mathbf{i}}, n_{A,\mathbf{i}}, r_{C,\mathbf{i}}, r_{A,\mathbf{i}}, s_{C,\mathbf{i}}, s_{A,\mathbf{i}} \geq 0 \} \end{aligned}$$

Note that the definition of the state space excludes degenerated transitions. So, in case $n_{C,\mathbf{i}} + n_{A,\mathbf{i}} < k$ it is impossible to have items in stock and a failed item always has to be replaced immediately. Overall, the state space grows polynomially with the installed base size k and maximum number of spare parts S . More precisely, the state space size is equal to:

$$\binom{S+4-1}{4-1} (k+1) + \sum_{a=1}^k a(k+S+2-a)$$

where the first term subsumes states with no backlog and the summation states with backlog.

Upon failure of an item, one has to take two decisions. First, whether to use an AM or a CM item from stock (if possible) to replace the failed item (maintenance decision). Note that the item which failed does not need to be replaced with the same item version. Second, we have to decide whether to reorder an AM or a CM item to replenish the stock (sourcing decision). Optimal decisions are dependent on the state \mathbf{i} .

We define a matrix \mathbf{X} to represent the decisions for all states, where each column corresponds to a certain state \mathbf{i} , and each row to a decision option c , defined as follows:

- $c=1$: take AM version from stock (if possible), and order AM version.
- $c=2$: take CM version from stock (if possible), and order AM version.
- $c=3$: take AM version from stock (if possible), and order CM version
- $c=4$: take CM version from stock (if possible), and order CM version

where each matrix element $x_c(\mathbf{i})$ represents a binary variable that is equal to 1 if we choose decision option c in state \mathbf{i} and 0 otherwise. Note that in cases where $s_{C,\mathbf{i}}+s_{A,\mathbf{i}}=0$, it is indifferent whether to choose $c=1$ or $c=2$. The same holds for the choice between $c=3$ and $c=4$. In Table 5.1, we summarize the notation.

Table 5.1: *Notation overview*

Notation	Explanation
k	Installed base size
λ_C	CM failure rate
λ_A	AM failure rate
μ_C	CM resupply rate
μ_A	AM resupply rate
c_C	CM piece price
c_A	AM piece price
o_A	Operational cost savings per AM part per time unit
m	Maintenance costs
b	Backorder costs per time unit
d	Depreciation costs per time unit
κ	Holding costs as fraction of the piece price per time unit
S	Maximum number of spare parts circulating in the inventory system
N	Total number of parts ($N=k+S$)
$\mathbf{i} \in \Omega$	Six-dimensional state vector, $(n_{C,\mathbf{i}}, n_{A,\mathbf{i}}, r_{C,\mathbf{i}}, r_{A,\mathbf{i}}, s_{C,\mathbf{i}}, s_{A,\mathbf{i}})$
$n_{C,\mathbf{i}}$	Number of CM parts in operation
$n_{A,\mathbf{i}}$	Number of AM parts in operation
$r_{C,\mathbf{i}}$	Number of CM parts in resupply
$r_{A,\mathbf{i}}$	Number of AM parts in resupply
$s_{C,\mathbf{i}}$	Number of CM parts in stock
$s_{A,\mathbf{i}}$	Number of AM parts in stock
$c \in \{1,2,3,4\}$	Maintenance and sourcing decision

5.3.3 Evaluation

Given \mathbf{X} , k and S , the model is evaluated using a continuous-time Markov Chain. In Figure 5.1 we show the ten possible transitions to and from state $(n_{C,\mathbf{i}}, n_{A,\mathbf{i}}, r_{C,\mathbf{i}}, r_{A,\mathbf{i}}, s_{C,\mathbf{i}}, s_{A,\mathbf{i}})$ for the case where an AM part arrives or fails. The ten transitions associated with the arrival or failure of a CM part are omitted but exhibit the same pattern. We use \mathbf{i}' to refer to the associated predecessor states of $(n_{C,\mathbf{i}}, n_{A,\mathbf{i}}, r_{C,\mathbf{i}}, r_{A,\mathbf{i}}, s_{C,\mathbf{i}}, s_{A,\mathbf{i}})$. Furthermore, if the maintenance decision is indifferent (i.e. no stock available), as we elaborated in the previous section, we use $c=2$ and $c=4$ as default in Figure 5.1.

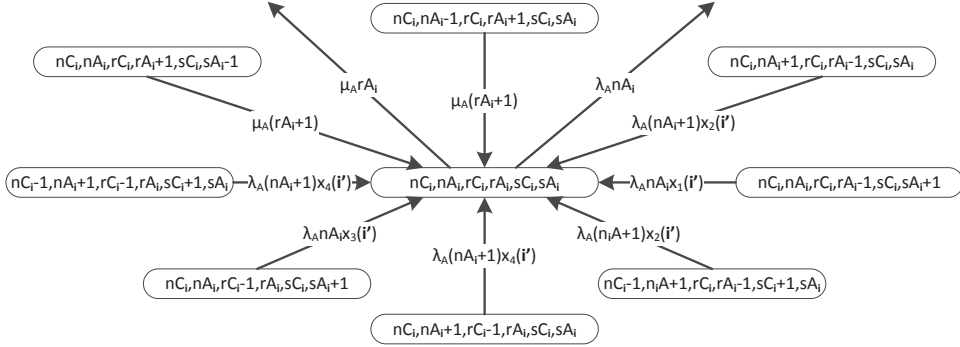


Figure 5.1: Possible transitions to and from state $(n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i}, s_{C,i}, s_{A,i})$ if an AM item arrives or fails

To explain the underlying logic of Figure 5.1, let us focus on the transition in the top-right corner, i.e. with a rate of $\lambda(n_{A,i}+1)x_2(i')$ we transition from state $(n_{C,i}, n_{A,i}+1, r_{C,i}, r_{A,i}-1, s_{C,i}, s_{A,i})$ to state $(n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i}, s_{C,i}, s_{A,i})$. This transition describes the situation where an AM item fails and we decide to order an AM part while no items are in stock. Hence, $n_{A,i}$ decreases by one unit while $r_{A,i}$ increases by one unit. The balance equations (cf. Appendix 5.B) directly follow from the transitions illustrated in Figure 5.1 and, in combination with the normalization equation, allow the computation of the state probabilities p_i with common methods. The result is captured by the column vector \mathbf{p} with elements p_i .

After \mathbf{p} was determined, the long-run average cost C are computable by $\mathbf{g}^T \mathbf{p}$ with \mathbf{g} representing an $|\Omega|$ -dimensional column vector of cost g_i in state \mathbf{i} . Given that we consider purchasing, maintenance, holding and backorder costs, the total expected cost, given the system is in state \mathbf{i} , are given by

$$g_i = \mu_C c_C r_{C,i} + \mu_A c_A r_{A,i} + m(\lambda_C n_{C,i} + \lambda_A n_{A,i}) + \kappa(c_C s_{C,i} + c_A s_{A,i}) + b(k - n_{C,i} - n_{A,i})$$

5.3.4 Optimization of the Sourcing and Maintenance Policy

Formally, the optimization problem can be expressed as Problem 5.1:

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} && \mathbf{g}^T \mathbf{p} \\ & \text{subject to} && \mathbf{Q}(\mathbf{X}) \mathbf{p} = \mathbf{0} \\ & && \mathbf{1}_{|\Omega|}^T \mathbf{p} = 1 \\ & && \mathbf{p} \geq \mathbf{0} \\ & && \mathbf{1}_{|C|}^T \mathbf{X} = \mathbf{1}_{|\Omega|}^T \\ & && x_c(\mathbf{i}) \in \{0, 1\} \quad \forall c, \mathbf{i} \end{aligned} \tag{5.1}$$

where $\mathbf{Q}(\mathbf{X}) \mathbf{p} = \mathbf{0}$ represents the balance equations with matrix $\mathbf{Q}(\mathbf{X})$ describing the generator of the Markov process and $\mathbf{0}$ denoting an $|\Omega|$ -dimensional vector of zeros. Furthermore,

$\mathbf{1}_{|\Omega|}^T \mathbf{p} = 1$ defines the normalization equation with $\mathbf{1}_m$ being an m -dimensional column vector of ones (here $m = |\Omega|$ or $m = |c|$).

Given the product $x_c(\mathbf{i})p_{\mathbf{i}}$ in the balance equations (cf. Appendix 5.B), Problem 5.1 is nonlinear and therefore less computationally tractable. Therefore, we transform (5.1) into an equivalent linear formulation. The key step of the transformation is the rearrangement of the balance equations and the substitution of the product $x_c(\mathbf{i})p_{\mathbf{i}}$ with the variable $y_c(\mathbf{i})$. One may interpret $y_c(\mathbf{i})$ as the long run fraction of the time that the system is in state \mathbf{i} and action c is chosen. This operation allows us to redefine Problem 5.1 as follows:

$$\begin{aligned}
 & \underset{\mathbf{Y}}{\text{minimize}} \quad \mathbf{g}^T \mathbf{p} \\
 & \text{subject to} \quad \sum_{c=1}^4 \mathbf{L}_c \mathbf{y}_c + \mathbf{M} \mathbf{p} = \mathbf{0} \\
 & \quad \mathbf{1}_{|\Omega|}^T \mathbf{p} = 1 \\
 & \quad \mathbf{p} \geq \mathbf{0} \\
 & \quad \mathbf{1}_{|c|}^T \mathbf{Y} = \mathbf{p}^T \\
 & \quad y_c(\mathbf{i}) \geq 0 \quad \forall c, \mathbf{i}
 \end{aligned} \tag{5.2}$$

where matrix \mathbf{L}_c contains all transition rates dependent on decision option c , and matrix \mathbf{M} contains all transition rates independent of decision option c . For example, as shown in Appendix B, \mathbf{M} contains the last four terms in the balance equations. Furthermore, we use \mathbf{y}_c to describe an $|\Omega|$ -dimensional column vector with elements $y_c(\mathbf{i})$ and \mathbf{Y} to describe a matrix with row vectors \mathbf{y}_c .

As a consequence, Problem (5.2) allows us to find the optimal values of \mathbf{Y} and \mathbf{p} with linear programming methods. Next, the optimal decisions \mathbf{X}^* are recovered with the relation $x_c(\mathbf{i}) = y_c(\mathbf{i})/p_{\mathbf{i}}$. Note that in case $0 < y_c(\mathbf{i}) < p_{\mathbf{i}}$ we would obtain a randomized policy which violates the constraint $x_c(\mathbf{i}) \in \{0, 1\} \quad \forall c, \mathbf{i}$ of Problem 5.1. Sleptchenko and Johnson (2015) however, show that, given a linear cost function, $x_c(\mathbf{i}) \in \{0, 1\} \quad \forall c, \mathbf{i}$ holds, i.e. we always obtain a deterministic policy with Problem (5.2).

Compared to an MDP approach, the described optimization procedure is more efficient: instead of four equations per state (one for each action), we can represent the problem with two equations per state (balance equation and policy constraint). Furthermore, we do not need to discretize time, which reduces the number of possible transitions in each state.

5.3.5 Optimization of the Stocking Policy

The optimal base stock levels in the $(S-1, S)$ spare parts supply systems are usually determined using a greedy approach using convexity properties of the cost function in S , cf. Van Houtum and Kranenburg (2015). However, in the dual sourcing case presented here, a changing sourcing strategy with increasing S causes a non-convex behavior of the individual costs terms. Figure 5.2 gives an example for this behavior, and in Figure 5.3 we show the corresponding failure rates. Observe that, with increasing S the CM failure rate increases

which is caused by the increasing preference to install CM rather than AM items. The switch from AM to CM produced components (with the increase of S), leads to the non-convexity of the purchasing, backorder and holding costs terms in the objective function.




Figure 5.2: Examples of cost changes when S increases




Figure 5.3: Examples of failure rate changes when S increases

At the same time, we were not able to find any example of non-convexity of the objective function. The fluctuations in the costs terms appear to compensate each other, and the total cost function remains convex in S . Therefore, we propose the following conjecture that will be used in the optimization procedure described below.

Conjecture 1 *The optimal total cost $C(S)$ with fixed base stock level S and optimized ordering policy \mathbf{X} has only one local minimum for all $S \geq 0$, which consequently will be the global minimum.*

To find the optimal number S^* of spare parts circulating in the inventory system, we use a greedy approach to determine the optimal long-run average cost $C^*(0)$ given that $S=0$ first. Next, we set $S=1$ and determine $C^*(1)$. In case $C^*(0) < C^*(1)$, $S=0$ is the optimal base stock level. Otherwise, we continue to increment S until $C^*(S) < C^*(S+1)$.

5.3.6 Joint Optimization

It is also possible to jointly optimize the stocking policy with the sourcing and maintenance policy (cf. Appendix 5.C). In the joint optimization approach, we set an upper bound on the number of spare parts required, and include the option not to order any part upon failure in state \mathbf{i} .

However, numerical experiments show that the joint optimization is computationally inferior to the greedy approach. The key reason for this characteristic is that it is difficult to find a tight upper bound on the number of spare parts, in combination with the fact that the computation times grow rapidly with the size of the state space. On the other hand,

an advantage of the joint optimization approach is that this model can be extended to a dynamic stocking policy (cf. Appendix 5.D). Given the different failure rates of the AM and CM version, it is likely that the optimal inventory level depends on the mixture of AM parts and CM parts in the installed base. So, the inventory level will be state-dependent. The impact of such a dynamic spare parts inventory policy is a topic for further research.

5.3.7 Design costs

Designing two instead of one part version causes additional design costs but holds the promise of a more flexible inventory system. The trade-off between more desirable system parameters and additional costs has not received much attention in the literature but is regarded as a promising direction for future research, cf. Hu et al. (2018). Often, it is assumed that the design costs are an increasing convex function of the part reliability, see Mettas (2000) and Öner et al. (2010). It remains unclear how other aspects such as AM resupply lead time and AM unit costs may also impact the design costs for the situation discussed in this chapter. More research is needed to explore this relationship.

In order to include design costs in our analysis, we assume that design costs are depreciated over the remaining life cycle length. Depending on the applied depreciation method, the per time unit expense d is then derived from the remaining life cycle length and design costs difference between single and dual sourcing.

5.3.8 Operational Cost Savings

Operational cost savings depend on the installed base composition. Thus, if we assume operational cost savings of o_A per time unit and installed AM item, we can easily include the savings by adding the term $-o_A n_{A,i}$ to the cost function g_i . We will include operational cost savings in the case study presented in Section 5.5, but do not consider this cost factor in the numerical experiments. Operational cost savings are highly case dependent and thus appear unsuitable for a generic assessment of the effects of dual sourcing with AM. Furthermore, business cases highly dominated by AM-specific operational cost savings diminish the value of dual sourcing as we will discuss in Section 5.5.

5.4 Numerical Experiments

To study the value of dual sourcing with two different production methods, we carry out numerical experiments. In Section 5.4.1, we elaborate on the experimental settings and justify the selected parameter ranges. In Section 5.4.2, we present the results while Section 5.4.3 summarizes the key findings.

5.4.1 Experimental Settings

The value of AM technology in the spare parts business is perceived highest for low-volume applications, i.e., for applications where economies of scale are of minor importance (Khajavi

et al., 2014) and (Liu et al., 2014). So, we focus on scenarios where the combination of failure rate and installed base size causes 1 to 7 demands with CM sourcing per year.

For expensive spare parts with infrequent demand, the replenishment lead time is often long. For example, Basten (2010) observes spare parts replenishment lead times of more than a year in the defense sector. In comparison, AM lead times are typically short which is exemplified by Liu et al. (2014) who report on data from the aerospace industry. For CM items they find lead times between 1.5 and 8 month, whereas, for AM parts, they typically assume lead times of less than a month. AM lead times in the order of magnitude smaller than a week appear unrealistic – at least for metal printing processes. Pre- and post-processing activities seem inevitable for industrial applications today. For instance, usually a preheating and lengthy cool-down phase is necessary to ensure an acceptable part quality (Shouche et al., 2016). Furthermore, most companies rely on external AM suppliers which causes additional transportation times. Accordingly, we study AM lead times of 1 to 4 weeks and CM lead times of 2 to 6.5 month.

The AM piece price is typically higher than for CM parts. Not only does this situation relate to the novelty of industrial AM process, but also to the high raw material costs. Additionally, AM processes remain labor and skill intensive (Weller et al., 2015). As we already discussed in Section 5.1, the reliability of printed parts is often inferior compared to their CM counterparts. Thus, for various applications customized AM designs become necessary to achieve a comparable or even superior quality compared to the CM version. As we will also exemplify in the case study, one possible approach is the utilization of higher quality raw material. For other applications, a lower reliability might be acceptable also in view of an otherwise significantly higher AM piece price. For instance, in Section 5.4.2, we discuss AM applications at remote locations where printed, low-reliability parts are used as emergency solution to avoid high holding and downtime costs. Overall, we evaluate scenarios where the AM piece price varies between 1 and 3 times of the CM piece price.

We consider three different CM items in the numerical experiments. Item 1 exhibits a failure rate of $\lambda_C = 0.02$ failures per month, i.e., a MTBF of about 4 years; a resupply rate of $\mu_C = 0.5$ per month, i.e, an average lead time of two months, and a unit costs of $c_C = 10$. Item 2 fails less frequently ($\lambda_C = 0.015$) and has a longer resupply lead time ($\mu_C = 0.25$). Item 3 exhibits an even lower failure rate ($\lambda_C = 0.01$) and longer resupply lead time ($\mu_C = 0.15$). All other parameters follow commonly observed conditions in the low-volume, high-variety spare parts business.

5.4.2 Sensitivity Analysis

We consider 6,561 problem instances following the experimental setup discussed in the previous section. Table 5.2 shows the parameter settings and results. In the column ‘*Cost savings compared to*’ we show the percentage of cost that can be saved with a dual sourcing approach if compared to single sourcing with CM, single sourcing with AM and the best instance specific single sourcing approach as average over associated instances. The column ‘*Best sourcing approach*’ shows the percentage of instances where single sourcing with CM,

single sourcing with AM or dual sourcing, respectively, is the best approach. The columns ‘Base stock level’ and ‘Inventory level’ show the average base stock and inventory level for the different sourcing concepts over associated instances.

Table 5.2: Numerical experiments as average over all instances, where single sourcing with only CM is denoted by CM, single sourcing with only AM by AM, single sourcing with the best instance specific approach by Best and dual sourcing by Dual.

Parameter	Value	Unit	Cost savings compared to			Best sourcing approach			Base stock level			Inventory level		
			CM	AM	Best	CM	AM	Dual	CM	AM	Dual	CM	AM	Dual
Item	1	-	7%	43%	4%	24%	11%	64%	3.37	1.97	2.83	2.59	1.59	1.94
	2	-	10%	48%	8%	11%	4%	85%	4.30	1.97	3.37	3.12	1.59	2.14
	3	-	14%	57%	13%	2%	0%	98%	4.48	1.97	3.48	3.17	1.59	2.16
k	10	installed parts	12%	46%	9%	11%	6%	82%	2.74	1.40	2.16	2.20	1.22	1.55
	20	Installed parts	10%	50%	8%	12%	5%	83%	4.15	2.01	3.27	3.06	1.63	2.13
	30	installed parts	9%	52%	7%	13%	4%	83%	5.26	2.50	4.25	3.62	1.93	2.57
λ_A	0.0175	failures/month	15%	27%	8%	4%	16%	18%	4.05	1.40	2.83	2.96	1.22	1.65
	0.035	failures/month	9%	54%	9%	13%	0%	87%	4.05	2.01	3.36	2.96	1.64	2.23
	0.0525	failures/month	7%	67%	7%	20%	0%	80%	4.05	2.49	3.49	2.96	1.92	2.37
μ_A	1	supplies/month	6%	51%	5%	19%	4%	77%	4.05	2.72	3.54	2.96	2.06	2.36
	2	supplies/month	10%	49%	8%	10%	6%	84%	4.05	1.88	3.21	2.96	1.56	2.05
	4	supplies/month	14%	48%	11%	7%	6%	86%	4.05	1.31	2.94	2.96	1.17	1.84
c_A	10	euro/part	15%	32%	9%	6%	16%	79%	4.05	2.25	2.91	2.96	1.86	1.65
	20	euro/part	9%	53%	9%	12%	0%	88%	4.05	1.93	3.32	2.96	1.56	2.23
	30	euro/part	7%	64%	7%	19%	0%	81%	4.05	1.72	3.46	2.96	1.36	2.37
b	20	euro/month	4%	50%	3%	34%	5%	61%	2.63	0.90	2.31	1.58	0.59	1.29
	200	euro/month	11%	49%	9%	2%	5%	92%	4.07	2.01	3.25	2.97	1.61	2.09
	2000	euro/month	15%	49%	12%	1%	5%	94%	5.44	3.00	4.13	4.33	2.59	2.87
m	2	euro/failure	13%	49%	10%	9%	6%	85%	4.05	1.97	3.16	2.96	1.60	2.00
	10	euro/failure	10%	50%	8%	12%	5%	85%	4.05	1.97	3.24	2.96	1.59	2.09
	18	euro/failure	8%	50%	6%	15%	4%	80%	4.05	1.96	3.29	2.96	1.59	2.15
κ	0.15	-	9%	51%	7%	13%	5%	83%	4.22	2.11	3.40	3.13	1.72	2.24
	0.2	-	10%	49%	8%	13%	5%	82%	4.04	1.95	3.22	2.95	1.57	2.08
	0.25	-	11%	48%	9%	11%	6%	83%	3.89	1.85	3.07	2.80	1.48	1.93

As an average over all instances, we obtain cost savings of more than 8% if we compare the dual sourcing cost with the cost of the best single sourcing option of each specific instance. In extreme cases, we even find instances with cost savings of more than 30% which is remarkable because also for these cases the AM failure rate and the AM piece price is often two to three times higher than for the CM part.

Single sourcing with AM is significantly outperformed by dual sourcing, see the column ‘Cost savings compared to AM’ in Table 5.2. This result is explained by the following observations. First, an installed base composed of AM items typically causes a higher demand rate, which increases the maintenance and the purchasing costs, cf. Figure 5.4. Second, given that we only order the usually more expensive AM item, the purchasing and the holding costs increase (the holding costs are a fixed fraction of the corresponding piece prices).

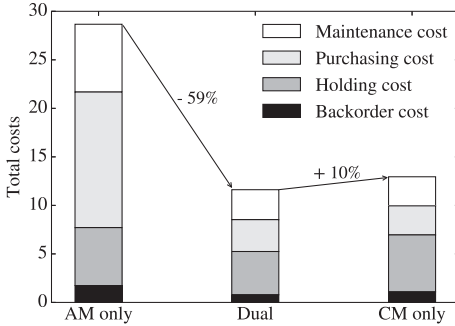


Figure 5.4: Cost composition for the different sourcing concepts as average over all instances




Figure 5.5: Number of dual sourcing instances where we first order from the AM source if $x\%$ of the total stock has been depleted

In this context, we also stress that printing spare parts on demand is no option for most instances. For example, even if the resupply lead time is equal to one week ($\mu_A = 4$), we find an average AM base stock level of 1.31 over associated instances, cf. Table 5.2. Hence, the common belief, see e.g. P eres and Noyes (2006), Campbell et al. (2012) and Thomas and Gilbert (2014), that the spare parts business may particularly benefit from a demand driven production approach has to be reconsidered in this respect.

Single sourcing with CM is usually also outperformed by dual sourcing, see the column ‘Cost savings compared to CM’ in Table 5.2 which is caused by a reduction of the holding and the backorder costs with dual sourcing in comparison to single sourcing with CM, cf. Figure 5.4. Also, the purchasing and the maintenance costs increase with dual sourcing is marginal compared to single sourcing with CM. This finding relates to the flexibility to order from the AM source if the stock is (nearly) depleted while, otherwise, the typically cheaper and more reliable CM part is ordered. In fact, as shown in Figure 5.5, the first order from the AM source usually only takes place if nearly no stock is left. Hence, in line with the dual sourcing literature, the AM source associated with a short resupply lead time usually functions as an emergency option.

The cost savings with dual sourcing are most sensitive to the AM resupply rate (μ_A) and the backorder costs (b), if we study the column ‘Cost savings compared to Best’ in Table 5.2. This finding can be explained by the observation that both parameters primarily influence the holding and the backorder costs. As we just discussed, only by reducing either of these cost factors dual sourcing may outperform single sourcing with CM methods.

In contrast, the cost savings potential with dual sourcing is less sensitive to the AM failure rate (λ_A) and the AM piece price (c_A), cf. column ‘Cost savings compared to Best’ in Table 5.2. So, dual sourcing appears to be a suitable vehicle to exploit the short AM resupply lead time. Negative effects of high AM unit costs and low AM part reliability turn out less cru-

cial for a successful implementation. As a consequence, various new service concepts are foreseeable. For instance, dual sourcing with AM may provide a good option to efficiently service downtime critical goods at remote locations. Today, long lead times and high inventory costs often cause high service cost. These may be reduced if locally printed – possibly less reliable spare parts – become available as an emergency solution. First applications can be found in defense organizations which experiment with mobile AM production facilities (McLearn, 2015).

A small sub-experiment may substantiate the potential further. Therefore we consider an AM resupply lead time of about one day ($\mu_A=25$) and a MTBF of 5 months ($\lambda_A=0.2$). The CM part follows the characteristics of Item 3 ($\lambda_C=0.01$, $\mu_C=0.15$, $c_C=10$). The other parameter values remain unchanged. Even under these conditions, dual sourcing leads to cost savings of about 17% on average. We conclude that the application of AM at remote locations deserves more attention in the literature.




Figure 5.6: Percentage point change of the cost factors if parameter value is increased from its minimum value to its maximum value




Figure 5.7: Effect of design costs on the cost savings with dual sourcing compared to single sourcing with CM

The cost savings with dual sourcing become larger when increasing the holding costs fraction (κ), decreasing the maintenance costs (m) and decreasing the installed base size (k), cf. column ‘Cost savings compared to Best’ in Table 5.2. If we study the impact of changing these values on the cost composition in Figure 5.6, the underlying reasons become clear. In case of parameter m , the holding and the backorder costs decrease as a percentage of the total costs while the maintenance costs increase significantly. Hence, the cost saving potential of dual sourcing reduces. We find a similar result for the installed base size (k), except that also the purchasing costs increase as a percentage of the total costs. Note that the installed base size primarily influences the demand rate and thus explains the increase in the purchasing costs. In contrast, an increase in the holding costs fraction (κ) increases the holding costs and the backorder costs while the maintenance and purchasing costs decrease.

As discussed in Section 5.3.7, designing one instead of two part versions typically causes

additional design costs. In Figure 5.7 we show the effect of additional design costs on the average cost savings with dual sourcing compared to single sourcing with CM. Here, the assessed values for the additional design costs are in the same order of magnitude as we find in the case study presented in Section 5.5. As expected, the value of a dual sourcing approach decreases with increasing design costs. Nevertheless, even in case of high design costs we still find various instances where dual sourcing offers cost savings of more than 20% compared to single sourcing with CM in Figure 5.7.

5.4.3 Summary of Key Findings

The numerical experiments lead to the following conclusions:

1. With dual sourcing, it is possible to exploit most benefits of a typically short AM resupply lead time while common drawbacks of high AM piece price and/or low AM part reliability have a significantly lower impact than with AM single sourcing. For example, even if the AM piece price (c_A) or the AM failure rate (λ_A) are up to three times higher than those of the CM part, we still find instances with cost savings of more than 30% compared to the best single sourcing approach.
2. Printing spare parts on demand does not appear to be suitable for downtime critical spare parts. Instead, stock remains necessary to reduce the risk of expensive downtime. As a result, dual sourcing largely outperforms single sourcing with AM. On average, we find cost savings of more than 48% even if the AM resupply lead time is more than twenty times shorter than the CM resupply lead time.
3. In line with the dual sourcing literature, AM typically functions as an emergency source. In most cases, more than 80% of the base stock has to be depleted before we order the first time from the AM source.
4. If single sourcing with CM leads to high holding and backorder costs relative to the other cost factors, dual sourcing is most promising. This type of conditions are typically obtained with a low CM resupply rate (μ_C), a high holding costs fraction (κ), a high backorder costs (b), low maintenance costs, and/or a low demand rate.

5.5 Case Study

To gain further insights into the practical implications of combining AM and CM sourcing modes, we conduct a case study at a service provider in the aerospace industry. More explicitly, we consider a hinge bracket that is used for connecting the rudder and the aircraft. According to the Federal Aviation Administration (FAA), the hinge is categorized as a Class 2 product. That is, a failure may jeopardize the safety of an aircraft and thus is considered critical. To avoid devastating consequences, aircraft manufacturers typically work with redundancies. Nevertheless, a failure of a hinge has to be corrected upon discovery. Hence, from an operational perspective, redundancies only exist to prevent failure risks during a flight but do not lead to a k-out-of-n maintenance system. The most common failure modes are fractures of the hinge that are caused by fatigue, tensile stress or corrosion.

Figure 5.8 illustrates the CM and AM design of the hinge. The service provider estimated the total design costs for the AM hinge at around 15,000 euro. While the CM hinge is an aluminum machined part, the AM hinge is built from titanium powder (Ti-6AL-4V) with Selective Laser Melting (SLM). Next to accounting for the different material properties, the AM design is topologically optimized to reduce the weight of the hinge. Overall, despite titanium being heavier than aluminum, the topology optimization leads to a weight reduction of about 25%, which leads to fuel savings of approximately $\sigma_A = 15$ euro per AM part per year.




Figure 5.8: CM (a) and AM hinge (b) design

According to further investigations reported by the service provider, cost reductions resulting from fuel savings though were not sufficient to justify a transition to AM. Today's unstandardized AM certification process in the aerospace industry for structural parts would require an investment of about 250,000 euro which renders any application of AM economically infeasible. However, the service provider is confident that fast regulatory process development may dismantle the certification costs barrier soon. For instance, the FAA just distributed a draft version of its Additive Manufacturing Strategic Roadmap to establish a standardized certification process (3ders.org, 2017).

Another concern constitutes the high AM piece price. The service provider argues that the AM piece price has to decrease significantly before the production with AM methods becomes economically feasible. We provide the input parameters for the hinge case in Table 5.3. If not mentioned otherwise below, the parameter values were obtained from company records. For the design costs, we assume that the costs are evenly spread over the remaining life cycle of 15 years and thus encounter depreciation costs $d = 1,000$ euro per year. Compared to the numerical experiments, we take operational cost savings into account by following the explanation in Section 5.3.7.

The failure rate λ_A is based on the following insights: mechanical tests revealed that the AM hinge exhibits superior static strength compared to the CM hinge. Also, given that the AM hinge is produced with titanium powder rather than with aluminum, failures caused by corrosion can be ruled out. On the other hand, the surface roughness and porosity expected from the SLM production process may concentrate more tensions in the hinge.

Table 5.3: *Model input data hinge case*

Parameter	Value	Unit
c_A	1,197	euro/unit
c_C	480	euro/unit
μ_A	12	supplies/year
μ_C	4	supplies/year
λ_A	0.027	failures/year
λ_C	0.036	failures/year
k	382	installed parts
b	32,500	euro/day
m	100	euro/failure
d	1,000	euro/year
o_A	15	euro/AM part/year
κ	0.2	-

Hence, it is likely that the AM hinge fails more often under cyclic load caused by fatigue. Combining these different observations, AM experts estimate $\lambda_A = 0.75\lambda_C$. Note that from a regulatory perspective even inferior reliability would be acceptable as long as it is possible to prove that the integrity of the system is not compromised and that airworthiness is not affected. For the hinge bracket, these criteria are met given the use of redundancy, regular inspections, and the favorable failure behavior.

The backorder costs b follow from the criticality of the hinge. According to company representatives, it is likely that a stock-out leads to additional downtime of the aircraft. Hence, depending on the aircraft type, backorder costs b vary between 15,000 to 50,000 euro per day. For our analysis, we use $b = 32,500$ as a base case but note that our results are not very sensitive to changes of b in that range: the long-run average service cost per year differ by less than 1% if we compare $b = 15,000$ and $b = 50,000$.

Currently, the case company serves an installed base of $k = 382$ hinges, each with a mean time between failure of 27.78 years. To obtain computationally tractable data, we need to aggregate demand streams. To that end, we consider a smaller k' instead and multiply the failure rates λ_A and λ_C with the fraction $382/k'$ in order to keep the total demand rate unchanged. Given that, as soon as k' is sufficiently large, the demand variability becomes nearly independent of the installed base size, this transformation leads to proportionally the same results. After the results are computed, we simply multiply the operational cost savings with $382/k'$ to obtain the savings for the case $k = 382$. We illustrate this effect by plotting the long-run average net cost (cost minus savings) per year with dual sourcing for different values of k' in Figure 5.9. In the remainder, we use $k' = 50$ for our analysis. Furthermore, in the next chapter, we revisit the hinge bracket case and are able to confirm that the choice to consolidate demand streams does not lead to a large approximation error.

Figure 5.10 shows that dual sourcing appears less valuable for the hinge case. In particular, we observe that dual sourcing leads to total costs of about 13,838 euro. Even if we were considering that with dual sourcing we can secure cost savings of about 4,581 euro from

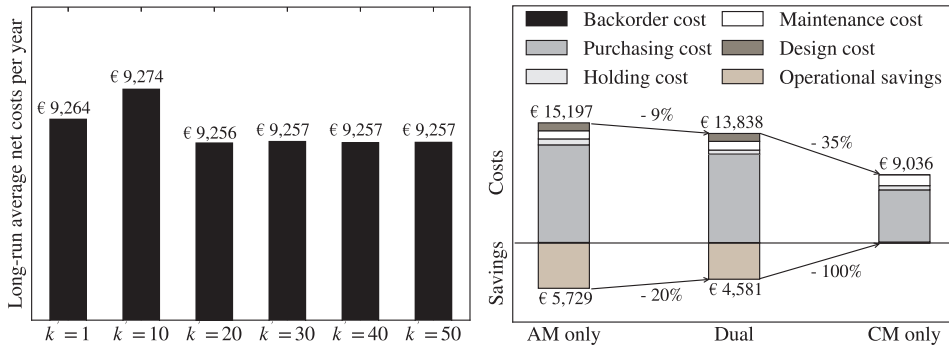


Figure 5.9: Effect of consolidating demand streams where $\lambda_A = 382/k\lambda_A$ and $\lambda_C = 382/k\lambda_C$

Figure 5.10: Total cost and savings for the three sourcing options excluding certification costs

weight reductions, using CM only remains the better choice. An explanation for this results offer the high purchasing costs. As we found in the numerical experiments (cf. Section 5.4), dual sourcing is most valuable if the holding or the backorder costs dominate. For the specific case this is clearly not the case since we only find 70 euro (AM only), 55 euro (Dual), and 120 euro (CM only) for the backorder costs, and 752 euro (AM only), 507 euro (Dual), and 534 euro (CM only) for the holding costs. To that end, it appears likely that the value of AM for the hinge bracket case (if any) would originate from operational rather than service cost savings. However, it remains an open research question if and how a dual sourcing approach is useful for such scenarios. The results clearly indicate a trade-off between operational cost savings and service cost. That is, in order to secure higher operational cost savings, we typically would install more AM parts than is optimal to minimize service cost.

Finally, and this represents a key learning for us from this case study, we emphasize that today's attention for AM technologies is mainly motivated by the prospect of design improvements/operational cost savings. While this certainly has its value, it dilutes the contribution of AM technologies to service cost savings. In our opinion there are two reasons for this: first, design improvements often scale with quantity. For example, if topology optimization leads to weight reductions of a few grams only, the effort is typically justified by a large installed base size. Accordingly, AM demonstrators frequently exhibit "high" demand rates (say, $>10/\text{year}$) rather than high holding and backorder costs. So, as also exemplified by the hinge case, service cost savings become less interesting. Second, the prospect of design improvements and technological complexity of AM methods typically put engineers in the lead to identify parts worthy of production with AM methods. Unfortunately, as a result, logistic opportunities may be underrated/overlooked. Instead, logisticians should become more involved in this task. To that end, we believe that it is important to carry out more case-based research in the low-volume, downtime-critical spare parts business to exemplify the merit of AM methods in this field to decision makers.

5.6 Conclusion

Demand fulfillment with a mix of AM and CM production methods has not been sufficiently discussed in the literature yet. In this chapter, we have addressed this gap by studying a dual sourcing concept where AM and CM methods are used in parallel to fulfill spare parts demand. A key aspect of our model is that we account for the different failure behavior of parts obtained from AM and CM supply sources.

Overall, dual sourcing turns out to reinforce the value of AM methods in the spare parts business. In particular, this holds if backorder costs and holding costs are high relative to purchasing and maintenance costs. Moreover, our results stipulate that the concept of printing spare parts on demand usually is not suitable for downtime-critical parts. Instead, stock remains necessary even if the AM resupply lead time is very short. In our opinion, the most remarkable finding, however, is that dual sourcing offers an approach to profit from the fast AM resupply lead time even if the AM piece price is high and/or AM part reliability is low compared to the CM part. Consequently, dual sourcing may extend the operating range of AM methods in the spare parts business significantly. In the light of this finding, new sourcing concepts are likely. For instance, our results indicate that supplementing CM supply with less reliable, but locally producible AM parts may reduce operating costs of capital goods at remote locations considerably.

It appears valuable to investigate the effect of this approach more in detail. To that end, one may consider scenarios where the availability of the CM source is restricted. Another model extension may consider a typically higher uncertainty about the AM failure behavior. While for the CM source, we typically (should) know historical failure records, the failure behavior of AM items is still rather unclear. Finally, in this chapter, we assumed a base stock policy. However, state-dependent demand rates may deem a dynamic inventory policy optimal. The model extension proposed in Section 5.3.6, may serve future research as a starting point for further investigations on this matter.

Appendices

5.A Exponentially Distributed Resupply Rate

In order to gain insights on how the replenishment variability influences the results, we conducted a simulation study. Figure 5.11 and Figure 5.12 show the effect of an increasing replenishment variability on the cost for two different instances. As we observe, the effect is limited and thus is in line with the findings in the literature, cf. Section 5.3.1.

The parameter settings for both cases are in the same range as for the numerical experiments discussed in Section 5.4. In particular, for the first instance (Figure 5.11) we use following parameter settings: $k=20$ installed parts, $b=2000$ euro/month, $\lambda_A=0.035$ failures/month, $\mu_C=0.15$ supplies/month, $c_C=10$ euro/part, $m=10$ euro/failure, $\kappa=0.25$, $c_A=20$ euro/part, $\lambda_C=0.01$ failures/month and $\mu_A=4$ supplies/month. For the second instance (Figure 5.12) we use: $k=30$ installed parts, $b=200$ euro/month, $\lambda_A=0.035$ failures/month, $\mu_C=0.15$ supplies/month, $c_C=10$ euro/part, $m=2$ euro/failure, $\kappa=0.25$, $c_A=20$ euro/part, $\lambda_C=0.01$ failures/month and $\mu_A=4$ supplies/month.




Figure 5.11: Simulation results for different variability of the replenishment lead time




Figure 5.12: Simulation results for different variability of the replenishment lead time

5.B Balance Equations Basic Model

Subsequently, we present the balance equations in which p_{\bullet} refers to the state probability under consideration. For example, in case $x_2(n_{C,i}, n_{A,i} + 1, r_{C,i}, r_{A,i} - 1, s_{C,i}, s_{A,i})$ then $p_{\bullet} = p(n_{C,i}, n_{A,i} + 1, r_{C,i}, r_{A,i} - 1, s_{C,i}, s_{A,i})$. Furthermore, if the maintenance decision is indifferent (i.e. no stock available) we use $x_2(\mathbf{i})$ and $x_4(\mathbf{i})$ as default.

$$\begin{aligned}
& (\lambda_C n_{C,i} + \lambda_A n_{A,i} + \mu_C r_{C,i} + \mu_A r_{A,i}) p_{\mathbf{i}} = \\
& \lambda_A (n_{A,i} + 1) x_2(n_{C,i}, n_{A,i} + 1, r_{C,i}, r_{A,i} - 1, s_{C,i}, s_{A,i}) p_{\bullet} + \\
& \lambda_A n_{A,i} x_1(n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i} - 1, s_{C,i}, s_{A,i} + 1) p_{\bullet} + \\
& \lambda_A (n_{A,i} + 1) x_2(n_{C,i} - 1, n_{A,i} + 1, r_{C,i}, r_{A,i} - 1, s_{C,i} + 1, s_{A,i}) p_{\bullet} + \\
& \lambda_A (n_{A,i} + 1) x_4(n_{C,i}, n_{A,i} + 1, r_{C,i} - 1, r_{A,i}, s_{C,i}, s_{A,i}) p_{\bullet} + \\
& \lambda_A n_{A,i} x_3(n_{C,i}, n_{A,i}, r_{C,i} - 1, r_{A,i}, s_{C,i}, s_{A,i} + 1) p_{\bullet} + \\
& \lambda_A (n_{A,i} + 1) x_4(n_{C,i} - 1, n_{A,i} + 1, r_{C,i} - 1, r_{A,i}, s_{C,i} + 1, s_{A,i}) p_{\bullet} + \\
& \lambda_C (n_{C,i} + 1) x_2(n_{C,i} + 1, n_{A,i}, r_{C,i}, r_{A,i} - 1, s_{C,i}, s_{A,i}) p_{\bullet} + \\
& \lambda_C (n_{C,i} + 1) x_1(n_{C,i} + 1, n_{A,i} - 1, r_{C,i}, r_{A,i} - 1, s_{C,i}, s_{A,i} + 1) p_{\bullet} + \\
& \lambda_C n_{C,i} x_2(n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i} - 1, s_{C,i} + 1, s_{A,i}) p_{\bullet} + \\
& \lambda_C (n_{C,i} + 1) x_4(n_{C,i} + 1, n_{A,i}, r_{C,i} - 1, r_{A,i}, s_{C,i}, s_{A,i}) p_{\bullet} + \\
& \lambda_C (n_{C,i} + 1) x_3(n_{C,i} + 1, n_{A,i} - 1, r_{C,i} - 1, r_{A,i}, s_{C,i}, s_{A,i} + 1) p_{\bullet} + \\
& \lambda_C n_{C,i} x_4(n_{C,i}, n_{A,i}, r_{C,i} - 1, r_{A,i}, s_{C,i} + 1, s_{A,i}) p_{\bullet} + \\
& \mu_A (r_{A,i} + 1) p(n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i} + 1, s_{C,i}, s_{A,i} - 1) + \\
& \mu_A (r_{A,i} + 1) p(n_{C,i} - 1, n_{A,i}, r_{C,i}, r_{A,i} + 1, s_{C,i}, s_{A,i}) + \\
& \mu_C (r_{C,i} + 1) p(n_{C,i}, n_{A,i}, r_{C,i} + 1, r_{A,i}, s_{C,i} - 1, s_{A,i}) + \\
& \mu_C (r_{C,i} + 1) p(n_{C,i} - 1, n_{A,i}, r_{C,i} + 1, r_{A,i}, s_{C,i}, s_{A,i})
\end{aligned}$$

5.C Joint Optimization

To jointly optimize the base stock level, maintenance policy and sourcing policy, we increase the decision space of the sourcing decision and add the possibility to order nothing upon failure of an item. This extension enables the transition to states where N decreases. Accordingly, we obtain the following updated decision space:

$$\begin{aligned}
\Omega = \{ & (n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i}, s_{C,i}, s_{A,i}) : \\
& LB \geq N \leq UB \\
& n_{C,i} + n_{A,i} + r_{C,i} + r_{A,i} + s_{C,i} + s_{A,i} = N \\
& s_{C,i} + s_{A,i} = \max\{N - k - r_{C,i} - r_{A,i}; 0\} \\
& n_{C,i} + n_{A,i} \leq k \\
& n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i}, s_{C,i}, s_{A,i} \geq 0\}
\end{aligned}$$

where LB and UB describe the lower and upper bound on N respectively. We set $LB=k$ which permits the extreme case to manufacture parts on demand only. In case of the UB , we compute the optimal base stock level \hat{S} of a single sourcing model with $\lambda = \max\{\lambda_C; \lambda_A\}$, $\mu = \min\{\mu_C; \mu_A\}$ and $c = \min\{c_C; c_A\}$. Next, we set $UB = k + \hat{S}$. As this is a worst case scenario, we always find an upper bound, even though it may not be very tight. An alternative is to use as a heuristic upper bound which is equal to the base stock level of the better performing single sourcing option. Unfortunately, our numerical experiments reveal that this bound is not always sufficient. In this case we set $S = UB - k$ and follow the greedy heuristic as explained in Section 5.3.5.

To include the decision option to order nothing we increase the action space of the sourcing decision. Accordingly, we add decision options:

- $c=5$: take AM version from stock (if possible), and order nothing.
- $c=6$: take CM version from stock (if possible), and order nothing.

The increase of the action space leads to six additional transitions. These are illustrated for state $(n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i}, s_{C,i}, s_{A,i})$ in Figure 5.13.

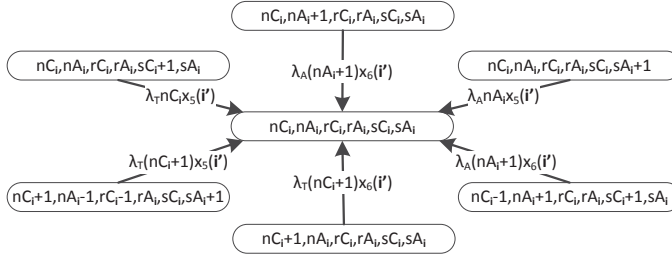


Figure 5.13: Additional transitions for state $(n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i}, s_{C,i}, s_{A,i})$ for joint optimization

5.D Dynamic Inventory Policy

Given that the part is available with two versions of different reliability, the expected number of failures is state dependent. Same holds for the expected number of arrivals given the difference in resupply rate. As a consequence, the stocking policy might depend on the state. For example, consider the situation where the installed base is mainly equipped with AM items. In this case, a higher failure frequency is likely, and thus we may want to increase the base stock level to avoid backorder costs. Then, if the number of CM components increases, we may find that it is cost efficient to reduce the base stock level to decrease holding costs.

Given that we already added the possibility to reduce N (cf. Appendix 5.C), we only need to include transitions that allow increasing N to facilitate a dynamic base stock level. Again, we realize this extension by increasing the action space to allow ordering more than

one part. Given that typically $\lambda_C < \lambda_A$, the possibility to order more than one part is most valuable if a CM part fails. Also, it is not reasonable to order more than 2 parts because this would imply that it would have been useful to already order 2 parts at a previous failure of a CM part. Accordingly, we add six decision options:

- $c=7$: take CM part from stock (if possible), and order AM and CM part.
- $c=8$: take CM part from stock (if possible), and order 2 AM parts.
- $c=9$: take AM part from stock (if possible), and order AM and CM part.
- $c=10$: take AM part from stock (if possible), and order 2 AM parts.
- $c=11$: take CM part from stock (if possible), and order 2 CM parts.
- $c=12$: take AM part from stock (if possible), and order 2 CM parts.

We can incorporate these actions by regarding nine additional transitions. These are illustrated for state $(n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i}, s_{C,i}, s_{A,i})$ in Figure 5.14.




Figure 5.14: Additional transitions for state $(n_{C,i}, n_{A,i}, r_{C,i}, r_{A,i}, s_{C,i}, s_{A,i})$ for dynamic inventory policy

Chapter 6

Large-Scale Dual Sourcing Problems with AM Supply

In the previous chapter, we analyzed the effects of dual sourcing of CM and AM parts. In contrast with existing dual sourcing studies, the item failure behavior may depend on the sourcing mode in our study. While it was possible to derive insights for smaller problem instances, the proposed approach becomes computationally burdensome for larger problems. Similar difficulties arise if we are considering other problem settings, such as different order policies or demand patterns.

In this chapter¹, we suggest and evaluate a procedure that facilitates the analysis of large problem instances. Furthermore, we discuss the option to model dual sourcing problems with AM as an Approximate Dynamic Programming (ADP) problem. This chapter concerns the fifth contribution as mentioned in Section 1.4.3.

6.1 Introduction

As shown by Whittemore and Saunders (1977), the optimal dual sourcing policy generally depends on delivery time and quantity of in-transit parts. The exact analysis quickly becomes intractable. As we reviewed in Section 5.2, many heuristics were proposed to overcome this problem. A common technique is the aggregation of state information to approximate ordering policies. For instance, various heuristics base order decisions on the inventory position (stock level plus in-transit items) rather than the delivery time and size of each order.

However, for dual sourcing problems with supply mode dependent failure behavior, aggregation of state information becomes more complicated. In our case, we have to

¹This chapter is based on the working paper Knofius et al. (2018a).

distinguish between CM and AM items that are installed, in stock, and in resupply. Otherwise, it would not be possible to account for the impact of the order decision on future demand. For instance, if we decide to order (and later to install) an AM item that may be characterized by a lower quality than the CM item installed thus far, the expected demand increases.

In this chapter, we propose an extension for dual sourcing methods to account for this effect, namely an iterative approach. The proposed extension is particularly appealing since it can be applied in combination with dual sourcing heuristics from the literature. To that end, we may use this approach for the analysis of a wide range of problem settings. In this chapter, we demonstrate its application and performance under two problem settings, i.e., the setting studied in Chapter 5 and a new one to be introduced in this chapter. Furthermore, we examine the option to formulate dual sourcing problems with AM as an ADP. As it turns out that the specific problem structure renders typical ADP designs to be unsuitable, we analyze the observed shortcoming and suggest future research directions to overcome them.

The chapter is structured as follows. In Section 6.2, we introduce the problem and point out the differences compared to the problem setting studied in Chapter 5. Next, in Section 6.3, we formulate the problem as an MDP model. Its exact solution serves as benchmark to evaluate the iterative procedure which we propose in Section 6.4 and analyze in Section 6.5. In Section 6.6, we discuss the ADP formulation, after which we conclude the chapter with Section 6.7.

6.2 Problem Formulation and Notation

Similar to the previous chapter, we consider a service provider that may order spare parts for a specific component using either a CM or an AM supply channel. Both supply modes are characterized by different part quality, replenishment lead time and order costs. Unlike the model in the previous chapter, holding costs are item type independent, which is a common model choice in the dual sourcing literature, see e.g., Veeraraghavan and Scheller-Wolf (2008) or Axsäter (2014). Furthermore, this time the CM supply mode imposes fixed order costs and batch ordering. To this end, the service provider has to decide how many *batches* to order from the CM supplier, and how many *items* from the AM supplier. In this chapter, we do not address the maintenance decision: *If both item types are in stock which type do we use first?* Instead, we assume that parts with higher reliability are consumed first. We use this simplification because there is no longer a direct monetary incentive to install either item version first which is caused by the identical holding costs. Without loss of generality, we assume that CM items exhibit a higher reliability throughout this chapter.

We define a discrete-time, periodic-review model in which we assume the following order of events. At the beginning of each period, we monitor the inventory positions, possibly place orders, and account for the costs which consist of purchasing costs (fixed and variable), holding costs, backorder costs and expected maintenance costs. During each period, items fail and are replaced by functioning components without delay if sufficient stock is available. Otherwise, we backorder the demand. At the end of each period, we may receive replenishment orders

which are put in stock or are used for backorder clearing without further delay. In addition, we make the following assumptions to obtain a tractable model with an exact solution procedure.

1. *Lead times are deterministic.*

In case of stochastic lead times, we would need to consider order-crossing effects. The study of these effects is not the focus in this chapter. Furthermore, we again assume external suppliers for which deterministic resupply lead times are usually part of contractual agreements. In case a supplier is unable to meet these agreements, delays are typically compensated which further supports this assumption.

2. *Lead times are a multiple of the review period.*

Since we chose the review period small compared to the lead times, this assumption will effect the results only marginally. Besides, this assumption simplifies the analysis because orders may only arrive at the end of a period.

3. *We allow for at most one outstanding order from each supply mode (AM and CM).*

In the dual sourcing literature, cf. Section 5.2, we often observe less restrictive assumptions on the number of outstanding orders. However, given that we need to differentiate between CM and AM items, common methods such as deriving order decisions from the inventory position or other aggregated measures, do not apply for our problem setting. Instead, we require two state variables for each outstanding order to track the relevant information, cf. Section 6.3.1. The assumption of at most one outstanding order from each supply mode is also justified by the observation that fixed order costs usually impose a sufficiently large CM batch size and hence the probability of a second order placed before the first one has arrived is negligible. For AM parts, the lead time is generally short, and therefore again the probability of a new order released before the first one has arrived, is very small.

4. *Demand per item is assumed to follow a logarithmic compound Poisson process and is i.i.d. throughout a period, conditional on the number of operating CM and AM parts at the beginning of the period.* The latter choice constitutes a model assumption, since failures and replacements may change the proportion of operating CM and AM items during a period. Yet, in case we chose a small review period length, the effect should be negligible. This assumption, however, allows for demand patterns having higher variability than Poisson demand and is thus more flexible.

We use τ to describe the review period length and t to indicate the period. The installed base size is equal to N , and we denote the number of CM and AM failures in a period by k_C and k_A . To describe the failure behavior of one CM (AM) part, we use the mean μ_C and variance V_C (μ_A and V_A) of failures per period, respectively. Orders for the CM supply channel are placed in integer multiples of a given batch size Q_C (input). Hence for each CM batch, we have variable CM order costs $c_C Q_C$. Furthermore, we use K_C to denote the fixed CM order costs (independent of the number of batches in the same order). AM supply is more flexible because of the low setup costs for AM technology. Accordingly, we may order any quantity with a piece price c_A . Fixed order costs do not apply for this supply mode. The CM replenishment lead time is denoted by l_C and the commonly shorter AM replenishment lead time by l_A .

For each failure, the service provider encounters maintenance costs m , which covers spare parts transportation costs, engineer time and materials for diagnosis and repair. Unfulfilled demand is backlogged and causes backorder cost at a rate b per period. Furthermore, we use a holding cost rate h per period independent of the item version. In Table 6.1, we summarize the notation that is used throughout the chapter.

6.3 Markov Decision Problem Formulation

In this section, we formulate the described model as a Markov Decision Problem (MDP). The section is organized as follows. In Section 6.3.1, we describe the state space. Next, in Section 6.3.2, we define the decision space and in Section 6.3.3, we discuss possible transitions and associated transition probabilities. We close with Section 6.3.4 in which we describe the evaluation of the expected costs, given a certain state and decision.

6.3.1 State Space

The state is described by the tuple $\mathbf{i}_t = (n_{C,t}, n_{A,t}, s_{C,t}, s_{A,t}, u_{C,t}, u_{A,t}, w_{C,t}, w_{A,t})$. The number of operating CM and AM items at the beginning of period t , are denoted by $n_{C,t}$ and $n_{A,t}$ respectively. Additionally, we use $s_{C,t}$ and $s_{A,t}$ to describe the CM and AM on-hand stock at the beginning of period t . The state variables $u_{C,t}$ and $u_{A,t}$ denote the number of review periods before $w_{C,t}$ CM batches and $w_{A,t}$ AM parts arrive. Hence, we may express the number of backorders at the start of period t by Equation 6.1 and the inventory position at the start of period t by Equation 6.2.

$$B_t = N - n_{C,t} - n_{A,t} \quad (6.1)$$

$$IP_t = s_{C,t} + s_{A,t} + w_{C,t}Q_C + w_{A,t} - B_t \quad (6.2)$$

We obtain the following definition of the state space in period t :

$$\mathcal{I} = \{(n_{C,t}, n_{A,t}, s_{C,t}, s_{A,t}, u_{C,t}, u_{A,t}, w_{C,t}, w_{A,t}) : \quad (6.3a)$$

$$0 \leq n_{C,t} + n_{A,t} \leq N \quad (6.3a)$$

$$IP_t \leq S \quad (6.3b)$$

$$B_t > 0 \implies s_{C,t} + s_{A,t} = 0 \quad (6.3c)$$

$$0 \leq u_{C,t} \leq l_C - \tau \quad (6.3d)$$

$$0 \leq u_{A,t} \leq l_A - \tau \quad (6.3e)$$

$$w_{C,t} = 0 \Leftrightarrow u_{C,t} = 0 \quad (6.3f)$$

$$w_{C,t} > 0 \Leftrightarrow u_{C,t} > 0 \quad (6.3g)$$

$$w_{A,t} = 0 \Leftrightarrow u_{A,t} = 0 \quad (6.3h)$$

$$w_{A,t} > 0 \Leftrightarrow u_{A,t} > 0 \quad (6.3i)$$

Constraint 6.3a ensures that the number of installed items cannot become negative nor exceed the installed base size. Constraint 6.3b limits the number of items in the inventory

Table 6.1: Notation overview

Input variables	
τ	Length of the review period
μ_C	Mean number of failures of one CM item per period
μ_A	Mean number of failures of one AM item per period
V_C	Variance of the number of failures of one CM item per period
V_A	Variance of the number of failures of one AM item per period
l_C	Deterministic CM replenishment lead time
l_A	Deterministic AM replenishment lead time
c_C	CM piece price
c_A	AM piece price
K_C	Fixed CM order costs
Q_C	Fixed CM batch size
m	Maintenance costs, i.e., costs to replace a failed part
h	Holding cost per item in stock at the beginning of the period
b	Backorder cost per item backlogged at the beginning of a period
N	Installed base size
S	Maximum number of spare parts circulating in the inventory system
$p(\mathbf{i}_{t+1} \mathbf{i}_t, \mathbf{a}(\mathbf{i}_t))$	Probability to transition to state $\mathbf{i}_{t+1} \in \mathcal{I}_{t+1}$ given $\mathbf{i}_t \in \mathcal{I}_t$ and $\mathbf{a}(\mathbf{i}_t) \in \mathcal{A}$
Auxiliary variables	
$\mathbf{i}_t \in \mathcal{I}_t$	System state tuple, $\mathbf{i}_t = (n_{C,t}, n_{A,t}, s_{C,t}, s_{A,t}, u_{C,t}, u_{A,t}, w_{C,t}, w_{A,t})$
$\mathbf{a}(\mathbf{i}_t) \in \mathcal{A}$	Action tuple, $\mathbf{a}(\mathbf{i}_t) = (x_C, x_A)$
$p_C(k_C n_{C,t})$	Probability of k_C failures of CM items in a review period, given $n_{C,t}$
$p_A(k_A n_{A,t})$	Probability of k_A failures of AM items in a review period, given $n_{A,t}$
B_t	Number of backorders at the beginning of period t
IP_t	Inventory position at the beginning of period t
State variables	
$n_{C,t}$	Number of operating CM parts at the beginning of period t
$n_{A,t}$	Number of operating AM parts at the beginning of period t
$s_{C,t}$	CM on-hand stock at the beginning of period t
$s_{A,t}$	AM on-hand stock at the beginning of period t
$u_{C,t}$	Remaining periods at the beginning of period t , until CM order arrival
$u_{A,t}$	Remaining periods at the beginning of period t , until AM order arrival
$w_{C,t}$	Number of CM batches Q_C in a running CM order at the beginning of period t
$w_{A,t}$	Number of AM parts in a running AM order at the beginning of period t
Decision variables	
x_C	Number of CM batches to order
x_A	Number of AM parts to order
Output variables	
π	Stationary policy which prescribes $\mathbf{a}(\mathbf{i}_t) \forall \mathbf{i}_t \in \mathcal{I}_t$
$g(\pi)$	Long-run average service cost if policy π is applied
$v(\mathbf{i}_t, \pi)$	Relative value if we are in state \mathbf{i}_t and follow policy π

system. That is, the inventory position IP_t at any period t may not exceed some predefined value S . This constraint is necessary to obtain a finite state space. We determine S numerically by increasing S until IP_t never exceeds $S - 1$. Constraint 6.3c implies that it is not possible to have on-hand stock if we observe a backlog at the beginning of period t . Constraint 6.3d and Constraint 6.3e restrict the number of review periods a CM and an AM order may be outstanding. The remaining constraints establish that if there is an outstanding order, the remaining lead time has to be positive. Otherwise, both the remaining lead time and order size are equal to zero.

6.3.2 Decision Space

The decision is defined by the tuple $\mathbf{a}(\mathbf{i}_t) = (x_C, x_A)$ where x_C describes the number CM batches and x_A the number of AM parts that are ordered at the beginning of period t . Using state information and the relationship between both decision variables, we limit the action space to

$$\mathcal{A} = \{(x_C, x_A) : \begin{aligned} &u_{C,t} > 0 \implies x_C = 0 \end{aligned} \tag{6.4a}$$

$$\begin{aligned} &u_{A,t} > 0 \implies x_A = 0 \end{aligned} \tag{6.4b}$$

$$0 \leq x_C Q_C + x_A \leq S - IP_t \} \tag{6.4c}$$

where Constraint 6.4a and Constraint 6.4b establish that there is at most one outstanding CM and AM order respectively (cf. Assumption 3 in Section 6.2). Furthermore, Constraint 6.4c ensures that the total order quantity cannot be negative nor raise the inventory position above the maximum number of spare parts in the system S .

6.3.3 Transitions

We first describe the transition function for each state variable and then how to compute the transition probability. Recall that CM items are consumed first if we have both item versions in stock and that the number of CM and AM failures in a period is denoted by k_C and k_A , respectively. The installed base composition at the beginning of period $t+1$ is obtained as follows:

$$n_{C,t+1} = n_{C,t} - k_C + y_C + z_C \tag{6.5}$$

$$n_{A,t+1} = n_{A,t} - k_A + y_A + z_A \tag{6.6}$$

where the auxiliary variable y_C (y_A) denotes the maximum number of CM (AM) items that may be installed *before* a possible order arrival at the end of period t , and z_C (z_A) the maximum number of CM (AM) items that may be installed *after* a possible order arrival. We need this distinction, since we do not assume a random order in which items are used. Accordingly, we have:

- $y_C = \min\{B_t + k_C + k_A, s_{C,t}\}$ where the first element describes the total demand (backorders + new demand arrivals in period t), and the second element the CM on-hand stock.
- $y_A = \min\{B_t + k_C + k_A - y_C, s_{A,t}\}$ where the first element describes the total demand not satisfied by CM before a possible order arrival, and the second element the AM on-hand stock.
- $z_C = \min\{B_t + k_C + k_A - y_C - y_A, A_{C,t}\}$ where the first element describes the backlog just before a possible order arrives, and the second element is equal to the number of CM units that may arrive at $t + \tau$, i.e., $A_{C,t} = w_{C,t}Q_C$ if $u_{C,t} = 1$ or $A_{C,t} = x_CQ_C$ if $l_C = \tau$.
- $z_A = \min\{B_t + k_C + k_A - y_C - y_A - z_C, A_{A,t}\}$ where the first element describes the remaining backlog after we accounted for the CM order arrival, and the second element is equal to the size of a potential AM order arrival, i.e., $A_{A,t} = w_{A,t}$ if $u_{A,t} = 1$ or $A_{A,t} = x_A$ if $l_A = \tau$.

For the on-hand stock in period $t + 1$, we have:

$$s_{C,t+1} = s_{C,t} + A_{C,t} - y_C - z_C \quad (6.7)$$

$$s_{A,t+1} = s_{A,t} + A_{A,t} - y_A - z_A \quad (6.8)$$

The timing of the CM replenishment orders is updated as follows:

$$u_{C,t+1} = \begin{cases} u_{C,t} - 1 & , \text{ if } u_{C,t} > 0 \\ (l_C - \tau) / \tau & , \text{ if } u_{C,t} = 0 \wedge x_C > 0 \\ 0 & , \text{ if } u_{C,t} = 0 \wedge x_C = 0 \end{cases} \quad (6.9)$$

The number of CM batches in the replenishment pipeline at the beginning of period $t + 1$ is equal to:

$$w_{C,t+1} = \begin{cases} w_{C,t} & , \text{ if } u_{C,t} > 1 \\ x_C & , \text{ if } u_{C,t} = 0 \wedge l_C > \tau \\ 0 & , \text{ if } u_{C,t} = 1 \vee l_C = \tau \end{cases} \quad (6.10)$$

The timing and size of AM orders is updated analogously to Equation 6.9 and Equation 6.10.

Next, we discuss the derivation of the transition probability $p(\mathbf{i}_{t+1} | \mathbf{i}_t, \mathbf{a}(\mathbf{i}_t))$ that we are in state \mathbf{i}_t and observe a transition to \mathbf{i}_{t+1} provided that we have chosen action $\mathbf{a}(\mathbf{i}_t)$. Since the failures are the only random events, we need to determine the number of CM item failures k_C and AM item failures k_A in a review period. Using Assumption 4 defined in Section 6.2 and the state information, the probability mass functions $p_C(k_C | n_{C,t})$ and $p_A(k_A | n_{A,t})$ are characterized by $E[k_x] = \mu_x n_{x,t}$ and $V[k_x] = V_x n_{x,t}$ for production method $x \in \{C, A\}$. Note that Assumption 4 may cause that $k_C > n_{C,t}$ and $k_A > n_{A,t}$ since we do not keep track of the installed base development during a review period. We modify this by assigning the associated probability masses to the case where $k_C = n_{C,t}$ and $k_A = n_{A,t}$, thereby limiting the number of failures to $n_{C,t}$ and $n_{A,t}$ respectively. This appears to be a reasonable simplification given that the probability that the number of parts that fails within one review period exceeds the size of the entire installed base is highly unlikely for realistic problem instances. As a result, the transition probabilities are defined as follows:

$$p(\mathbf{i}_{t+1} | \mathbf{i}_t, \mathbf{a}(\mathbf{i}_t)) = p_C(k_C | n_{C,t}) p_A(k_A | n_{A,t}) \quad (6.11)$$

6.3.4 Cost Function

We aim to minimize the average service costs per period, $C(\mathbf{i}_t, \mathbf{a}(\mathbf{i}_t))$, consisting of purchasing costs $P(\mathbf{a}(\mathbf{i}_t))$ (fixed and variable), holding costs $H(\mathbf{i}_t)$, backorder costs $BO(\mathbf{i}_t)$ and maintenance costs $M(\mathbf{i}_t)$ over an infinite horizon. The purchasing costs are equal to

$$P(\mathbf{a}(\mathbf{i}_t)) = \mathbf{1}_{x_C > 0} K + c_C x_C Q_C + c_A x_A \quad (6.12)$$

where $\mathbf{1}_{x_C > 0}$ is the indicator function, which is equal to 1 if $x_C > 0$ and 0 otherwise. The holding costs and the expected backorder costs are equal to

$$H(\mathbf{i}_t) = h(s_{C,t} + s_{A,t}) \quad (6.13)$$

$$BO(\mathbf{i}_t) = \max\{(\mu_C n_{C,t} + \mu_A n_{A,t}) + B_t - s_{C,t} - s_{A,t}, 0\} b \quad (6.14)$$

We encounter expected maintenance cost for each failure, thus we have

$$M(\mathbf{i}_t) = m(\mu_C n_{C,t} + \mu_A n_{A,t}) \quad (6.15)$$

If the policy π is optimal, the optimal long-run average service costs $g(\pi)$ and the relative value function $v(\mathbf{i}_t)$ satisfy the Bellman equation:

$$v(\mathbf{i}_t) = \min_{\mathbf{a}(\mathbf{i}_t) \in \mathcal{A}} \{C(\mathbf{i}_t, \mathbf{a}(\mathbf{i}_t)) - g(\pi) + \sum_{\mathbf{i}_{t+1} \in \mathcal{I}_{t+1}} p(\mathbf{i}_{t+1} | \mathbf{i}_t, \mathbf{a}(\mathbf{i}_t)) v(\mathbf{i}_{t+1})\}, \forall \mathbf{i}_t \in \mathcal{I} \quad (6.16)$$

In Appendix 6.A, we outline an exact solution procedure to find the optimal policy.

6.4 Iterative Procedure

Solving MDP problems often becomes computationally intractable, because of a rapid growth in (i) the state space, (ii) the action space, and (iii) the number of stochastic events that have to be evaluated. Powell (2011) refers to these problems as the three curses of dimensionality.

Also for the problem considered in this chapter, we deal with the curses of dimensionality. In particular, the size of the state space renders an exact analysis impossible for any realistic problem size, due to its problem specific high dimension. To give an example, already for the simple model proposed in Section 6.3, we have to distinguish between 350,000 states for a medium sized problem instance ($N=50$, $l_C=10\tau$, $l_A=2\tau$ and $S=10$). Crucial for the fast state space growth is that we have to differentiate between CM and AM items in order to respect the effect of a different failure behavior.

In this section, we elaborate how an iterative procedure may be used to separate this effect from the solution procedure which significantly reduces the state space growth. As our discussions will highlight, the iterative procedure may be used in combination with most dual sourcing solution procedures and therefore allows the extension of various powerful dual sourcing heuristics (cf. Section 5.2).

The key idea is the following. If we assume a certain average ratio of CM and AM items in operation, we can compute the expectation and the variance of the number of failures per period of an arbitrary item in the installed base. Using these results as input for a common dual sourcing solution approach, then allows the computation of an improved estimate for the average ratio of CM and AM items in operation. We continue with this procedure until the average ratio of CM and AM items in operation converges. Subsequently, we discuss each step in greater detail.

Suppose, we have an estimate for the average ratio γ_j of CM and AM items in operation in iteration j . We then can compute the expectation $E[X]$ and the variance $Var[X]$ of a random variable X that measures the number of failures per period of an arbitrary item in the installed base. We define the indicator variable $\mathbf{1}_A$ that is equal to 1 if the arbitrary item is an AM item and 0 otherwise. Using this variable, we derive the mean and variance of X using the standard conditioning rules $E[X] = E[E[X|\mathbf{1}_A]]$ and $Var[X] = E[Var[X|\mathbf{1}_A]] + Var[E[X|\mathbf{1}_A]]$, and obtain:

$$E[X] = \gamma_j \mu_A + (1 - \gamma_j) \mu_C \quad (6.17)$$

$$Var[X] = \gamma_j V_A + (1 - \gamma_j) V_C + \gamma_j (1 - \gamma_j) (\mu_A - \mu_C)^2 \quad (6.18)$$

Next, to find the new estimate of the average ratio of CM and AM items in operation γ_{j+1} for iteration $j+1$, we measure the fraction ρ_j of sourced AM items when using a specified dual sourcing solution procedure using the previously computed $E[X]$ and $Var[X]$. Note that ρ_j has to be equal to the fraction of AM items entering or leaving the installed base, as shown in Equation 6.19. Finally, we solve Equation 6.19 for γ_{j+1} and obtain Equation 6.20.

$$\rho_j = \frac{\gamma_{j+1} \mu_A}{\gamma_{j+1} \mu_A + (1 - \gamma_{j+1}) \mu_C} \quad (6.19)$$

$$\gamma_{j+1} = \frac{\rho_j \mu_C}{(1 - \rho_j) \mu_A + \rho_j \mu_C} \quad (6.20)$$

The procedure continuous, until the difference Δ between two subsequent observations of γ is less than α . We summarize the iterative procedure in Algorithm 6.1.

Algorithm 6.1: Iterative procedure

Set $j := 1$, $\gamma_1 := 0$ and $\Delta := \alpha$;

repeat

Compute $E[X]$ and $Var[X]$ with Equation 6.17 and 6.18;
 Measure ρ_j and derive π_j based on $DualSourcingMethod(E[X], Var[X])$;
 Compute γ_{j+1} using ρ_j and Equation 6.20;
 $j = j + 1$;

until $|\gamma_{j-1} - \gamma_j| < \Delta$;

return Heuristic policy $\pi_h = \pi_j$ and long-run average cost C_{i_j, a_j}

6.5 Performance Analysis

As we stressed, the iterative procedure may be applied in combination with most dual sourcing methods since it only relies on an estimate of a fraction of items ordered from either supply mode. However, in this section we apply the iterative approach in combination with exact methods that do *not* distinguish between the supply mode specific failure behavior. By using exact methods, we avoid approximation errors unrelated to the iterative procedure, and therefore obtain a more precise measure of its performance.

The section is structured as follows. In Section 6.5.1, we describe the experimental design used to assess the performance of the iterative procedure. Next, in Section 6.5.2 we benchmark its performance with the exact solution. Finally, in Section 6.5.3 we will revisit the hinge bracket case we discussed in Chapter 5 (cf. Section 5.5) to demonstrate the application of the iterative procedure with both a large size problem instance and a different dual sourcing model.

6.5.1 Experimental Design

We limit our experiments to the case in which the installed base size $N=7$ items and the period length τ equals one week. This restriction is necessary to obtain an exact benchmark. However, in Section 5.5, we also study the performance for a larger problem instance. Next, we motivate the chosen parameter settings in the test bed.

We study two different scenarios regarding the CM item version. In the first scenario, the CM item is expensive, fails rarely and has a long replenishment lead time. More specifically, we assume that $c_C=5000$ euro with fixed order cost $K_C=2000$ euro, $\mu_C=0.01$ failures per week (i.e., an item has to be replaced approximately every 2 years on average) with a variance of $V_C=2\mu_C$. Furthermore, we assume a lead time of $l_C=8$ weeks and a minimum order quantity of 5 parts, i.e. $Q_C=5$. In the second scenario, the CM item is cheaper, fails more often, but has a shorter replenishment lead time. Accordingly, we assume that $c_C=1000$ euro with fixed order cost $K_C=750$ euro, $\mu_C=0.025$ failures per week, $V_C=2\mu_C$ and a replenishment lead time of $l_C=4$ weeks. For this scenario, we set $Q_C=7$.

To determine the parameters of the corresponding AM item, we rely on common characteristics of AM items which we discussed in Section 5.4.1. Again, we consider two scenarios. In the first scenario, the AM item is substantially more expensive, i.e., $c_A=2c_C$, exhibits a substantially lower reliability with $\mu_A=2\mu_C$ and $V_A=2\mu_A$, but has a shorter lead time $4l_A=l_C$. The second scenario is more moderate with $c_A=1.5c_C$, $\mu_A=1.5\mu_C$, $V_A=3\mu_A$ and $2l_A=l_C$.

For the remaining parameters, we also distinguish between two cases. In the first case, the maintenance cost and holding cost are high, i.e., $m=c_C$ euro and $h=7(0.3c_C/365)$ euro/week. Note that we expressed the maintenance costs in terms of c_C only for orientation. So, we do not intend to suggest that maintenance costs depend on the CM unit costs in real-life. Furthermore, we aim at a high fill rate (99.5% of demand has to be satisfied from stock). Here, we approximate the fill rate with $b/(b+h)$. Hence, if we consider the expensive CM item, we have $b=5724.66$ euro/week. For the second case, we reduce the

maintenance and holding costs, i.e., $m=0.25c_C$ euro and $h=7(0.2c_C/365)$ euro/week. Also, we aim at a lower fill rate (99% of demand can be satisfied from stock).

The experimental setup leads to eight instances. However, to also regard more extreme cases, we added (although maybe less realistic) Instance 9 and 10. We motivate this choice as for realistic problem settings, we typically find 100% AM usage or a small fraction of AM usage. Instance 9, on the other hand, has been designed such that it leads to a relatively large usage of AM as emergency source. Instance 10 was selected to show the performance of the iterative procedure in case of substantial failure rate differences between CM and AM parts. The entire set of instances is summarized in Table 6.2.

Table 6.2: *Input parameters for performance analyses*

Parameter	Instance									
	1	2	3	4	5	6	7	8	9	10
c_C	5000	5000	5000	5000	1000	1000	1000	1000	2000	2000
K_C	2000	2000	2000	2000	750	750	750	750	2000	1000
μ_C	0.01	0.01	0.01	0.01	0.025	0.025	0.025	0.025	0.01	0.025
Var_C	0.02	0.02	0.02	0.02	0.05	0.05	0.05	0.05	0.02	0.05
l_C	8	8	8	8	4	4	4	4	10	6
Q_C	5	5	5	5	7	7	7	7	4	7
c_A	10000	10000	7500	7500	2000	2000	1500	1500	3000	2400
μ_A	0.02	0.02	0.015	0.015	0.05	0.05	0.0375	0.0375	0.015	0.125
Var_A	0.04	0.04	0.045	0.045	0.1	0.1	0.1125	0.1125	0.03	0.375
l_A	2	2	4	4	1	1	2	2	1	1
m	5000	1250	5000	1250	1000	250	1000	250	500	2000
h	29	19	29	19	6	4	6	4	12	12
b	5725	1899	5725	1899	1145	380	1145	380	2290	2290

6.5.2 Results

In this section, we apply the iterative procedure in combination with a *simplified version* of the MDP model described in Section 6.3. As outlined in Appendix 6.B, the simplified MDP model no longer distinguishes between the supply mode specific failure behavior since this aspect is taken care of by the iterative procedure. Next, we examine how well the iterative procedure performs compared to the exact solution obtained with the *regular version* of the MDP model described in Section 6.3.

Table 6.3 shows the results for each problem instance specified in Table 6.2. We observe that the iterative procedure performs well across all instances and, as shown in the last column, converges after only a view iterations. However, Instance 10 is an exception. Here we observed cycling between two nearly identical solutions (cost difference of about 0.6%). In particular, the solution procedure iterated between a solution with an inventory position of up to 11 items (i.e. $S=11$) and a small AM usage ($\sim 1\%$), and a solution with an inventory position of up to 10 items (i.e. $S=10$) but a higher AM usage ($\sim 3\%$).

Table 6.3: Comparison of exact method with iterative procedure

Instance	Exact method			Iterative procedure				
	g^*	S^*	γ^*	g	Δg	S	γ	j
1	914	8	1%	930	1.75%	7	3%	4
2	579	7	1%	580	0.17%	7	1%	3
3	917	8	1%	922	0.55%	8	2%	2
4	576	7	2%	581	0.87%	7	2%	3
5	414	10	1%	415	0.24%	10	1%	3
6	265	9	1%	266	0.38%	9	1%	3
7	413	10	1%	414	0.24%	10	1%	2
8	264	9	1%	265	0.38%	9	1%	3
9	271	6	7%	275	1.47%	6	8%	3
10	823	11	1%	850	3.28%	10	2%	N/A

The cycling can be explained by two characteristics. First, we encounter discretization effects. That is, a slight change in input parameters (here: expectation and variance of the failure rate) may lead to an increase or decrease of S by one. Second, due to the large difference in failure behavior of CM and AM parts for this instance, a small change in γ substantially changes the expected demand which increases the occurrence of discretization effects. We solve this issue by fixing S equal to the solution with the best performance thus far. As a result, discretization effects are avoided and the algorithm converges to the solution as shown in Table 6.3 after 4 iterations.

6.5.3 Case Study Revisited

In this section, we return to the hinge bracket case discussed in Chapter 5, cf. Section 5.5 to apply the iterative procedure in combination with a different dual sourcing model. To that end, we again create a *simplified version* of the dual sourcing model discussed in Chapter 5. As we outline in Appendix 6.C, the simplified MDP model no longer distinguishes between the supply mode specific failure behavior since this aspect is taken care of by the iterative procedure.

In line with the analysis in Chapter 5, we consider both service costs and operational savings. Recall that the latter are a result of fuel savings which are caused by the lighter AM part design. Compared to the analysis in Chapter 5, we no longer require the consolidation of demand streams to obtain computationally tractable results. Hence, we can directly apply the iterative procedure without refinements to the parameter settings shown in Table 5.3.

As we observe in Figure 6.1, both the total service costs and the individual cost factors only differ slightly. Furthermore, the iterative procedure correctly identified that sourcing with CM only is the best sourcing approach. Hence, the result give a first indication that the iterative procedure performs well for larger problem instance. Besides, we obtain further evidence that the choice to consolidate demand streams in Chapter 5 appears suitable.

However, we stress that the model in Chapter 5 was motivated for a specific application domain – namely the low-volume, high-downtime costs spare parts business. To assess




Figure 6.1: *Hinge bracket case evaluation comparison*

the performance of the iterative procedure under broader conditions for large problem instance (e.g. batch ordering, fixed ordering costs and more general lead time and failure distributions), it would be required to design a dual sourcing benchmark heuristic that distinguishes between CM and AM failure behavior. A suitable starting point might be the solution framework discussed in the next section.

6.6 Approximate Dynamic Programming Formulation

An even more general approach to overcome the curses of dimensionality is Approximate Dynamic Programming (ADP). In this section, we study the problem structure of the MDP model proposed in Section 6.3, and use the gained insights to derive an ADP design. The design relies on a typical value iteration scheme with an approximation of the value function. It turns out that our approximation of the value function works well, but unfortunately the value iteration scheme is not able to cope with the specific problem structure. We analyze the underlying reasons and suggest alternative modelling choices.

ADP is based on a solution framework that combines simulation with approximate analytical techniques to solve stochastic decision problems such as the one introduced in Section 6.3. Here, we approximate (learn) the relative value function $v(\cdot)$ (cf. Equation 6.16), which is a powerful technique to handle the computational burden associated with a large state space, cf. Powell (2011). In particular, we continuously improve our estimate of the relative value function $v(\cdot)$, by observing the value of visited states during the simulation. Ultimately, the estimate of the relative value function can be used to derive an approximation of the optimal average cost policy.

There exist several methods to approximate the relative value function $v(\cdot)$, cf. (Bertsekas, 2012) or (Powell, 2011). Here, we approximate $v(\cdot)$ by a linear combination of so-called

basis functions $\phi_f(\cdot)$. In analogy to linear regression, the basis functions may be viewed as independent variables that describe a so-called state feature $f \in \mathcal{F}$, such as the inventory position or the inventory level. Accordingly, we obtain the following approximation $\bar{v}(\cdot)$ for the relative value function:

$$v(\cdot) \approx \bar{v}(\cdot) = \sum_{f \in \mathcal{F}} \theta_f \phi_f(\cdot) \quad (6.21)$$

where θ_f denotes the weight of each basis function. Compared to other approximation techniques, the benefit of the basis function approach is that each observation obtained with the simulation affects the estimate of the relative value function $v(\cdot)$. Hence, we may predict the value of a certain state based on its features without ever having visited it before. We determine a set of suitable features in the next section.

In general, we may approximate the relative value function based on the features of the “normal” states. However, we use the features of so-called post-decision states instead. Conceptually, the post-decision state separates random from deterministic factors of the transition which, as we explain below, leads to significant computational benefits.

In our case, the post-decision is equal to the state we would transition to, if we apply a certain action and disregard any failure arrivals, i.e., $k_C, k_A = 0$. For example, suppose at the beginning of some period there are ten CM and two AM items in operation while there are no items in stock or in the replenishment pipeline, i.e., $\mathbf{i}_t = (10, 2, 0, 0, 0, 0, 0)$. If we decide to order two AM items, i.e., $\mathbf{a}(\mathbf{i}_t) = (0, 2)$, then the post-decision state is equal to $\mathbf{i}_t^{\mathbf{a}(\mathbf{i}_t)} = (10, 2, 0, 0, 0, l_A - \tau, 0, 2)$. To complete the transition to state \mathbf{i}_{t+1} , we then need to consider possible failures during the period.

Using this approach, we eliminate the need to learn the value of each state-action pair for all possible failure realizations (k_C, k_A) . Instead, we only have to learn an approximation for the different failure realizations of each post-decision state. Clearly, if the number of possible failure realizations (outcome space) becomes large, the post-decision concept significantly reduces the computational burden.

6.6.1 Problem Features

In this section, we derive the set of features that are being used in the ADP method. The section is structured as follows. First we derive a set of possible features using domain knowledge. Next, we use statistical analyses to select the most promising features and to confirm their predictive power to approximate the relative value function $v(\cdot)$.

Several features may qualify for the description of a post-decision state. As is further motivated below, we consider each state variable, and the following characteristics relevant:

- IP_t : the inventory position as defined by Equation 6.2.
- $IP_{t,A}$: the inventory position if we only take outstanding orders into account that arrive within the AM lead time, i.e., $IP_{t,A} = s_{C,t} + s_{A,t} + \mathbf{1}_{u_{C,t} \leq l_A/\tau} w_{C,t} Q_C + w_{A,t} - B_t$, with $\mathbf{1}_{u_{C,t} \leq l_A/\tau}$ denoting an indicator function equal to 1 if $u_{C,t} \leq l_A/\tau$ and 0 otherwise.

- $E[D_t]$: the expected demand in period t , i.e., $E[D_t] = n_{C,t}\mu_C + n_{A,t}\mu_A$.
- $V[D_t]$: the demand variance in period t , i.e., $V[D_t] = n_{C,t}V_C + n_{A,t}V_A$.
- IL_t : the inventory level in period t , i.e., $IL_t = s_{C,t} + s_{A,t} - B_t$.
- L : the time until the next replenishment order arrives, i.e., $L = \min(u_A, u_C)\tau$. If no AM (CM) order is outstanding, we define $u_A = l_C/\tau$ ($u_C = l_C/\tau$). We motivate this choice based on the post-decision state in which no outstanding order implies that it is likely that we have sufficient stock.
- *Constant*: a feature independent of the post-decision state.

IP_t and $IP_{t,A}$ are commonly used in dual sourcing heuristics, cf. Section 5.2. Hence, it appears reasonable to assume that they carry high predictive power. The features $E[D_t]$ and $V[D_t]$ provide insight into the state dependent failure behavior, and IL_t and L give an indication about the stock-out risk/severity. Finally, the constant ensures that costs which cannot be explained by the linear combination (cf. Equation 6.21) do not lead to over-fitting.

Next, we evaluate how well the proposed features may serve as predictors for the actual value of each state. For that purpose we use LASSO regression (Tibshirani, 1996). LASSO regression holds the benefit that it allows the further refinement of the set of features. The underlying principle is that it solves an ordinary least squares problem, which simultaneously penalizes non-zero regression coefficients. Therefore, coefficients of features with low predictive value become equal to zero. Note that we computed the actual state values with the exact solution procedure described in Appendix 6.A and evaluated the problem instances defined in Table 6.2. Furthermore, since we study an average cost MDP formulation, we normalized the relative values and considered each instance separately. The results are shown in Table 6.4.

Table 6.4: LASSO regression models for the instance defined in Table 6.2

Feature	Instance									
	1	2	3	4	5	6	7	8	9	10
n_C	-0.03	-0.02	-0.03	-0.01	-0.02	-0.02	-0.02	-0.01	-0.03	-0.04
n_A	-0.01	0	-0.02	0	0	0	0	0	-0.02	0
s_C	0	0	0	0	0	-0.01	0	-0.01	0	0
s_A	0.01	0.01	0	0	0.02	0.01	0.01	0	0	0.04
u_C	-0.01	-0.02	-0.01	-0.02	-0.05	-0.05	-0.04	-0.05	0.00	-0.01
u_A	0	0	0	0	0	0	0	0	0	0
w_C	0	-0.01	0	0	0	0	0	0	0	-0.08
w_A	0	0	0	0	0	0	0	0	0	0
IP_t	-0.01	-0.01	0.00	-0.01	0	-0.01	0	-0.01	0	0
$IP_{t,A}$	-0.01	-0.02	-0.01	-0.02	-0.01	-0.01	-0.02	-0.02	-0.01	-0.01
IL_t	-0.01	-0.01	-0.01	-0.01	-0.01	0	-0.01	0	-0.02	-0.02
L	0	0	0	0	0	0.01	0	0	0.01	0
$E[D_t]$	0	0	0	0	0	0	0	0	0	0
$V[D_t]$	0	0	0	0	0	0	0	0	0	0
<i>Constant</i>	0.45	0.44	0.36	0.40	0.48	0.57	0.39	0.53	0.30	0.54
R²	0.75	0.73	0.78	0.81	0.78	0.86	0.84	0.88	0.77	0.83

Overall, the set of chosen features achieves a high predictive power (on average $R^2 = 80\%$). However, we observe that some features only carry a low degree of predictive value and thus are less suitable for the ADP design. We exclude features for which the coefficient is unequal to zero in less than three instances. As a result, the final feature set does not contain s_C , u_A , w_C , w_A , L , $E[D_t]$ and $V[D_t]$.

6.6.2 Performance Analysis

In this section, we analyze the performance of the ADP design. As outlined in Appendix 6.D, the solution procedure is based on the so-called ϵ -greedy approach in combination with recursive least square to update the basis function weights after each iteration.

In Table 6.5, we show the performance of the ADP design compared to the exact method for the set of instances, which we defined earlier in Table 6.2. As we observe, the performance of the ADP design is unsatisfactory for most instances. Next to a large cost difference for most instances (Δg), the policy obtained from the ADP algorithm suggest to never order items from the AM source, i.e., $\gamma = 0\%$ for each instance.

Table 6.5: Comparison of exact method with ADP method

Instance	Exact method			ADP method		
	g^*	S^*	γ^*	g	Δg	γ
1	914	8	1%	2090	129%	0%
2	579	7	1%	583	0.69%	0%
3	917	8	1%	1394	52%	0%
4	576	7	2%	13290	2207%	0%
5	414	10	1%	419	1.2%	0%
6	265	9	1%	2658	903%	0%
7	413	10	1%	8015	1841%	0%
8	264	9	1%	278	0.53%	0%
9	271	6	7%	325	20%	0%
10	823	11	1%	825	0.24%	0%

Since the ADP design performed well for some instances, we first examined whether we did not allow for a sufficient number of iterations for the other instances. It turned out that this was not the reason for the observed results as we discuss in Appendix 6.E.

Next, we evaluated for the least accurate Instance 4 whether the basis function design and chosen algorithm parameters (cf. Appendix 6.D) would not allow for an adequate approximation of the relative value function. To that end, we guided the weight learning procedure by choosing the actions within the ADP algorithm according to the optimal policy (which we obtained from the exact solution).

Figure 6.2 plots the state values obtained from the resulting value function approximation against the exact values for each state of Instance 4. As we observe, the basis function design and chosen algorithm parameters lead to a good approximation, i.e., states with

a low value in the MDP solution also receive a low value with the ADP solution (and vice versa). Note that the less accurate predictions for some high values are uncritical, since these would be avoided in the optimal policy anyways. For example, the least precise approximation was obtained for the state $(0, 0, 0, 0, 7, 0, 1, 0)$, i.e., no items in operation nor in stock while we decide to order one batch of CM items.




Figure 6.2: *ADP design estimates compared to exact values for all states, Instance 4 – using optimal policy*

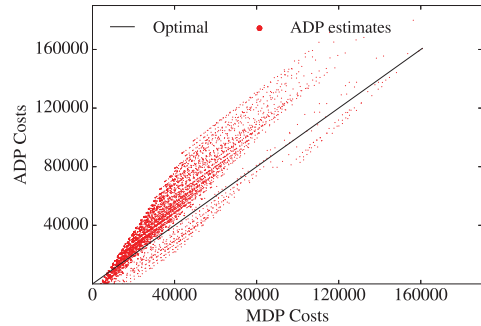


Figure 6.3: *ADP design estimates compared to exact values for all states, Instance 4*

In comparison, Figure 6.3 shows the results if we do *not* follow the optimal policy for Instance 4, but the ϵ -greedy approach. As we observe, the approximation quality is poor. By analyzing the simulation trajectory of the ADP, we found that this result is caused by the specific problem structure. Due to the low demand rates, the optimal decision for most states is to order nothing. Therefore, the order decision of only a few states determines the solution quality. For example, for Instance 4, we find that the optimal action is to order nothing in more than 93% of all states. Learning the weights under such conditions is problematic since our ADP design is biased by relatively more frequently observed states (in which nothing is ordered), thereby largely influencing the value function approximation. Since AM orders are even less common, this finding may also explain why the ADP results suggest to never order from the AM source across all instances (cf. Table 6.5). We conclude that, despite a suitable approximation framework (cf. Figure 6.2), the problem structure distorts the learning phase and therefore leads to inaccurate state value approximations (cf. Figure 6.3).

An alternative to cope with the described problem may be an ADP approach in which we aggregate states. Aggregation refers to a concept in which we group states according to their features. With aggregation it becomes possible to isolate states with a high likelihood of ordering and therefore to obtain a more specific value function approximation. Future research has to show whether aggregation is a fruitful approach in combination with a rare event ADP problem, such as ours.

6.7 Conclusion

In Chapters 4 and 5, we demonstrated the value of applying AM in combination with CM methods. However, computational limitations only supported the assessment of modest problem instances. In this chapter, we developed an iterative procedure that may overcome this limitation. The approach is designed to extend dual sourcing methods which do not distinguish between CM and AM failure behavior. We demonstrated its application in combination with two dual sourcing models. In both cases, the performance of the iterative procedure turned out to be promising. The optimality gap was less than 4% across all instances (typically way less), and the iterative procedure confirmed the sourcing approach determined for the case study in Chapter 5. For future research, it seems worthwhile to study the performance of the iterative procedure under more general conditions. For that purpose, it is necessary to develop a benchmark heuristic first, since exact solution methods are limited by computational tractability.

To obtain a benchmark heuristic, we examined the option to model the problem with ADP. Although we identified suitable problem features, it turned out that the application of ADP is not straightforward for the observed problem structure. In particular, we found that the optimal policy typically orders parts in only a very limited number of states. Learning the value of these states was difficult for our ADP design since it is biased towards relatively more frequently observed states in which nothing is ordered. A design based on state aggregation appears promising to model such rare event ADP problems and hence is a good candidate for future research.

Appendices

6.A Exact Solution Procedure

In Algorithm 6.2, we describe how the policy iteration algorithm may be used to find the optimal policy π for the Bellman Equations 6.16.

Algorithm 6.2: Policy iteration

```

Set  $j=1$  and chose a stationary policy  $\pi_j$ ;
repeat
  Determine  $g(\pi_j), v(\mathbf{i}_t, \pi_j)$  with the system of linear equations;
   $v(\mathbf{i}_t, \pi_j) = C(\mathbf{i}_t, \pi_j) - g(\pi_j) + \sum_{\mathbf{i}_{t+1} \in \mathcal{I}_{t+1}} p(\mathbf{i}_{t+1} | \mathbf{i}_t, \pi_j) v(\mathbf{i}_{t+1}, \pi_j), \forall \mathbf{i}_t \in \mathcal{I}$ ;
   $v(\mathbf{s}_t, \pi_j) = 0$  with  $\mathbf{s}_t$  an arbitrary state;
  Find the new policy  $\pi_{j+1}$ ;
   $\mathbf{a}(\mathbf{i}_t)_j = \operatorname{argmin}_{\mathbf{a}(\mathbf{i}_t) \in \mathcal{A}} \{C(\mathbf{i}_t, \pi_j) - g(\pi_j) + \sum_{\mathbf{i}_{t+1} \in \mathcal{I}_{t+1}} p(\mathbf{i}_{t+1} | \mathbf{i}_t, \pi_j) v(\mathbf{i}_{t+1}, \pi_j)\}$ ;
   $\pi_{j+1}(\mathbf{i}_t) = \mathbf{a}(\mathbf{i}_t)_j, \forall \mathbf{i}_t \in \mathcal{I}$ ;
  Set  $j = j + 1$ ;
until  $\pi_j = \pi_{j-1}$ ;
return the optimal policy  $\pi_j$ 

```

6.B Markov Decision Problem Changes

Here, we elaborate how the MDP model described in Section 6.3 may be simplified to obtain a model in which we do not distinguish between CM and AM items.

Any combination of $n_{C,t}$ and $n_{A,t}$ is replaced by n_t . The same holds for $s_{C,t}$ and $s_{A,t}$ and thus we define $\mathbf{i}_t = (n_t, s_t, u_{C,t}, u_{A,t}, w_{C,t}, w_{A,t})$. The decision space and cost calculations remain unchanged, except that we do not differentiate between AM and CM failures any longer. Hence, the expected demand in period t is equal to the “general” failure rate (μ) multiplied with the number of operating parts (n_t), and we have to change the calculation of the maintenance and backorder costs accordingly. The new transition probability is equal to $p(\mathbf{i}_{t+1} | \mathbf{i}_t, \mathbf{a}(\mathbf{i}_t)) = p(k | n_t)$, where k describes the total number of failures in period t . The number of operating systems at the beginning of period $t+1$ simplifies to

$$n_{t+1} = n_t - k + x \quad (6.22)$$

where x is equal to the total number of items that may be installed *before* and *after* a possible replenishment order arrives in period t , i.e., $x = \min\{B_t + k, s_t + A_t\}$, where the first element describes the total demand (backorders + new demand arrival in period t), and the second element the on-hand stock plus the total number of items A_t arriving at

$t+\tau$. Finally, we simplify the computation of the on-hand stock in period $t+1$ to

$$s_{t+1} = s_t + A_t - x \quad (6.23)$$

To measure the fraction ρ of sourced AM items, we multiply the steady state probabilities with the corresponding order quantities.

6.C Simplified Version of the Hinge Bracket Model

Since we no longer differentiate between the CM and AM failure rate, we can reduce the dimension of the state space and therefore describe the state \mathbf{i} of the inventory system by the tuple $(n_{G,\mathbf{i}}, r_{C,\mathbf{i}}, r_{A,\mathbf{i}}, s_{G,\mathbf{i}})$. Note that we still differentiate between $r_{C,\mathbf{i}}$ and $r_{A,\mathbf{i}}$ because these state variables define the resupply rates of both sourcing options.

The action space reduces as well. The maintenance decision, i.e., which item type to take from stock, becomes superfluous since both items types show the same failure behavior and we approximate the holding costs based on the average ratio γ_{j-1} of CM and AM items in operation in iteration $j-1$. To that end, we only have to decide which item type we order after each failure. Furthermore, the calculation of the total expected cost simplifies and we rewrite Equation 5.1 for iteration j

$$g_{\mathbf{i}} = \mu_C c_C r_{C,\mathbf{i}} + \mu_A c_A r_{A,\mathbf{i}} + m(\lambda_G n_{G,\mathbf{i}}) + \kappa((1-\gamma_{j-1})c_C + \gamma_{j-1}c_A)s_{G,\mathbf{i}} + b(k - n_{G,\mathbf{i}}; 0)$$

The final alteration concerns the operational savings (o_A) which we secure with each installed AM item. Given the simplified state space, we also use γ_{j-1} to find an approximation of the operational savings. That is, the operation costs savings are equal to $-\gamma_{j-1}o_A n_{G,\mathbf{i}}$ for state \mathbf{i} . For the optimization, we apply the same approach as described in Chapter 5.

6.D Approximate Dynamic Programming Algorithm

Before we outline the ADP algorithm, we first highlight its key elements. To learn the value of different state action pairs, we apply the so-called ϵ -greedy approach, in which random actions are chosen with a probability ϵ . Otherwise, we select actions that appear most valuable according to the current value function approximation. Furthermore, after each iteration, we chose the next state at random with equal probability which leads to a uniform exploration of the state space.

For updating the basis function weights θ_f after each iteration, we use the recursive least squares method. In comparison to ordinary least squares, with recursive least squares it is not necessary to reconsider all past information after each iteration (by solving a matrix inversion). We only need to update the weights with the new information. Furthermore, by using a recursive updating procedure, it becomes possible to attribute a higher importance to more recent observations using a so-called *forgetting factor* which is introduced to give less weight to older observations. For a more detailed discussion of recursive updating procedures, we refer to Powell (2011).

Finally, we use a step-function to update the average cost estimate \bar{g} after each iteration. The larger we choose the step size, the more weight is given to the more recent observation. Here, we apply a step size which is decreasing after each iteration.

Algorithm 6.3 outlines the ADP approach for one iteration j . Note that algorithm parameters such as the step size, the number of iterations or the probability with which we choose a random action are tunable. As we discuss, in Section 6.6.2, our settings appear reasonable since they lead to a good value function approximation if actions are chosen according to the optimal policy.

Algorithm 6.3: ADP procedure for iteration j

1. Choose a random state $\mathbf{i}_j \in \mathcal{I}_t$
 2. Determine an action $\mathbf{a}(\mathbf{i}_j)$ with the ϵ -greedy approach based on the average cost estimate \bar{g}_{j-1} and the value function approximation $\bar{v}_{j-1}(\cdot)$
 3. Determine the post-decision state $\mathbf{i}_{j,\mathbf{a}(\cdot)}$
 4. Sample a failure realization k_C and k_A leading to the next state \mathbf{i}_j^{t+1}
 5. Determine the best action $\hat{\mathbf{a}}(\mathbf{i}_j^{t+1})$ based on \bar{g}_{j-1} and $\bar{v}_{j-1}(\cdot)$
 6. Determine the new post-decision state $\mathbf{i}_{j,\hat{\mathbf{a}}(\cdot)}^{t+1}$
 7. Determine the error $\Delta = \bar{v}_{j-1}(\mathbf{i}_{j,\mathbf{a}(\cdot)}) - (C(\mathbf{i}_j^{t+1}, \hat{\mathbf{a}}(\mathbf{i}_j^{t+1})) + \bar{v}_{j-1}(\mathbf{i}_{j,\hat{\mathbf{a}}(\cdot)}^{t+1}))$
 8. Find $\bar{v}_j(\cdot)$ with recursive least square based on the observed error Δ
 9. Update the average cost estimate with $\bar{g}_j = (1 - \beta)\bar{g}_{j-1} + \beta C(\mathbf{i}_j^{t+1}, \hat{\mathbf{a}}(\mathbf{i}_j^{t+1}))$
-

6.E ADP Convergence Analysis

To determine the number of iterations that are required before the value function approximation $\bar{v}(\cdot)$ converges, we study the change of the basis function weights θ_f . Their development is shown for Instance 4 and 8 (cf. Table 6.2) in Figure 6.4. The graphs indicate




Figure 6.4: Convergence of weights $\theta_f | f \in \mathcal{F}$, Instance 4 (left) and Instance 8 (right)

that the weights converge after about 80,000 iterations (less than 1 minute computation time). Remaining fluctuations are a consequence of the random trajectory in which states are visited. Non-stationary updating procedures for the basis function weights may even lead to stronger convergence. However, these procedures demand more extensive parameter tuning and therefore appear more suitable to fine-tune the algorithm after the current approach leads to satisfactory results.

Chapter 7

Conclusions

In the previous chapters, we investigated the possible impact of the introduction of AM technology on after-sales service supply chains. In this chapter, we draw the main conclusions. In Section 7.1, we present our results by responding to the research questions presented in Chapter 1. Furthermore, we discuss possible model extensions and suggest an implementation framework in Section 7.2. Subsequently, in Section 7.3, we outline research areas that may contribute to the dissemination of AM in after-sales service supply chains. Section 7.4 concludes this thesis.

7.1 The Research Agenda Revisited

We formulated the research objective of this thesis as follows:

To offer decision support for actors in after-sales service supply chains to identify and understand the value of AM technology for their organization, and to provide quantitative insights into both when and how AM technology may be used or combined with conventional manufacturing methods to improve the efficiency of service logistics.

To address this research objective, we formulated five research questions. In Section 7.1.1, we review the results for each research question and elaborate on the possible extensions thereafter in Section 7.1.2.

7.1.1 Results

The first research question was motivated by the observation that organizations struggle to identify high-impact spare parts for the application of AM in spare parts supply chains. The corresponding research question was formulated as follows:

RQ 1. *How can organizations identify spare parts that are economically viable and technologically feasible for the application of AM technology?*

By studying the literature and observing the practices in the organizations that participated in this research project, we recognized that the selection of parts is solely based on expert judgment. While this, as we call it, bottom-up procedure certainly has value, we have identified various limitations. For instance, practitioners are often not aware of the spare parts characteristics that influence the value of AM technology from a supply chain perspective and tend to focus on the technological features of the product. In addition, a bottom-up procedure only allows the consideration of a small subset of parts. Both limitations may direct unnecessary attention towards less interesting spare parts, which, in turn, causes organizations to incur avoidable exploration costs and, moreover, may diminish the enthusiasm for AM technology.

We propose a top-down approach that is based on the Analytic Hierarchy Process (AHP) and relies on spare part information that is retrievable from ERP systems. This has the following two advantages: first, the approach can be customized according to specific company characteristics, and second, a very large number of spare parts may be assessed simultaneously. A case study in the aerospace industry has demonstrated this approach. Furthermore, sensitivity analyses have been performed to evaluate the robustness of the method. Our findings provided evidence that the method facilitates a useful prioritization of a large spare part assortment. Moreover, sensitivity analyses demonstrate both the robustness of the approach and the method's application flexibility. Today, the procedure has been applied in more than five organizations and has led to various positive business cases. In addition, we recognized that the application of the top-down approach prompts organizations, and in particular, their higher management, to reflect on the goals associated with AM for their spare part operations. This stage proved to be crucial to re-adjust expectations and create a high level of engagement within the organizations.

The second research question was motivated by the observation that AM technology often leads to substantial design changes. While the effects of design changes are quite well understood from a technological and operational perspective, the implications for service logistics are commonly ignored or oversimplified. Since consolidation of parts was recognized as the most promising application of AM technology for operations (Wohlers Report, 2014), we formulated the following research question:

RQ 2. *How does the consolidation of spare parts with AM technology affect the total life cycle cost of capital goods and when is it valuable?*

We primarily found that authors emphasized the benefits of consolidating parts with AM technology in the literature. For instance, it was argued that the avoidance of the assembly steps and the reduced organizational complexity will cause substantial reductions in lead time. We have attempted to draw a more balanced picture by also highlighting the possible

drawbacks. Furthermore, we conceptually discussed how the total life cycle costs may be affected if we compare the operation of a CM-produced segmented spare part design with an AM-produced consolidated design.

Based on these insights, we identified the VARI-METRIC framework to be a suitable approach to assess the effects of consolidation with AM quantitatively. Our results suggest that consolidation with AM often increases service costs. Thus, in about only 13% of our analyzed instances, consolidation leads to cost reductions. This finding mainly stems from a loss of flexibility. For example, in case of a failure, a consolidated component must be replaced in its entirety, whereas the conventional assembled component may only require the replacement of a defective sub-component. Furthermore, short re-supply lead times for the consolidated spare part turn out to be less beneficial than perceived, therefore putting into perspective one conjectured benefit of consolidation with AM.

On the other hand, we found that the reduction in purchasing costs and improvement in reliability are indicative for a high value of consolidation with AM. Furthermore, parts with comparable characteristics in terms of lead time, reliability and piece price typically yield the highest potential for consolidation. This finding is intuitively explained by the consideration that the integration of comparable parts causes only marginal reductions in flexibility. Overall, our results exemplified the necessity to adopt a broader perspective when judging the effects of design improvements with AM. Otherwise, design improvements may lead to unforeseen effects that may render its application debatable, despite the substantial functionality improvements.

The third research question was driven by the wish to extend the results from the investigation pertaining to the first research question. While the top-down approach proposed in Chapter 2 simplifies the identification of high-impact spare parts, it remains unclear when and how organizations should start the transition to the AM-produced spare part version. Therefore, we posed the following research question:

RQ 3. *When and how does a transition to AM technology become profitable for the low-volume spare parts business?*

Based on a scenario observed at an OEM of radar systems, we developed a finite horizon stochastic dynamic programming model. We selected this approach because we identified that the change of characteristics through time, such as AM production cost reductions and installed base size changes, are deterrents to the transition strategy. Furthermore, we wanted to understand how the existing CM production capacities may affect the investment decision in AM technology.

Utilizing a case from the OEM and numerical experiments, we were able to derive various crucial insights. First, a transition to AM technology for spare parts supply is often valuable during the service period. In other words, we found average costs savings of nearly 35% across all instances. To a certain extent, this finding contrasts the common belief that investment in AM technology only pays off if initiated at the beginning of an asset's life cycle. Second,

(initially) high AM unit and setup costs may incorrectly bias organizations against the value of moving to AM technology. In particular, the high inventory cost savings – in terms of holding and backorder cost reductions – may offset these costs. Finally, the investment in AM technology should not be postponed. We observed that it is usually better to prepare the AM method as soon as possible. Then, during the transition period, AM may supplement the regular supply source. By applying this dual sourcing strategy, companies can benefit from the short AM replenishment lead times in stock-out situations, and on the other hand, they can avoid high AM unit costs by ordering the majority of parts from a cheaper CM source.

The fourth research question builds on the findings from the study of the previous research question and investigates the dual-sourcing phase in greater detail. In particular, we relaxed the assumption that the CM and AM versions exhibit the same failure behavior. Arguments pertaining to design changes, different production procedures and different materials motivate this extension. Therefore, we formulated the following research questions:

RQ 4. *Under what conditions does the sourcing of low-volume spare parts with a combination of AM and CM methods pay off if we acknowledge that part quality is largely influenced by the production method?*

To explore this question, we developed an infinite horizon continuous-time Markov Decision Model, in which we explicitly considered the dependency between the sourcing decision and future demand. This dependency has not been addressed in the dual sourcing literature yet but appears to be essential if both supply modes use different production techniques.

From numerical experiments as well as a case study conducted in the aerospace industry, we derived the conditions under which a dual sourcing approach with AM is most valuable. In general, we observed that dual sourcing reduces the purchasing and maintenance costs as compared to single sourcing with AM methods; and it reduces the backorder and holding costs as compared to single sourcing with CM methods. In addition, we showed that dual sourcing regularly leads to cost savings of more than 30%, even if the AM part reliability or AM unit costs were three times greater than that of the corresponding CM part. In fact, even higher cost savings were possible if we considered scenarios where the backorder and holding costs were high relative to the purchasing and maintenance costs. Finally, our results exemplified that printing spare parts on demand is typically not a good sourcing strategy for downtime critical spare parts.

The fifth research question was motivated by the computational limitations of the model proposed in Chapter 5. Computational limitations may arise for problem instances in which we regard factors, such as higher demand rates, more complex demand patterns or different order policies. Thus, we formulated the following research question:

RQ 5. *How can we analyze large problem instances of the problem discussed in RQ 4, given the computational limitations of the exact analytical methods?*

We addressed this research question by developing an iterative procedure that supports the separation of the effects of the supply mode specific failure behavior from other aspects of the model. This decomposition significantly reduces the model's complexity and, therefore, enables the approximate analysis of larger problem instances. Furthermore, the iterative procedure facilitates the extension of dual sourcing heuristics from the literature (cf. Section 5.2) through a supply mode specific failure behavior. This extension is easily implemented, as the iterative procedure only relies on an estimate of a fraction of items ordered from either of the supply modes. Therefore, combining the existing dual sourcing heuristics with the iterative procedure provides a flexible solution framework to assess a wide range of problems with AM as a dual sourcing option.

Employing small problem instances, we showed that the iterative procedure achieved an optimality gap of less than 4% (typically even smaller) across all instances. Furthermore, we re-evaluated the case study presented in Chapter 5, which was possible this time without any further simplifications to obtain a computationally solvable model. The results obtained with the iterative procedure confirmed the results of Chapter 5. Hence, we find first evidence that the iterative procedure performs well in addressing large problem instances.

Furthermore, we examined the option to model the problem with Approximate Dynamic Programming (ADP). Thus, we studied an ADP design that uses state features to find a value function approximation. While it turned out that the selected features were suitable to approximate the value function, the studied ADP design did not perform satisfactorily. In particular, we found that the optimal policy typically orders parts in only a very limited number of states. Learning the value of these states was difficult for our ADP design since it is biased towards relatively more frequently observed states in which nothing is ordered. We concluded that the proposed ADP design was unsuitable for the encountered problem structure and that an alternative design should be explored.

7.1.2 Model Extensions and Improvements

Here, we elaborate on the possible extensions of the models introduced and analyzed in this thesis.

- *Refinements to the top-down approach.*

Based on the experience we gained while applying the top-down approach proposed in Chapter 2, we found that the data collection step is usually the most challenging. In particular, organization-specific terminology or decentralized data storage often complicates this step. Furthermore, extensive investigations are usually necessary to uncover the implications of certain spare part attributes available in the databases. So far, these time consuming steps were undertaken in graduation projects. However, we believe that statistical analyses or data mining techniques, such as clustering, may have the potential to simplify these steps. Furthermore, it may be worthwhile to extend the proposed method by also considering part attributes, which present options for design improvements. In the current approach, we ignored these opportunities while acknowledging that most applications of AM are motivated by design improvements at present.

- *Reduced supplier dependency through consolidation.*

In Chapter 3, we assumed a static system and, therefore, did not consider the effects of supply disruptions. However, if AM is applied to consolidate items in particular, the dependency on suppliers is significantly reduced. To that end, it may be worthwhile to study consolidation with AM in a dynamic system in which the supply channels may become unavailable.

- *Stochastic AM piece price development.*

In Chapter 4, we based our investigation on the assumption that AM piece price developments are known upfront (or can at least be reliably estimated). Future research may relax this assumption and, thus, study models with a stochastic AM piece price development. We suspect that this additional source of uncertainty increases the transition period in which both CM and AM methods are applied in parallel.

- *Mitigating supply risks with an early preparation of AM.*

We believe that it may be worthwhile to consider scenarios where the CM tooling may “get lost”. In a graduation project, which was conducted within the scope of the SINTAS project, we learned that such situations are not uncommon, particularly if the demand is low. Avoiding the associated risk of supply discontinuation with an (additional) supply source may further increase the value of an early investment in AM technology.

- *Joint optimization of service and operational costs.*

The application of the model presented in Chapter 5 to a case in the aerospace industry revealed that the joint optimization of service cost and operational cost savings may be a beneficial direction for future research. In particular, we observed that if AM offers the prospect of design improvements (leading to operational cost savings), the optimal sourcing strategy can change significantly. Furthermore, we suspect that the parallel consideration of operational and service cost savings may justify the application of AM technology in higher demand environments. This hypothesis is based on the observation that design improvements typically scale with quantity and are, therefore, more interesting for high demand environments.

- *Developing an ADP design based on state aggregation.*

The ADP design proposed in Chapter 6 was deemed unsuitable for the encountered problem structure. Other ADP designs where states are aggregated according to their features may be a better fit for modelling such type of problems. For future research, we believe that the grouping of states with a high likelihood of ordering will be a good starting point.

7.2 Roadmap for Implementation

Each organization has different requirements and follows different objectives with the implementation of AM in their after-sales service supply chains. For example, we observed organizations that aimed at improving system availability at remote locations while others were primarily concerned with the diversification of their supply options. To that end, any implementation strategy has to be customized for the specific firm. But while the outcomes

may be specific, the steps to run through are rather generic and hence it is possible to outline a framework that may guide the implementation process. In addition, the framework discussed below offers an overview at which stages during the implementation process the conducted research in this thesis becomes applicable. In the following, we will describe the implementation framework which is illustrated in Figure 7.1.

For most organizations the *spare parts assortment* size rules out an individual assessment of each spare parts. Hence, instead of considering spare parts at random, it is advisable to follow a *prioritization* routine. As we elaborated in Chapter 2, prioritization may be achieved by different means. For example, one may follow a bottom-up procedure where company representatives are encouraged to suggest promising spare parts for AM. Yet, as we discussed in Section 2.1, a top-down procedure based on *company goals* and easily retrievable *spare parts attributes* may be preferable for various conditions. The resulting *prioritized spare part list* directs the order in which the spare parts are assessed in the next step. Typically, the *assessment and selection* step requires the consultation of an AM technology expert. AM method specific process knowledge is required in order to determine the technical feasibility. Also less easily accessible information may rule out a further consideration of the assessed part. For example, in Section 2.3, we discussed a case for which a supplier offered an inexpensive final order of the conventional manufactured spare part version. Ultimately, this new insight led to the decision not to consider the spare part any further.

Next, we analyze the value of AM for each *selected spare part* in further depth. This process step may either be performed sequentially or in parallel for the selected parts. In any case, we propose to *generate feasible design options* first since cost and logistical parameters often largely depend on the specific design. During this step, company or industry specific *regulations and design constraints* usually have to be taken into account. The *design option evaluation* should be based on a total life cycle costs estimation. As our results in Chapter 3 demonstrated, hidden cost may quickly lead to a negative business case even despite substantial design improvements compared to the CM counter part. Also, at this stage it seems advisable to obtain a first overview about possible *AM supply options*. As we discussed in Section 1.2.2, the supplier (internal or external) is often highly involved in the design process. Furthermore, the assessment of supply options gives a first indication about the expected cost and lead times.

Based on the *selected AM spare part design*, we then analyze the value of introducing AM. Obviously, the assessment depends on the specific case. Section 1.4 gives an overview which situations have been already considered in the literature. However, for the implementation framework, we position the research conducted in this thesis. A good indicator for which research is applicable provides the *system life cycle stage* of the associated capital good. During the initial life cycle stage, our work in Chapter 4 is most relevant since it supports a *service strategy definition* with AM. For other systems that already reached a more mature stage, AM may offer a possibility to *improve the service approach*. We suggest to first regard the work of Chapter 4 as it reveals whether AM becomes relevant during the remaining service horizon. Depending on the result, Chapter 5 & 6 provide the means to evaluate the value of changing the service approach in further detail. During the final life cycle stage, preliminary research in the form of Bachelor and Master projects indicated that the value of AM primarily is derived

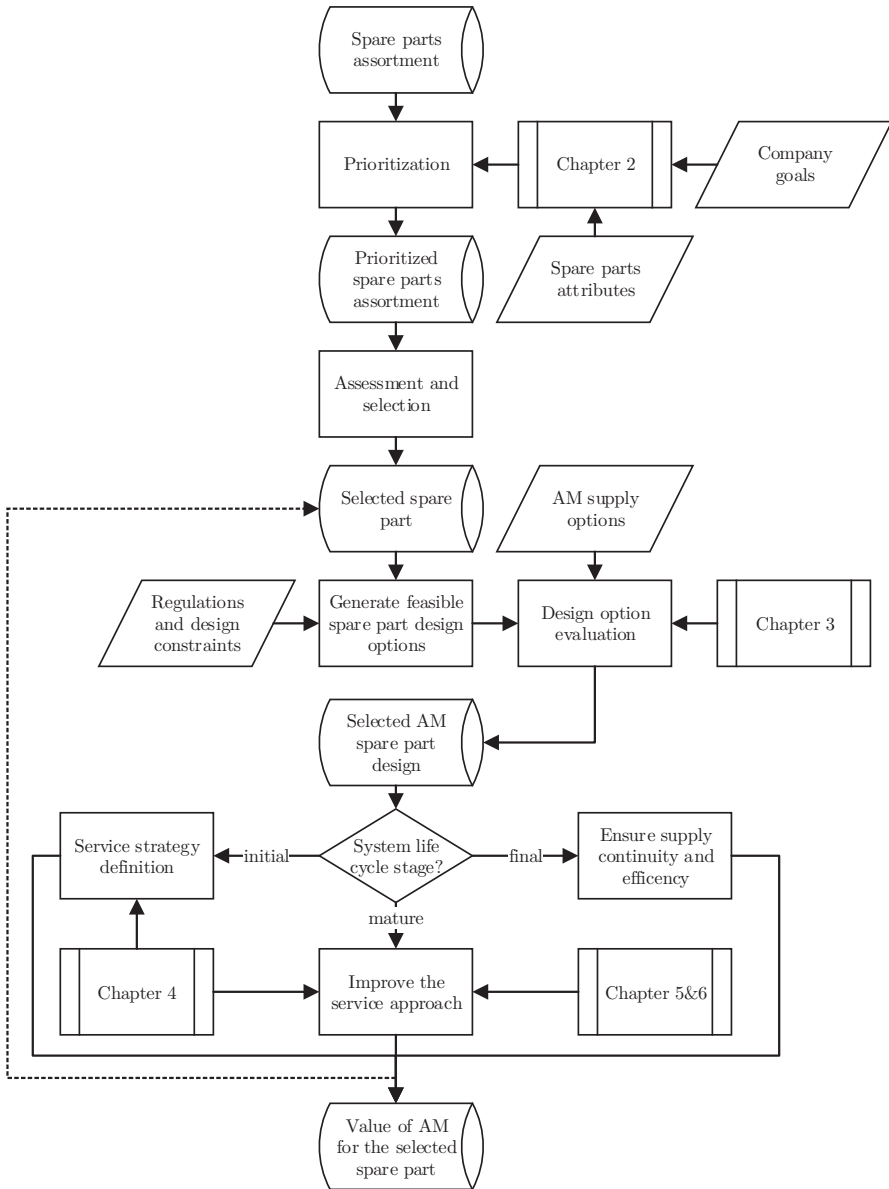


Figure 7.1: Implementation Framework

from the flexibility of AM as opposed to other supply options such as a final order, aftermarket supply or refurbished parts. However, future research is required to reveal to which extent AM may contribute to *ensure supply continuity and efficiency* during the final life cycle stage.

The estimated *value of AM for the selected spare part* then should provide sufficient evidence whether an implementation is profitable for the organization. Finally, note that we discussed the implementation framework for the scenario of an existing spare parts assortment. Yet, the suggested approach may also be applied during the design phase of a capital good. For such a scenario, it is likely that the prioritization step is based on expert judgment. Furthermore, it may even offer additional flexibility to base the assessment on functional descriptions rather than predefined parts. As we also discuss in the next section, the combination of design decisions with service considerations offers a fruitful direction for future research.

7.3 Discussion of Further Research Areas

In this section, we outline the research areas related to AM application, which were not addressed in this thesis. For this purpose, we first recall Section 1.2.2 where we described the possible effects of AM in after-sales service supply chains. Today, most effects still require further exploration. Here, we will focus our discussion on a few selected research areas that we consider particularly promising for the further dissemination of AM technology in after-sales service supply chains.

- *Circumvention of certification cost.*

Our results revealed the value of printing critical parts for the spare parts business. However, regulations often prevent the application of AM in this domain. Methods in which the part is not printed directly may bypass these problems. For example, by printing the sub-components, the component re-supply/repair lead time may be reduced significantly, while regulations become way less restrictive. We suspect that such printing options may also affect the level of repair decisions and, thereby, may alter maintenance concepts in general.

Another example is the option to produce (disposable) AM tooling and then use it for the production of spare parts. Again, such indirect methods may (partially) help avoid certification costs while replenishment lead times and unit costs are likely to reduce. In a graduation project carried out within the scope of the SINTAS project, it was found that the reduced unit costs are particularly promising in low demand environments.

- *AM in combination with condition based monitoring.*

Our investigations revealed that printing spare parts on demand is typically not an option for downtime critical capital goods. Despite the short AM production lead times, the total throughput time is still too long. However, the shorter lead time may be valuable in combination with condition based monitoring. At present, long CM production lead times may force organizations to order spare parts even when the information available about the monitored part is insufficient. Yet this approach increases stock levels and may cause unnecessary production costs. With AM technology, the production decision may be postponed until sufficient evidence concerning the part's condition has been gathered.

Future research may show whether the postponement of the production decision can justify investment in AM technology.

- *Mitigating the risks of supply disruptions.*

As we discussed in Chapter 1, we often observe problems with supply disruptions during the service period of capital goods. The flexibility of AM technology to produce spare parts with general purpose equipment raises the question: For what asset components will it be beneficial to proactively design AM backup solutions? This may be a particularly valuable extension to a recent research that utilized company data to predict possible supply disruptions, e.g., Li et al. (2016).

- *Restructuring of the spare parts supply chain.*

In Chapter 1, we elaborated that, at present, most service providers rely on complex manufacturing supply chains to source spare parts. With the maturation of AM technology, this dependency may reduce and, thus, allow different sourcing concepts. For example, service providers may change the supply source, depending on the proximity to asset location. Moreover, service providers may seek to minimize import duties or transportation costs. As elaborated in Section 1.2.2, logistic service providers already experiment with such opportunities in the consumer market.

- *Investment of service providers in reliability improvements.*

As discussed in Chapter 1, AM simplifies design changes. Therefore, it becomes perceivable that not only an OEM but also a service provider (co-)invests in the reliability increases of components. The arising investment decision may be studied in the context of a spare parts system approach (cf. Section 1.5), and relates to the literature pertaining to the value of system parameter improvements, cf. Hu et al. (2018). Furthermore, the described scenario may give rise to models that study investment decisions from a cooperative game theory perspective.

7.4 Concluding Remarks

The application of AM technology in after-sales service supply chains remains fairly unexplored. With this thesis, we addressed various research directions that were primarily derived from observations made at our industrial and governmental project partners. While we believe that we have contributed to a better understanding regarding the potentials and risks of AM technology for after-sales service logistics, we are also aware of the wide range of unanswered questions that require further attention from the scientific community.

Bibliography

- 3ders.org (2017). *FAA's 8-year 'Additive Manufacturing Strategic Roadmap' lays groundwork for 3D printed aircraft development*. URL: <https://www.3ders.org/articles/20171023-faas-8-year-additive-manufacturing-strategic-roadmap-lays-groundwork-for-3d-printed-aircraft-development.html>.
- Achillas, Ch., D. Aidonis, E. Iakovou, M. Thymianidis, and D. Tzetzis (2015). "A methodological framework for the inclusion of modern additive manufacturing into the production portfolio of a focused factory". In: *Journal of Manufacturing Systems* 37, pp. 328–339. DOI: 10.1016/j.jmsy.2014.07.014.
- Additive Industries (2018). *Additive Industries - MetalFAB1 brochure*. URL: <https://additiveindustries.com/systems/metalfab1>.
- Airbus (2016). *Innovative 3D printing solutions are "taking shape" within Airbus*. URL: <http://www.airbus.com/newsevents/news-events-single/detail/innovative-3d-printing-solutions-are-taking-shape-within-airbus/>.
- Alberth, Stephan (2008). "Forecasting technology costs via the experience curve — Myth or magic?" In: *Technological Forecasting and Social Change* 75.7, pp. 952–983. DOI: 10.1016/j.techfore.2007.09.003.
- Alfredsson, Patrik and Jos Verrijdt (1999). "Modeling Emergency Supply Flexibility in a Two-Echelon Inventory System". In: *Management Science* 45.10, pp. 1416–1431. DOI: 10.1287/mnsc.45.10.1416.
- Allon, Gad and Jan A. Van Mieghem (2010). "Global Dual Sourcing: Tailored Base-Surge Allocation to Near- and Offshore Production". In: *Management Science* 56.1, pp. 110–124. DOI: 10.1287/mnsc.1090.1099.
- Alvarez, E.M., M.C. van der Heijden, and W.H.M. Zijm (2013). "The selective use of emergency shipments for service-contract differentiation". In: *International Journal of Production Economics* 143.2, pp. 518–526. DOI: 10.1016/j.ijpe.2012.02.019.
- AM Sub - Platform (2014). *Strategic Research Agenda*. URL: <http://www.rm-platform.com/linkdoc/AM%20SRA%20-%20February%202014.pdf>.
- Andersson, Olov, Andreas Graichen, Hakan Brodin, and Vladimir Navrotsky (2016).

- “Developing Additive Manufacturing Technology for Burner Repair”. In: *Volume 6: Ceramics; Controls, Diagnostics and Instrumentation; Education; Manufacturing Materials and Metallurgy*. ASME, V006T21A003. DOI: 10.1115/GT2016-56594.
- Arts, Joachim, Marcel van Vuuren, and G.P. Kiesmüller (2011). “Efficient optimization of the dual-index policy using Markov chains”. In: *IIE Transactions* 43.8, pp. 604–620. DOI: 10.1080/0740817X.2010.550908.
- Axsäter, Sven (2007). “A heuristic for triggering emergency orders in an inventory system”. In: *European Journal of Operational Research* 176.2, pp. 880–891. DOI: 10.1016/j.ejor.2005.09.002.
- (2014). “An improved decision rule for emergency replenishments”. In: *International Journal of Production Economics* 157.1, pp. 313–317. DOI: 10.1016/j.ijpe.2014.01.021.
- Barankin, E. W. (1961). “A delivery-lag inventory model with an emergency provision (the single-period case)”. In: *Naval Research Logistics Quarterly* 8.3, pp. 285–311. DOI: 10.1002/nav.3800080310.
- Barz, Andreas, Tobias Buer, and Hans-Dietrich Haasis (2016). “Quantifying the effects of additive manufacturing on supply networks by means of a facility location-allocation model”. In: *Logistics Research* 9.1, p. 13. DOI: 10.1007/s12159-016-0140-0.
- Basten, R. J.I. (2010). “Designing logistics support systems: level of repair analysis and spare parts inventories”. PhD thesis. Enschede, The Netherlands: University of Twente, p. 173. DOI: 10.3990/1.9789036529679.
- Basten, R. J.I. and Geert-Jan van Houtum (2014). *System-oriented inventory models for spare parts*. DOI: 10.1016/j.sorms.2014.05.002.
- Baumers, Martin, Phill Dickens, Chris Tuck, and Richard Hague (2016). “The cost of additive manufacturing: machine productivity, economies of scale and technology-push”. In: *Technological Forecasting and Social Change* 102, pp. 193–201. DOI: 10.1016/J.TECHFORE.2015.02.015.
- Bechthold, Laura, Veronika Fischer, Andre Hainzmaier, Daniel Hugenroth, Kristina Kroth, Benedikt Römer, Edyta Sikorska, and Vincent Sitzmann (2015). *3D Printing A Qualitative Assessment of Applications , Recent Trends and the Technology ’ s Future Potential*. Tech. rep. 17, p. 119. URL: http://www.e-fi.de/fileadmin/Innovationsstudien_2015/StuDIS_17_2015.pdf.
- Behfard, S., M.C. van der Heijden, A. Al Hanbali, and W.H.M. Zijm (2015). “Last time buy and repair decisions for spare parts”. In: *European Journal of Operational Research* 244.2, pp. 498–510. DOI: 10.1016/J.EJOR.2015.02.003.
- Bertsekas, DP (2012). *Dynamic programming and optimal control*. 4th ed. Vol. II. Belmont, MA.: Athena Scientific, p. 712.
- Campbell, Thomas, Christopher Williams, Olga Ivanova, and Banning Garrett (2012). *Could 3D Printing Change the World?* Tech. rep. Atlantic Council, pp. 3–7.
- Cattani, Kyle D. and Gilvan C. Souza (2003). “Good buy? Delaying end-of-life purchases”. In: *European Journal of Operational Research* 146.1, pp. 216–228. DOI: 10.1016/S0377-2217(02)00212-6.
- Chopra, Sunil and Peter Meindl (2016). *Supply chain management : strategy, planning, and operation*. 6th ed. Harlow: Pearson Education, p. 516.

-
- Cohen, Morris A., Narendra Agrawal, and Vipul Agrawal (2006). "Winning in the Aftermarket". In: *Harvard Business Review* May, pp. 129–138. URL: <https://hbr.org/2006/05/winning-in-the-aftermarket>.
- Craeghs, Tom, Stijn Clijsters, Jean-Pierre Kruth, Florian Bechmann, and Marie-Christin Ebert (2012). "Detection of Process Failures in Layerwise Laser Melting with Optical Process Monitoring". In: *Physics Procedia* 39.1, pp. 753–759. DOI: 10.1016/j.phpro.2012.10.097.
- Cunningham, Ross, Andrea Nicolas, John Madsen, Eric Fodran, Elias Anagnostou, Michael D Sangid, and Anthony D Rollett (2017). "Analyzing the effects of powder and post-processing on porosity and properties of electron beam melted Ti-6Al-4V". In: *Materials Research Letters* 5.7, pp. 516–525. DOI: 10.1080/21663831.2017.1340911.
- D'Aveni, Richard (2015). "The 3-D Printing Revolution". In: *Harvard Business Review* May, pp. 40–48. URL: <https://hbr.org/2015/05/the-3-d-printing-revolution>.
- Dekker, Rommert, Çerağ Pınç, Rob Zuidwijk, and Muhammad Naiman Jalil (2013). "On the use of installed base information for spare parts logistics: A review of ideas and industry practice". In: *International Journal of Production Economics* 143.2, pp. 536–545. DOI: 10.1016/j.ijpe.2011.11.025.
- Elmakis, David and Anatoly Lisnianski (2006). "Life cycle cost analysis: Actual problem in industrial management". In: *Journal of Business Economics and Management* 7.1, pp. 5–8. DOI: 10.1080/16111699.2006.9636115.
- EOS (2015). *EOS Industrial 3D printing - Process, method and benefits*. URL: https://www.eos.info/additive_manufacturing/for_technology_interested.
- Everton, Sarah K, Matthias Hirsch, Petros Stravroulakis, Richard K Leach, and Adam T Clare (2016). "Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing". In: *Materials & Design* 95.6, pp. 431–445. DOI: 10.1016/j.matdes.2016.01.099.
- Fleiss, Joseph L., Bruce Levin, and Myunghee Cho. Paik (2003). *Statistical methods for rates and proportions*. Ed. by Joseph L. Fleiss, Bruce Levin, and Myunghee Cho. Paik. 3rd. New York, NY: J. Wiley, p. 760.
- Ford (2017). *Ford Tests Large-Scale 3D Printing with Light-Weighting and Personalization in Mind*. URL: <https://media.ford.com/content/fordmedia/fna/us/en/news/2017/03/06/ford-tests-large-scale-3d-printing.html>.
- Frazier, William E (2014). "Metal Additive Manufacturing: A Review". In: *Journal of Materials Engineering and Performance* 23.6, pp. 1917–1928. DOI: 10.1007/s11665-014-0958-z.
- GE Aviation (2015). *3D Printing Aircraft Parts GE Global Research*. URL: <http://www.geglobalresearch.com/innovation/3d-printing-creates-new-parts-aircraft-engines>.
- Gibson, Ian, David W Rosen, and Brent Stucker (2010). *Additive Manufacturing Technologies*. Boston, MA: Springer US, pp. 437–446. DOI: 10.1007/978-1-4419-1120-9.
- Graham, Ronald L (1972). "An efficient algorithm for determining the convex hull of a finite planar set". In: *Information processing letters* 1.4, pp. 132–133.
- Graves, Stephen C (1985). "A Multi-Echelon Inventory Model for a Repairable Item with One-for-One Replenishment". In: *Management Science* 31.10, pp. 1247–1256. DOI:

- 10.1287/mnsc.31.10.1247.
- Hopkinson, N and P Dicknes (2003). “Analysis of rapid manufacturing—using layer manufacturing processes for production”. In: *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* 217.1, pp. 31–39. DOI: 10.1243/095440603762554596.
- Hu, Qiwei, John E Boylan, Huijing Chen, and Ashraf Labib (2018). “OR in spare parts management: A review”. In: *European Journal of Operational Research* 266.2, pp. 395–414. DOI: 10.1016/j.ejor.2017.07.058.
- Huang, Shuo, Sven Axsäter, Yifan Dou, and Jian Chen (2011). “A real-time decision rule for an inventory system with committed service time and emergency orders”. In: *European Journal of Operational Research* 215.1, pp. 70–79. DOI: 10.1016/j.ejor.2011.05.029.
- Inderfurth, Karl and Kampan Mukherjee (2008). “Decision support for spare parts acquisition in post product life cycle”. In: *Central European Journal of Operations Research* 16.1, pp. 17–42. DOI: 10.1007/s10100-007-0041-z.
- Johnson, Michael and Randolph Kirchain (2009). “Quantifying the effects of parts consolidation and development costs on material selection decisions: A process-based costing approach”. In: *International Journal of Production Economics* 119.1, pp. 174–186. DOI: 10.1016/j.ijpe.2009.02.003.
- Kalpakjian, Serope and Steven R. Schmid (2014). *Manufacturing Engineering and Technology*. 7th ed. New Jersey, USA: Pearson, p. 1216.
- Khajavi, Siavash H., Jouni Partanen, and Jan Holmström (2014). “Additive manufacturing in the spare parts supply chain”. In: *Computers in Industry* 65.1.
- Kiesmüller, G.P. (2003). “A new approach for controlling a hybrid stochastic manufacturing/remufacturing system with inventories and different leadtimes”. In: *European Journal of Operational Research* 147.1, pp. 62–71. DOI: 10.1016/S0377-2217(02)00351-X.
- Kleber, Rainer, Tobias Schulz, and Guido Voigt (2012). “Dynamic buy-back for product recovery in end-of-life spare parts procurement”. In: *International Journal of Production Research* 50.6, pp. 1476–1488. DOI: 10.1080/00207543.2011.560202.
- Knofius, N., M.C. van der Heijden, M.R.K. Mes, and W.H.M. Zijm (2018a). “Large-Scale Dual Sourcing Problems with an Additive Manufacturing Supply Mode”. In: *Working Paper*.
- Knofius, N., M.C. van der Heijden, A. Slepchenko, and W.H.M. Zijm (2017). “Improving effectiveness of spare part supply by additive manufacturing as dual sourcing option.” In: *Beta Working Paper series 530* 530.May.
- Knofius, N., M.C. van der Heijden, and W.H.M. Zijm (2016). “Selecting parts for additive manufacturing in service logistics”. In: *Journal of Manufacturing Technology Management* 27.7, pp. 915–931. DOI: 10.1108/JMTM-02-2016-0025.
- (2018b). “Consolidating spare parts for asset maintenance with additive manufacturing”. In: *International Journal of Production Economics*. DOI: 10.1016/j.ijpe.2018.11.007.
- (2018c). “Moving to additive manufacturing for spare parts supply”. In: *Working Paper*.
- Kornbrot, Diana (2005). “Pearson Product Moment Correlation”. In: *Encyclopedia of Statistics in Behavioral Science*. Ed. by Brian S. Everitt and David C. Howell. Chichester: John Wiley & Sons, Ltd. DOI: 10.1002/0470013192.

-
- Kranenburg, Bram (2006). "Spare parts inventory control under system availability constraints". PhD thesis. DOI: 10.6100/IR616052.
- Landis, J. Richard and Gary G. Koch (1977). "The Measurement of Observer Agreement for Categorical Data". In: *Biometrics* 33.1, p. 159. DOI: 10.2307/2529310.
- Law, Averill M. (2007). *Simulation modeling and analysis*. 4th ed. Boston, MA: McGraw-Hill, p. 800.
- Li, Xishu, Rommert Dekker, Christiaan Heij, and Mustafa Hekimoğlu (2016). "Assessing End-Of-Supply Risk of Spare Parts Using the Proportional Hazard Model". In: *Decision Sciences* 47.2, pp. 373–394. DOI: 10.1111/deci.12192.
- Li, Yao, Guozhu Jia, Yang Cheng, and Yuchen Hu (2017). "Additive manufacturing technology in spare parts supply chain: a comparative study". In: *International Journal of Production Research* 55.5, pp. 1498–1515. DOI: 10.1080/00207543.2016.1231433.
- Lindemann, Christian, Thomas Reiher, Ulrich Jahnke, and Rainer Koch (2015). "Towards a sustainable and economic selection of part candidates for additive manufacturing". In: *Rapid Prototyping Journal* 21.2, pp. 216–227. DOI: 10.1108/RPJ-12-2014-0179.
- Lindemann, Christiane, Uwe Jahnke, Matthias Moi, and Richard Koch (2012). "Analyzing Product Lifecycle Costs for a Better Understanding of Cost Drivers in Additive Manufacturing". In: *Proceedings of International Solid Freeform Fabrication Symposium 2012*. Austin: Solid Freeform Fabrication Symposium, pp. 177–188.
- (2013). "Impact and Influence Factors of Additive Manufacturing on Product Lifecycle Costs". In: *Proceedings of International Solid Freeform Fabrication Symposium 2013*. Austin: Solid Freeform Fabrication Symposium, pp. 998–1009.
- Liu, Peng, Samuel H Huang, Abhiram Mokasdar, Heng Zhou, and Liang Hou (2014). "The impact of additive manufacturing in the aircraft spare parts supply chain: supply chain operation reference (scor) model based analysis". In: *Production Planning & Control* 25.13-14, pp. 1169–1181. DOI: 10.1080/09537287.2013.808835.
- Löber, Lukas, Christoph Flache, Romy Petters, Uta Kühn, and Jürgen Eckert (2013). "Comparison of different post processing technologies for SLM generated 316l steel parts". In: *Rapid Prototyping Journal* 19.3, pp. 173–179. DOI: 10.1108/13552541311312166.
- Magee, C.L., S Basnet, J.L. Funk, and C.L. Benson (2016). "Quantitative empirical trends in technical performance". In: *Technological Forecasting and Social Change* 104, pp. 237–246. DOI: 10.1016/j.techfore.2015.12.011.
- McCue, T. J. (2015). "3D Printing Is Changing the Way We Think". In: *Harvard Business Review* July, pp. 2–5. URL: <https://hbr.org/2015/07/3d-printing-is-changing-the-way-we-think>.
- McLearn, Luke J. (2015). "Additive manufacturing in the Marine Corps". PhD thesis. Naval Postgraduate School. URL: <https://calhoun.nps.edu/handle/10945/45903>.
- Mettas, A (2000). "Reliability allocation and optimization for complex systems". In: *Annual Reliability and Maintainability Symposium. 2000 Proceedings. International Symposium on Product Quality and Integrity (Cat. No.00CH37055)*. IEEE, pp. 216–221. DOI: 10.1109/RAMS.2000.816310.
- Minner, Stefan (2003). "Multiple-supplier inventory models in supply chain management: A review". In: *International Journal of Production Economics* 81-82, pp. 265–279. DOI:

- 10.1016/S0925-5273(02)00288-8.
- Moinzadeh, Kamran and Charles P Schmidt (1991). “An (S–1, S) Inventory System with Emergency Orders”. In: *Operations Research* 39.2, pp. 308–321. DOI: 10.1287/opre.39.2.308.
- Montero, Joaquin, Kristin Paetzold, Matthias Bleckmann, and Jens Holtmannspoetter (2018). “Re-design and re-manufacturing of discontinued spare parts implementing additive manufacturing in the military field”. In: International design conference - design 2018, pp. 1269–1278. DOI: 10.21278/idc.2018.0444.
- Muckstadt, John A (1973). “A Model for a Multi-Item, Multi-Echelon, Multi-Indenture Inventory System”. In: *Management Science* 20.4-part-i, pp. 472–481. DOI: 10.1287/mnsc.20.4.472.
- (2005). *Analysis and Algorithms for Service Parts Supply Chains*. Springer Series in Operations Research and Financial Engineering. New York, NY: Springer New York. DOI: 10.1007/b138879.
- Nagy, Béla, J. Doyne Farmer, Quan M. Bui, and Jessika E. Trancik (2013). “Statistical Basis for Predicting Technological Progress”. In: *PLoS ONE* 8.2. Ed. by Luís A. Nunes Amaral, e52669. DOI: 10.1371/journal.pone.0052669.
- Öner, K.B., R. Franssen, G.P. Kiesmüller, and Geert-Jan van Houtum (2007). “Life cycle costs measurement of complex systems manufactured by an engineer-to-order company”. In: *The 17th International Conference on Flexible Automation and Intelligent Manufacturing*, pp. 589–596.
- Öner, K.B., G.P. Kiesmüller, and Geert-Jan van Houtum (2010). “Optimization of component reliability in the design phase of capital goods”. In: *European Journal of Operational Research* 205.3, pp. 615–624. DOI: 10.1016/j.ejor.2010.01.030.
- Pérès, François and Daniel Noyes (2006). “Envisioning e-logistics developments: Making spare parts in situ and on demand”. In: *Computers in Industry* 57.6, pp. 490–503. DOI: 10.1016/j.compind.2006.02.010.
- Powell, Warren B. (2011). *Approximate dynamic programming: solving the curses of dimensionality*. 2nd ed. New Jersey, USA: Wiley, p. 656.
- Puterman, Martin L. (2014). *Markov Decision Processes*. Ed. by Martin L. Puterman. Wiley Series in Probability and Statistics. Hoboken, NJ, USA: John Wiley & Sons, Inc., p. 684.
- Roland Berger (2013). *Additive manufacturing: a game changer for the manufacturing industry?* URL: https://www.rolandberger.com/en/Publications/pub_additive_manufacturing_2013.html.
- Saaty, Thomas L. (1980). *The analytic hierarchy process : planning, priority setting, resource allocation*. Pittsburgh: McGraw-Hill International Book Co, p. 287.
- (2008). “Decision making with the analytic hierarchy process”. In: *International Journal of Services Sciences* 1.1, pp. 83–98. DOI: 10.1504/IJSSci.2008.01759.
- Sasson, Amir and John Chandler Johnson (2016). “The 3D printing order: variability, supercenters and supply chain reconfigurations”. In: *International Journal of Physical Distribution & Logistics Management* 46.1, pp. 82–94. DOI: 10.1108/IJPDLM-10-2015-0257.

-
- Savastano, Marco, Carlo Amendola, Fabrizio D'Ascenzo, and Enrico Massaroni (2016). "3-D Printing in the Spare Parts Supply Chain: An Explorative Study in the Automotive Industry". In: *Lecture Notes in Information Systems and Organisation*. Ed. by Leonardo Caporarello, Fabrizio Cesaroni, Raphael Giesecke, and Michele Missikoff. Vol. 18. Cham: Springer International Publishing, pp. 153–170. DOI: 10.1007/978-3-319-40265-9{_}_}11.
- Scheller-Wolf, Alan, Senthil Veeraraghavan, and Geert-Jan van Houtum (2006). *Inventory policies with expedited ordering: Single index policies*. Philadelphia.
- Schröder, Malte, Björn Falk, and Robert Schmitt (2015). "Evaluation of Cost Structures of Additive Manufacturing Processes Using a New Business Model". In: *Procedia CIRP* 30, pp. 311–316. DOI: 10.1016/J.PROCIR.2015.02.144.
- Sherbrooke, Craig C (1968). "Metric: A Multi-Echelon Technique for Recoverable Item Control". In: *Operations Research* 16.1, pp. 122–141. DOI: 10.1287/opre.16.1.122.
- (1971). "An Evaluator for the Number of Operationally Ready Aircraft in a Multilevel Supply System". In: *Operations Research* 19.3, pp. 618–635. DOI: 10.1287/opre.19.3.618.
- (1986). "VARI-METRIC : Improved Approximations for Multi-Indenture, Multi-Echelon Availability Models". In: *Operations Research* 34.2, pp. 311–319. DOI: 10.1287/opre.34.2.311.
- (2004). *Optimal Inventory Modeling of Systems*. Vol. 72. International Series in Operations Research & Management Science. Boston, MA: Springer US. DOI: 10.1007/b109856.
- Shouche, Surabhi, Richard A Wysk, Russell E King, and Ola L.A. Harrysson (2016). "Supply Chain Operations Reference model for U.S. based powder bed metal additive manufacturing processes". In: *2016 Winter Simulation Conference (WSC)*. IEEE, pp. 1158–1169. DOI: 10.1109/WSC.2016.7822173.
- Siemens (2014a). *Additive Manufacturing: Facts and Forecasts*. URL: <https://www.siemens.com/innovation/en/home/pictures-of-the-future/industry-and-automation/Additive-manufacturing-facts-and-forecasts.html>.
- (2014b). "Using AM for gas turbine repair". In: *Metal Powder Report* 69.6, pp. 36–37. DOI: 10.1016/S0026-0657(14)70278-4.
- Simkin, Z. and A. Wang (2014). "Cost-benefit Analyses". In: *Wohlers Report 2014*. Ed. by Terry T. Wohlers and Tim Caffrey. Fort Collins: Wohlers Associates. Chap. 6, pp. 178–189.
- Sirichakwal, Issariya and Brett Conner (2016). "Implications of Additive Manufacturing for Spare Parts Inventory". In: *3D Printing and Additive Manufacturing* 3.1, pp. 56–63. DOI: 10.1089/3dp.2015.0035.
- Sleptchenko, A., M.C. van der Heijden, and A. van Harten (2003). "Trade-off between inventory and repair capacity in spare part networks". In: *Journal of the Operational Research Society* 54.3, pp. 263–272. DOI: 10.1057/palgrave.jors.2601511.
- Sleptchenko, A. and M.E. Johnson (2015). "Maintaining Secure and Reliable Distributed Control Systems". In: *INFORMS Journal on Computing* 27.1, pp. 103–117. DOI: 10.1287/ijoc.2014.0613.
- Sleptchenko, A. and M.C. Van der Heijden (2016). "Joint optimization of redundancy level and spare part inventories". In: *Reliability Engineering & System Safety* 153, pp. 64–74. DOI: 10.1016/j.res.2016.04.006.

- Song, Jing-Sheng, Li Xiao, Hanqin Zhang, and Paul Zipkin (2017). “Optimal Policies for a Dual-Sourcing Inventory Problem with Endogenous Stochastic Lead Times”. In: *Operations Research* 65.2, pp. 379–395. DOI: 10.1287/opre.2016.1557.
- Song, Jing-Sheng and Yue Zhang (2016). *Stock or Print? Impact of 3D Printing on Spare Parts Logistics*. DOI: 10.2139/ssrn.2884459.
- Song, Jing-Sheng and Paul Zipkin (2009). “Inventories with Multiple Supply Sources and Networks of Queues with Overflow Bypasses”. In: *Management Science* 55.3, pp. 362–372. DOI: 10.1287/mnsc.1080.0941.
- Stein, Walter (2012). *Spare Parts Planning at ASML*. Ed. by A. Van Weele and Geert-Jan Van Houtum. The European Supply Chain Forum, pp. 1–14.
- Tapia, Gustavo and Alaa Elwany (2014). “A Review on Process Monitoring and Control in Metal-Based Additive Manufacturing”. In: *Journal of Manufacturing Science and Engineering* 136.6. DOI: 10.1115/1.4028540.
- TechCrunch (2016). *NASA is sending a 3D printer to space that you can use*. URL: <https://techcrunch.com/2016/03/19/nasa-is-sending-a-3d-printer-to-space-that-you-can-use/?guccounter=1>.
- Thomas, Douglas S (2016). “Costs, benefits, and adoption of additive manufacturing: a supply chain perspective”. In: *The International Journal of Advanced Manufacturing Technology* 85.5-8, pp. 1857–1876. DOI: 10.1007/s00170-015-7973-6.
- Thomas, Douglas S and Stanley W Gilbert (2014). *Costs and Cost Effectiveness of Additive Manufacturing*. Tech. rep. 6. Gaithersburg, MD: National Institute of Standards and Technology, pp. 1–89. DOI: 10.6028/NIST.SP.1176.
- Tibshirani, Robert (1996). “Regression Shrinkage and Selection via the Lasso”. In: *Journal of the Royal Statistical Society, Series B (Methodological)* 58.1, pp. 267–288.
- Tijms, Henk C. (2003). *A First Course in Stochastic Models*. Chichester, UK: John Wiley & Sons, Ltd, pp. 1–491. DOI: 10.1002/047001363X. URL: <http://doi.wiley.com/10.1002/047001363X>.
- Törnqvist, Leo, Pentti Vartia, and Yrjö O Vartia (1985). “How Should Relative Changes be Measured?” In: *The American Statistician* 39.1, pp. 43–46. DOI: 10.1080/00031305.1985.10479385.
- Tsai, Chih-Yang (2017). “The impact of cost structure on supply chain cash flow risk”. In: *International Journal of Production Research* 55.22, pp. 6624–6637. DOI: 10.1080/00207543.2017.1330568.
- UPS (2016). *UPS On-Demand 3D Printing Network Expanded To Asia*. URL: pressroom.ups.com/pressroom/ContentDetailsViewer.page?ConceptType=PressReleases&id=1474287337483-522.
- Van der Heijden, M.C., E.M. Alvarez, and J.M.J. Schutten (2013). “Inventory reduction in spare part networks by selective throughput time reduction”. In: *International Journal of Production Economics* 143.2, pp. 509–517. DOI: 10.1016/j.ijpe.2012.03.020.
- Van Houtum, Geert-Jan and Bram Kranenburg (2015). *Spare Parts Inventory Control under System Availability Constraints*. Vol. 227. International Series in Operations Research & Management Science. Boston, MA: Springer US. DOI: 10.1007/978-1-4899-7609-3.
- Van Houtum, Geert-Jan and W.H.M. Zijm (2000). “On the relationship between cost and

- service models for general inventory systems”. In: *Statistica Neerlandica* 54.2, pp. 127–147. DOI: 10.1111/1467-9574.00132.
- Veeraraghavan, Senthil and Alan Scheller-Wolf (2008). “Now or Later: A Simple Policy for Effective Dual Sourcing in Capacitated Systems”. In: *Operations Research* 56.4, pp. 850–864. DOI: 10.1287/opre.1080.0552.
- Walter, Manfred, Jan Holmström, and Hannu Yrjölä (2004). “Rapid manufacturing and its impact on supply chain management”. In: *Proceedings of the Logistics Research Network Annual Conference*. November, p. 12.
- Weller, Christian, Robin Kleer, and Frank T. Piller (2015). “Economic implications of 3D printing: Market structure models in light of additive manufacturing revisited”. In: *International Journal of Production Economics* 164, pp. 43–56. DOI: 10.1016/j.ijpe.2015.02.020.
- Westerweel, Bram, Rob J.I. Basten, and Geert-Jan van Houtum (2018a). “Printing Spare Parts at Remote Locations: Fulfilling the Promise of Additive Manufacturing”. Eindhoven. URL: https://www.researchgate.net/publication/324562056_Printing_Spare_Parts_at_Remote_Locations_Fulfilling_the_Promise_of_Additive_Manufacturing.
- (2018b). “Traditional or Additive Manufacturing? Assessing Component Design Options through Lifecycle Cost Analysis”. In: *European Journal of Operational Research* 270.2, pp. 570–585. DOI: 10.1016/j.ejor.2018.04.015. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0377221718303187>.
- Whittmore, A S and S C Saunders (1977). “Optimal Inventory Under Stochastic Demand with Two Supply Options”. In: *SIAM Journal on Applied Mathematics* 32.2, pp. 293–305. DOI: 10.1137/0132023.
- Wits, Wessel W, J Roberto Reyes García, and Juan M Jauregui Becker (2016). “How Additive Manufacturing Enables more Sustainable End-user Maintenance, Repair and Overhaul (MRO) Strategies”. In: *Procedia CIRP* 40, pp. 693–698. DOI: 10.1016/j.procir.2016.01.156.
- Wohlers Report (2014). *3D Printing and Additive Manufacturing State of the Industry*. Ed. by Terry T. Wohlers and Tim Caffrey. Fort Collins: Wohlers Associates.
- Yang, Sheng, Yunlong Tang, and Yaoyao Fiona Zhao (2015). “A new part consolidation method to embrace the design freedom of additive manufacturing”. In: *Journal of Manufacturing Processes* 20.3, pp. 444–449. DOI: 10.1016/j.jmapro.2015.06.024.
- Zhou, Sean X and Chaolin Yang (2016). “Continuous-Review (R, nQ) Policies for Inventory Systems with Dual Delivery Modes”. In: *Operations Research* 64.6, pp. 1302–1319. DOI: 10.1287/opre.2016.1538.
- Zijm, W.H.M., N. Knofius, and M.C. van der Heijden (2019). “Additive manufacturing and its impact on the supply chain”. In: *Operations, Logistics and Supply Chain Management*. Ed. by W.H.M. Zijm, M. Klumpp, A. Regattieri, and S. Heragu. SPRINGER.
- Zopounidis, Constantin and Michael Doumpos (2002). “Multicriteria classification and sorting methods: A literature review”. In: *European Journal of Operational Research* 138.2, pp. 229–246. DOI: 10.1016/S0377-2217(01)00243-0.

Summary

Additive Manufacturing (AM, also known as 3D printing) is developing into a powerful complement to more conventional manufacturing (CM) methods. In comparison to CM methods such as milling, drilling, casting and forging, AM technologies build complete parts by adding materials layer upon layer without using any dedicated tooling.

The resulting ability to produce complex structures without lengthy and expensive setup procedures could turn out particularly valuable for the low-volume spare parts business. Short AM lead times are likely to significantly improve the balance between spare parts inventory investment and system downtime. Generic AM processes could relax the dependence on suppliers and therefore decrease risks and costs associated with supply disruptions. Ultimately, AM could even enable the implementation of a decentralized production concept that holds the promise of increased supply chain responsiveness at low costs.

However, it is necessary to deconstruct these concepts and to separate the hype from reality to leverage the potentials of AM technology in after-sales service supply chains. In this dissertation, we aim to contribute to this undertaking by offering a scientific perspective on how and to what extent after-sales service supply chains can benefit from AM technology. To that end, we develop and apply techniques from the field of Operations Research to learn from the various case studies that were conducted at different organizations throughout this research.

After an introduction to AM technology, its potential application in after-sales service supply chains, and a broad overview of the related literature, we identified the following research objective:

To offer decision support for actors in after-sales service supply chains to identify and understand the value of AM technology for their organization, and to provide quantitative insights into both when and how AM technology may be used or combined with conventional manufacturing methods to improve the efficiency of service logistics.

In the following sections, we summarize how we addressed the research objective. The first study was motivated by the observation that organizations struggle to identify high-impact spare parts for the application of AM. Particularly, it turned out that the current practice to primarily rely on expert judgment was not suitable for the spare parts industry. Large spare parts assortment only allowed the assessment of a very selected subset of the part population while the evaluated spare parts often turned out uninteresting for the application of AM. To overcome these shortcomings, we developed a top-down approach that solely relies on spare parts information that is retrievable from ERP systems. Furthermore, this approach adopts results from an analytical hierarchy process to customize the underlying ranking procedure according to the specific company preferences. Our findings and applications of the method in more than five organizations provided evidence for the value of a top-down approach during the initial screening for interesting business cases. We further recognized that the application of the top-down approach prompts organizations, especially their higher management, to reflect on the goals associated with AM for their spare part operations. This stage proved to be crucial to re-adjust expectations and create a high level of engagement within the organizations.

In the next part, we considered the common practice (or requirement) of altering the design with AM technology. While the effects of design changes appear to be quite well understood from a technological and operational perspective, we found that the implications for service logistics were largely unexplored. Since the consolidation of parts was recognized as the most promising application of AM technology for operations, we explored the associated effects in further detail. Thus, we first discussed how consolidation with AM could affect the total life cycle costs. Subsequently, we used a model from the spare parts optimization literature to quantify the associated effects. It was observed that consolidation using AM often leads to higher costs throughout the entire life cycle than with the conventional part design. The key reason behind this being that replacing a consolidated AM part is often more expensive than repairing a conventional part by replacing a failed subcomponent. Overall, our results for consolidation with AM provides an effective example for the necessity to adopt a broader perspective when judging the effects of design changes with AM. Otherwise, in the long-run, even design improvements may lead to unforeseen effects that could render its application debatable despite possible functionality improvements.

The next part of this thesis was driven by the observation that even if organizations identified high-impact spare parts for the application of AM, it was often unclear to them when or if at all to start the implementation. Particularly, the anticipation of AM unit cost reductions in the near future and the cost associated with setting up a new production method caused skepticism regarding investment in AM being a good choice. To address this problem, we developed a finite horizon stochastic dynamic programming model. Using this model, we were able to show under which conditions and when during the service horizon an investment in AM is advisable. Interestingly, it was observed that postponing the investment in AM is typically not a good decision. For example, the case study conducted during this part of our thesis revealed that the best choice would be an immediate investment in AM technology. Before the company is advised to completely rely on AM technology, however, we observed a transition period of up to nine years. During this transition period, our results suggest that

parallel use of AM and CM technology would be the best approach. In particular, CM would be used for regular supply and AM for emergency shipments. Next to direct cost savings, this so-called dual sourcing approach also holds the indirect benefit of gaining knowledge regarding the application of AM while the conventional supply source is still available.

In the following part of this thesis, we investigate the underlying mechanisms of the dual sourcing approach with AM and CM methods in further detail. In particular, we studied how a production source dependent part quality would influence the sourcing decisions. For us, it seemed evident that a fundamentally different production paradigm, a frequently altered design, and even sometimes the use of different materials would more often than not lead to a different quality for both AM and CM parts. Since this aspect has not been addressed in the dual sourcing literature yet, we proposed an infinite horizon continuous-time Markov Decision Model. A novel aspect of this model is that sourcing decisions influence future demands, which is an immediate consequence of the different part quality. Using a case study in the aerospace industry and numerical experiments, we confirmed that dual sourcing appears to be a suitable approach to profit from AM at an earlier point in time. For example, we found that dual sourcing reduces the negative effects of high AM unit costs or low AM part quality while the benefits of a short AM replenishment lead time are exploitable.

The next part of the thesis was motivated by the computational limitations of the method proposed in the previous part. For example, it would have been computationally cumbersome to study higher demand rates, more complex demand patterns or problems for which spare parts are ordered in batches. To overcome this problem, we investigated the performance of approximate solution procedures. Therefore, we designed a framework to extend dual sourcing heuristics from the literature with a feature to also regard production method dependent part quality. Since the extension is easily implementable, the approach would provide a flexible solution framework to assess a wide range of problems with AM as a dual sourcing option. Using small problem instances, we were able to show that the approach offers high solution quality. However, future research is required to also assess the solution quality for larger problems.

In the final part of this thesis, we outline our findings and suggest a roadmap for the implementation of AM technology in after-sales service supply chains. With the roadmap, we aim to provide guidance for practitioners who want to exploit AM in their spare parts operations. Furthermore, the roadmap offers an overview of the stages of the implementation process in which the research conducted in this thesis would be valuable. We conclude the final part of this thesis with suggestions for future research that could bring forward the discovery and implementation of new solutions with AM technology in after-sales service supply chains.

Acknowledgements

It is the end of 2018, I am speaking more English than German in my daily life, and now I am about to complete my PhD thesis. My younger self would have been surprised about this outcome. And true, it was a twisted path that led me to this stage in my life. With the completion of my PhD thesis, I would like to take the opportunity to take one step back and express my gratitude to all those that supported me during this journey.

Particularly, I am grateful to Henk Zijm and Matthieu van der Heijden. Having you as my promotor and co-promotor was certainly the best combination I could have imagined. Matthieu, your detailed comments and advice always gave me the feeling to have a companion and mentor on my side. No matter when, you took your time to discuss or just listen to my ideas. Henk, your enthusiasm for my research and your determination to delve into every tiny detail of my models was truly inspiring. With your unique ability to break down complex concepts into its essential pieces, you gave my results an even more clear-cut interpretation many times. Thank you to both of you for all the time – often also during weekends – you have spent to help me improving this thesis.

I would not have reached the decision to pursue a PhD without the advice and encouragements of many people. Especially, I would like to thank Frau Siegel, Jürgen Lohmann, Karsten Schmidt, and Leo van der Wegen. I am also grateful for the opportunity to work in such a great environment during my PhD. Both, the colleagues of the IEBIS department and of the SINTAS project offered the support and the friendship to endure also the most stressful times. Without Kaveh Alizadeh, Vincent van Vliet, Aad van den Berg, Jeroen van der Wel, and Jelmar den Boer, many of the presented case studies would not have been feasible. The possibility to ask for advice or to discuss modeling ideas with Andrei Sleptchenko, Arturo Perez Rivera, Bram Westerweel, Engine Topan, Martijn Mes, Rob Basten, Sajjad Rahimi Ghahroodi and Sina Behfard certainly has left its mark in this thesis. Also, I would like to thank Elke van der Veen, Hilde Meijer and Jasper de Graaf for the organization of events, conference visits or just a short chat in the coffee corner.

My family and friends were the anchor stone throughout my PhD. The time in The Netherlands would have been only half as fun without Abhishta, Letizia, Luca, Arturo, Wouter,

Wenyi, Sajjad, Sina, Devrim, Piotr, Tyar, Dimitris, Nestoras, Vova, Xavier, Bahruz, Aytaj, Fabian, Bram, Guido, Vahid and Andreij. The frequent reunions and trips with Achim, Felix, Marissa, Nick, Rico, Simon, Steph and Tim always gave time for reflection and extraordinary experiences. It also would not have been possible without Henrik & Verena, Philip, Jonas, Tim, Ingrid und Magrit who always gave me the feeling of staying connected with the north.

It would be probably clear without mentioning, but certainly I would not have achieved all this without my parents and Laura. Mutti, Paps, Danke für die uneingeschränkte Unterstützung, den Glauben an meine Talente und die unglaubliche Vorbildfunktion während all dieser Jahre. Diese Arbeit ist euch gewidmet. Laura, thank you for the empathy, joyfulness and curiosity. No matter where, you give me the feeling of being home.

About the Author

Nils Knofius was born in Henstedt-Ulzburg, Germany, on January 18th, 1988. In 2008, he completed his high school education at the Dietrich-Bonhoeffer-Gymnasium in Quickborn, Germany and started his four year undergraduate education in industrial engineering and business administration at the private university of applied science, NORDAKADEMIE, in Elmshorn, Germany. His studies were financed and accompanied by a four year traineeship program at the gravure printing division of arvarto Bertelsmann, Prinovis Ltd. & Co. KG, in Ahrensburg, Germany. His primary tasks during the traineeship ranged from production planning and scheduling to the optimization of material flows and setup processes of gravure printing machines. During his Bachelor project, Nils optimized the material flows in the graver printing form production.

After Nils obtained his Bachelor of Science in 2012, he pursued a Master of Science degree in industrial engineering and management at the University of Twente, Enschede, The Netherlands. During his studies, he received funding from the Twente Graduate School in order to prepare for his Ph.D. project. Furthermore, he had the opportunity to study one semester at the National Chiao Tung University, Hsinchu, Taiwan. While conducting his graduation assignment, he worked for nine months at ASML, a supplier of photolithography systems for the semiconductor industry. In his master project, he investigated how installed base size and system life cycle information can be used to improve the final order decision for spare parts. In 2014, Nils defended his master thesis cum laude and started his Ph.D. project in the department of industrial engineering and business information systems at the University of Twente, under the supervision of prof. dr. W.H.M. Zijm and dr. M.C. van der Heijden. The results of his studies are presented in this dissertation.

Scientific Output

Publications

Knofius, N., M.C. van der Heijden, and W.H.M. Zijm (2016). “Selecting parts for additive manufacturing in service logistics”. In: *Journal of Manufacturing Technology Management* 27.7, pp. 915–931. doi: 10.1108/JMTM-02-2016-0025.

Knofius, N., M.C. van der Heijden, and W.H.M. Zijm (2018b). “Consolidating spare parts for asset maintenance with additive manufacturing”. To appear in: *International Journal of Production Economics*. doi: 10.1016/j.ijpe.2018.11.007.

Zijm, W.H.M., N. Knofius, and M.C. van der Heijden (2019). “Additive manufacturing and its impact on the supply chain”. In: *Operations, Logistics and Supply Chain Management*. Ed. by W.H.M. Zijm, M. Klumpp, A. Regattieri, and S. Heragu. Springer.

Working Papers

Knofius, N., M.C. van der Heijden, A. Slepchenko, and W.H.M. Zijm (2017). “Improving effectiveness of spare part supply by additive manufacturing as dual sourcing option.” In: *Beta Working Paper series 530. Under review.*

Knofius, N., M.C. van der Heijden, M.R.K. Mes, and W.H.M. Zijm (2018a). “Large-Scale Dual Sourcing Problems with an Additive Manufacturing Supply Mode”, University of Twente.

Knofius, N., M.C. van der Heijden, and W.H.M. Zijm (2018c). “Moving to additive manufacturing for spare parts supply”, University of Twente. *Under review.*

