Chapter 1

Introduction

1.1 Motivation

Efficient inventory management is a core element of running a successful business in many industries, but it is a complex and challenging task. Holding inventory reduces stockout cost, facilitates smooth operations, and improves service levels and customer experience. However, it ties up capital and goes along with costs for storage, obsolescence, handling, and other. Looking at the largest economy in the world, the cost for carrying inventories were \$454.6 billion in 2019 and accounted for 28% of the total US business logistics costs (Zimmermann et al. 2020). For the European market, numbers are even higher in both absolute and relative terms: in 2018, 47% (converted \$624,2 billion) of the total logistics costs were related to inventory carrying activities (Schwemmer 2019). Until a drop in 2020 related to the COVID19 pandemic, for both Europe and the US, inventories held by companies have constantly increased (Eurostat 2021, Zimmermann et al. 2020). In that regard, the US total carried inventory has grown by as much as 40% since 2010 and accounts for more than 13% of the US gross domestic product (Zimmermann et al. 2020).

As inventories grow, so does their improvement potential. For decades, inventory research has developed models with increasing complexity that capture more details of the real-word situation and find optimal inventory policies for challenges companies are facing. The approaches promise improvements such as lower inventories, reduced cost, or higher service levels.

To achieve these improvements, inventory models build on assumptions. For example, they typically assume direct implementation of new inventory policies and inventory systems in steady state, full adherence to the model's recommendations, and perfect data. In practice, these assumptions seldomly fully hold. When a new inventory policy is implemented, it takes time before the new steady state is reached, decision makers modify and overrule model recommendations, or the data is inaccurate and does not reflect reality.

As a consequence, the actual performance of the inventory system deviates from the projected performance and the full potential of the models cannot be exploited. Despite pure financial implications, this can lead to disappointment and disbelieve in the inventory models in general and can jeopardize the successful implementation of state-of-the-art inventory research in practice.

This thesis aims to improve the application of inventory research in practice with different methods of supply chain analytics. We address three main challenges that we have encountered in our work with companies: the introduction of new inventory policies in existing inventory systems, the use of algorithmic advice by human planners, and the accuracy of master data on which inventory models rely.

1.2 Outline

In the following, we present the structure of this thesis. The research presented in Chapters 2 to 5 are independent research projects but share the common goal of improving the application of inventory models in practice. In each chapter, we address an issue related to inventory management that we observed at real-world companies.

Chapter 2 improves the transition of running inventory systems when the inventory policy changes.[1] It was motivated by a company that implemented new base stock levels for its spare part division. A few months after implementation, the company observed a substantial increase in inventory cost. We quickly identified that the reason for the increase was the transient behavior of the system during the transition to the new target state. To address this issue, we model the inventory transition as a finite-horizon optimization problem and determine transient base stock levels for the parts. We solve the problem with a column generation approach and a heuristic that is based on marginal analysis. Using data of the company that motivated our research, we illustrate how the transition can be controlled to quickly improve fill rates without exceeding the initial inventory budget.

Chapter 3 addresses the interaction of human decision makers and algorithmic decision support.[2] Such decision support is omnipresent in many managerial tasks, such as inventory management, and it can substantially improve decision quality. However, in practice, the recommendations of algorithms are oftentimes

[1]This chapter is based on the paper by Haubitz and Thonemann (2021) that was published in Production and Operations Management. The problem definition, model formalization, development and implementation of solution approaches, numerical study, and writing of the paper was done by Christiane Haubitz. Professor Thonemann gave input for the problem definition, modeling, and solution approaches and proofread the paper. Further, it benefited from the comments of two anonymous referees and the editors of Production and Operations Management and from participants' feedback at the *Supply Chain Research Seminar* at the University of Cologne, January 30, 2019, and the *ISIR Summer school 2019* at University KU Leuven, Belgium, August 30, 2019.

[2]This chapter is based on the paper by Lehmann et al. (2020) that was published in the conference proceedings of the International Conference on Information Systems (ICIS) 2020 and was presented at ICIS 2020 on December 14, 2020. It benefited from the comments of three anonymous referees and an associate editor of the conference proceedings. It is joint work with Cedric Lehmann. We split the work in different tasks: problem identification, pre-study, literature review, development of hypotheses, experimental design, programming and execution of the experiment, analysis of the experimental results. To the problem identification, pre-study, and development of hypotheses both Cedric Lehmann and Christiane Haubitz contributed with the same share. The literature review and analysis of the experimental results was mainly done by Cedric Lehmann. The experimental design, and the programming and execution of the experiment was mainly done by Christiane Haubitz. Overall, the work was divided fairly between Cedric Lehmann and Christiane Haubitz. Professor Andreas Fügener and Professor Ulrich Thonemann participated in discussions about the experimental design and gave input for the modeling approach, the design and analysis of the experiments, and the positioning of the paper.

overruled. One reason that is often stated as barrier of successful human-machine collaboration is the lack of algorithm transparency. In this chapter, we analyze the effects of algorithm transparency on the perceived value of algorithmic advice and its resulting utilization for an easy advice-giving algorithm. In a laboratory experiment, we simulate a task that many inventory planners face in practice, that is, demand forecasting. We present algorithmic advice to the participants and only inform the treatment group about the underlying principles of the simple yet optimal advice-giving algorithm. While the explanation increases the understanding of the algorithmic procedure, it reduces the perceived value of the algorithmic advice, its utilization and the participants' performance.

Chapter 4 builds on the insights of Chapter 3 and adds the dimension of algorithm complexity.[3] In a new experimental study, we consider the moderating effect of algorithm complexity on the use of algorithmic advice. We provide advice from algorithms with different degrees of transparency and complexity. Aligned with our results in Chapter 3, we find that increasing the transparency of a simple algorithm reduces the use of advice. However, increasing the transparency of a complex algorithm increases the use of advice.

Chapter 5 addresses the challenge of improving the quality of master data. We focus on external supplier lead times, a type of master data that is oftentimes highly inaccurate and that is critical for inventory planning. We suggest a machine-learning based approach to build a model for lead time prediction. Our approach

[3]This chapter is joint work with Cedric Lehmann. We split the work in different tasks: problem identification, literature review, development of hypotheses, experimental design, programming and execution of the experiment, analysis of the experimental results. To the problem identification and development of hypotheses both Cedric Lehmann and Christiane Haubitz contributed with the same share. The literature review and analysis of the experimental results was mainly done by Christiane Haubitz. The experimental design, and the programming and execution of the experiment was mainly done by Cedric Lehmann. Overall, the work was divided fairly between Cedric Lehmann and Christiane Haubitz. Professor Andreas Fügener and Professor Ulrich Thonemann participated in discussions about the experimental design and gave input for the modeling approach, the design and analysis of the experiments, and the positioning of the paper.

allows to predict both individual order lead times and general planned lead times that can be used, for example, in inventory planning. We test our approach on historical purchase orders of a large German provider of healthcare equipment. Our approach outperforms the currently used planned lead times from the company's enterprise resource planning (ERP) system and other statistically derived values. We demonstrate the practical relevance and monetary impact on an inventory system with a periodic-review base stock policy.

Chapter 6 concludes this thesis. We summarize and critically review our key insights and outline promising directions for future research.

1.3 Contribution

This thesis contributes to the research on inventory control with supply chain analytics and the successful application of inventory research in practice. We apply three different methodologies to address three challenges that we have experienced at companies we have been working with. The impact of all our approaches is validated with real-world data.

In **Chapter 2**, we contribute to the literature of inventory control by considering the transition of a multi-item inventory system from a current state to an optimized state. We formally introduce the problem as a multi-period optimization problem and develop two solution approaches that rely on column generation and on marginal analysis. Compared to upper bounds on the objective function value, both approaches generate solutions that are close to optimal. We demonstrate the value of controlling the inventory transition in an extended case study that is based on data from a global equipment manufacturer. We also analyze under which circumstances controlling the transition is particularly valuable and discuss managerial implications.

In **Chapters 3 and 4**, we contribute to better understanding the interaction of human planners and planning algorithms. More generally, we contribute to the research streams of human-machine interaction and advice taking. We analyze the effects of algorithm transparency on the use of algorithmic advice and contribute to a better understanding of when algorithmic principles should be made transparent and when not. Chapter 3 is the initial study and focuses on simple, easy-to-understand algorithms for which an easy explanation can be provided. In laboratory experiments, we find that the explanation of the algorithm increases understanding of the algorithmic principles, but the higher understanding does not translate to a higher use of advice. In Chapter 4, we add the dimension of algorithm complexity to our analyses. In a new laboratory experiment, we find that the negative effects of algorithm transparency on use of advice diminishes when the algorithm complexity increases. We provide important insights and managerial implications in the area of human-machine collaboration. We show that making advice-giving algorithms transparent is not always beneficial as it bears the risk of disappointing people's expectations, which can backfire and eventually reduce the use of advice.

In **Chapter 5**, we contribute to supporting the process of improving the quality of master data using machine learning algorithms. We focus on external supplier lead times. Our machine-learning based approach allows to predict both individual order lead times and general planned lead times that can be used, for example, in inventory planning. It is particularly valuable for a setting with a large variety of products and suppliers but only few purchase orders per product and supplier. It enables a prediction of external supplier lead times solely based on available past purchase order data without the need for external supplier information. Moreover, we provide a way to predict planned lead times of new products, which is not possible with other statistical methods.

Chapter 2

How to change a running system – Controlling the transition to optimized spare parts inventory policies

Inventory optimization approaches typically optimize steady-state performance, but do not consider the transition of an initial state to the optimized state. In this study, we address this transition. Our research is motivated by a company that implemented an improved inventory policy for its spare parts division. The improved policy suggested new base stock levels for the majority of the parts. For parts with increased base stock levels, inventory increases were realized after the part lead times, but for low-demand parts with decreased base stock levels, inventory reductions were slow. As a result, inventory cost increased over the first months after the new inventory policy had been introduced and exceeded the inventory budget substantially. To avoid such undesirable effects, base stock level changes must be phased in. We consider a multi-item spare parts inventory system, initially operating under an item approach policy that achieves identical fill rates for all parts. Our approach addresses the transition to a superior system approach policy that maximizes the system fill rate. We model the inventory transition as a finite-horizon optimization problem and apply column generation and a marginal analysis heuristic to determine transient base stock levels for all parts. Using data from the company that motivated our research, we illustrate how the transition can be controlled to quickly improve fill rates without exceeding the inventory budget.

2.1 Introduction

The trigger of inventory optimization projects is often a sub-optimal performance of the current inventory system. In spare parts management, for instance, companies can set base stock levels such that a certain fill rate is achieved by every individual part. This approach is referred to as *item approach*. Instead of optimizing each part individually, companies can reduce inventory or increase the system fill rate by considering all parts in the inventory system jointly when making decisions about base stock levels. This approach is referred to as *system approach* (Sherbrooke 2004). For a spare parts inventory system for high-end computer servers, for example, Thonemann et al. (2002) show improvements in inventory investment of up to 25% when applying a system approach instead of an item approach. For a spare parts inventory system at the Royal Netherlands Navy, Rustenburg et al. (2003) demonstrate an increase in spare parts availability by 34 percentage points, while simultaneously reducing the inventory investment by about 10%.

Implementing a system approach requires the adjustment of inventory control policy parameters. For example, base stock levels of inexpensive fast movers are increased and base stock levels of expensive slow movers are reduced. Overall inventory performance is improved once the inventory system has reached its new steady state. However, during the transition to the new steady state, system fill rate or system inventory holding cost can deteriorate temporarily, particularly if lead times are long or demand rates are low.

We experienced this in an inventory optimization project with the service division of a global B2B (business-to-business) equipment manufacturer. The company generates annual multi-billion euros turnover and operates in more than 50 countries. Its service division offers repair and maintenance services for the specialized and expensive equipment. The division wanted to improve its inventory