Data-driven Inventory Optimization

Inauguraldissertation
zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften
der Universität Mannheim

vorgelegt von

Sebastian Müller
Limburg an der Lahn

Dekan: *Joachim Lutz*
Referent: *Prof. Dr. Moritz Fleischmann*
Korreferent: *Prof. Dr. Cornelia Schön*

Tag der mündlichen Prüfung: 27. November 2020 to Oskar and Paul

Acknowledgements

This dissertation would not have been possible without the inspiration, encouragement, and support of many people that have accompanied me during the last five years.

First and foremost, I would like to thank Moritz Fleischmann, my doctoral advisor. I am very grateful for his support and motivating words whenever needed. His very detailed and specific feedback about the work that I presented to him was invaluable, improved my research, and hopefully made me a better researcher. His way of thinking about the relevance of research has had a lasting influence on me. As a boss, his trusting way of leading the Chair has always created a productive and pleasant working environment. I would like to thank Cornelia Schön who was so kind to take the effort of reviewing this dissertation.

I am also grateful to Jakob Huber, from whom I learned a lot about how to apply state-of-the-art machine learning to real-world data. Our discussions have always been productive in the end, even if he sometimes felt that the decision-making process could have been faster. I would also like to thank Heiner Stuckenschmidt for his support.

I was very fortunate to meet many bright and inspiring people during the last few years, especially my fellow doctoral students, Luca Biscaro, Christian von Falkenhausen, Hendrik Guhlich, Hossein Nikpayam, Maryam Nouri Roozbahani, Kilian Seifried, Gerrit Schumacher, Jonas Schwamberger, postdocs Jochen Schlapp, and Danja Sonntag, and secretaries Judith Fuhrmann, Helke Naujok, and Ruth Pfitzmann. It was you who made the long train ride worth it.

Finally, I am greatly indebted to my family. The unconditional support of my parents throughout my life has enabled me to achieve my goals and to write this dissertation. Most importantly, I thank my wife Ronja, for always being there for me, her love, and support. Without her, this dissertation would not have been possible. Concerning this dissertation, I would *not* like to thank my sons, Oskar and Paul. They made my nights short, my days busy, and I had less time and energy to focus on my research. Yet, they make my life full of surprises, happiness, and love. I dedicate this dissertation to them.

Summary

The recent explosion of data availability opens up opportunities for companies to make better decisions. However, it is not clear, in general, how to get from data to a good decision. Exploiting these data for improved decision making requires adequate methodologies. Inventory management decisions are a particularity important set of decision problems for virtually every company that buys, produces, distributes, or sells physical products. In this dissertation, we investigate the question of how to get from data to a good decision in inventory management problems. To this end, we revisit three fundamental inventory management problems, propose new data-driven methodologies, and measure their impact on inventory performance. Chapter II covers the newsvendor problem. To investigate how to exploit the available data, we propose a framework that distinguishes three levels on which data can generate value. Furthermore, we present a novel solution method that integrates the traditionally separate steps of demand estimation and inventory optimization into a single optimization problem. In our empirical analysis with real-world data, we find that data-driven methods outperform traditional approaches in most cases and that the benefit of improved forecasting dominates other potential benefits of data-driven methodologies. Chapter III is concerned with managing inventories for multiple products in a product category. We present a novel data-driven solution approach based on machine learning that integrates the estimation and optimization steps and takes complex substitution effects into account. We evaluate our approach on two real-world datasets. We find that our data-driven approach outperforms the benchmark on the first dataset and performs competitively on the second. Chapter IV focuses on dynamic inventory problems. We propose a novel solution approach that leverages auxiliary data. Our approach divides the problem into multiple single-stage problems using dynamic programming and uses machine learning methods in each stage to improve inventory decisions. In a computational study, we find that our method performs close to the optimal decision and significantly outperforms the benchmark.

Table of Contents

List of Figures

List of Tables

Chapter I

Introduction

The International Data Corporation estimates the total amount of data that is created, captured, or replicated worldwide to grow from 33 zettabytes [1] in 2018 to 175 zettabytes by 2025 (Reinsel et al. 2018). This data is generated e.g. by mobile phones and computers that track user behavior, sensor-equipped machinery that reports its current status and condition, and digitized vehicles that transmit their positions and mileage. Parallel to this development, the processing capabilities of modern computers have also improved significantly and enabled the handling of the emerging *Big Data*.

More and better data open up opportunities for companies to make better decisions. The explosion of data availability and computing power has led to increased interest and adoption of data-driven decision making across industries (Brynjolfsson and McElheran 2016a,b). Exploiting these data for better decision making requires new methodologies and algorithms. One of these methodologies that has attracted a lot of attention in recent years is machine learning. Examples of machine learning applications include approaches that leverage large data sets to predict crime (Kadar and Pletikosa 2018), predict results of football matches (Baboota and Kaur 2019), or predict retail store sales (Kaneko and Yada 2016). Deriving optimal decisions, however, is usually not considered in machine learning. The lack of established methodologies for data-driven decision making is one of the reasons why many companies are still struggling to leverage the available data.

Inventory management decisions are a particularity important set of decision problems for virtually every company that buys, produces, distributes, or sells physical products. Controlling inventories is especially challenging when customer demand is uncertain. In

[1] One zettabyte is equal to one trillion gigabytes.

competitive environments, good service is a key factor for success. Thus, companies attach great importance to product availability to satisfy their existing customers and attract new ones. On the one hand, holding too little inventory leads to stock-outs, unsatisfied customers, and finally to financial disadvantages. On the other hand, holding too much inventory is associated with excessive costs, such as warehousing costs, handling costs, costs of capital, and depreciation. Depending on the company and its business model, the value and therefore also the costs of inventory are substantial. For example, the world's largest chemical company, BASF, holds inventories with a value of over 11.2 billion Euro, corresponding to approximately 13% of total assets (BASF 2020); one of Germany's largest retailers, REWE Group, holds inventories with a value of over 4.1 billion Euro, corresponding to approximately 13% of total assets (REWE 2020). Thus, managing the trade-off between too little and too much inventory and taking appropriate inventory decisions contributes significantly to a company's economic success.

It is at the heart of operations research and management science to provide models, methods, and insights to help make better inventory decisions. However, most traditional inventory management approaches do not provide ways to get from data to decision in this new data-rich world. They either make strong and unrealistic assumptions about the data or concentrate only on a limited part of the data and neglect the remaining information. Thus, new methods are needed that leverage the available data to the fullest.

Motivated by the recent explosion of data availability, the main research question of this dissertation is

How to get from data to good decisions in inventory management?

To answer this question, we revisit three fundamental inventory management problems in Chapters II to IV and investigate the potential benefits of data-driven decision making therein.

Chapter II covers the most fundamental stochastic inventory problem, the newsvendor problem. In the newsvendor problem, the decision-maker has to manage the trade-off between ordering too much and ordering too little of a single product with uncertain demand for a single selling period. Traditionally, the newsvendor problem is solved based on a demand distribution assumption. However, in reality, the true demand distribution is hardly ever known to the decision-maker. Instead, large datasets of point-of-sales data as well as auxiliary feature data are available that enable the use of empirical distributions. Chapter II investigates how to exploit this data for improved decision making.

To address this question, we identify three levels on which data can generate value. The first level is *demand estimation*, where data can help to improve the demand forecast. The second level is *inventory optimization*, where data can replace demand distribution assumptions. The third level is integrated estimation and optimization, where the inventory decision is estimated directly from data. While most traditional methods use the first two levels, we introduce new methods for integrated estimation and optimization based on machine learning and quantile regression that do not require the assumption of a specific demand distribution. We empirically evaluate the impact of data-driven decision making on the three levels with point-of-sales data of a large German bakery chain. We find that data-driven methods outperform traditional approaches in most cases. Furthermore, we find that the benefit of improved forecasting (first level) dominates other potential benefits of data-driven methodologies.

Chapter III is concerned with managing inventories for multiple products in a product category. Managing inventories for these products is particularly challenging due to substitution effects within the category. Substitution effects occur if a customer cannot find his or her preferred product, due to a stock-out, and substitutes for a similar product within the same category. This substitution behavior makes the inventory decisions for individual products interdependent, and the resulting optimization problem notoriously hard. Another difficulty is that, again, the true demand distributions of products are usually unknown to the decision-maker. In this chapter, we investigate the question of how to leverage the available data for the multi-product newsvendor problem under customer substitution, and we measure its value for a real-world problem.

To answer this question, we present a novel solution approach for the multi-product newsvendor problem. Our method integrates the demand estimation step and the inventory optimization step into a single optimization problem. The method is based on modern machine learning techniques that leverage large available datasets, including, data on historical sales, weather, store location, and special days, and are able to take complex substitution effects into account. We empirically evaluate our approach on two real-world datasets of a large German bakery chain. Furthermore, we evaluate the effects of demand estimation accuracy, feature data, and substitution on overall performance. We find that our data-driven approach outperforms the model-based benchmark on the first dataset and performs competitively on the second dataset. Our forecast accuracy analysis reveals that the performance difference is mainly due to the additional data that can be leveraged by our method compared to the benchmark approach. Another benefit of our approach is that it prescribes a decision for any combination of feature data values once it is trained. There is no need to solve the notoriously hard multi-product newsvendor problem in each period.

In Chapter IV, we focus on another fundamental class of inventory management problems, namely dynamic inventory problems. Dynamic inventory problems occur in a large number of industrial, distribution, and service applications. Most models in the literature assume that the demand distribution is stationary and known a priori in these contexts. In reality, this is rarely the case. Instead, the actual demand distribution is unknown and may change over time. The recent explosion of data availability may help to avoid such assumptions. We revisit the classical stochastic dynamic inventory problem of Scarf (1959), where the demand distribution is unknown, and the decision-maker has access to historical demand data and associated feature data instead. The main question that we address in is how to incorporate feature data into the inventory decision.

To answer this question, we combine concepts of classical dynamic inventory management and machine learning to develop a data-driven solution approach that leverages auxiliary data such as weather, location, and calendar data for the stochastic dynamic inventory problem. More specifically, we built the idea of *predictive prescriptions* of Bertsimas and Kallus (2020). Predictive prescriptions are functions that prescribe a decision for a stochastic optimization problem given some feature data. The approach is similar to sample average approximation. However, instead of assigning the same weight to every historical demand realization, these realizations are reweighed according to their importance for the decision, using machine learning. The concept can be applied to general single-stage optimization problems but it is not immediately clear how to transfer it to multi-stage problems. We use the dynamic programming formulation of the dynamic inventory problem to split it into successive single-stage problems and apply the concept of predictive prescriptions to each of these problems. In a computational study, we compare our approach to an extant non-parametric method and to the optimal dynamic programming solution. In our computational, we find that our method results in an optimality gap of only 0.75% to 2.92% whereas the benchmark results an optimality gap of of 6.63% to 7.58%.

Chapter II

A Data-driven Newsvendor Problem: From Data to Decision

with Jakob Huber, Moritz Fleischmann,
and Heiner Stuckenschmidt[1]

Abstract

Retailers that offer perishable items are required to make ordering decisions for hundreds of products on a daily basis. This task is non-trivial because the risk of ordering too much or too little is associated with overstocking costs and unsatisfied customers. The well-known newsvendor model captures the essence of this trade-off. Traditionally, this newsvendor problem is solved based on a demand distribution assumption. However, in reality, the true demand distribution is hardly ever known to the decision maker. Instead, large datasets are available that enable the use of empirical distributions. In this paper, we investigate how to exploit this data for making better decisions. We identify three levels on which data can generate value, and we assess their potential. To this end, we present data-driven solution methods based on Machine Learning and Quantile Regression that do not require the assumption of a specific demand distribution. We provide an empirical evaluation of these methods with point-of-sales data for a large German bakery chain. We find that Machine Learning approaches substantially outperform traditional methods if the dataset is large enough. We also find that the

[1]The research presented in this chapter is based on a paper entitled “A Data-driven Newsvendor Problem: From Data to Decision” coauthored with Jakob Huber, Moritz Fleischmann and Heiner Stuckenschmidt.

benefit of improved forecasting dominates other potential benefits of data-driven solution methods.

2.1. Introduction

Demand uncertainty is a major challenge in supply chain management practice and research. An important remedy for demand risk is the deployment of safety stock. In order to set appropriate stock levels, many inventory models assume a specific demand distribution (Silver et al. 2017). These problems are then solved in a two-step procedure. First, the parameters of a given demand distribution are estimated, and second, an optimization problem based on this distribution is solved. Despite the theoretical insights generated, the distribution assumption is problematic in real-world applications, as the actual demand distribution and its parameters are not known to the decision maker in reality and may even change over time (Scarf 1958).

The growing availability of large datasets ("Big Data") may help overcome this issue and improve the performance of inventory models in real-world situations. Data that are indicative of future demand provide an opportunity to make better-informed decisions. These data include external information that is available through the Internet and data from internal IT systems. While this potential is widely recognized (see e.g. Bertsimas and Kallus (2020)), it is unclear how to best exploit it. Extant literature is rather fragmented in that regard and proposes multiple alternative directions. Our paper intends to contribute to a more wholistic understanding of the potential of data-driven inventory management. To this end, we distinguish three levels on which data can be used to revise the traditional decision process (see Figure 2.1). We discuss how these levels are interrelated, and we quantify their respective impact in a real-life application.

The first level on which data can be exploited is demand estimation. The available data may contain information about future demand that can be extracted by suitable forecasting methods. These methods use historical demand data and other feature data (e.g. weekdays, prices, weather, and product ratings) to estimate future demand. The output of these models is a demand estimate together with historical forecast errors. If additional information can be extracted, the reduced demand risk results in more accurate decisions. Machine Learning (ML) has attracted a great deal of attention in the past decade. ML methods are able to process large datasets and have been successfully

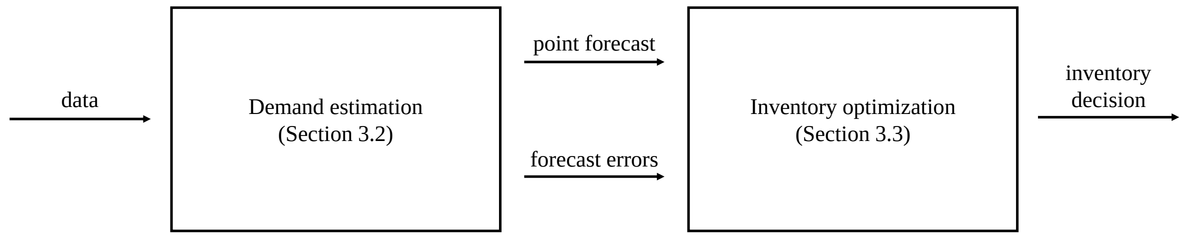

(a) Separate estimation and optimization

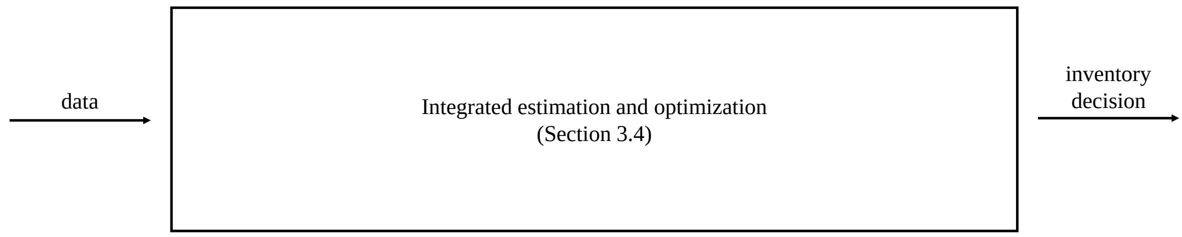

(b) Integrated estimation and optimization

Figure 2.1.: The three levels of data-driven inventory management

applied to numerous forecasting problems (Barrow and Kourentzes 2018, Crone et al. 2011, Carbonneau et al. 2008, Thomassey and Fiordaliso 2006).

On the second level, the inventory decision is optimized based on the demand forecast and the historical forecast errors. To this end, it is necessary to incorporate the remaining uncertainty associated with the forecast. Traditionally, uncertainty is modeled through a demand distribution assumption (Silver et al. 2017). We call this approach *model-based* since it explicitly models a demand distribution. However, this assumption might be misspecified and leads to suboptimal inventory policies (Ban and Rudin 2019). Instead of speculating about a parametric demand distribution, the assumption can be replaced by empirical data that are now available on large scale. This approach is called Sample Average Approximation (SAA) (Kleywegt et al. 2002, Shapiro 2003) and we call it *data-driven* as it does not rely on a distribution assumption.

On the third level, demand estimation and optimization are integrated into a single model that directly predicts the optimal decision from historical demand data and feature data, as depicted in Figure 2.1b (Beutel and Minner 2012, Sachs and Minner 2014, Ban and Rudin 2019, Bertsimas and Kallus 2020). This approach is also *data-driven*, as it does not require the assumption of a demand distribution and works directly with data.

From the existing literature, it is not yet clear whether and under which circumstances data-driven approaches are preferred to model-based approaches. Furthermore, the question of the conditions under which separate or integrated estimation and optimization is superior remains open.

To shed light on these questions, we focus on the newsvendor problem as the basic inventory problem with stochastic demand. We empirically analyze the effects of data-driven approaches on overall costs on the three levels. Moreover, we develop novel data-driven solution methods that combine modern ML approaches with optimization and empirically compare them to well-established methods.

In our approaches, we integrate Artifical Neural Networks (ANNs) and Decision Trees (DTs) into an optimization model. Most previous work on integrated estimation and optimization assumed the inventory decision to be linear in the explanatory features (Beutel and Minner 2012, Sachs and Minner 2014, Ban and Rudin 2019). This assumption poses many restrictions on the underlying functional relationships. We extend this literature by integrating multiple alternative ML methods and optimization in order to avoid these strong assumptions and incorporate unknown seasonality, breaks, thresholds, and other non-linear relationships. Recently, Oroojlooyjadid et al. (2018) and Zhang and Gao (2017) also used ANNs in this context.

We evaluate our solution approaches with real-world data from a large bakery chain in Germany. The company produces and sells a variety of baked goods. It operates a central production facility and over 150 retail stores. Every evening, each store must order products that are delivered the next morning. Reordering during the day is not possible. Most of the goods have a shelf life of only one day. Thus, leftover product at the end of the day is wasted, while stock-outs lead to lost sales and unsatisfied customers.

From an optimization perspective, the problem can be represented by a newsvendor model, and the available point-of-sales data can be used to calculate forecasts. We apply our data-driven methods to the problem and compare their performance to the performance of well-established approaches. To summarize, our key contributions include the following:

- We identify and conceptualize three levels of data-driven approaches in inventory management.
- We investigate the impact of the three levels on overall performance in a newsvendor problem.

- We present novel data-driven solution approaches to the newsvendor problem based on Machine Learning.
- We compare our method to well-established approaches on the three levels and show that data-driven methods outperform their model-based counterparts on our real-world dataset in most cases.

The remainder of this paper is organized as follows. In the next section, we provide an overview of related literature. In Section 2.3, we describe the problem and introduce the methodology, including the data-driven ML approaches. Section 2.4 contains an introduction to the reference models, an empirical evaluation, and a discussion of the results. In Section 2.5, we summarize our findings and outline opportunities for further research.

2.2. Related Literature

Most inventory management textbooks assume that the relevant demand distribution and its parameters are exogenously given and known (Silver et al. 2017). For a review of newsvendor-type problems, see Qin et al. (2011). In this section, we review the literature on inventory problems in which the demand distribution is unknown. More specifically, we focus on Robust Optimization, Sample Average Approximation (SAA), and Quantile Regression (QR).

One approach that needs only partial information on demand distributions is robust optimization (Ben-Tal et al. 2009). Scarf (1958) studies a single period problem in which only the mean and the standard deviation of the demand distribution are known. He then optimized for the maximum minimum (max-min) profit for all distributions with this property. Gallego and Moon (1993) further analyzed and extended it to a setting where reordering is possible. Bertsimas and Thiele (2006) and Perakis and Roels (2008) provide more insights into the structure of robust inventory problems. The main drawback of robust optimization is its limitation to settings with very risk-averse decision makers. For most real-world applications, robust optimization is overly conservative. For our analysis, we focus on methods that minimize expected costs instead of the max-min objective.

A data-driven method with a wider range of applications is Sample Average Approximation (SAA) (Kleywegt et al. 2002, Shapiro 2003). Here, the demand distribution assumptions are replaced by empirical data. Levi et al. (2007) analyze the SAA solution of a newsvendor model and its multi-period extensions. The authors calculate bounds on the number of observations that are needed to achieve similar results compared to the case with full knowledge of the true demand distribution. These bounds are independent of the actual demand distribution. More recently, Levi et al. (2015) showed that the established bound is overly conservative and does not match the accuracy of SAA obtained in simulation studies. Therefore, they develop a tighter bound that is distribution specific. In this paper, we provide empirical support for the good performance of SAA and compare the results of diverse methods.

Instead of using sequential estimation and optimization, integrating both steps into a single optimization model has been suggested (Bertsimas and Kallus 2020). Beutel and Minner (2012) incorporate a linear regression function for demand into their newsvendor model. The authors test their approach on simulated data and actual retail data. The model was later extended to situations with censored demand observations (Sachs and Minner 2014). Ban and Rudin (2019) propose an algorithm that is equivalent to the one in Beutel and Minner (2012), in addition to a kernel optimization method. Furthermore, the authors show several properties of the algorithm and test it with empirical data in a newsvendor-type nurse staffing problem. Oroojlooyjadid et al. (2018) and Zhang and Gao (2017) integrate a neural network into a newsvendor model and compare it to several other approaches from the literature. However, they do not distinguish the effects of estimation, optimization, and integrated estimation and optimization. A drawback of extant research on integrated estimation and optimization is that non-linear relationships between inventory decision and feature data remain understudied. By using ML instead of a linear decision rule, our approaches can detect a priori unknown non-linear relationships between the optimal decision and the input features. Furthermore, we disentangle the effects of the three different levels of data usage highlighted in Figure 2.1.

It is well known that the optimal solution to the standard newsvendor model corresponds with a certain quantile of the demand distribution (Silver et al. 2017). Estimating a certain quantile of a distribution is known as Quantile Regression (QR) in the statistics and ML literature (Koenker 2005). A very general approach to QR is presented by Takeuchi et al. (2006). The authors derive a quadratic programming problem and provide bounds and convergence statements of the estimator. Taylor (2000) use an ANN for QR in order to estimate conditional densities of financial returns. Similarly, Cannon

(2011) describes an implementation of ANNs for QR and gives recommendations on solution approaches with gradient algorithms. More related to our application, Taylor (2007) applies QR to forecast daily supermarket sales. The proposed method can be interpreted as an adaption of exponential smoothing to QR. In the empirical evaluation, the author tests three implementations of the method: one with no regressors, one with a linear trend term, and one with sinusoidal terms to account for seasonality. None of the papers on QR we found uses QR to evaluate the costs of an inventory decision. For our solution approach, we build on the existing literature on QR by integrating ML methods into the optimization model and and evaluate the resulting costs of the newsvendor decision.

The challenge of incorporating demand uncertainty in inventory models without demand distribution assumptions is most recently also discussed by Trapero et al. (2019). They argue that the typical assumption of normal i.i.d. forecast errors should be questioned and suggest using a non-parametric kernel density approach for short lead times. Prak and Teunter (2019) propose a framework for incorporating demand uncertainty in inventory models that mitigates the parameter estimation uncertainty.

To summarize, we empirically evaluate the impact of data-driven approaches on the three levels (1) estimation, (2) optimization, and (3) integrated estimation and optimization. To this end, we extend the literature by proposing novel data-driven approaches to the newsvendor problem that are based on ML and build on the existing knowledge on QR in order to leverage existing big data and computation power for inventory optimization. We also illustrate the connection between QR and integrated estimation and optimization in the newsvendor context. Finally, we empirically compare the data-driven methods to their model-based counterparts and other well-established approaches.

2.3. Methodology

2.3.1. Problem Description

We consider a classical newsvendor problem with an unknown demand distribution: a company sells perishable products over a finite selling season with uncertain demand. The company must choose the number of products to order prior to the selling season. If the order is too high and not all products can be sold, the company bears a cost of c_o for each unit of overage. If the order is too low and more units could have been sold, the company bears costs of c_u for each unit of underage. Thus, the objective is to minimize the total expected costs according to

$$\min_{q\geq 0} \mathbb{E}\left[c_u(D-q)^+ + c_o(q-D)^+\right], \tag{2.1}$$

where q is the order quantity and D is the random demand. The well-known optimal solution to this problem is to choose as the order quantity the quantile of the cumulative demand distribution function F that satisfies

$$q^* = \inf\left\{p : F(p) \geq \frac{c_u}{c_u + c_o}\right\}, \tag{2.2}$$

where $\frac{c_u}{c_u+c_o}$ is the optimal service level. The service level represents the probability of satisfying demand in a given period.

The problem that we address is that in most real-world cases, the actual demand distribution F is unknown. However, historical data $S_n = \{(d_1, \mathbf{x_1}), ..., (d_n, \mathbf{x_n})\}$ are available, where d_t is the demand and $\mathbf{x_t}$ is a vector of covariates or features (e.g. weekday, historical demand, and price) in period t. These data can be leveraged in different ways to reduce demand risk.

In the following sections, we present approaches that use the data on the three levels introduced in Section 2.1. First, we introduce forecasting models based on ML that we use throughout our analysis. Next, we describe a data-driven optimization approach that leverages the empirical distribution of forecast errors. Finally, we present novel data-driven models that integrate ML and the optimization model.

2.3.2. Demand Estimation

If the underlying structure of the demand data is unknown, it is reasonable to consider very general forecasting models. ML methods have been applied to numerous forecasting tasks. Compared to traditional forecasting methods, ML is able to "learn" non-linear relationships between inputs and outputs. The most widely and successfully used methods are Artificial Neural Networks (ANNs) and Gradient Boosted Decision Trees (DTs).

ANNs are data-driven models that can approximate any continuous function (Hornik 1991), making them suitable for forecasting if enough data are available and it is difficult to specify the underlying data generation process. An overview of time series forecasting with ANNs is provided by Zhang et al. (1998). The multilayer perceptron with a single hidden layer is commonly used for time series forecasting (Zhang et al. 1998):

$$\hat{y}(\mathbf{x}) = o\left(\mathbf{W^{(2)}}a\left(\mathbf{W^{(1)}}\mathbf{x} + \mathbf{b^{(1)}}\right) + \mathbf{b^{(2)}}\right) \tag{2.3}$$

Equation (2.3) specifies a fully-connected feed-forward ANN. All input nodes $\mathbf{x}$ are connected to the nodes in the hidden layer, which is represented by the weight matrix $\mathbf{W^{(1)}}$. The activated output of the hidden layer is connected to the output layer by $\mathbf{W^{(2)}}$. The vectors $\mathbf{b^{(1)}}$, $\mathbf{b^{(2)}}$ describe the bias for each node. The functions $a(\cdot)$ and $o(\cdot)$ are the activation functions of the hidden layer and output layer, respectively.

Decision trees (DTs) are simple binary trees that map an input to the corresponding leaf node. Since the introduction of Classification and Regression Trees (CART) several approaches have been developed that combine multiple DTs for one prediction (e.g. Random Forrest Breiman (2001)). Gradient boosted DTs are tree ensemble models, that use K additive functions to predict the output $\hat{y}$ (Friedman 2001):

$$\hat{y}(\mathbf{x}) = \sum_{k=1}^{K} f_k(\mathbf{x}), \tag{2.4}$$

where each function f_k represents a decision tree that maps the input $\mathbf{x}$ to the corresponding leaf in the tree.

2.3.3. Optimization

Recall that the true demand distribution F is unknown to the decision maker. In the following sections, we present two different ways to deal with this problem: traditional model-based optimization and data-driven optimization based on SAA. Both approaches use the point forecast and the historical estimation errors as inputs to determine an inventory decision.

Model-based Optimization

The model-based approach assumes a certain forecast error distribution $\bar{F}$ (e.g. normal distribution) whose parameters θ (e.g. mean and standard deviation) are estimated based on historical forecast errors. The order quantity is then optimized by evaluating the function at the service level quantile and adding it to the forecast:

$$q(\mathbf{x}) = \hat{y}(\mathbf{x}) + \inf\left\{p : \bar{F}(p, \hat{\theta}) \geq \frac{c_u}{c_u + c_o}\right\}, \tag{2.5}$$

where $\hat{y}(\mathbf{x})$ is the mean forecast, given that the features $\mathbf{x}$, and $\hat{\theta}$ are the parameters of the error distribution estimated from the resulting forecast errors. In our evaluation, we adopt normally distributed errors for the model-based approaches.

Of course, this approach yields the optimal decision if the distribution assumption is true. However, in reality, the distribution is unknown and may even change over time. The observed forecast errors depend on the model chosen to produce the forecast. A misspecified model leads to errors that are not distributed as assumed. If the demand distribution is misspecified, highly distorted decisions may result. Ban and Rudin (2019) show this for the example of a normal distribution assumption where the actual demand is exponentially distributed.

Data-driven Optimization with Sample Average Approximation

A data-driven method to optimize the inventory decision is SAA. Here, the error distribution $\bar{F}$ is determined by the empirical forecast errors $\epsilon_1, ..., \epsilon_n$. A distribution assumption is not needed. Thus,

$$\bar{F}(p) = \frac{1}{n}\sum_{t=1}^{n} \mathbb{I}(\epsilon_t \leq p). \tag{2.6}$$

To optimize the order quantity, the service level quantile of the empirical distribution is selected and added to the point forecast. Thus, the resulting order quantity given the features $\mathbf{x}$ is

$$q(\mathbf{x}) = \hat{y}(\mathbf{x}) + \inf\left\{p : \frac{1}{n}\sum_{t=1}^{n} \mathbb{I}(\epsilon_t \leq p) \geq \frac{c_u}{c_u + c_o}\right\}. \tag{2.7}$$

The performance of the optimization highly depends on the quality of the forecast, the number of available data points, and the target service level. Levi et al. (2007, 2015) provide worst-case bounds for a given number of observations. An important and intuitive result is that if the optimal service level is close to 0 or 1, i.e., extreme quantiles need to be estimated, the required sample size is much higher than for service levels close to 0.5, as extreme observations are rare.

2.3.4. Integrated Estimation and Optimization with Quantile Regression

Instead of sequentially forecasting demand and optimizing inventory levels, one can also directly optimize the order quantity by integrating the forecasting model into the optimization problem. The optimal order quantity q of the standard newsvendor model (2.1) is then a function of the feature data $\mathbf{x}$. Instead of first estimating the mean demand and the error distribution and then solving the newsvendor problem, we can now directly estimate the optimal order quantity from the feature data. Beutel and Minner (2012) and Ban and Rudin (2019) formulate this problem as a linear program. This implies that the optimal order quantity is a linear function of the features. We extend these approaches by incorporating ML and thus also allowing for non-linear relationships:

$$\min_{\mathbf{\Phi}} \frac{1}{n} \sum_{t=1}^{n} \left[c_u (d_t - q_t(\mathbf{\Phi}, \mathbf{x_t}))^+ + c_o (q_t(\mathbf{\Phi}, \mathbf{x_t}) - d_t)^+ \right], \tag{2.8}$$

where $q_t(\mathbf{\Phi}, \mathbf{x_t})$ is the output of the ML method in period t with parameters $\mathbf{\Phi}$ (e.g. weight matrix of an ANN) and input variables $\mathbf{x_t}$.

By introducing dummy variables u_t and o_t for the underage and overage in period t, the problem can be reformulated as a non-linear program:

$$\min_{\mathbf{\Phi}} \frac{1}{n} \sum_{t=1}^{n} (c_u u_t + c_o o_t) \tag{2.9}$$

subject to:

$$u_t \geq d_t - q_t(\mathbf{\Phi}, \mathbf{x_t}) \qquad \forall t = \{1, ..., n\}, \tag{2.10}$$

$$o_t \geq q_t(\mathbf{\Phi}, \mathbf{x_t}) - d_t \qquad \forall t = \{1, ..., n\}, \tag{2.11}$$

$$u_t, o_t \geq 0 \qquad \forall t = \{1, ..., n\}. \tag{2.12}$$

The objective function (2.9) minimizes the empirical underage and overage costs, while the constraints (2.10) to (2.12) ensure that deviations of the estimate from the actual demand are correctly assigned to underages and overages. By solving the problem for the empirical data $S_n = \{(d_1, \mathbf{x_1}), ..., (d_n, \mathbf{x_n})\}$, we obtain parameters $\mathbf{\Phi}^*$ for the ML method that minimize the empirical costs with respect to these data. Once the model has been trained, the resulting order quantity for period p is the quantile forecast with

$q_p(\mathbf{\Phi}^*, \mathbf{x}_p)$.

Bertsimas and Kallus (2020) and Ban and Rudin (2019) showed that integrating forecasting in the optimization model is equivalent to the more general QR problem in Takeuchi et al. (2006). For a better understanding, we elaborate on this relation in more detail. The basic idea of QR is to estimate the unobservable quantile by modifying the loss function of a standard regression model. Minimizing the sum of squared errors $\sum_{t=1}^{n}(y_t - \hat{y}_t)^2$ yields the mean, while minimizing the sum of absolute errors $\sum_{t=1}^{n}|y_t - \hat{y}_t|$ yields the median. By weighting the underages with the quantile $\tau \in (0, 1)$ and overages with $(1-\tau)$, thus $\sum_{t=1}^{n}\tau(y_t - \hat{y}_t)^+ + (1-\tau)(\hat{y}_t - y_t)^+$, we obtain an estimate for the quantile (Koenker 2005). The optimal solution of the newsvendor model is the quantile $\tau = \frac{c_u}{c_u+c_o}$ of the demand distribution; thus, $(1-\tau) = \frac{c_o}{c_u+c_o}$. Inserting these values of τ and $(1-\tau)$ into the objective function of the quantile regression yields the optimization problem (2.9).

The main advantage of QR over the model-based approach and SAA is its ability to model conditional quantiles under heteroscedasticity and for unknown error distributions. However, the performance of the approach depends crucially on the underlying model q. On the one hand, if q is too simplistic (e.g. linear), the model might not be able to capture the structure in the training data. On the other hand, if q is too complex, there is a risk of overfitting the model.

2.3.5. Summarizing the Three Levels of Data-driven Inventory Management

We conclude this chapter by linking our methodology explained in Subsections 2.3.2 - 2.3.4 to our framework of data-driven inventory management introduced in Figure 2.1. To this end, Figure 2.2 positions each piece of our methodology in the framework.

On the first level (demand estimation), we choose a parametric forecasting model (e.g. ETS or ANN). For the ML models, this includes the selection and optimization of hyper-parameters (e.g. number of layers of ANNs). We then use the data to fit the model by optimizing its parameters in order to minimize a certain objective function (i.e. MSE). The outputs of the first level of data-driven inventory management are a point demand forecast and the resulting empirical error distribution.

On the second level (inventory optimization), we operationalize a model-based approach by fitting a normal distribution and distinguish it from a data-driven (SAA)

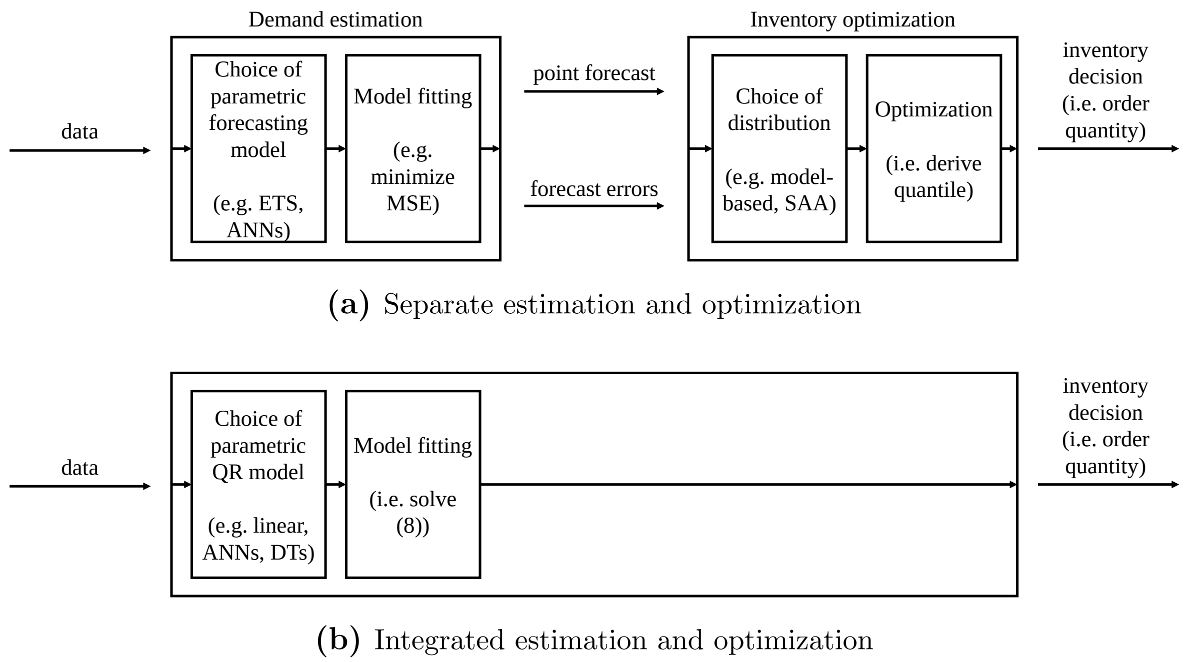

Figure 2.2.: Relating our methodology to the three levels of data-driven inventory management

approach. We then optimize by selecting a certain quantile of the respective demand distribution. This gives us the resulting order quantity.

On the third level (integrated estimation and optimization), we choose a parametric QR model (e.g. ANNs) and fit its parameters by solving problem (2.8) instead of minimizing the MSE.

From the existing literature, it is not yet clear how the choices on each of the three levels affect performance. In the following, we investigate this question empirically.

2.4. Empirical Evaluation

Our empirical evaluation aims to assess the impact of data-driven approaches for the three levels – (1) demand estimation, (2) optimization, and (3) integrated estimation and optimization – on average costs for the newsvendor problem. To this end, we evaluate the performance of the methods with respect to costs by using a real-world dataset to compare it to various standard approaches.

2.4.1. Data

We evaluate the proposed approaches using daily demand data of a German bakery chain. The observed sales are not necessarily equal to demand, as stock-outs occur and lead to censored demand information (Conrad 1976). In order to estimate the daily demand in the case of a stock-out, we leverage intra-day sales patterns of point-of-sales data (Lau and Lau 1996). In particular, for each product and weekday, we determine the average demand proportion of each hour in relation to the total demand on days on which the product was not sold out. This process allows us to interpolate the sales when a stock-out occurs and obtain an estimate for historical demand. The approach is feasible because we have access to point-of-sales data and information on the overage for each product per day. Figure 2.3 shows the strong weekly seasonality of demand for (a) a representative product and (b) a box plot that confirms this pattern for all time series. While median demand on Tuesdays and Thursdays is the lowest, it is slightly higher on Mondays, Wednesdays and Fridays. The median demand on Saturday is higher than it is for all other days. The standard deviation of demand does not vary strongly across the weekdays.

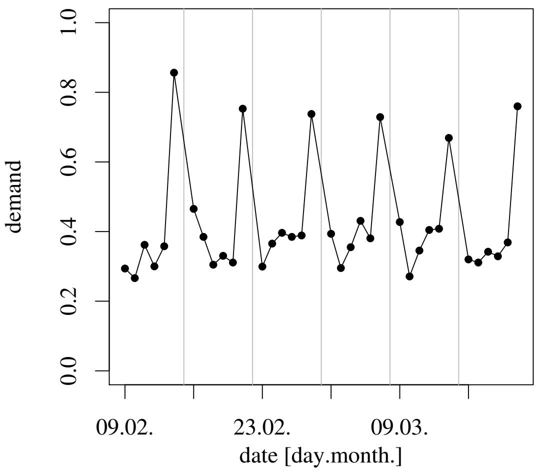

(a) Exemplary demand time series of six weeks

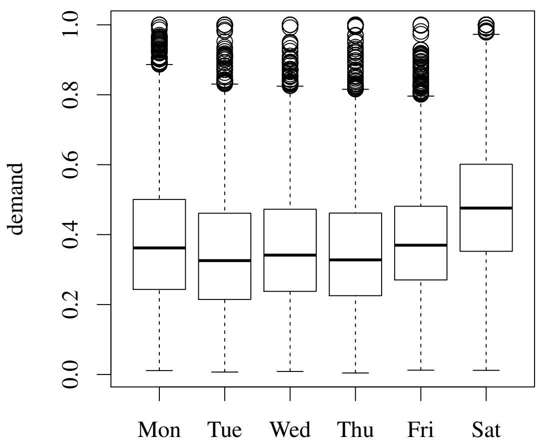

(b) Boxplot of demand per weekday

Figure 2.3.: The demand shows a strong weekly seasonality. The demand levels for working days (Mon-Fri) are comparable, while the demand level on the weekend (Sat) is noticeably higher.

The dataset comprises eleven stock-keeping units, namely, six breads and five buns, for five stores over a period of 88 weeks, where each store is open from Monday to

Data Source	Features
Master Data	store class, product category, opening times (day, hours/duration)
Transactional Data	lagged sales, rolling median of sales, binary promotional information
Calendar	day of year, month, day of month, weekday, public holiday, day type, bridge day, nonworking day, indicators for each special day, school holidays
Weather	temperature (minimum, mean, maximum) and cloud cover of target day
Location	general location (city, suburb, town); in proximity to the store: shops (bakeries, butcher, grocery, kiosk, fast-food, car repair), amenities (worship, medical doctors, hospitals), leisure (playground, sport facility, park), education (kindergarden, school, university)

Table 2.1.: Features used in the machine learning methods

Saturday. This configuration amounts to 55 ordering decisions per day. Additionally, we enrich the dataset with external explanatory features related to calendar, weather, and location of the store (see Table 2.1). We split the dataset into a training set containing up to 63 weeks and a test set containing the remaining 25 weeks (see Table 2.2). We perform a rolling 1-step-ahead prediction evaluation on the test set in order to assess the performance of the methods. We fit the models and distribution parameters every 10 days on a rolling training dataset with constant size. Due to computational constraints, we fit the parameters of the ANNs every 50 days only. To evaluate the effect of the amount of available data, we use different sample sizes for the training set. The full training set (sample size 1.0) covers 63 weeks, while the smallest training set (sample size 0.1) contains only 6 weeks (see Table 2.2).

sample	1.0	0.8	0.6	0.4	0.2	0.1
train length (days)	378	300	228	150	78	36
test length (days)	150	150	150	150	150	150

Table 2.2.: Training & test periods for different sample sizes.

While traditional time series methods such as exponential smoothing or ARIMA are able to process only a single times series at a time, a major advantage of the ML methods is their ability to deal with a large number and variety of features. In order to leverage this advantage, we do not only train them with a single time series per product but alternatively also across products and stores. In the latter case, we also include the features listed in Table 2.1.

2.4.2. Experimental Design

In our experiment, we evaluate the impact of different (1) estimation, (2) optimization, and (3) integrated estimation and optimization approaches on the costs of the newsvendor model. We start by assessing the impact of forecast performance. In addition to the ANNs and DTs introduced in the previous section, we evaluate six different reference forecasting methods, which we outline in the next section. For each forecasting method, we measure the forecast accuracy (Section 2.4.4) and then investigate its impact on costs (Section 2.4.5). Second, we compare the model-based optimization assuming a normal distribution (*Norm*) with the data-driven optimization using *SAA*. To this end, we calculate the average costs for different target service levels (Section 2.4.5). Third, we assess the performance of the integrated estimation and optimization approach with QR and compare it to the separate approaches (Section 2.4.5). Fourth, we evaluate the sensitivity to the sample size in order to assess the value of a large training set (Section 2.4.5). Overall, the database of the evaluation results comprises more than 9.1 million entries, i.e., close to 0.6 million point forecasts and approximately 8.6 million order quantities. We employ the Wilcoxon signed-rank test to test the statistical significance of our results at the 5% significance level.

2.4.3. Reference Methods and ML Setup

In order to evaluate the ML approaches, we compare them to well-established forecasting methods. With the exception of the first approach (*Median*), we rely on methods that are explicitly able to model seasonal time series because the demand for baked goods exhibits a strong weekly seasonality (see Figure 2.3).

Reference Methods

Median and Seasonal-Median

The first benchmark forecast is the median of the entire training set (*Median*); it does not consider seasonality. Nonetheless, we include it in our comparison in order to evaluate the benefit of seasonal demand models. Its seasonal variant estimates the median by weekday (*S-Median*).

Seasonal-Naïve

A popular benchmark method for forecasting is the Naïve method and its seasonal variant (*S-Naïve*). The forecast is set to the last observed value from the same part of the season: $\hat{y}_{t+h} = y_{t+h-m}$. Hence, we need to specify only the frequency of the seasonality m, which we set to 6 for the considered time series.

Seasonal Moving Average

The seasonal moving average method (*S-MA*) sets the forecast to an average of the last observations from the same part of the season: $\hat{y}_{t+h} = \frac{1}{k}\sum_{i=1}^{k} y_{t+h-mk}$. Besides setting the frequency of the seasonality m, we must set k, which controls the number of considered values. We determine k in the range from 3 to 12 based on the last 20% of the training set for each time series. We choose the value of k that minimizes the sum of squared errors.

Seasonal Autoregressive Integrated Moving Average

Autoregressive integrated moving average (ARIMA) and its seasonal variant *S-ARIMA* represent a widely used forecasting method. The autoregressive part of ARIMA represents a linear combination of past values, while the moving average part is a linear combination of past forecast errors. The time series must be stationary, which can be achieved by differencing. We employ the method `auto.arima()` function from the `forecast` package (Hyndman and Khandakar 2008) for the statistical software R (R Core Team 2017) in order to identify the most suitable model per time series. The `auto.arima()` function selects a suitable model using a step-wise approach that traverses the space of possible models in an efficient way until the best model is found.

Exponential Smoothing

Exponential smoothing methods calculate the forecast by computing a weighted average of past observations. The weights decay as the observations get older. Hyndman et al. (2002, 2008) propose innovation space models that generalize exponential smoothing methods (*ETS*). These models include a family of 30 models that cover different types of errors, seasonal effects and trends (none, additive, multiplicative). We use the `ets()`

function from the `forecast` package (Hyndman and Khandakar 2008) for the statistical software R (R Core Team 2017).

ML Setup

In this subsection, we introduce methods that take multiple time series and additional features (see Table 2.1) into account. For these methods, we also evaluate the integrated estimation and optimization approach introduced in Section 2.3.4.

Linear Regression

The linear regression model uses lagged demand data (lags: 1, 2, ..., 6, 12, 18) which are linearly scaled between 0 and 0.75 as input. The weekly seasonality is modeled through binary variables. When all time series across stores and products and the extended feature set are used for the prediction, further variables are introduced. In order to avoid overfitting, we include a regularization term in the objective function. The integrated linear approach is equivalent to the models in Beutel and Minner (2012) and Ban and Rudin (2019).

ANNs

We apply *ANNs* as described in Section 2.3.2. Several hyper-parameters (learning rate, batch size, number of hidden nodes, activation function of hidden layer) are optimized by a random search (Bergstra and Bengio 2012) in combination with cross-validation on the training set. As activation function for the output layer we use a linear function, which is reasonable for regression with ANNs (Zhang et al. 1998).

In order to encode deterministic seasonality, we use trigonometric functions as features, as proposed by Crone and Kourentzes (2009). This is a parsimonious approach which requires only two additional input variables. Additionally, the approach is non-parametric, as no seasonal indices need to be estimated. The two variables are $x_{t,1}$ and $x_{t,2}$ in period t, with m representing the frequency of the seasonality:

$$x_{t,1} = sin(2\pi t/m) \tag{2.13}$$

$$x_{t,2} = cos(2\pi t/m) \tag{2.14}$$

The input consists of lagged demand information (lags: 1, 2, ..., 6, 12, 18), which is linearly scaled between 0 and 0.75, as this is similar to what other seasonal methods consider. When all time series across products and stores are considered, we enrich the dataset with further explanatory features (see Table 2.1).

The performance of an ANN depends on its initial weights, which are randomly set. Therefore, we employ an ensemble of ANNs with the *median* ensemble operator, as this approach is robust to the initial weights and provides reliable results (Barrow et al. 2010, Kourentzes et al. 2014). Another crucial aspect is the training of ANNs. We use the stochastic gradient-based algorithm ADAM proposed by Kingma and Ba (2015) to optimize the weights of the ANN. We also employ early stopping to avoid overfitting and train an ensemble of 50 *ANNs* in order to obtain more reliable and accurate results (Barrow et al. 2010, Kourentzes et al. 2014).

DTs

The DT approach is a tree-based ensemble model as described in Section 2.3.2. We use Microsoft's LightGBM implementation (Ke et al. 2017). Similar to the ANNs, several hyper-parameters (learning rate, number of leaves, minimum amount of data in one leaf, maximum number of bins, maximum depth of tree) are selected based on a random search within the training data (Bergstra and Bengio 2012). The number of trees is controlled by early stopping, which also reduces the risk of overfitting. We consider the same features as in the other ML methods.

2.4.4. Point Forecast Analysis

The relevant performance measure of the newsvendor model is overall costs (overage and underage). Before evaluating the impact of the different estimation and optimization approaches on cost in Section 2.4.5, we separately measure the accuracy of the point forecasts in order to relate it to overall costs in the subsequent analysis.

For each forecasting method introduced in the previous section, we compute a set of common accuracy measures, including the Mean Percentage Error (MPE), Symmetric Mean Absolute Percentage Error (SMAPE), Mean Absolute Percentage Error (MAPE), Mean Absolute Scaled Error (MASE) (Hyndman and Koehler 2006), Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and Relative Absolute Error (RAE). We provide more than one measure because each of them has its strengths and weaknesses. For instance, RMSE and MAE are scale-dependent error measures and do not allow for comparisons between time series at different scales, while percentage-based error measures (SMAPE, MAPE) are not always defined and may result in misleading outcomes if demand is low. Table 2.3 shows the average forecast accuracy over all time series by method.

Method	MPE	SMAPE	MAPE	MASE	RMSE	MAE	RAE
Median	-22.34	29.71	39.43	1.01	39.89	15.70	1.72
S-Median	-21.45	24.74	33.73	0.82	28.42	11.99	1.31
S-Naïve	**-11.84**	28.71	34.86	0.92	27.80	12.56	1.37
S-MA	-14.61	23.32	30.15	0.75	22.27	10.14	1.11
ETS	-12.47	22.19	28.47	0.71	21.83	9.66	1.06
S-ARIMA	-14.35	22.88	29.71	0.73	21.40	9.87	1.08
Linear	-18.73	23.75	32.07	0.77	23.43	10.54	1.15
DT-LGBM	-18.80	22.88	31.13	0.73	21.98	9.92	1.08
ANN-MLP	-14.73	22.63	29.59	0.72	21.28	9.75	1.07
Linear (all)	-14.33	22.14	29.18	0.71	21.23	9.63	1.05
DT-LGBM (all)	-13.44	21.51	28.34	**0.68**	**20.06**	**9.15**	**1.00**
ANN-MLP (all)	**-12.62**	**21.42**	**27.87**	**0.68**	**20.09**	**9.16**	**1.00**

Table 2.3.: Forecast performance of the point predictions (sample size: 1.0). The best performance for each metric is underlined. Results that do not differ from the one of the best method at a significance level of 5% for each metric are printed in bold face.

Not surprisingly, the worst accuracy is achieved by the *Median* forecast, which is the only method that does not incorporate the weekly seasonality pattern. The results improve noticeably (more than 5 percentage points in MAPE) when the weekly seasonality is considered (*S-Median*). *S-Median* is also more robust against sudden changes in demand and provides more reliable results than *S-Naïve*. *S-MA* outperforms all baseline methods (*Median*, *S-Median*, *S-Naïve*) and its accuracy is even competitive to more sophisticated approaches. It is not as prone to outliers but follows minor level shifts. Overall, *ETS* is the best method compared to models that are trained on a single time series as it captures the main characteristics of the time series by computing the weighted average of past observations. Even the more complex ML approaches cannot improve the forecast. However, when trained across stores and products with additional features, the ML methods further improve significantly. *ANN-MLP* and *DT-LGBM* also outperform ETS. The information contained in the features and supplementary time series has additional explanatory potential that is effectively extracted by all three ML approaches.

We note that the negative MPE throughout all methods indicates that in the test data, there are low-demand events that cannot be foreseen by the models based on historical demand. These low-demand events are more frequent, more extreme, or both during the test period than events of unexpectedly high demand. This observation might be due to the fact that situations with very low demand (e.g. supply disruption, partial shop closing, and construction) are more likely than situations with extremely high demand.

2.4.5. Inventory Performance Analysis

The purpose of the newsvendor model is to determine the cost-minimal order quantity by considering demand uncertainty and underage and overage costs. In order to perform a comprehensive analysis of the introduced methods, we calculate the order quantities and compute the resulting average costs for each approach. As underage and overage cost may vary among products and stores, we analyze multiple target service levels. The target service level $c_u/(c_u + c_o)$ is the optimal probability of having no stock-out during the day. In the repeated newsvendor model, this corresponds to the long run fraction of periods in which demand is fully satisfied. By setting the unit price and the sum of underage and overage costs $(c_u + c_o)$ to 1.00 and varying their relative share, we obtain six different target service levels. This process allows us to interpret c_u as the profit margin and c_o as the unit costs (e.g. material and production costs) of an item. In order to compare the different methods, we measure the performance relative to the best method for each target service level. Additionally, we report the realized average service level for each approach. We calculate the realized service level as the relative share of days on which total demand was met. A large deviation of the realized service level from the target service level indicates that a method tends to overestimate or underestimate the optimal order quantity. Note that the reported service level just serves to characterize the solution by relating it to the newsvendor solution. It does not reflect a cost-service trade-off since costs include both overage and underage costs. The results are reported in Table 2.4.

In the following sections, we analyze the effects of (1) demand estimation, (2) optimization, and (3) integrated estimation and optimization on average costs and observed service levels. Furthermore, we evaluate the sensitivity of the results to the size of the available sample.

	Method		TSL = 0.5		TSL = 0.6		TSL = 0.7		TSL = 0.8		TSL = 0.9		TSL = 0.95	
	Estimation	Optimization	Δ Cost	SL	Δ Cost	SL	Δ Cost	SL	Δ Cost	SL	Δ Cost	SL	Δ Cost	SL
Benchmarks	Median	Norm	72.5%	0.61	83.9%	0.72	93.2%	0.79	99.4%	0.86	97.8%	0.92	92.0%	0.95
		SAA	72.5%	0.61	79.4%	0.70	87.9%	0.79	99.6%	0.87	109.5%	0.94	101.9%	0.98
	S-Median	Norm	31.8%	0.64	33.5%	0.74	34.7%	0.83	34.9%	0.89	32.2%	0.95	29.6%	0.97
		SAA	31.8%	0.64	30.3%	0.72	30.4%	0.80	29.5%	0.88	27.4%	0.95	31.2%	0.98
	S-Naive	Norm	38.0%	0.51	37.5%	0.63	37.0%	0.75	37.3%	0.85	37.2%	0.93	37.2%	0.96
		SAA	38.4%	0.51	37.6%	0.61	35.6%	0.71	34.3%	0.81	32.2%	0.91	33.3%	0.96
	S-MA	Norm	11.5%	0.56	13.6%	0.68	16.0%	0.78	17.6%	0.86	18.3%	0.94	16.6%	0.97
		SAA	10.5%	0.52	11.0%	0.62	11.4%	0.73	11.7%	0.82	12.2%	0.92	13.9%	0.96
	ETS	Norm	6.1%	0.53	6.7%	0.64	7.0%	0.74	7.1%	0.83	5.6%	0.91	5.7%	0.95
		SAA	6.2%	0.50	6.5%	0.61	6.7%	0.71	6.7%	0.80	5.6%	0.90	5.9%	0.95
	S-ARIMA	Norm	8.5%	0.55	8.9%	0.65	8.8%	0.75	8.3%	0.84	7.5%	0.92	7.2%	0.95
		SAA	8.0%	0.52	8.1%	0.62	8.0%	0.71	7.7%	0.81	6.5%	0.91	7.2%	0.95
ML single time series	Linear	Norm	15.8%	0.58	17.7%	0.69	19.6%	0.78	20.8%	0.85	20.2%	0.93	20.9%	0.95
		SAA	15.6%	0.56	17.2%	0.66	18.9%	0.75	20.1%	0.84	20.7%	0.93	21.8%	0.96
		QR	10.6%	0.54	10.7%	0.64	11.4%	0.73	11.2%	0.82	11.8%	0.91	18.9%	0.96
	DT-LGBM	Norm	9.0%	0.60	8.6%	0.68	8.5%	0.76	8.8%	0.83	10.2%	0.89	**15.2%**	0.93
		SAA	7.9%	0.57	7.7%	0.65	7.8%	0.73	8.4%	0.81	10.0%	0.89	14.4%	0.94
		QR	11.1%	0.59	10.8%	0.68	12.1%	0.78	15.3%	0.85	20.8%	0.93	29.0%	0.96
	ANN-MLP	Norm	7.2%	0.55	8.4%	0.66	9.0%	0.75	9.6%	0.83	9.4%	0.91	10.5%	0.95
		SAA	6.6%	0.52	7.6%	0.63	8.2%	0.72	8.6%	0.82	8.6%	0.91	10.2%	0.95
		QR	7.5%	0.53	7.9%	0.64	8.6%	0.73	9.8%	0.82	13.0%	0.91	18.1%	0.95
ML pooled time series + features	Linear (all)	Norm	5.9%	0.53	5.5%	0.64	5.6%	0.75	6.1%	0.84	4.9%	0.91	4.0%	0.95
		SAA	5.4%	0.51	5.3%	0.62	5.0%	0.72	5.3%	0.82	5.2%	0.91	4.9%	0.95
		QR	5.1%	0.52	4.5%	0.62	5.2%	0.72	7.2%	0.81	10.0%	0.90	12.8%	0.95
	DT-LGBM (all)	Norm	**0.6%**	0.53	**0.4%**	0.62	0.0%	0.71	0.1%	0.80	0.4%	0.87	2.1%	0.92
		SAA	0.9%	0.51	**0.4%**	0.61	**0.0%**	0.69	**0.0%**	0.79	0.2%	0.88	1.7%	0.92
		QR	1.6%	0.52	1.7%	0.61	1.6%	0.71	3.1%	0.80	6.4%	0.90	11.4%	0.94
	ANN-MLP (all)	Norm	0.7%	0.52	0.7%	0.63	0.7%	0.73	0.7%	0.82	**0.0%**	0.90	**0.0%**	0.95
		SAA	**0.3%**	0.51	**0.2%**	0.61	0.3%	0.72	0.4%	0.81	0.4%	0.90	1.5%	0.95
		QR	**0.0%**	0.50	**0.0%**	0.61	0.9%	0.72	3.3%	0.82	6.8%	0.91	11.2%	0.95

Table 2.4.: Inventory performance analysis: Average cost increase relative to the best approach and average service level (SL) for various target service levels (TSLs) and a sample size of 1.0. Methods denoted with *all* are trained on data across all products and stores. The best approach for each target service level is underlined. Results that do not differ from the one of the best method at a significance level of 5% for each service level are printed in bold face.

The Effect of Demand Estimation

To evaluate the effect of demand estimation on costs, we compare the average cost of the different estimation approaches for each target service level in Table 2.4. The best approach for each target service level is underlined. We see that the approaches based on the ML forecasts that use data across stores and products and additional features (*all*) provide the lowest average costs for all target service levels. The performance of *ANN-MLP* and *DT-LGBM* is very similar, while methods based on the *Linear* forecast yield higher costs. An interesting result is that *ETS* performs best when training is restricted to single time series. This is particularly noteworthy when considering its computational efficiency compared to the ML methods. Overall, we observe that approaches based on accurate estimation methods achieve significantly lower costs, independent of the optimization approach. Thus, the level of demand estimation has a substantial impact on overall performance.

In order to further substantiate this statement, we conduct a correlation analysis. We compute the Spearman's rank correlation coefficient ρ between costs and forecast accuracy (SMAPE and RMSE) for each store-article-service level combination. The results are depicted in Table 2.5.

	Costs	SMAPE	RMSE
Costs	-	0.8799 (± 0.1211)	0.9406 (± 0.0481)
SL	0.4202 (± 0.2879)	0.4253 (± 0.2834)	0.3034 (± 0.2678)

Table 2.5.: Median of Spearman's Correlations (± standard deviation) between absolute service level deviation (SL), costs, and forecast accuracy (SMAPE, RMSE).

The analysis supports the claim that the general ranking of methods with respect to costs is similar to the ranking with respect to forecast accuracy, with a median ρ of 0.8799 for the rank correlation of costs and SMAPE and 0.9406 for the median rank correlation of costs and RMSE. The reason for this observation is that more accurate point predictions lead to more precise demand distribution estimates, which make the succeeding optimization phase less crucial.

We complement the above cost analysis by looking at the realized service levels which provide further insights into the order quantities obtained from the different methods.

Table 2.5 also shows the Spearman Correlations between the absolute service level deviations (i.e. difference between average observed service level and the newsvendor target service level) and costs and forecast errors, respectively.

From Table 2.4, we can see that all methods overachieve the target service level on average. This matches our observation of Section 2.4.4, that all forecasting methods overestimate the demand on average, due to events with unexpectedly low demand in the test data.

We further see that the correlation between the absolute service level deviation and costs is relatively low (0.4202). This shows that the ability of a method to achieve a desired service level on average is not a very good indicator for the cost performance of that method. The service level measures only whether or not there was a stock-out and thus indicates the direction of the deviation from the optimal order quantity on average. It does not take into account the order of magnitude of overages and underages. The low correlation between the forecast accuracy measures and the service level deviation confirms this conclusion.

The Effect of Optimization

To assess the impact of model-based vs. data-driven optimization on costs, we compare the average cost of *Norm* and *SAA* for each estimation method and target service level. We perform a Shapiro-Wilk test on the residuals of the forecasts of S-ARIMA and ETS and find that for approximately one quarter of the time series the residuals are normally distributed at 95% confidence level. Thus, the normal distribution assumption can be justified, although one cannot expect that all residuals follow the distribution assumption in a real-world data set. We observe that the performance differences between *SAA* and *Norm* are relatively small and the effect of accurate demand estimation clearly outweighs the effect of data-driven optimization. However, for the majority of estimation methods, *SAA* leads to lower costs than *Norm* for target service levels up to 0.9, while the normal distribution assumption can be beneficial for higher service levels.

The good performance of *SAA* and its weaknesses for higher service levels are in line with the theoretical results of Levi et al. (2015). The authors provide a bound on the accuracy of *SAA* for the newsvendor model (Theorem 2 *Improved LRS Bound*) that does not rely on assumptions on the demand distribution. The bound has an exponential rate that is proportional to the sample size and $\min(c_u, c_o)/(c_u + c_o)$. In our case, the bound implies that using *SAA*, in order to obtain the same accuracy for a service level of 0.9 (0.95) as for a service level of 0.8, we would need 1.5 (4) times more data. However, in the bakery industry, such high service levels are not common, and our dataset is sufficient to let *SAA* outperform *Norm* for service levels up to 0.9 for most approaches.

The Effect of Integrated Estimation and Optimization

We also employ the QR approach that integrates the demand estimation into the optimization model for the *linear* approach and the ML methods *DT-LGBM* and *ANN-MLP*. In order to focus on the effect of integrated estimation and optimization, we compare *QR* to *SAA* for the respective approaches. For *DT-LGBM* and *ANN-MLP* trained on single time series, *QR* performs worse than *SAA*, while *Linear QR* outperforms *SAA*. For high service levels *QR* generally performs relatively poor for all three estimation approaches. When trained on data across stores and products and including features, integration of estimation and optimization improves the performance of *Linear (all)* and *ANN-MLP (all)* for low service levels. However, for high target service levels, *SAA* and *Norm* perform better than *QR* for all estimation approaches.

The theoretical advantage of the *QR* approach is its ability to estimate *conditional* quantiles that depend on the features (see Figure 2.4). The observation that for the approaches trained only on single time series, *QR* is not beneficial, might be explained by the fact that too little features are available to leverage the feature-dependency of the quantile. The previous statement is supported by the fact that *Linear (all)* and *DT-LGBM (all)* improve through integration at low service levels as more data are available and feature-dependent variance can be estimated more accurately. However, this theoretical advantage cannot be observed for higher service levels. We suspect that more extensive hyper-parameter optimization in combination with alternative scaling of the input data for each individual target service level might improve the performance.

Our results for the single time series case are in line with the outcome of the empirical analysis of Ban and Rudin (2019) who also report that separate estimation and optimization outperforms the linear integrated approach on their relatively small dataset of one year. We observe that this effect gets smaller when the models are trained with pooled time series and features.

The effect of learning across products and external features

Our dataset comprises sales data of several breads and buns across multiple stores. These products are relatively similar to one another and therefore one time series might

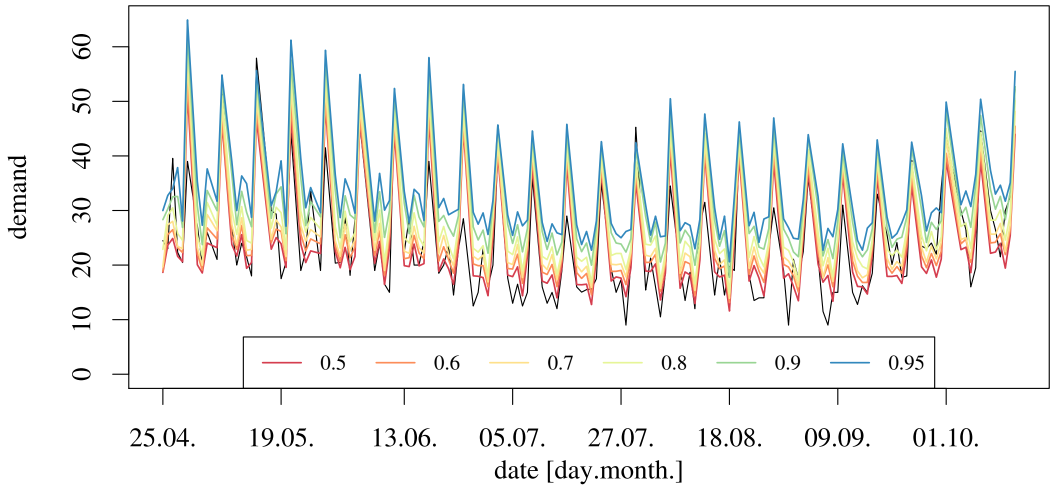

Figure 2.4.: Forecasts for different service levels using ANN QR.

contain information about the other. Univariate time series models can only consider a single product at time, while ML methods are able to process a large number of inputs. Therefore, we train *linear (all)*, *DT-LGBM (all)*, and *ANN-MLP (all)* across all products and stores. The pooling of training data also makes it possible to enhance the data set with a large number of additional features that cannot be employed if the models are trained per time series.

From Table 2.4 we observe that indeed all ML methods benefit from the additional data and improve significantly. *DT-LGBM (all)* and *ANN-MLP (all)* perform similarly and outperform all other methods. We note that a similarity of time series is not specific to our case but can be found in many retail settings.

Sensitivity to Sample Size

The power of the data-driven approaches lies in their ability to leverage large amounts of available data, which makes them very flexible but may limit their deployability if not enough data is available. In order to determine the dependency of the different approaches on data availability, we vary the size of the training data and compare the results on a fixed test set (see Table 2.2). The results of this experiment are given in Table 2.6 and depicted in Figure 2.5 for the data-driven approaches. We present only the results for target service level 0.7, noting that the qualitative results also apply to the other service levels.

Based on our results, the methods can be divided into three groups: The first group

	Method		S = 0.1		S = 0.2		S = 0.4		S = 0.6		S = 0.8		S = 1.0	
	Estimation	Optimization	Δ Cost	SL	Δ Cost	SL	Δ Cost	SL	Δ Cost	SL	Δ Cost	SL	Δ Cost	SL
Benchmarks	Median	Norm	79.5%	0.72	82.0%	0.75	87.1%	0.79	92.7%	0.80	92.7%	0.80	93.2%	0.79
		SAA	77.2%	0.70	78.2%	0.74	82.6%	0.78	88.0%	0.79	87.7%	0.79	87.9%	0.79
	S-Median	Norm	13.5%	0.70	15.5%	0.76	23.2%	0.81	32.5%	0.84	33.6%	0.83	34.7%	0.83
		SAA	13.3%	0.65	12.1%	0.72	18.5%	0.79	26.8%	0.81	29.1%	0.80	30.4%	0.80
	S-Naive	Norm	36.7%	0.72	36.5%	0.73	36.8%	0.75	37.5%	0.75	37.1%	0.75	37.0%	0.75
		SAA	40.2%	0.68	37.4%	0.69	36.2%	0.70	35.9%	0.71	35.7%	0.71	35.6%	0.71
	S-MA	Norm	15.8%	0.73	14.6%	0.75	15.1%	0.77	15.9%	0.78	15.6%	0.78	16.0%	0.78
		SAA	19.2%	0.67	15.4%	0.67	12.1%	0.70	11.0%	0.72	11.8%	0.72	11.4%	0.73
	ETS	Norm	**12.1%**	0.70	**8.8%**	0.73	7.9%	0.74	7.7%	0.74	6.6%	0.74	7.0%	0.74
		SAA	14.0%	0.66	9.6%	0.68	8.1%	0.69	7.6%	0.70	6.5%	0.71	6.7%	0.71
	S-ARIMA	Norm	25.8%	0.69	13.8%	0.72	10.7%	0.74	10.3%	0.75	9.2%	0.75	8.8%	0.75
		SAA	29.2%	0.64	**14.5%**	0.68	10.5%	0.68	9.8%	0.70	8.6%	0.72	8.0%	0.71
ML single time series	Linear	Norm	29.8%	0.70	18.0%	0.71	18.9%	0.76	21.1%	0.79	20.4%	0.78	19.6%	0.78
		SAA	30.2%	0.68	18.3%	0.70	18.1%	0.74	20.1%	0.76	19.3%	0.76	18.9%	0.75
		QR	32.1%	0.67	17.4%	0.68	14.3%	0.71	13.2%	0.73	12.0%	0.73	11.4%	0.73
	DT-LGBM	Norm	50.4%	0.72	22.6%	0.73	13.6%	0.77	12.5%	0.79	10.4%	0.78	8.5%	0.76
		SAA	48.7%	0.68	22.6%	0.68	11.1%	0.73	10.2%	0.75	8.5%	0.75	7.8%	0.73
		QR	52.4%	0.73	27.3%	0.75	17.7%	0.78	16.8%	0.80	13.6%	0.79	12.1%	0.78
	ANN-MLP	Norm	33.2%	0.69	14.3%	0.73	13.9%	0.78	12.0%	0.78	9.9%	0.78	9.0%	0.75
		SAA	33.5%	0.70	14.9%	0.72	12.4%	0.74	10.2%	0.75	8.4%	0.75	8.2%	0.72
		QR	33.1%	0.67	17.8%	0.71	12.5%	0.75	10.7%	0.75	9.2%	0.75	8.6%	0.73
ML pooled time series + features	Linear (all)	Norm	**16.0%**	0.69	**14.2%**	0.70	14.8%	0.73	8.8%	0.74	8.2%	0.75	5.6%	0.75
		SAA	**17.2%**	0.64	14.5%	0.66	13.5%	0.69	7.3%	0.71	7.6%	0.72	5.0%	0.72
		QR	**15.3%**	0.65	12.7%	0.67	13.3%	0.69	7.7%	0.71	7.1%	0.72	5.2%	0.72
	DT-LGBM (all)	Norm	**21.1%**	0.67	10.4%	0.69	**8.0%**	0.69	**5.5%**	0.71	2.5%	0.70	0.0%	0.71
		SAA	**20.3%**	0.65	10.6%	0.66	**7.5%**	0.67	**4.7%**	0.69	**2.4%**	0.68	**0.0%**	0.69
		QR	18.5%	0.71	10.8%	0.69	7.7%	0.70	**4.8%**	0.71	3.4%	0.70	1.6%	0.71
	ANN-MLP (all)	Norm	33.5%	0.60	12.4%	0.69	6.1%	0.69	3.6%	0.73	0.7%	0.71	0.7%	0.73
		SAA	36.3%	0.58	12.6%	0.66	**5.4%**	0.67	**2.8%**	0.71	**0.3%**	0.69	0.3%	0.72
		QR	**18.7%**	0.63	12.4%	0.65	6.2%	0.68	**3.1%**	0.70	0.3%	0.70	0.9%	0.72

Table 2.6.: The effect of the sample size: Average cost increase relative to the best approach (over all sample sizes) and average service level (SL) for the target service level 0.7 and various sample sizes (S). Methods denoted with *all* are trained on data across all products and stores. The best approach for each sample size is underlined. Results that do not differ from the one of the best method at a significance level of 5% for each sample size are printed in bold face.

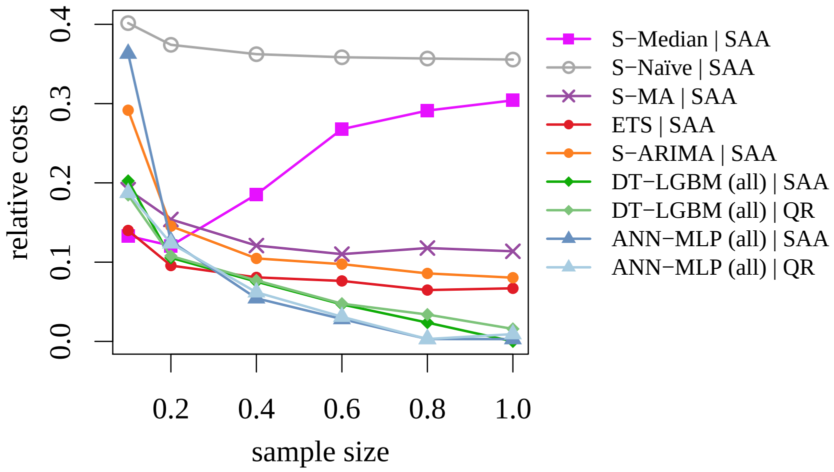

Figure 2.5.: Effect of the sample size (TSL = 0.7).

consists of methods whose performance hardly depends on the sample size. In our case this includes methods based on the *S-Naïve* forecast. The *S-Naïve* approaches simply forecast the demand of the same weekday of the week before. Thus, it does not improve as more data becomes available. The second group consists of methods whose performance diminishes as more training data become available. The approaches with a *Median* (not depicted in Figure 2.5, see Table 2.6) and *S-Median* forecast are part of this group. The costs increase as more training data are available and as more "outdated" data are included. In our real-world case, this observation implies that, for example, demand data from Winter is used to estimate the median forecast for Summer although these data are not representative of this season. The third group consists of methods whose performance improves as more data become available. This group comprises the ML methods proposed in this paper. We also include methods based on *S-ARIMA*, *ETS*, and *linear* forecast in this group. However, the performance of *S-ARIMA* and *ETS* stagnates for sample sizes larger than 0.6. This effect might be due to the fact that we use a little over one year of training data and consequently some months are included twice. It seems that the ML approaches can account for this matter. Thus, in the present application, the purely data-driven approaches benefit most from a large training set.

Comparing the different optimization methods, we find that with a sample size of $S = 0.4$ (150 days) and larger, the data-driven *SAA* method yields lower costs than its model-based counterpart *Norm* for most forecasting methods at a service level of 0.7.

This observation implies that a normal distribution assumption is beneficial in our case only if a very limited dataset is available or if the target service level is very high (see Section 2.4.5).

The performance and the ranking of the methods varies depending on the sample size. However, if more data are available, it is possible to employ a method that reduces the costs compared to the best method on the smaller dataset. For sample size 0.1, *ETS Norm* is the best approach, while costs can be reduced by 17.4% using an *DT-LGBM Norm* with a sample size of 1.0.

2.5. Conclusion

In this study, we propose a framework for how data can be leveraged in inventory problems on three different levels: demand estimation, optimization, and integrated estimation and optimization. We highlight that integrated estimation and optimization in the newsvendor problem is equivalent to the Quantile Regression problem, and we introduce novel data-driven methods for the newsvendor problem based on Machine Learning and Quantile Regression. Moreover, we empirically compare the methods to well-established standard approaches on a real-world dataset. We are specifically interested in the effect of data-driven approaches on the three levels on the overall performance.

The key result of our evaluation is that data-driven approaches outperform their model-based counterparts in most cases. In our evaluation, this finding already holds for a demand history of beyond 25 weeks (i.e. 150 data points). However, overall performance depends heavily on the demand estimation method employed. We found that poor forecasts cannot be compensated for by the choice of the subsequent optimization approach. Thus, the selection of the forecast model is the most crucial decision in the case of separated estimation and optimization.

The empirical evaluation of the Quantile Regression approaches revealed that integrating forecasting and optimization is beneficial only if enough data are available to estimate the conditional quantiles and limited to target service levels smaller than 0.8. When working with single time series, separate estimation and optimization yields superior results. This finding is in line with the empirical analysis of Ban and Rudin (2019).

More sophisticated estimation methods such as ANNs and Gradient Boosted Decision Trees require more training data in order to produce reliable results. However, these methods are also the only methods that constantly improve as more data becomes available. In our example, the demand history should contain more than six months of training data before employing Machine Learning. If a limited amount of data is available, simple methods such as the seasonal moving average can be suitable alternatives.

The major advantage of ML methods is that they are very flexible with respect to the input and that they are naturally able to process large datasets. The ability of ML methods to leverage similarities of time series across products and stores significantly improved their performance in our case. Additionally, they do not require restrictive assumptions on the demand process. Hence, they can identify patterns that traditional time series methods cannot detect. For instance, they can model multiple seasonalities (e.g. week and year), special days (e.g. public holidays), promotional activities and outliers (Barrow and Kourentzes 2018). A drawback of these approaches is that they are a black box, which makes it more difficult to justify the resulting predictions. However, when the improvements in forecast accuracy can be easily measured, as in the case of baked goods, the advantage of accurate predictions should outweigh the issue of interpretability.

Data-driven inventory management is an active field of research with a variety of opportunities for future work. Our analysis is based on a particular data set of bakery products. It would be interesting to repeat the analysis on other data sets, including other products. The methodology is applicable to perishable products with repetitive sales (bread, fresh produce, newsprint,...). In other newsvendor situations, little or no historical sales data may be available (fashion, electronics, sport events,...). In that case, forecasting requires other leading indicators than historical sales. It will be interesting to investigate the performance of alternative approaches to derive decisions from data under those circumstances.

We presented a data-driven approach for the single-item newsvendor model. It seems natural to explore the multi-product case as well. Particularly in the bakery domain, it is a common practice to plan safety stocks on the product category level. This step is reasonable because the substitution rates within a category in the case of stock-outs are high for perishable goods (Van Woensel et al. 2007). Thus, it could be possible to leverage hierarchical demand forecasts (Huber et al. 2017) in order to optimize inventory and to make globally optimal decisions. Especially for the multi-product case, joint capacity restrictions and lot sizes should also be considered.

Some bakery products can be sold over multiple days. Thus, expanding the model to a multi-period inventory model is reasonable. It would widen the application of the model to many other grocery products that can be reordered during the selling season. There are several papers that deal with the the multi-period problem with unknown demand distribution (e.g. Godfrey and Powell (2001), Levi et al. (2007)). Given the inherent similarity between reorder point calculations and newsvendor trade-offs, one may expect machine learning approaches to also be beneficial in that context.

In our application, there is no lead time. However, in other problem settings lead time plays an important role. Prak et al. (2017) show that using one-period-ahead forecast errors to optimize inventories leads to insufficient safety stock levels in case of a positive lead time.

In addition to the problem specific extensions, the methodology of the presented approaches may also be adjusted. Other machine learning approaches can be used for integrated forecasting and optimization, e.g., random forest or kernel methods.

Chapter III

Data-driven Inventory Management under Customer Substitution

with Jakob Huber, Moritz Fleischmann, and Heiner Stuckenschmidt[1]

Abstract

Most retailers that sell perishable goods offer multiple products in a product category (e.g., fresh food or fashion). Managing the inventories of these products is especially challenging due to frequent stock-outs and resulting substitution effects within the category. Furthermore, the true demand distributions of products are usually unknown to the decision maker. New digital technologies have enormously expanded the availability of data, storage capacity, and computing power and may thereby help improve inventory decisions. In this paper, we present a novel solution approach for the multi-product newsvendor problem. Our method is based on modern machine learning techniques that leverage large available datasets (e.g., data on historical sales, weather, store location, and special days) and are able to take complex substitution effects into account. We empirically evaluate our approach on two real-world datasets of a large German bakery chain. We find that our data-driven approach outperforms the model-based benchmark on the first dataset and performs competitively on the second dataset.

[1]The research presented in this chapter is based on a paper entitled "Data-driven Inventory Management under Customer Substitution" coauthored with Jakob Huber, Moritz Fleischmann and Heiner Stuckenschmidt.

3.1. Introduction

Stock-outs are common in many retail settings. When customers cannot find their preferred product in stock, some might choose a similar product instead (Gruen et al. 2002, Van Woensel et al. 2007). This substitution behavior makes inventory optimization for multi-product portfolios challenging due to the resulting interdependencies of stocking decisions. A second major challenge is that demand distributions are usually unknown and have to be estimated from the available data. Recent advances in data storage and processing technologies have led researchers to develop decision models that work with "big data" (Feng and Shanthikumar 2018). Although the adoption of data-based decision making has increased in practice, many companies are still struggling to turn data into better decisions (Brynjolfsson and McElheran 2016b).

In this paper, we address a multi-product newsvendor problem, where demand distributions are unknown. We present and test a novel solution approach based on machine learning (ML) that prescribes ordering decisions for substitutable products. Our method leverages the available data directly and without demand distribution assumptions and takes the complex substitution effects into account. Once the ML model has been trained, i.e., optimized, ordering decisions can be obtained instantaneously without solving the notoriously difficult multi-product newsvendor problem in each period.

Traditionally, solving the multi-product newsvendor problem with unknown demand distributions involves two steps: *demand estimation* and *inventory optimization*. The demand estimation problem and the inventory optimization problem are usually addressed separately in the literature and in practice.

The objective in the *estimation* problem is to find a demand model with a "good fit" to the data. There is a large set of well-established and refined estimation methods that predict demand distributions. Time-series approaches such as moving averages, exponential smoothing, and ARIMA have been widely used in practice and implemented in most forecasting software programs. More recently, ML approaches have been applied to forecasting tasks, as they are able to leverage large data sets due to advances in information technology (Carbonneau et al. 2008, Crone et al. 2011). The problem of optimal inventory levels is usually not addressed in the forecasting and ML literature.

The objective of the inventory *optimization* problem is to maximize the overall profit by setting appropriate inventory levels for each product *given the demand distributions*. While the classical single-product newsvendor problem is well-solved, the multi-product version is known to be hard to solve, due to the interdependencies between individual stocking decisions (Netessine and Rudi 2003, Schlapp and Fleischmann 2018). The operations research (OR) literature on this problem is mainly concerned with establishing theoretical properties (Parlar and Goyal 1984, Netessine and Rudi 2003, Schlapp and Fleischmann 2018) and developing efficient optimization algorithms (Hübner et al. 2016, Farahat and Lee 2018, Zhang et al. 2018). These studies assume that the demand distributions are known to the decision maker. Few papers study the interaction between estimation and optimization in a multi-product newsvendor setting (Kök and Fisher 2007, Sachs 2015).

Despite its widespread application in practice and research, the strict separation of estimation and optimization may lead to suboptimal solutions to the original inventory problem. Choosing a demand model based on “traditional” criteria (e.g. mean squared error) in the estimation step does not necessarily lead to a model that produces good decisions in the optimization step (Prak et al. 2017, den Boer and Sierag 2020, Bertsimas and Kallus 2020). Additionally, changes in the demand distribution estimates require a de novo execution of the optimization algorithm which is computationally expensive, especially in dynamic retail environments with many products and categories.

While den Boer and Sierag (2020) propose a model selection approach that selects models based on their corresponding decision quality rather than their “goodness of fit”, the need for a demand model can be questioned more fundamentally, the more data gets available to replace these model assumptions.

Instead of first estimating a demand model and then optimizing inventory decisions, integrated estimation and optimization approaches have been proposed for the single-product newsvendor problem (Liyanage and Shanthikumar 2005, Beutel and Minner 2012, Ban and Rudin 2019, Huber et al. 2019). We refer to these as *data-driven* approaches compared to the traditional *model-based* approaches because they do not require prespecified demand model assumptions. The multi-product version of the newsvendor problem is known to be much harder than its single-period counterpart due to the interdependencies of stocking decisions. To the best of our knowledge, no data-driven approach exists for the multi-product version for more than two products. Thus, the question as to what such an approach could look like remains to be answered.

We apply the idea of *data-driven* optimization to a multi-product newsvendor problem with unknown demand distributions. Our method builds on an artificial neural network (ANN) that is able to process large datasets. The ANN is trained on a problem-specific loss function that reflects the complex interactions between substitutable products. Once trained, the ANN prescribes inventory decisions without the need to solve an optimization problem in each period. We evaluate our approach based on two datasets of a large German bakery chain for two different product categories (breads and buns), each consisting of 6 products. The datasets comprise historical sales data and data on additional features, such as store location, opening hours, and weather data. From our analysis of the real-world problem, we derive insights for retailers on the importance of data and substitution for inventory decisions.

To summarize, our main contributions in this paper are as follows:

- We present a novel data-driven solution approach for the multi-product newsvendor problem.
- We document the viability of our method with real-world data.
- We derive insights for retailers with respect to considering data and substitution in inventory management decisions.

The remainder of this paper is structured as follows. In the next section, we summarize the related literature. Section 3.3 contains the formal problem description of the multi-product newsvendor problem with unknown demand distributions. In Section 3.4, we introduce our solution approach. In Section 3.5, we report and discuss the results of the empirical evaluation. Finally, we summarize our findings and outline opportunities for further research in Section 3.6.

3.2. Related Literature

In this section, we review the ML literature on demand forecasting, the OR literature on the multi-product newsvendor problem under customer substitution, and empirical studies on substitution rates.

A large part of ML is concerned with the prediction of a certain quantity (e.g. demand), given a large dataset of features. For an overview of ML methods, see Hastie et al. (2017). We restrict our review to references that use ML for demand forecasting. Carbonneau et al. (2008) use several ML methods to forecast demand in a supply chain. They find that recurrent neural networks and support vector machines are the most accurate forecasting techniques tested on their dataset. Barrow et al. (2010) study the performance of different neural network setups in a time-series forecasting task. Crone et al. (2011) report the results of the NN3 forecasting competition with a focus on ANNs. The results suggest that ANNs perform competitively with traditional statistical approaches but cannot outperform them. Barrow and Kourentzes (2018) use ANNs to predict call center arrivals from time-series data with a focus on outlying periods. In general, ML is concerned with prediction alone and does not consider optimal decision making. In our approach, we use ML, namely ANNs, to get from data to decisions directly, in a multi-product newsvendor problem.

There are many studies in the OR literature on the structural properties of the multi-product newsvendor problem under customer substitution and on the related solution algorithms. Kök et al. (2015) provide a broad review of the topic. Two main modeling approaches for substitution can be distinguished. Models in the first approach assume a specific customer choice model, e.g., a multinomial logit model (van Ryzin and Mahajan 1999, Musalem et al. 2010, Vulcano et al. 2010, 2012, Topaloglu 2013, Farahat and Lee 2018). Models in the second approach represent substitution by exogenous substitution rates. We follow the latter approach which is prevalent in the inventory management literature. For the two-product case, Parlar and Goyal (1984) show that the objective function of the maximization problem is concave under mild conditions and provide necessary optimality conditions. This work is extended for more than two products by Netessine and Rudi (2003), who derive necessary optimality conditions for the more general case. The authors study the centralized case and competition. Schlapp and Fleischmann (2018) include capacity restrictions in addition to substitution.

Another stream of the extant research focuses on efficient optimization algorithms to solve real-world inventory problems. To this end, Zhang et al. (2018) develop two mixed-integer linear programs that are able to solve problems of realistic sizes for many applications. For very large problems, they provide approximation algorithms. Closely related to our work is Kök and Fisher (2007), who describe a step-by-step approach from the estimation of substitution rates and demand distributions to the final inventory decision. They develop a heuristic for the problem, apply their approach to a large supermarket chain and are able to gain a large increase in profit relative to the current practice. Based on the work of Kök and Fisher (2007), Hübner et al. (2016) develop an optimal solution procedure and heuristics to increase solution quality and speed. The work on the multi-product newsvendor problem relies on the separate estimation of demand distributions and the optimization of inventory levels, where an optimization problem needs to be solved in each period if the distribution estimate changes. We integrate both problems into a single optimization problem that needs to be solved only once. While estimation and optimization have been integrated for the single-product problem (Liyanage and Shanthikumar 2005, Beutel and Minner 2012, Ban and Rudin 2019, Huber et al. 2019) and for the two-product case with a linear decision rule (Sachs 2015), we are not aware of any approach for more than two products and with a nonlinear ML approach.

To estimate exogenous substitution rates, Anupindi et al. (1998) propose an approach based on maximum likelihood estimation that works with inventory transaction data. They test their approach with data from vending machines. Kök and Fisher (2007) generalize this approach to dynamic choice processes. The estimation method in Fisher and Vaidyanathan (2014) is also based on MLE. Wan et al. (2018) compare the accuracy of a customer choice model (nested logit) to exogenous substitution rates in a multi-store environment. For the estimation of substitution rates in our empirical evaluation, we adapt the methodology of Anupindi et al. (1998), as all the necessary data are available.

There are several empirical studies that measure substitution rates for diverse product categories. For ground coffee, orange juice, peanut butter tomato sauce, and toothpaste, Emmelhainz et al. (1991) found that between 65% and 83% of customers substitute in response to a stock-out. Campo et al. (2000) find substitution rates of 44% and 51% for cereals and margarine, respectively. The most extensive study by Gruen et al. (2002) found that substitution rates vary significantly by category and are approximately 45% on average. Most related to our research is the work of Van Woensel et al. (2007). The authors investigate consumer responses to stock-outs of bakery bread and find that approximately 82% of customers are willing to substitute for another product if their first choice is unavailable. From these empirical studies, we conclude that substitution rates are high across different product categories and that the multi-product nature of the problem should therefore be taken into account when making inventory decisions.

To summarize, the literature on ML and OR is still relatively disjointed. The ML literature in the retail context is mostly concerned with estimating demand distributions from data and does not consider optimal decision making. The OR literature on the multi-product newsvendor problem focuses on theoretical properties and efficient optimization algorithms with strong distributional assumptions. To avoid these strong assumptions, we propose a method that integrates demand estimation and inventory optimization into a single optimization problem. Our approach can leverage large datasets, is able to reflect the inherently challenging substitution effects, and prescribes feature-dependent inventory decisions without having to solve an optimization problem in each period.

3.3. Problem Formulation

We study the multi-product newsvendor problem with stock-out-based substitution, where demand distributions are unknown. As a starting point, we follow the prevalent model formulation in the inventory management literature (Netessine and Rudi 2003, Kök et al. 2015, Schlapp and Fleischmann 2018).

Consider a retailer selling n partially substitutable products with uncertain demand D_i of product i over a finite selling season. The retailer must choose the order quantity q_i of product i before the selling season, such that the expected total profit Π is maximized. The unit sales price of product i is p_i, and the unit cost is c_i. Unsold units of product i that are left over at the end of the season have a unit salvage value of s_i. Naturally, $p_i > c_i > s_i \geq 0$.

To model substitution, we assume that a fraction $\alpha_{ji} \in [0, 1]$ of customers who cannot find their preferred product j in stock (i.e. when $D_j > q_j$) will substitute for product i, where $\sum_{i \neq j} \alpha_{ji} \leq 1$. This substitution behavior results in an inflation of the *initial demand* D_i of product i. The *substitution demand* of product i is $D_i^s = D_i + \sum_{j \neq i} \alpha_{ji}(D_j - q_j)^+$.

Thus, the retailer's objective is to maximize the expected total profit according to

$$\max_{q_i \geq 0} \Pi = \mathbb{E} \sum_i \left[u_i D_i^s - u_i (D_i^s - q_i)^+ - o_i (q_i - D_i^s)^+ \right], \tag{3.1}$$

$$= \sum_i \left(u_i q_i - (u_i + o_i) \mathbb{E} \left[q_i - D_i^s \right]^+ \right), \tag{3.2}$$

where $u_i = p_i - c_i$ and $o_i = c_i - s_i$ are product i's underage and overage costs, respectively.

The major difficulty in solving problem (3.1) is that D_i^s depends on the order quantities of the other products $j \neq i$. Netessine and Rudi (2003) show that the objective function with more than two products is not necessarily concave or quasiconcave, which makes the problem particularly hard to solve.

In contrast to the traditional problem, we assume that the probability distributions of the initial demand D_i for each product i are not known to the decision maker a priori. Instead, historical data $S_T = \{(\mathbf{d}_1, \mathbf{x}_1), ..., (\mathbf{d}_T, \mathbf{x}_T)\}$ are available, where $\mathbf{d}_t = [d_{1,t}, ..., d_{n,t}]$ is a vector of historical demand realizations of all n products, and $\mathbf{x}_t = [x_{1,t}, ..., x_{m,t}]$ is a vector of m covariates or *features* (e.g., store location, opening hours, and weather data) in period t.

The traditional approach to solving this problem would be first to estimate a model for the distribution of D_i for every product i from these data and then to optimize the actual inventory problem (3.1). Each time the distribution estimate changes (e.g., due to seasonality or the influence of other features), the problem needs to be solved again. We propose an alternative solution approach that integrates the estimation into the optimization problem and is optimized only once. We elaborate on our data-driven method and the traditional model-based approach in the following section and empirically compare them in Section 3.5.

3.4. Solution Approaches

3.4.1. "Traditional" Model-based Approach

Most of the inventory management literature on the multi-product newsvendor problem neglects the fact that the demand distributions of products are unknown to the decision maker (Netessine and Rudi 2003, Kök et al. 2015, Schlapp and Fleischmann 2018). The extant papers that address both the inventory problem and the demand estimation problem use a two-step procedure. First, estimating a demand model. Second, optimizing the inventory decisions based on the estimated demand distributions (Kök and Fisher 2007).

Estimation

If only historical demand data $\mathbf{d}_1, ..., \mathbf{d}_T$ are available, one can approximate the actual demand distributions with a parametric distribution or the empirical distribution. If additional feature data $\mathbf{x}_1, ..., \mathbf{x}_T$ (e.g. weekdays, opening hours, weather data) are available that are correlated with demand, the estimates might be improved because the *conditional* forecast can be more accurate (Ban and Rudin 2019, Huber et al. 2019).

Based on a specific measure for "goodness of fit" (e.g., mean squared error), one would choose a demand model that "fits" the data $S_T = \{(\mathbf{d}_1, \mathbf{x}_1), ..., (\mathbf{d}_T, \mathbf{x}_T)\}$ well. The estimate for the demand distributions is then constructed from the point forecast of the demand model that has the best "fit" and the historical forecast errors of that model, i.e. residuals.

For our analysis, we choose an exponential smoothing approach as it has been shown to perform well in a single product setting, even compared to more complex ML methods that leverage more data (Huber et al. 2019). In our implementation, we use the `ets()` function from the `forecast` package (Hyndman and Khandakar 2008) for the statistical software R (R Core Team 2017). It is based on Hyndman et al. (2002, 2008) who propose innovation space models that generalize exponential smoothing methods (*ETS*). These models include a family of 30 models that cover different types of errors, seasonal effects and trends (none, additive, multiplicative).

Optimization

Given an estimate $\hat{D}_i$ for the demand distribution of each product i from a demand estimation procedure, the next step is to solve the original multi-product newsvendor problem (3.1). This problem is known to be notoriously hard to solve due to the nonconvexity of the objective function (Netessine and Rudi 2003). Obtaining optimal analytical solutions is intractable, and there are only a few efficient solution algorithms that provide near-optimal solutions in reasonable time for real-world problem sizes (Hübner et al. 2016, Zhang et al. 2018). Zhang et al. (2018) developed two mixed-integer linear program (MILP) formulations of the problem that we use throughout the paper for the optimization part of the model-based solution approach. For expositional purposes, we introduce only the first formulation and refer the reader to Zhang et al. (2018) for more details.

$$\max_{q_i \geq 0} \sum_i \left(u_i q_i - (u_i + o_i) \frac{1}{n} \sum_t y_{i,t} \right) \tag{3.3}$$

subject to:

$$y_{i,t} \geq q_i - \hat{d}_{i,t} - \sum_{j \neq i} \alpha_{ji} v_{i,t} \qquad \forall i, t \tag{3.4}$$

$$v_{i,t} \leq \hat{d}_{i,t} - q_i + M_i z_{i,t} \qquad \forall i, t \tag{3.5}$$

$$v_{i,t} \geq \hat{d}_{i,t} - q_i - M_i z_{i,t} \qquad \forall i, t \tag{3.6}$$

$$v_{i,t} \leq \hat{d}_{i,t}(1 - z_{i,t}) \qquad \forall i, t \tag{3.7}$$

$$v_{i,t}, y_{i,t} \geq 0 \qquad \forall i, t \tag{3.8}$$

$$z_{i,t} \in \{0, 1\} \qquad \forall i, t \tag{3.9}$$

The authors reformulate the expectation in the objective function of the original problem (3.1) as a finite summation over the discrete demand estimates $\left(\hat{d}_{i,t}\right)_{t=1,\ldots,T}$. If $\hat{D}_i$ is assumed to be continuous, then one might also generate i.i.d. samples from $\hat{D}_i$.

$y_{i,t} = q_i - \hat{d}_{i,t} - \sum_{j \neq i} \alpha_{ji}(\hat{d}_{j,t} - q_j)^+$ and $v_{i,t} = (\hat{d}_{i,t} - q_i)^+$ represent overages and underages, respectively. Constraints (3.4) to (3.8) ensure that these equations hold. To linearize the $(\cdot)^+$ functions, the formulation uses binary variables $z_{i,t}$, where $z_{i,t} = 1$ if $q_i \geq \hat{d}_{i,t}$ (i.e., $v_{i,t} = 0$) and $z_{i,t} = 0$ if $q_i < \hat{d}_{i,t}$ (i.e., $v_{i,t} = \hat{d}_{i,t} - q_i$). M_i is an upper bound of the order quantity q_i of product i.

3.4.2. Data-driven Approach

Although the demand model might have a “good fit” to the data based on the loss minimization problem, choosing the model according to a criterion that does not incorporate the quality of the decisions can lead to suboptimal results for the overall problem if the model is misspecified (Liyanage and Shanthikumar 2005, den Boer and Sierag 2020). Additionally, if the distribution estimate of the demand model changes (e.g., through changes in feature values), then the optimization model needs to be re-solved.

We propose the estimation of the optimal order quantities directly from data with an ML approach (e.g., ANN) without prespecifying a demand model. To this end, we express the order quantities $q_{i,t}$ as a function of the feature vector $\mathbf{x}_t$ and parameters $\mathbf{W}$

of the ML approach (e.g., the weight matrix of an ANN). Instead of minimizing the L_2 norm, we replace the loss function with the objective function of the actual multi-product newsvendor problem.

$$\max_{\mathbf{W}} \frac{1}{n} \sum_{t} \sum_{i} \left(u_i q_{i,t}(\mathbf{W}, \mathbf{x}_t) - (u_i + o_i) \left(q_{i,t}(\mathbf{W}, \mathbf{x}_t) - d^s_{i,t} \right)^+ \right), \tag{3.10}$$

where $d^s_{i,t} = d_{i,t} + \sum_{j \neq i} \alpha_{ji}(d_{j,t} - q_{j,t}(\mathbf{W}, \mathbf{x}_t))^+$.

Given the historical dataset of demand and features $S_T = \{(\mathbf{d}_1, \mathbf{x}_1), ..., (\mathbf{d}_T, \mathbf{x}_T)\}$, we can train the network, i.e., optimize the parameters $\mathbf{W}$ in (3.10). Note that compared to the model-based approach, the decision variables of the optimization problem are no longer the inventory decisions but rather the parameters of a regression model. This regression model is then used to estimate the optimal decisions.

Our data-driven approach has several theoretical advantages compared to the traditional approach. First, it does not rely on a specific demand model that has been selected on the basis of a "goodness-of-fit" criterion. Second, demand uncertainty is feature-dependent, while in the model-based approach, only the mean of the demand distribution is feature-dependent. Third, once the model is trained, it prescribes a decision for any combination of features, without solving an additional optimization problem in each period. The latter leads to less computational effort compared to the model-based approach, where the optimization problem (3.3) - (3.9) has to be solved in each period.

While (3.10) is generally indepentent of a specific regression model, we choose an ANN to represent the order quantity in our approach because it is able to approximate any continuous function (Hornik 1991), which makes it suitable for forecasting if enough data are available, and it is difficult to specify the underlying data generation process. We rely on feed-forward neural networks, i.e., multilayer perceptrons. In a feed-forward neural network with L hidden layers ($L \geq 1$), the output $\mathbf{h}^{(k)}(\mathbf{x})$ of each layer k gets passed to the next layer ($1 \leq k \leq L+1$):

$$\mathbf{h}^{(k+1)}(\mathbf{W}, \mathbf{x}_t) = \sigma^{(k+1)}(\mathbf{b}^{(k+1)} + \mathbf{W}^{(k+1)}\mathbf{h}^{(k)}(\mathbf{x})) \tag{3.11}$$

The output of the input layer is defined as $\mathbf{h}^0(\mathbf{x}) = \mathbf{x}$, while the output of the last layer represents the prediction of the network, i.e., the order quantities, $\mathbf{q}(\mathbf{W}, \mathbf{x}) =$

$\mathbf{h}^{(L+1)}(\mathbf{W}, \mathbf{x})$. The output of each layer is connected with a fully connected weight matrix $\mathbf{W}^{(k)}$ to the next layer. The input of a layer is adjusted with the biases $\mathbf{b}^{(k)}$ of each neuron before it passes an activation function $\sigma^{(k)}$.

The performance of an ANN depends on its architecture and the hyperparameter settings. We optimize these hyperparameters (learning rate, patience for early stopping, batch size, number of hidden nodes, number of hidden layers, and activation function of hidden layers) with tree of parzen estimators (Bergstra et al. 2011, 2013) in combination with cross-validation of the training set. We use rectified linear units (RELUs) as activation functions at the input layer and the hidden layers and a linear activation function at the output layer (Zhang et al. 1998). Moreover, we train and employ an ensemble of multiple *ANNs* with the *median* ensemble operator, as this approach is robust to the initial weights (Barrow et al. 2010, Kourentzes et al. 2014). To optimize the weights of an ANN, we use the stochastic gradient-based algorithm ADAM proposed by Kingma and Ba (2015) in combination with early stopping to avoid overfitting. The input consists of the lagged demand information of each product and further explanatory features (see Section 3.5).

3.5. Empirical Evaluation

3.5.1. Data Description and Preparation

We evaluate our proposed approach using data from a large German bakery chain. In this section, we describe the datasets and the data preparation process. The datasets comprise the hourly sales data of the six most frequently sold stock-keeping units from the product categories buns and breads for nine stores over a period of 987 days. We enrich the datasets with additional master data, transactional data, and external explanatory features related to the calendar, weather, and location of the store (see Table 3.1).

Data Source	Features
Master Data	store class, product category, opening times (day, hours/duration)
Transactional Data	lagged sales, rolling median of sales, binary promotional information
Calendar	day of year, month of year, day of month, day of week, public holiday, day type, bridge day, nonworking day, indicators for each special day, school holidays
Weather	temperature (minimum, mean, and maximum) and cloud cover of target day
Location	general location (city, suburb, and town); in proximity to the store: shops (bakeries, butcher, grocery, kiosk, fast food, and car repair), amenities (worship, medical doctors, and hospitals), leisure (playground, sports facility, and park), education (kindergarten, school, and university)

Table 3.1.: Features used in the machine learning methods.

To apply and compare our optimization approaches, we need price and cost parameters, substitution rates, and daily demand data. We report price and cost parameters for both product categories in Tables 3.2 and 3.3. Unit prices can be directly observed, whereas unambiguous cost parameters cannot be obtained due to varying cost accounting methods and parameters. The unit costs for each category are based on expert judgment.

Buns	P1	P2	P3	P4	P5	P6
Unit price [EUR]	0.30	0.50	0.60	0.50	0.50	0.50
Unit cost [EUR]	0.06	0.06	0.06	0.06	0.06	0.06

Table 3.2.: Price and cost parameters for buns.

Breads	P1	P2	P3	P4	P5	P6
Unit price [EUR]	1.75	2.45	2.45	1.70	2.85	2.95
Unit cost [EUR]	0.60	0.60	0.60	0.60	0.60	0.60

Table 3.3.: Price and cost parameters for breads.

We estimate the substitution probabilities similar to Anupindi et al. (1998). Based on the transaction data, the method measures the spillover demand from stocked-out products to the products that are still available. From the magnitude of spillover demand in stock-out periods compared to non-stock-out periods, substitution rates can be estimated. As the assumption of stationary demand during the day does not hold in our case, due to a strong intraday sales pattern, we apply the approach to each hour of the day. We obtain an estimate for the substitution matrices for each hour of the day and compute the average over all hours. The results are shown in Tables 3.4 and 3.5.

↗	P1	P2	P3	P4	P5	P6	Total
P1	-	0.11	0.15	0.14	0.14	0.15	0.69
P2	0.22	-	0.08	0.11	0.12	0.12	0.65
P3	0.24	0.07	-	0.07	0.08	0.07	0.53
P4	0.26	0.09	0.07	-	0.12	0.12	0.66
P5	0.19	0.10	0.10	0.12	-	0.14	0.65
P6	0.17	0.13	0.11	0.11	0.13	-	0.65

Table 3.4.: Estimated substitution rates of buns (from row to column).

↗	P1	P2	P3	P4	P5	P6	Total
P1	-	0.14	0.09	0.25	0.12	0.12	0.72
P2	0.19	-	0.09	0.19	0.14	0.14	0.75
P3	0.07	0.09	-	0.07	0.10	0.10	0.43
P4	0.17	0.16	0.09	-	0.20	0.10	0.72
P5	0.18	0.12	0.09	0.24	-	0.12	0.75
P6	0.16	0.13	0.08	0.22	0.15	-	0.74

Table 3.5.: Estimated substitution rates of breads (from row to column).

Overall, we estimate that 43 % to 75 % of customers are willing to substitute for another product if their first choice is not available. To validate our results, we compare our estimates to earlier empirical work on substitution. Van Woensel et al. (2007) found rates of 75% to 82% for bakery products. However, the product portfolio in their study was much larger (208 products); therefore, it is also more likely that a substitution with a more similar product takes place and that substitution rates are higher. We note that in the category buns, the substitution rates *to* product 1 are relatively high. The same is true for product 4 in the category breads. These products stand out, as they have the lowest price within their respective categories.

A main drawback of point-of-sale data (e.g., our dataset) is that lost demand is unobservable in cases of stock-outs. Frequent stock-outs distort the sales data in two ways. First, the sales data of the out-of-stock product are *censored* if customers cannot find their preferred product and choose to substitute another product or leave the store without buying anything at all. Second, the sales data of the substitute products is *inflated* by the demand of the out-of-stock products. To address these problems, we *decensor* and *deflate* historical sales data by applying the following procedure to each product category on each day:

Define the set of all products of the category as $L = \{P_1, ..., P_n\}$;
while *at least one product in L goes out of stock* **do**
  find product $\tilde{P}$ that goes out of stock first;
  delete $\tilde{P}$ from L;
  decensor the demand of product $\tilde{P}$ based on Lau and Lau (1996);
  deflate the demand of all products in L based on substitution rates;
end

The *decensoring* of sales in Step 4 is based on Lau and Lau (1996). In particular, we calculate the average hourly share of demand for product $\tilde{P}$ in relation to the total demand on days on which the product was not sold out. Based on this averaged intraday demand pattern, we extrapolate the sales data when a stock-out occurs to estimate uncensored demand. To *deflate* the demand of the products in L, we subtract the sales in each hour of stock-out of product $\tilde{P}$ that are due to substitution demand.

Table 3.6 shows the average share of the daily demand of each product within each category. While this proportion is relatively homogeneous for breads (8.5% to 23.8%), product 1 dominates buns with a share of 64.3%.

Category	P1	P2	P3	P4	P5	P6
Buns	0.643	0.107	0.070	0.068	0.057	0.055
Breads	0.220	0.139	0.205	0.238	0.113	0.085

Table 3.6.: Average share of the demand of each product within each category.

We split the dataset into a training set, containing 110 weeks, and a test set, containing the remaining 31 weeks, and perform a rolling 1-step-ahead prediction evaluation on the test set to assess the performance of the methods. We fit the models and determine the error distributions every 14 days on a rolling training dataset of constant size.

3.5.2. Results

In this section, we report the results of our data-driven approach relative to the ex post optimal decision, i.e., relative to the optimal decision of the deterministic problem after demand has realized. To put these results into perspective and compare them to a traditional approach that does not leverage additional feature data, we employ a model-based approach. To this end, we use exponential smoothing as discussed in 3.4.1 in the estimation step and the MILP discussed in Section 3.4.1 in the optimization step. We solve the MILP to near optimality (optimality gap $\leq 0.01\%$) with Gurobi 8.1.

Figure 3.1 depicts the relative profit to the ex post maximum profit (= 1) of the data-driven and model-based approaches.

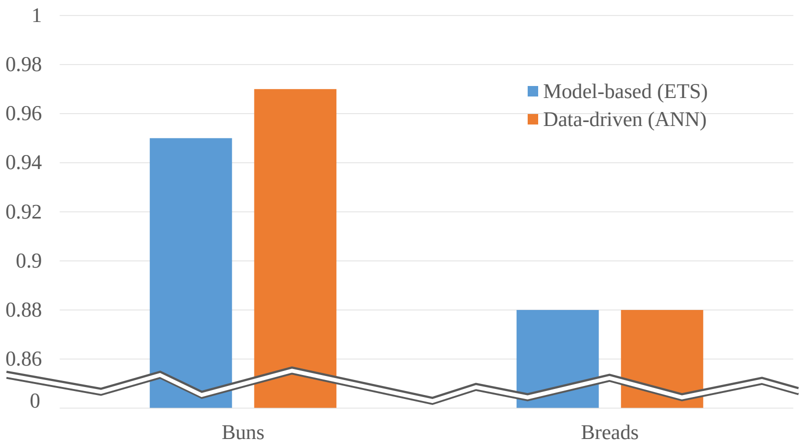

Figure 3.1.: Average profit relative to ex post maximum profit.

We observe that the traditional model-based approach with the exponential smoothing forecast (ETS) and MILP optimization achieves 95% and 88% of the maximum profit for the categories buns and breads, respectively. Our data-driven method outperforms the benchmark by 2 percentage points in the category buns and achieves the same average profit in the category breads. Note that the exponential smoothing benchmark already performs quite well, even without additional feature data. This confirms previous results, where seasonal exponential smoothing approaches have proven to perform well in sales forecasting tasks (Huber et al. 2019). As we discuss in Section 3.5.2 in more detail, calendar information, especially days of the week, is a very important feature in our setting due to the weekly seasonal demand pattern. This information is also available to the exponential smoothing method.

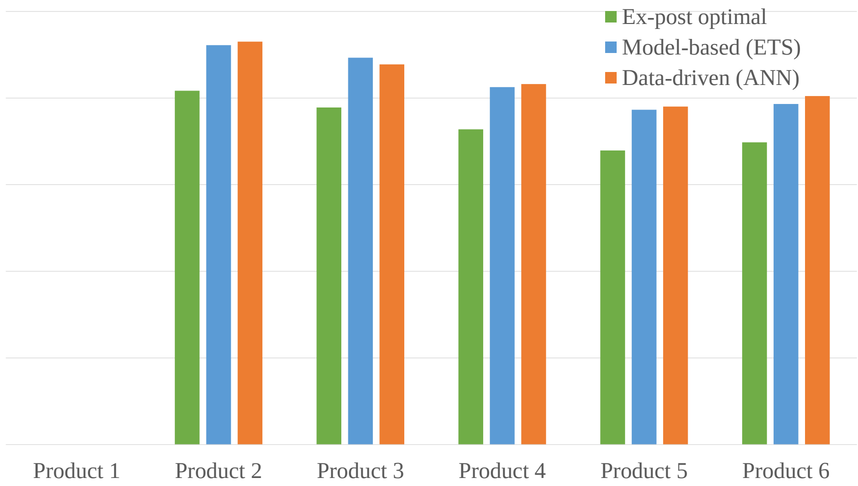

(a) Category: buns.

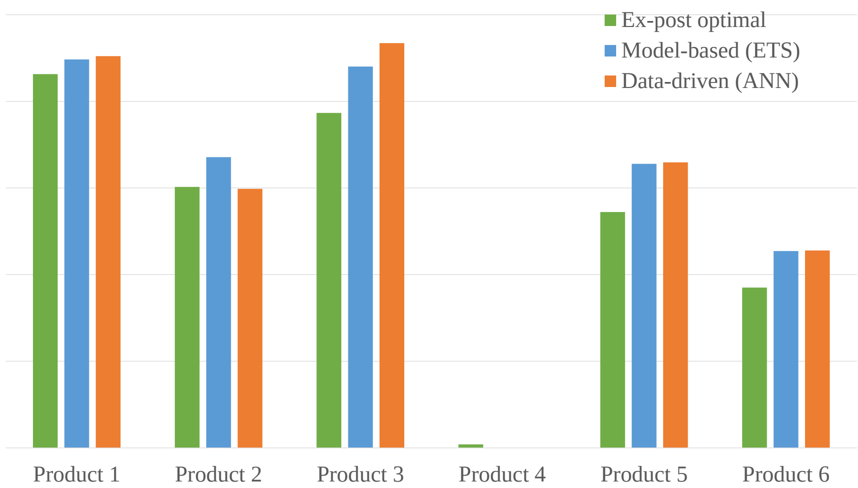

(b) Category: breads.

Figure 3.2.: Average order quantities per product and product category. The scale is not provided for reasons of confidentiality.

In Figure 3.2, we report the resulting average order quantities of our data-driven approach, the model-based benchmark, and the ex post optimal decisions. For buns and breads, it is optimal not to order product 1 and product 4, respectively. Both products have the lowest profit margin in their respective category (see Tables 3.2 and 3.3). Additionally, the substitution rates from these products to other products are relatively high. Therefore, it is optimal to "force" customers to substitute for higher-margin products by not including the lower-margin products in the assortment. We observe that both the model-based and the data-driven approaches follow this logic. Naturally, the order quantities of the data-driven and model-based approaches are generally higher than the optimal decision for the ex post deterministic problem, as they include safety stock to buffer against uncertainty.

To summarize, the data-driven method slightly outperforms the model-based benchmark on the first dataset, while it is competitive on the second dataset. In the following, we provide more insights into the effects that influence the decisions and performance of our approach.

The Effect of Demand Estimation Accuracy

A reason for the superior performance of our data-driven approach might be that it is able to leverage more data than the model-based exponential smoothing method. Therefore, we employ a model-based approach, where an ANN represents the demand model instead of ETS. Thus, the model is able to process the same data as our data-driven approach.

To measure the effect of demand estimation accuracy, we compute a set of common accuracy measures, including symmetric mean absolute percentage error (SMAPE), mean absolute percentage error (MAPE), mean absolute scaled Error (MASE) (Hyndman and Koehler 2006), root mean square error (RMSE), mean absolute error (MAE), and relative absolute error (RAE) for both the ETS and ANN forecast. We provide more than one measure because each has its strengths and weaknesses. For instance, RMSE and MAE are scale-dependent error measures and do not allow for comparisons between time series at different scales, while percentage-based error measures (SMAPE and MAPE) are not always defined and may result in misleading outcomes if demand is low. Table 3.7 shows the average forecast accuracy over all time series by method.

Category	Method	RMSE	MAE	MAPE	SMAPE	MASE
Buns	ETS	75.15	28.50	20.39	18.91	0.84
	ANN	44.13	22.00	16.75	15.59	0.67
Breads	ETS	4.77	3.33	36.47	28.04	0.83
	ANN	4.70	3.25	35.34	27.61	0.82

Table 3.7.: Forecast performance of the point predictions.

We observe that the ANN forecast outperforms the ETS forecast in all accuracy measures across both categories. The difference is relatively small for the category breads, while it is relatively large for the category buns. Although the ANN forecast is more accurate, its impact on the overall performance in terms of profit is not immediately clear. Thus, we use a two-step method, where we first forecast a demand distribution with ETS or ANN and then optimize based on the MILP (3.3) to (3.9). We report the resulting profits in Figure 3.3.

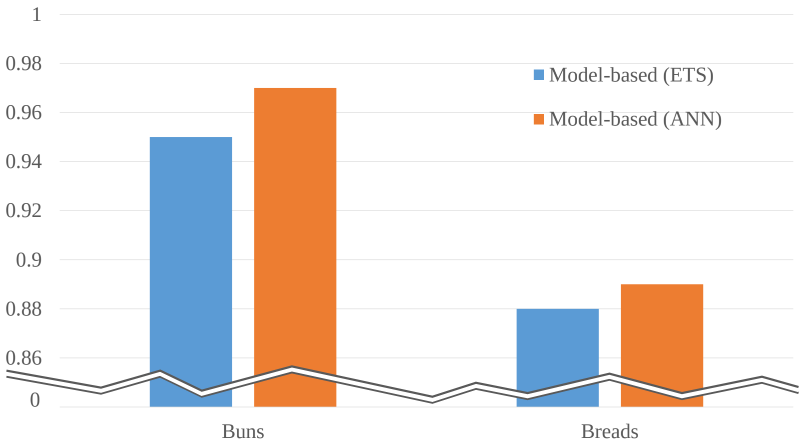

Figure 3.3.: Average profit relative to ex post maximum profit.

Not surprisingly, higher forecast accuracy leads to higher profits. Using the *ANN* as the forecast method in the model-based approach yields better results compared to the *ETS* forecast, with 2 percentage points in the category buns and 1 percentage point in the category breads. Recall, however, that the optimization step of the model-based approach has to be re-solved in each period, whereas the data-driven approach is trained only once.

When comparing the model-based ANN approach to the data-driven ANN approach, we observe that the performance is similar for both approaches, while the model-based approach even outperforms the data-driven approach by 1 percentage point in the category breads. We conclude that the empirical distribution estimate of the *ANN* is more accurate than the estimate of the *ETS* approach, which results in higher profits, and the profit increase of our data-driven approach is likely due to its ability to leverage additional feature data.

The Effect of Feature Data

We investigate the importance of the individual groups of feature data that are leveraged by our data-driven method. We use permutation feature importance (PFI), which measures the profit loss of the model after the feature values are permuted. PFI was originally introduced for random forest prediction models by Breiman (2001). We adapt the methodology and measure the profit loss after permutation instead in the increase of the prediction error after permutation. To this end, we randomly shuffle one set of features (see Table 3.1) at a time and therefore suppress its contribution to overall profit. For each set of features, we repeat the permutation 1,000 times. The permutation of an important feature with high prediction value will result in a high profit loss. The permutation of an unimportant feature will result in a low or no profit loss. The results are depicted in Figure 3.4.

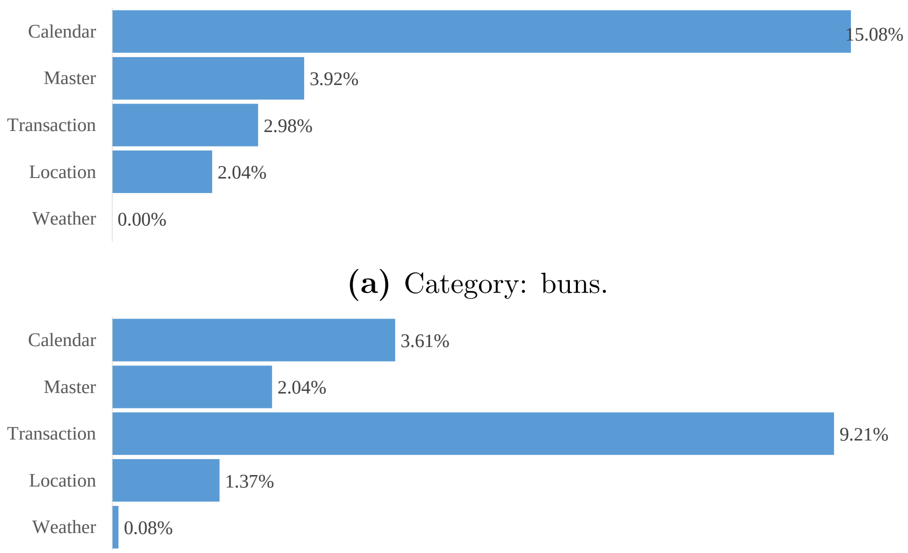

(a) Category: buns.

(b) Category: breads.

Figure 3.4.: Profit loss through permutation feature importance (PFI).

We observe that for both categories, calendar data, master data, and transaction data are the most important feature sources, although to a varying extent. Surprisingly, weather data have almost no value for optimizing inventory decisions in our datasets, even though they are used in many models for food retail (e.g., Beutel and Minner (2012), Kök and Fisher (2007)). For the category buns, excluding calendar data results in a substantial profit loss of 15.08%, while it is only 3.61% for the category breads.

Similarly, transaction data have different importance in both datasets. The impact of location data is relatively small (1.37 and 2.04%) in both datasets.

From the feature analysis, we conclude that calendar data, master data, and transaction data are important features in our dataset. Although breads and buns are relatively similar (fresh bakery products), the extent of profit loss through PFI varies greatly. Therefore, retailers should carefully select features for inventory optimization.

The Effect of Substitution

Next, we investigate the effect of substitution on the decisions and resulting average profits. To this end, we compare our data-driven approach and the model-based benchmark with their single-product newsvendor counterparts that ignore substitution effects. We report the percentage of lost profit by ignoring substitution in Figure 3.5.

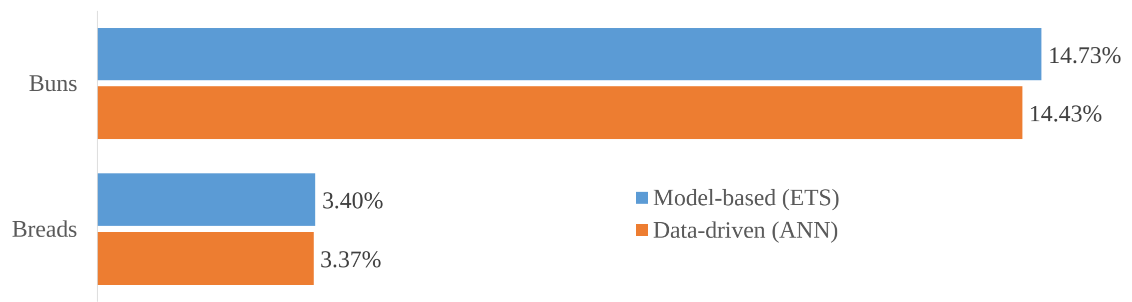

Figure 3.5.: Profit loss through ignoring substitution.

We observe that ignoring substitution leads to a profit loss of approximately 14.5%, independent of the decision method used, in the category buns. For the category breads the loss through ignoring substitution is much lower (approximately 3.4%), also independent of the decision method. The importance of substitution effects in the category buns is mainly due to the specific characteristics of product 1. Product 1 has by far the largest share of demand within the category (64.3%), the lowest profit margin, and the highest rates of substitution for other products (69%). Due to these properties, it is optimal to not order product 1 at all (fill rate = 0), as enough customers substitute for higher-margin products. We see a similar but smaller effect with product 4 in the category breads. It is also the lowest-margin product within the category, but it is not as dominant in terms of demand share (23.8%).

Overall, we find that accounting for substitution is a very important aspect that drives the performance of the methods considered. The large demand shifts away from product 1 in the category buns and from product 4 in the category breads result in large profit gains. However, they might be in conflict with other strategic objectives of the retailer. We discuss this issue in the following section and propose the implementation of fill-rate constraints.

The Effect of Fill-rate Constraints

In the previous section, we observed large demand shifts from low-margin products to substitutes due to very low ordering decisions for these products, which even resulted in an abandonment of these products from the assortment. We see similar effects in other applications of multi-product inventory models to real-world problems. Kök and Fisher (2007) find that "[p]roducts with low profit are dropped from the assortment, the number of facings of products with low marginal return are reduced, and the number of facings of those with higher returns are increased." They state that their recommendation suggests a more than 50% increase in profits compared to the current state. However, these decisions optimize the short-term profit. The long-term effects of stock-outs (e.g., dissatisfied customers and future demand losses) are not captured in these models, although they might be important for customers' store choice (Briesch et al. 2009). Including these long-term effects in the underage costs u_i of each product i is difficult, as long-term shortage costs are extremely difficult to estimate. Therefore, we propose adding target fill-rate constraints to the MILP (3.3) to (3.9) that account for strategic service level requirements.

We set a target fill-rate per category as

$$\beta^{cat} \leq \frac{\frac{1}{n}\sum_t \sum_i (q_i - y_{i,t})}{\frac{1}{n}\sum_t \sum_i D_{i,t}}. \tag{3.12}$$

Additionally, we introduce target fill-rate constraints per product as

$$\beta_i^{prod} \leq 1 - \frac{\frac{1}{n}\sum_t v_{i,t}}{\frac{1}{n}\sum_t D_{i,t}} \quad \forall i. \tag{3.13}$$

We set the category fill rate to values between 0.80 and 0.99 and the product fill rate to values between 0.80 and 0.95 and optimize the MILP based on the *ETS* and *ANN* distribution forecast. The resulting profits and fill rates are shown in Table 3.8 and Figure 3.6.

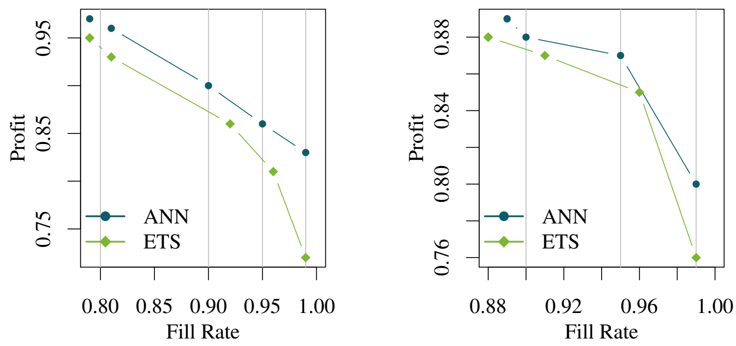

(a) Category: buns. (b) Category: breads.

Figure 3.6.: Average profit relative to ex post maximum profit of model-based approaches with fill-rate constraints at the category level.

Across all methods and categories, we can observe a trade-off between short-term profit and service (fill rate), which is illustrated in Figure 3.6. As soon as one of the fill-rate constraints is effective, the profit decreases. In the category buns, a product fill rate of 0.80 already drastically reduces the profit as the demand shifting away from product 1 to more profitable products is suppressed. In the category breads, this effect is smaller, as the products are more homogeneous with respect to volume and margin. Low category fill-rate constraints do not harm the profit too much, as they still allow for substitution.

When the long-term objectives of the retailer (e.g., attracting customers by means of a large assortment) are in conflict with the short-term objective of optimized inventories, fill-rate constraints can help achieve desired trade-off.

3.6. Conclusions

In this paper, we study the multi-product newsvendor problem with unknown demand distributions. We introduce a data-driven solution approach that takes the complex substitution effects into account and goes from data to decision in a single optimization problem. We empirically compare our method to a traditional model-based approach on a real-world dataset of a large German bakery chain. Furthermore, we evaluate

	Method	Target Fill Rate		Profit	Fill Rate						
		per Cat.	per Prod.		Prod. 1	Prod. 2	Prod. 3	Prod. 4	Prod. 5	Prod. 6	Cat.
Category Buns	Model-based (ANN)	-	-	0.97	0.00	1.00	1.00	1.00	1.00	1.00	0.79
		-	0.80	0.86	0.82	1.00	1.00	1.00	1.00	1.00	0.96
		-	0.90	0.84	0.92	1.00	1.00	1.00	1.00	1.00	0.98
		-	0.95	0.84	0.96	0.99	1.00	0.99	1.00	0.99	0.99
		0.80	-	0.96	0.05	1.00	1.00	1.00	1.00	1.00	0.81
		0.90	-	0.90	0.54	1.00	1.00	1.00	1.00	1.00	0.90
		0.95	-	0.86	0.80	1.00	1.00	1.00	1.00	1.00	0.95
		0.99	-	0.83	0.98	1.00	1.00	1.00	1.00	1.00	0.99
	Model-based (ETS)	-	-	0.95	0.00	0.99	1.00	1.00	1.00	1.00	0.79
		-	0.80	0.84	0.83	0.99	0.99	0.99	0.99	0.99	0.95
		-	0.90	0.82	0.93	0.99	0.99	0.99	0.99	0.99	0.97
		-	0.95	0.80	0.96	0.99	0.99	0.99	0.99	0.99	0.98
		0.80	-	0.93	0.13	1.00	1.00	1.00	1.00	1.00	0.81
		0.90	-	0.86	0.64	0.99	0.99	0.99	1.00	0.99	0.92
		0.95	-	0.81	0.87	0.99	0.99	0.99	0.99	0.99	0.96
		0.99	-	0.72	0.98	0.99	0.99	0.99	0.99	0.99	0.99
Category Breads	Model-based (ANN)	-	-	0.89	0.97	0.96	0.99	0.00	0.99	0.98	0.89
		-	0.80	0.87	0.91	0.93	0.98	0.85	0.98	0.96	0.94
		-	0.90	0.86	0.93	0.93	0.97	0.93	0.97	0.97	0.96
		-	0.95	0.82	0.97	0.95	0.97	0.97	0.98	0.98	0.98
		0.80	-	0.89	0.97	0.96	0.99	0.00	0.99	0.98	0.89
		0.90	-	0.88	0.98	0.96	0.99	0.12	0.99	0.98	0.90
		0.95	-	0.87	0.93	0.94	0.98	0.83	0.98	0.96	0.95
		0.99	-	0.80	0.98	0.96	0.99	0.99	0.98	0.98	0.99
	Model-based (ETS)	-	-	0.88	0.97	0.97	0.98	0.00	0.99	0.98	0.88
		-	0.80	0.86	0.90	0.95	0.96	0.84	0.98	0.97	0.94
		-	0.90	0.85	0.93	0.95	0.96	0.93	0.97	0.97	0.96
		-	0.95	0.81	0.97	0.96	0.97	0.96	0.98	0.98	0.98
		0.80	-	0.88	0.97	0.97	0.98	0.00	0.99	0.98	0.88
		0.90	-	0.87	0.97	0.97	0.98	0.23	0.99	0.98	0.91
		0.95	-	0.85	0.94	0.95	0.97	0.87	0.98	0.97	0.96
		0.99	-	0.76	0.99	0.97	0.99	0.99	0.99	0.98	0.99

Table 3.8.: Average profit relative to ex post maximum profit and resulting fill-rates under fill-rate constraints.

the effects of demand estimation accuracy, feature data, and substitution on overall performance in our dataset.

Our key result is that our approach outperforms a state-of-the-art method on the first dataset and performs competitively on the second dataset. From our forecast accuracy analysis, we conclude that the performance difference is mainly due to the additional data that can be leveraged by our method compared to the benchmark approach. In addition to the benefit of slightly better performance, our approach prescribes a decision for any combination of feature data values once it is trained. There is no need to solve an optimization problem in each period; only infrequent updates of the parameters (in our case, every two weeks) are needed. Furthermore, it does not rely on a demand distribution assumption and thus mitigates the risk of a misspecified demand model. Our approach can easily integrate additional explanatory feature data if available. Similar to previous research (Kök and Fisher 2007), we find in our empirical analysis that it is optimal for the maximization of short-term profits to delete a low-margin product from the assortment of each category. However, this might conflict with the long-term objectives (e.g., attracting customers through a large assortment) of the retailer. Therefore, we propose fill-rate constraints that account for this trade-off.

In addition to the methodological findings, we also provide empirical insights that might be valuable for inventory managers in fresh food retail. First, we estimate substitution rates from real-world sales data. Our estimates are in line with the previous research on bakery products, which directly asked customers for their substitution behavior (Van Woensel et al. 2007). Thus, we provide evidence for the applicability of data-based approaches for substitution estimation. Second, we find that ignoring substitution effects can result in significant profit reductions depending on product characteristics. In our case, ignoring substitution in the category buns resulted in a profit loss of approximately 14.5% and in the category breads of approximately 3.4%. Third, we analyze the impact of different feature data sources (master data, transactional data, calendar, weather, and location) on overall profits. While the impact depends on the product category, fresh food retailers should especially focus on calendar data, master data, and transaction data.

The interface between ML and OR is an active field of research that provides many opportunities for future work. In this paper, we focus on the bakery industry. It would be interesting to apply our method to other data-rich areas of inventory management, e.g., e-commerce, and analyze the impact of more feature data. To incrementally improve our method, we could apply more advanced ML methods to further improve forecast accuracy (e.g., recurrent neural networks). A more subnational extension of the analysis could include multi-period considerations, where products can be reordered during the selling season.

Chapter IV

Prescriptive Analytics for Dynamic Inventory Management[1]

Abstract

Dynamic inventory problems occur in a large number of industrial, distribution, and service applications. Most models in the literature assume that the demand distribution is stationary and known a priori in these contexts. In reality, this is rarely the case. Instead, the actual demand distribution is unknown and may change over time. The recent explosion of data availability may help to avoid unrealistic assumptions. In this paper, we combine concepts of classical dynamic inventory management and machine learning to develop a data-driven solution approach that leverages auxiliary data, including, weather, location, and calendar data, for the stochastic dynamic inventory problem. In a computational study, we compare our approach to an extant non-parametric method and to the optimal dynamic programming solution. We find that our approach performs close to the optimal solution with a gap of only 0.75% to 2.92%, whereas the benchmark without feature data achieves an optimality gap of 6.63% to 7.58%.

4.1. Introduction

One of the main challenges in inventory management is demand uncertainty. In the literature, there are two fundamentally different ways of capturing demand uncertainty

[1] The research presented in this chapter is based on a paper entitled “Prescriptive Analytics for Dynamic Inventory Management”.

in stochastic inventory models. *Parametric* approaches assume that the demand uncertainty is characterized by a specific probability distribution that is known to the decision-maker. *Non-parametric* or *data-driven* approaches rely directly on historical data without any assumptions on the underlying demand distribution. While most inventory models belong to the first group, the rising availability of data in practice has led to an increased interest in data-driven approaches (e.g., Ban and Rudin (2019), Huber et al. (2019), Ban (2020)). These data not only comprise information about historical demand but also additional feature data such as weather, location, and calendar data that might be correlated with demand and thus can help to improve inventory decisions. For dynamic inventory problems, however, feature data that have explanatory value for demand have largely been ignored so far.

It is at the heart of machine learning (ML) to use available feature data to predict uncertain quantities (e.g., demand). The explosion of available data and developments in ML algorithms have enabled ML applications such as predicting crime (Kadar and Pletikosa 2018), predicting results of football matches (Baboota and Kaur 2019), or predicting retail store sales (Kaneko and Yada 2016). Deriving optimal decisions under uncertainty, however, is usually not considered in ML.

In this paper, we combine ideas of classical inventory management and ML to develop a non-parametric solution approach that leverages feature data for the dynamic inventory problem under non-stationary demand.

Bertsimas and Kallus (2020) present a general solution approach for stochastic optimization problems that leverages feature data and call it *predictive prescriptions*. Predictive prescriptions are functions that prescribe a decision for a stochastic optimization problem given some feature data. One way of constructing predictive prescriptions is to assign weights to the historically observed quantities (e.g., demand) with ML according to their value for the decision and then solve the *weighted* sample average approximation problem. The approach can be applied to general single-stage optimization problems but it is not immediately clear how to transfer it to multi-stage problems.

The main idea of our paper is to use a dynamic programming (DP) formulation of the dynamic inventory problem that splits the problem into successive single-stage problems and apply the concept of predictive prescriptions to each of these problems.

We test our method in a computational study and compare it to a data-driven approach without feature data (Ban 2020) and to the optimal dynamic programming solution. We find that our approach performs close to the optimal policy with an average gap of only 0.75% to 2.92%, while the approach ignoring feature data results in an average optimality gap of 6.63% to 7.58%.

To summarize, our main contributions are:

- We combine ideas from classical inventory management and machine learning to leverage feature data for the dynamic inventory problem under non-stationary demand.
- We show how to apply the idea of *predictive prescriptions* to a multi-period problem using dynamic programming.
- We compare our method in a computational study to the optimal solution and to a less data-rich benchmark and find that our approach performs close to optimal.

The remainder of this paper is structured as follows. In the next section, we review the related literature. Section 4.3 contains the problem definition and model formulation. In Section 4.4, we test our approach in a computational study. Finally, we summarize our findings and outline opportunities for further research in Section 4.5.

4.2. Related Literature

In this section, we review the literature on inventory models with non-stationary demand and data-driven inventory models that do not require demand model assumptions.

Most of the literature on stochastic dynamic inventory management builds on the seminal work of Scarf (1959), who developed the theory of K-convexity to prove the optimality of (s, S)-policies for the multi-period problem. Karlin (1960a) investigates the optimal policy for the problem with non-stationary demand and different cost structures. Further theoretical results for special cases can be found in Karlin (1960b) for seasonal demand, and (Zipkin 1989) for cyclic demand. Our dynamic programming formulation builds on the theoretical results of Scarf (1959), and we exploit K-convexity. It is applicable to problems without specific assumptions on the demand process (e.g., trends or periodicity).

While the aforementioned literature is concerned with the structure of the optimal policy, another stream of research focuses on computing the actual values for a given policy. Due to the computational complexity of the problem, these methods solve the problem heuristically. Silver (1978) present a heuristic approach that sequentially determines if it is time to order, how many periods the order should cover, and how much to order. It can be interpreted as a stochastic version of the Silver-Meal heuristic. Askin (1981) present an alternative and use an order-up-to level in combination with a least period cost approach. Morton and Pentico (1995) and Bollapragada and Morton (1999) present myopic heuristics that perform well if the demand variation across periods is relatively small. Bollapragada and Rao (2006) include supply uncertainty and capacity limits and derive a heuristic based on the first two moments of the uncertain quantity. The discussed heuristics assume a non-stationary but known demand distribution which is rarely accurate in real-world situations.

In a departure from the previous work, the literature on non-parametric inventory management does not require demand distribution assumptions. Liyanage and Shanthikumar (2005), Beutel and Minner (2012), Ban and Rudin (2019) and Huber et al. (2019) consider the single-period newsvendor problem and incorporate feature data into the optimization problem. Non-parametric multi-period models include Kunnumkal and Topaloglu (2008) and Huh and Rusmevichientong (2009). These models do not require assumptions on the demand distribution and leverage historical demand/sales data directly. However, they do not consider additional feature data. Most related to our work is Ban (2020) who develops a data-driven solution approach for the stationary problem that uses historical sales data. The author derives finite sample properties and bounds of her approach and investigates the influence of censored data. Feature data is not considered. To the best of our knowledge, no data-driven approach that leverages feature data is currently available for dynamic inventory problems with non-stationary demand.

Motivated by this gap in the literature and the increase in data availability, we develop a DP for the dynamic inventory problem, that exploits feature data to improve the inventory decisions. To this end, we use the concept of *predictive prescriptions* developed by Bertsimas and Kallus (2020). Predictive prescriptions prescribe a decision for a problem, given some feature data. Bertsimas and Kallus (2020) describe various ML approaches that can be used to construct these decisions and study their asymptotics. They apply their approach to single-period problems, a two-stage shipment problem, and a real-world single-period inventory problem of a large media company. We transfer the idea of predictive prescriptions to a multi-period inventory problem and combine it with a non-parametric dynamic programming approach.

4.3. Problem Definition and Model Formulation

4.3.1. Problem Definition

We study the stochastic dynamic inventory problem of Scarf (1959). However, we drop the assumption of full knowledge of the demand distribution.

Thus, the problem is defined as follows. For each period $t = 1, ..., T$, the decision-maker has to decide about the order quantity q of a single product to minimize discounted costs, given a discount factor α_t. Demand D is uncertain, and demand that cannot be met is backlogged. The backordering costs and inventory holding costs in period t are b_t and h_t, respectively. If an order is placed in period t, it entails a fixed ordering cost K_t and a per-unit ordering cost c_t.

Instead of full knowledge of the demand distribution, the decision maker has access to a data set $\{(d_1, x_1), ..., (d_N, x_N)\}$. That is, for N historical periods, there is a set of tuples (d_n, x_n) with $n = 1, ..., N$, where d_n is the historical demand in period n and x_n is the historical feature vector in period n.

4.3.2. Model Formulation

To solve the above problem without assumptions on the demand distribution and to leverage the data set $\{(d_1, x_1), ..., (d_N, x_N)\}$, we use the idea of *predictive prescriptions* from Bertsimas and Kallus (2020).

Bertsimas and Kallus (2020) develop a general framework to include auxiliary feature data X that is correlated with the uncertain quantity D (in our case demand) into stochastic optimization problems. They use the term *predictive prescriptions* for functions $z(x)$ that prescribe a decision, given an observation of feature data x, i.e.,

$$z(x) = \arg\min_{z \in \mathcal{Z}} \mathbb{E}\left[c(z; D) | X = x\right]. \tag{4.1}$$

They propose to construct these prescriptions by solving the following weighted sample average approximation problem

$$\hat{z}(x) = \arg\min_{z \in \mathcal{Z}} \sum_{n=1}^{N} w_{N,n}(x) c(z; d_n). \tag{4.2}$$

where the weight functions $w_{N,n}(x)$ are based on ML approaches, such as kNN, kernel methods, local linear methods, and decision trees that leverage the observation of feature data x. This is a very generic formulation that applies to any single-stage stochastic program.

However, it is not immediately clear, how predictive prescriptions can be applied to multi-period problems because the uncertainty unfolds in multiple stages and the decisions in different periods are interdependent. In the following, we show how to use a DP formulation of the problem to apply the concept of predictive prescriptions to our multi-period inventory problem.

The main reasoning behind dynamic programming is to divide a complex multi-stage problem into multiple interrelated single-stage problems and solve those iteratively. We use this idea to apply predictive prescriptions to each single-stage problem of the DP formulation of the dynamic inventory problem.

As a starting point, we recap the classical DP formulation of Scarf (1959) with the notation of Ban (2020) and apply predictive prescriptions to each of the single-stage problems.

The total cost of ordering the quantity q in period t are

$$O_t(q) = \begin{cases} K_t + c_t q, & \text{if } q > 0 \\ 0, & \text{otherwise.} \end{cases} \tag{4.3}$$

Denote the stock level after deliveries in period t by y. The single-period newsvendor cost function is therefore

$$C(y, d) = b_t(d - y)^+ + h_t(y - d)^+, \tag{4.4}$$

and the Bellman equations can be formulated as

$$V_{T+1}(x) = 0, \qquad \forall x \geq 0 \tag{4.5}$$

$$V_t(I_t) = \min_{y \geq I_t} \{O_t(y - I_t) + \mathbb{E}\left[C_t(y, D_t) + \alpha_t V_{t+1}(y - D_t)\right]\}, \qquad 1 \leq t \leq T, \tag{4.6}$$

where I_t is the inventory level at the beginning of period t, I_1 is known, and $\alpha_t \in [0, 1)$ is the discount factor in period t. Note that I_t can be negative due to backordering.

Scarf (1959) shows that the optimal solution to the problem is a $(s_t, S_t)_{t=1}^{T}$-policy, i.e., if the inventory level I_t in period t drops to or below the reorder level s_t, it is optimal to order up to the level S_t. Given this policy, (Scarf 1959) shows that the optimal order-up-to level S_t in each period t can be found by minimizing

$$G_t(y) = c_t y + \mathbb{E}_t\left[C_t(y, D_t)\right] + \alpha_t \mathbb{E}_t\left[V_{t+1}(y - D_t)\right], \tag{4.7}$$

where

$$\begin{aligned} &V_{T+1}(x) = 0, && \forall x \geq 0 \\ &V_{t+1}(x) = \begin{cases} G_{t+1}(S_{t+1}) + K_t - c_t x, & \text{if } x < s_{t+1}, \\ G_{t+1}(x) - c_t x, & \text{if } x \geq s_{t+1}, \end{cases} && 1 \leq t \leq T-1 \end{aligned} \tag{4.8}$$

and optimal reorder point for each period $t = 1, ..., T$ is

$$s_t = \min_s \left\{ \underline{D} \leq s \leq S_t | G_t(s) - G_t(S_t) - K_t = 0 \right\}. \tag{4.9}$$

With (4.8), (4.7) can be rewritten as

$$G_T(y) = c_T y + \mathbb{E}_T\left[C_T(y, D_T)\right], \tag{4.10}$$

$$\begin{aligned} G_t(y) = {} & (1-\alpha_t) c_t y + \mathbb{E}_t\left[C_t(y, D_t)\right] + \alpha_t c_t \mathbb{E}_t\left[D_t\right] + \alpha_t G_{t+1}(s_{t+1}) \mathbb{E}_t\left[\mathbb{I}_t(y - s_{t+1})\right] \\ & + \alpha_t \mathbb{E}_t\left[G_{t+1}(y - D_t) \mathbb{I}_t^c(y - s_{t+1})\right], \ t = 1, ..., T-1, \end{aligned} \tag{4.11}$$

where $\mathbb{I}_t(x) := \mathbb{I}(x < D_t)$ and $\mathbb{I}_t^c(x)$ is its compliment.

To evaluate the expectations in the specified model and calculate optimal $(s_t, S_t)_{t=1}^T$ values, one would need perfect knowledge of the (non-stationary) demand distribution of D_t. This is not the case in most real-world settings.

We propose to construct predictive prescriptions instead. To do so, we follow the approach of Ban (2020) who replace the expectations in (4.10) and (4.11) by sample averages. We proceed analogously but use *weighted* sample averages, in line with (Bertsimas and Kallus 2020). Thus, we define

$$\hat{g}_T(y, d) = c_T y + C_T(y, d) \tag{4.12}$$

$$\begin{aligned} \hat{g}_t(y, d) = {} & (1-\alpha_t) c_t y + C_t(y, d) + \alpha_t c_t d + \alpha_t \hat{G}_{t+1}(\hat{s}_{t+1}, x_{t+1}) \mathbb{I}(y - \hat{s}_{t+1} \leq d) \\ & + \alpha_t \hat{G}_{t+1}(y - d, x_{t+1}) \mathbb{I}(y - \hat{s}_{t+1} > d), \ t = 1, ..., T-1, \end{aligned} \tag{4.13}$$

and set

$$\hat{G}_t(y, x_t) = \sum_{n=1}^{N} w_{N,n}(x_t)\hat{g}_t(y, d_n), \tag{4.14}$$

where we use ML and the feature data to construct weight functions $w_{N,n}(x_t)$ to reweigh the historical demand data.

We then choose the inventory control parameters as

$$\hat{S}_t = \arg\min_{y \in [\underline{D}, \overline{D}]} \hat{G}_t(y, x_t) \tag{4.15}$$

and

$$\hat{s}_t = \min_s \left\{ \underline{D} \leq s \leq \hat{S}_t | \hat{G}_t(s, x_t) - \hat{G}_t(\hat{S}_t, x_t) - K_t = 0 \right\}. \tag{4.16}$$

In our analysis below, we use two alternative ML approaches to design the weight functions $w_{N,n}(x_t)$, namely k Nearest Neighbor (kNN) regression (Hastie et al. 2017) and Classification and Regression Trees (CARTs) (Breiman et al. 1984). Both approaches are non-parametric ML algorithms that do not make strong prior assumptions about the functional relationship between input and output. kNN is a relatively simple, effective, and popular ML method that has proven to perform well in single-stage problems if the feature space is not too large (Bertsimas and Kallus 2020). kNN is an unsupervised learning method and if the feature space is large it might have problems to distinguish between valuable data and noise. CARTs are supervised learning methods that generally perform well, even if the feature space is very large.

The kNN weighting assigns a weight of $\frac{1}{k}$ to all kNNs and a weight of 0 to all other data points. To measure the distance, we use the L_2-norm $|| \cdot ||$ such that the neighborhood containing the k nearest neighbors is $\mathcal{N}_k(x) = \{x_i | \sum_{j=1}^{N} \mathbb{I}\left[||x - x_i|| \geq ||x - x_j||\right] \leq k, i = 1, ..., N\}$. Thus, the kNN weight function is

$$w_{N,n}^{k\text{NN}}(x_t) = \frac{1}{k}\mathbb{I}[x_n \in \mathcal{N}_k(x_t)]. \tag{4.17}$$

Note that kNN is an unsupervised learning approach, i.e., whether a data point is a kNN depends only on the feature values, not on the actual demand observation.

CARTs (Breiman et al. 1984) are binary trees that split a dataset into disjoint regions $R(\cdot)$, called "leaves". The trees are constructed recursively by splitting the data in each node into two child nodes such that the "impurity" (e.g., mean squared error) of the target variable in each node is minimized. The splitting stops as soon as a pre-specified stopping criterion is met (e.g., the maximum number of splits). The final nodes of the tree are the leaves. While the splitting is done on the feature values, the splitting rule is designed according to the target values (in our case demand). Once the tree has been constructed, the weight function is

$$w_{N,n}^{\text{CART}}(x_t) = \frac{\mathbb{I}\left[R(x_t) = R(x_n)\right]}{|\{j|R(x_j) = R(x_t)\}|}. \tag{4.18}$$

Each historical data point that resides in the same leaf as the current data point receives a weight of 1/(total number of data points in the leaf), all other data points have a weight of 0. Note that CART is a supervised learning algorithm that requires knowledge of the historical demand data to minimize impurity.

4.4. Computational Study

4.4.1. Study Design

In the previous chapter, we developed a method to incorporate feature data into multi-period inventory problems. We now test the viability of our approach and measure the value of additional feature data for inventory performance in a numerical study.To this end, we compare our data-driven approach with kNN weighting (*DD-kNN*) and CART weighting (*DD-CART*) with the optimal solution obtained by the dynamic program with full information of the (non-stationary) demand distribution and the data-driven approach of Ban (2020) that does not consider feature data (*DD-NoFeatures*). We compare both the cost incurred for each approach and the obtained policy values $(s_t, S_t)_{t=1}^{T}$.

We simulate 50 training data sets $\{(d_1, x_1), ..., (d_N, x_N)\}$ and corresponding test sets $\{(d_1, x_1),\ ..., (d_T, x_T)\}$. Each training set consists of $N = 50, 100, 200$ periods and each test set has a length of $T = 100$. For all simulations, we use the following cost parameters: $b_t = 10$, $h_t = 1$, $c_t = 0.1$, $K_t = 1280$, $\alpha_t = 1$, for all periods $t = 1, ..., T$.

As a basis for our study, we adapt the demand model from the simulation experiment of Bertsimas and Kallus (2020) which we find to be sufficiently complex to challenge our approach. In this model, the demand depends on three features. The feature data evolves according to the following 3-dimensional ARMA(2, 2) process:

$$x_t - \Phi_1 x_{t-1} - \Phi_2 x_{t-2} = u_t + \Theta_1 u_{t-1} + \Theta_2 u_{t-2}, \tag{4.19}$$

where $x_0 = x_1 = u_0 = u_1 = 0$, $u \sim \mathcal{N}(0, \Sigma_u)$ and

$$\Phi_1 = \begin{pmatrix} 0.5 & -0.9 & 0 \\ 1.1 & -0.7 & 0 \\ 0 & 0 & 0.5 \end{pmatrix}, \Phi_2 = \begin{pmatrix} 0 & -0.5 & 0 \\ -0.5 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{4.20}$$

$$\Theta_1 = \begin{pmatrix} 0.4 & 0.8 & 0 \\ -1.1 & -0.3 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \Theta_2 = \begin{pmatrix} 0 & -0.8 & 0 \\ -1.1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \Sigma_u = \begin{pmatrix} 1 & 0.5 & 0 \\ 0.5 & 1.2 & 0.5 \\ 0 & 0.5 & 0.8 \end{pmatrix}. \tag{4.21}$$

Demand is then generated according to the following model:

$$d_t = \max\{0, 150 + A(x_t + \delta/4) + (Bx_t)\epsilon\}, \tag{4.22}$$

where $\epsilon \sim \mathcal{N}(0, 1)$, $\delta \sim \mathcal{N}(0, 1)$, and

$$A = \begin{pmatrix} 2.8 \\ 2.8 \\ 0.35 \end{pmatrix}, B = \begin{pmatrix} -6 \\ 0.75 \\ -6 \end{pmatrix}. \tag{4.23}$$

Figure 4.1 shows the dependence of the demand on the three features and the resulting demand process over time for one exemplary selling horizon. We observe that demand fluctuates strongly over time, the features have different degrees of predictive value, and feature 1 exhibits significant heteroscedasticity.

A major focus in current ML research is on feature engineering. In reality, it is hard to identify or engineer the features that have high predictive power and some selected features might just add noise to the problem. To address this issue, we add additional feature data to our model that have no predictive power to the 50 runs with $N = 100$. We use $u = 0, 3, 6, 12, 24, 48, 96$ additional features. The additional features are normally distributed with 0 mean and sets of 3 features have the same covariance as the original features x.

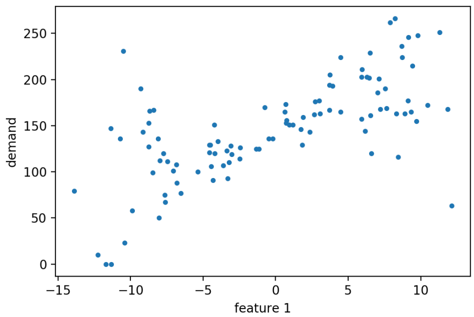

(a) Dependence of demand on feature 1.

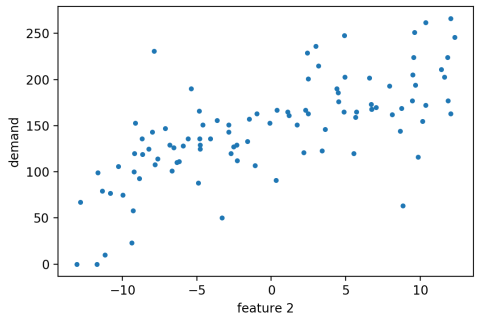

(b) Dependence of demand on feature 2.

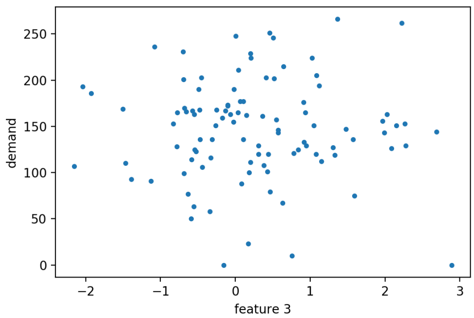

(c) Dependence of demand on feature 3.

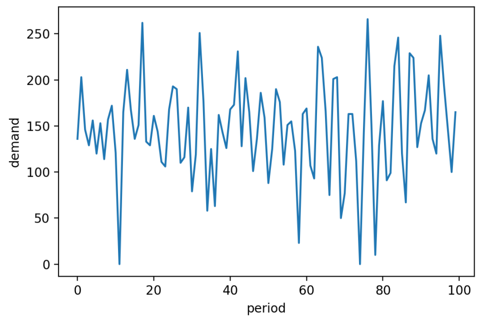

(d) Demand over time.

Figure 4.1.: Feature dependent demand.

To obtain the weights for our approach, we implement a kNN algorithm and use the CART algorithm from the Scikit-learn package in Python (Pedregosa et al. 2011).

To obtain the $(s_t, S_t)_{t=1}^T$ values for a given demand model (known distribution, *DD-NoFeatures*, *DD-kNN* and *DD-CART*), we use the following algorithm that exploits K-convexity (Bollapragada and Morton 1999): Starting from an upper bound on the order-up-to level, we evaluate (4.7) respectively (4.14) and then interatively decrease y by one unit. The lowest function value is stored. Once, the value of the function is greater than the lowest function value so far plus K, the loop terminates. This point of the function is the reorder level s_t (or $\hat{s}_t$) and the point where the lowest function value is obtained is the order-up-to level S_t (or $\hat{S}_t$).

4.4.2. Results

To evaluate our data-driven approaches *DD-kNN* and *DD-CART*, we compare them to the optimal solution with known demand distribution and to the data-driven approach without feature data *DD-NoFeature*, in terms of resulting costs and decisions.

Figure 4.2 shows the relative cost increase of *DD-kNN*, *DD-CART*, and *DD-NoFeature* compared to the optimal approach for the tested sample sizes $N = 50, 100, 200$ averaged over the 50 simulation runs. We observe that our approach *DD-kNN* performs best with an average optimality gap of only 0.75% ($N = 200$), 1.21% ($N = 100$), and 2.66% ($N = 50$) and relatively small variation. *DD-CART* performs second best with an average optimality gap of 2.93% ($N = 200$), 3.41% ($N = 100$), and 5.58% ($N = 50$). The benchmark approach *DD-NoFeature* achieves an average optimality gap of 7.03% ($N = 200$), 6.63% ($N = 100$), and 7.58% ($N = 50$). All approaches benefit from larger samples through lower costs and lower cost variation with the exception of *DD-NoFeatures* for $N = 200$. The results illustrate the viability of our approach and show its ability to extract the relevant information from the feature data to come up with near-optimal decisions. Especially our approach with kNN regression performs well with little modelling and implementation effort.

The lower costs of our approaches relative to the benchmark without feature data is a result of better decisions. To get an intuition for the differences of the resulting policies, Figure 4.3 shows the sample path of the $(s_t, S_t)_{t=1}^T$ values of the different methods for a selected sample season. The volatile pattern of the optimal $(s_t, S_t)_{t=1}^T$ values is a direct result of the simulated ARMA(2, 2) demand process (see Figure 4.1d). As feature data

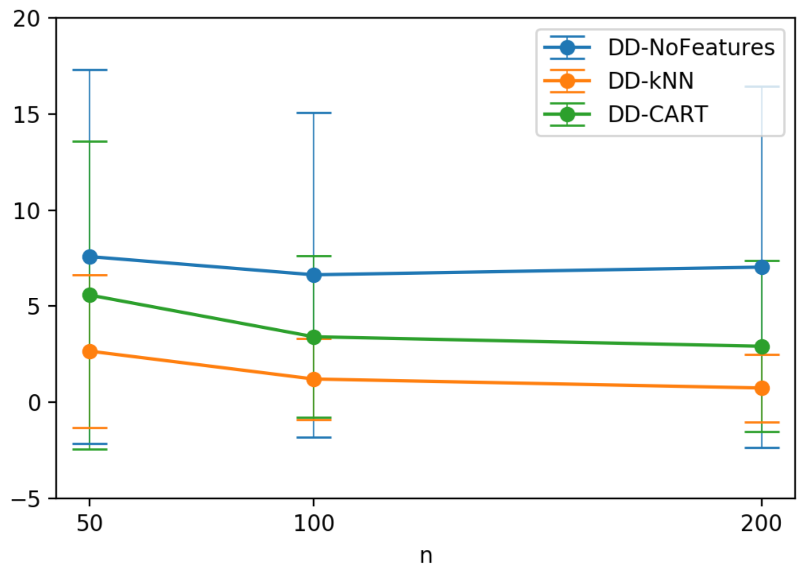

Figure 4.2.: Average relative cost increase (%) with 75% confidence intervals for varying sample sizes.

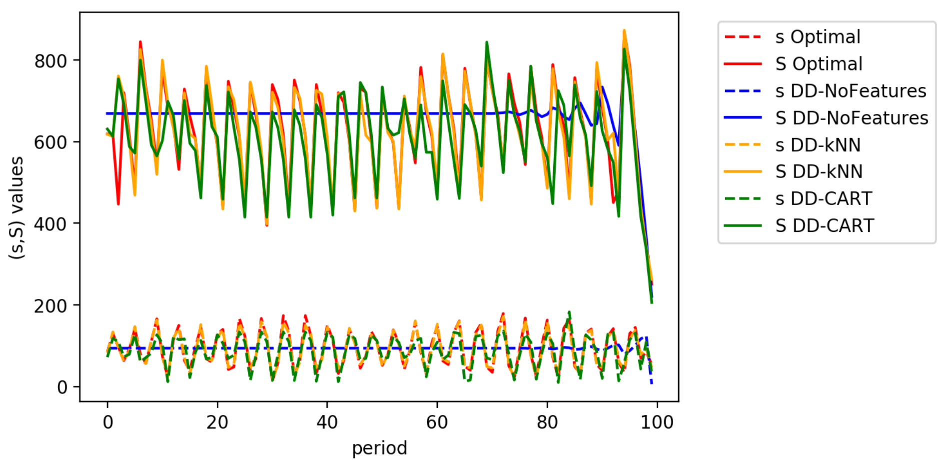

Figure 4.3.: Example of $(s_t, S_t)_{t=1}^{T}$ values for one selling horizon.

is ignored by *DD-NoFeature*, we observe that the $(s_t, S_t)_{t=1}^T$ values of the approach do not follow the optimal pattern, are constant over the main part of the planning horizon, and show an "end-of-horizon-effect". The $(s_t, S_t)_{t=1}^T$ values of our approaches follow the optimal pattern rather closely.

To analyse the policies more generally, Figure 4.4 shows the average relative error of the reorder level and the order-up-to level, respectively, of all approaches, compared to the optimal parameters for the tested sample sizes $N = 50, 100, 200$. We observe that

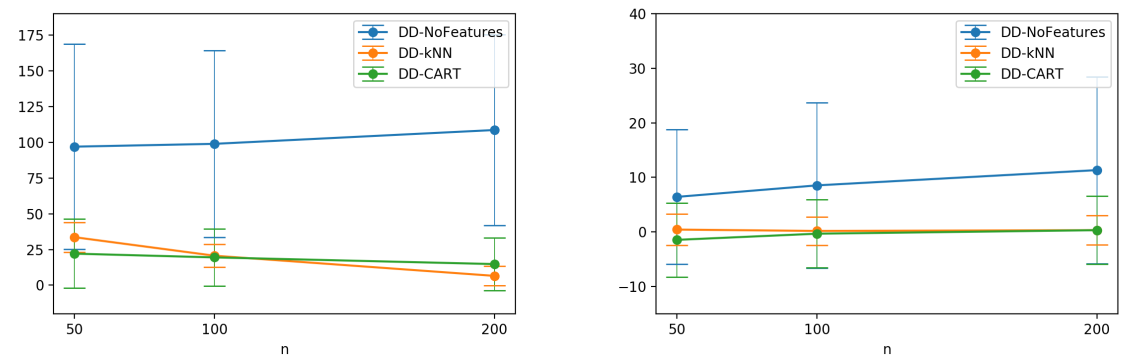

(a) Reorder level s_t.

(b) Order-up-to level S_t.

Figure 4.4.: Average relative errors (%) of policy parameters and 75% confidence intervals for varying sample sizes.

both approaches with feature data perform similar in terms of average relative errors for the policy parameters. Note that the larger relative errors for s_t, compared to S_t, are due to the difference in absolute values between both parameters. *DD-kNN* has significantly lower variation in the error, which also explains its superior cost performance. For the data-driven approaches, we see a positive sample size effect in the accuracy of the reorder level. In the benchmark approach *DD-NoFeature*, the relative error gets even larger with a larger sample size.

Figure 4.5 shows the average runtimes of the tested approaches for different sample sizes. The optimal policy parameters are calculated with a sampling approach with 400 samples. The runtime of the optimal approach depends mainly on this sample size. It takes on average 348 seconds to solve the problem optimally with this approach for one selling season with 100 periods on a standard PC (Intel Core i5, 2.3 GHz, 8 GB RAM). The runtime of the data-driven approaches depends linearly on the sample size and is very similar across the three data-driven methods. Both approaches with feature data

(*DD-kNN*, *DD-CART*) take 2 to 8 seconds longer than *DD-NoFeatures* due to the upfront reweighting of the data points. In relative terms, this amounts to a runtime increase of 4% to 6%. Overall, the runtime increase is small compared to the performance increase of our feature based approaches compared to the method without features.

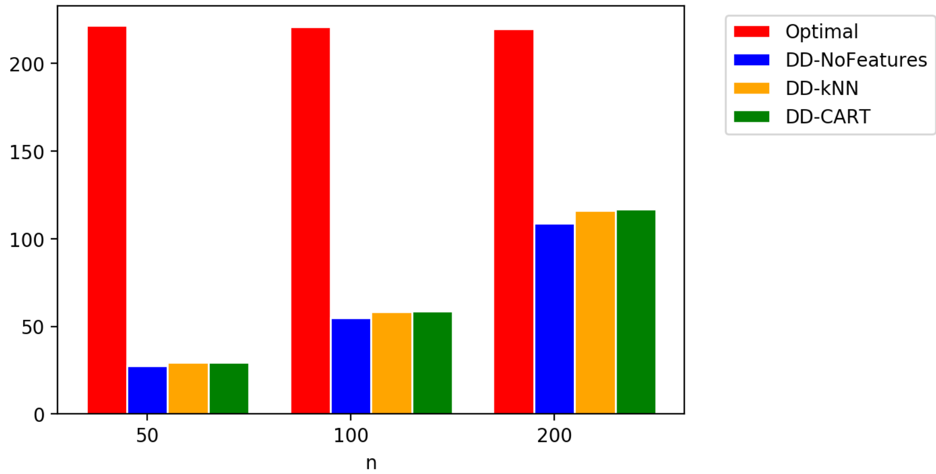

Figure 4.5.: Average runtimes (in seconds) depending on the sample size.

In our sampling approach, we assumed so far that the decision-maker can identify and incorporate those features that have predictive value for the inventory decision. In reality, however, identifying and engineering the most relevant features can be difficult. Therefore, Figure 4.6 depicts the average relative cost increase compared to the optimal approach of the tested methods if, in addition to the three useful features, several "useless" features are included, that have no predictive value for the inventory decision.

As the *DD-NoFeature* does not consider feature data, there is no impact of the larger feature space on costs. Our approach based on kNN suffers from a larger feature space, and the performance decreases with more features. As kNN takes into account the distance of all features without considering their impact on the demand, the average costs increase significantly with a larger feature space. On the other hand, the average performance of our CART approach is more stable. CARTs split the data according to the "impurity" of the target variable (in our case demand) in each node. Therefore, they can identify the features with predictive value. Nevertheless, kNN is better, except for a very large number of irrelevant features.

In summary, for the considered setting, both our approaches perform close to the optimal solution in terms of costs and parameter accuracy. Especially applying the

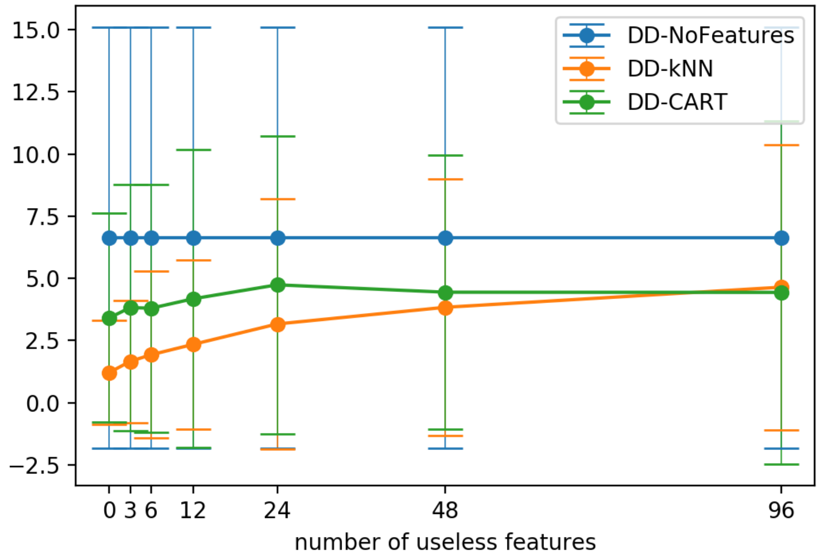

Figure 4.6.: Average relative cost increase (%) with 75% confidence intervals for varying dimensions of the feature space.

data-driven approach with kNN reweighing can save up to 5.87% of costs compared to the benchmark approach that does not consider feature data if enough data is available ($N = 200$). If the feature space is very large, the performance of our kNN approach deteriorates and *DD-CART* outperforms it.

4.5. Conclusions

In this paper, we develop a methodology to leverage the ever-increasing amount of available data for dynamic inventory decisions. To this end, we incorporate feature data into a data-driven dynamic inventory model by combining ideas from traditional inventory management and machine learning. We build on the idea of *predictive presciptions* (Bertsimas and Kallus 2020) that have been developed for single-stage stochastic problems and show how to apply the concept to a multi-stage problem by means of dynamic programming. Our method estimates a non-stationary (s, S)-policy by reweighing the historical demand samples with machine learning and feature data in a dynamic program. In a numerical study, it performs close to the optimal decision with a gap of only 0.75% to 2.92% whereas the benchmark without feature data achieves an optimality gap of 6.63% to 7.58% in our computational study. These results are in line with the results of similar predictive prescriptions for single-stage problems (Bertsimas and Kallus 2020). Thus, we conclude that our approach for multi-stage problems is viable and performs well for our dataset.

Our method is very flexible and can incorporate different machine learning approaches to reassign weights to the historical data. In particular, we show that the size of the feature space should be taken into account when choosing an estimation method. If the feature space is small, unsupervised learning methods like kNN perform well. If the feature space is very large, supervised learning approaches like CART should be preferred as they can filter irrelevant features. The runtime increase for incorporating feature data is only 4% to 6%. We conclude that by ignoring feature data, many traditional inventory models miss out on significant optimization potential, and combining traditional inventory management approaches with machine learning to data-driven methods provides near-optimal decisions with very little computational effort. If the noise in the data is small, already relatively simple methods like kNN perform well.

Our study considers a specific feature-driven demand model. Future studies should verify the generalizability of our results for other demand models, including, both other theoretical models and real-world data. Furthermore, other machine learning approaches such as kernel methods, artificial neural networks, or random forests could be applied to assign the weights to the data.

For given data, our approach searches for the optimal parameters by means of enumeration. This requires a relatively high computational effort. For time-critical applications, one could alternatively solve the parameter optimization heuristically.

Our approach is a multi-period version of the method by Bertsimas and Kallus (2020) and could be applied in future research projects to other operations management problems that are usually solved by dynamic programming.

Bibliography

Anupindi, R., Dada, M., and Gupta, S. (1998). Estimation of Consumer Demand with Stock-Out Based Substitution: An Application to Vending Machine Products. *Marketing Science*, 17(4):406–423.

Askin, R. G. (1981). A procedure for production lot sizing with probabilistic dynamic demand. *AIIE Transactions*, 13(2):132–137.

Baboota, R. and Kaur, H. (2019). Predictive analysis and modelling football results using machine learning approach for English Premier League. *International Journal of Forecasting*, 35(2):741–755.

Ban, G.-Y. (2020). Confidence Intervals for Data-Driven Inventory Policies with Demand Censoring. *Operations Research*, 68(2):309–326.

Ban, G. Y. and Rudin, C. (2019). The big Data newsvendor: Practical insights from machine learning. *Operations Research*, 67(1):90–108.

Barrow, D., Crone, S. F., and Kourentzes, N. (2010). An evaluation of neural network ensembles and model selection for time series prediction. In *The 2010 International Joint Conference on Neural Networks (IJCNN)*, Barcelona, Spain.

Barrow, D. and Kourentzes, N. (2018). The impact of special days in call arrivals forecasting: A neural network approach to modelling special days. *European Journal of Operational Research*, 264(3):967–977.

BASF (2020). Consolidated Financial Statements 2019. Technical report, BASF SE, Ludwigshafen am Rhein.

Ben-Tal, A., Ghaoui, L. E., and Nemirovski, A. (2009). *Robust optimization.* Princeton University Press, New Jersey.

Bergstra, J., Bardenet, R., Bengio, Y., and Kégl, B. (2011). Algorithms for Hyper-parameter Optimization. In *Proceedings of the 24th International Conference on Neural Information Processing Systems*, pages 2546–2554. Curran Associates Inc.

Bergstra, J. and Bengio, Y. (2012). Random Search for Hyper-Parameter Optimization. *Journal of Machine Learning Research*, 13:281–305.

Bergstra, J., Yamins, D., and Cox, D. D. (2013). Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures. In *Proceedings of the 30th International Conference on Machine Learning - Volume 28*, pages 115–123. JMLR.org.

Bertsimas, D. and Kallus, N. (2020). From Predictive to Prescriptive Analytics. *Management Science*, 66(3):1025–1044.

Bertsimas, D. and Thiele, A. (2006). A robust optimization approach to inventory theory. *Operations Research*, 54(1):150–168.

Beutel, A. L. and Minner, S. (2012). Safety stock planning under causal demand forecasting. *International Journal of Production Economics*, 140(2):637–645.

Bollapragada, R. and Rao, U. S. (2006). Replenishment planning in discrete-time, capacitated, non-stationary, stochastic inventory systems. *IIE Transactions (Institute of Industrial Engineers)*, 38(7):605–617.

Bollapragada, S. and Morton, T. E. (1999). A Simple Heuristic for Computing Nonstationary (s, S) Policies. *Operations Research*, 47(4):576–584.

Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.

Breiman, L., Friedman, J., Stone, C., and Olshen, R. (1984). *Classification and regression trees*. CRC Press, New York.

Briesch, R. A., Chintagunta, P. K., and Fox, E. J. (2009). How Does Assortment Affect Grocery Store Choice? *Journal of Marketing Research*, 46(2):176–189.

Brynjolfsson, E. and McElheran, K. (2016a). Data in Action: Data-Driven Decision Making in U.S. Manufacturing. *US Census Bureau Center for Economic Studies*, CES-WP-16-.

Brynjolfsson, E. and McElheran, K. (2016b). The Rapid Adoption of Data-Driven Decision-Making. *American Economic Review*, 106(5):133–139.

Campo, K., Gijsbrechts, E., and Nisol, P. (2000). Towards understanding consumer response to stock-outs. *Journal of Retailing*, 76(2):219–242.

Cannon, A. J. (2011). Quantile regression neural networks: Implementation in R and application to precipitation downscaling. *Computers and Geosciences*, 37(9):1277–1284.

Carbonneau, R., Laframboise, K., and Vahidov, R. (2008). Application of machine learning techniques for supply chain demand forecasting. *European Journal of Operational Research*, 184(3):1140–1154.

Conrad, S. A. (1976). Sales data and the estimation of demand. *Operational Research Quarterly*, 27(1):123–127.

Crone, S. F., Hibon, M., and Nikolopoulos, K. (2011). Advances in forecasting with neural networks? Empirical evidence from the NN3 competition on time series prediction. *International Journal of Forecasting*, 27(3):635–660.

Crone, S. F. and Kourentzes, N. (2009). Forecasting Seasonal Time Series with Multilayer Perceptrons-An Empirical Evaluation of Input Vector Specifications for Deterministic Seasonality. In *Proceedings of the 2009 International Conference on Data Mining*, pages 232–238, Las Vegas, USA.

den Boer, A. V. and Sierag, D. D. (2020). Decision-based model selection. *European Journal of Operational Research*, In Press.

Emmelhainz, L. W., Emmelhainz, M. A., and Stock, J. R. (1991). Logistics Implications of Retail Stockouts. *Journal of Business Logistics*, 12(2):129–142.

Farahat, A. and Lee, J. (2018). The Multiproduct Newsvendor Problem with Customer Choice. *Operations Research*, 66(1):123–136.

Feng, Q. and Shanthikumar, J. G. (2018). How Research in Production and Operations Management May Evolve in the Era of Big Data. *Production and Operations Management*, 27(9):1670–1684.

Fisher, M. and Vaidyanathan, R. (2014). A Demand Estimation Procedure for Retail Assortment Optimization with Results from Implementations. *Management Science*, 60(10):2401–2415.

Friedman, J. H. (2001). Greedy Function Approximation: A Gradient Boosting Machine. *The Annals of Statistics*, 29(5):1189–1232.

Gallego, G. and Moon, I. (1993). The distribution free newsboy problem: Review and extensions. *The Journal of the Operational Research Society*, 44(8):825–834.

Godfrey, G. A. and Powell, W. B. (2001). An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Science*, 47(8):1101–1112.

Gruen, T. W., Corsten, D. S., and Bharadwai, S. (2002). Retail Out-of-Stocks: A Worldwide Examination of Extent, Causes and Consumer Response. Technical report.

Hastie, T., Tibshirani, R., and Friedman, J. (2017). *The Elements of Statistical Learning*. Springer Series in Statistics. Springer, New York, NY, USA.

Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257.

Huber, J., Gossmann, A., and Stuckenschmidt, H. (2017). Cluster-based hierarchical demand forecasting for perishable goods. *Expert Systems with Applications*, 76:140–151.

Huber, J., Müller, S., Fleischmann, M., and Stuckenschmidt, H. (2019). A data-driven newsvendor problem: From data to decision. *European Journal of Operational Research*, 278(3):904–915.

Hübner, A., Kuhn, H., and Kühn, S. (2016). An efficient algorithm for capacitated assortment planning with stochastic demand and substitution. *European Journal of Operational Research*, 250(2):505–520.

Huh, W. T. and Rusmevichientong, P. (2009). A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research*, 34(1):103–123.

Hyndman, R. J. and Khandakar, Y. (2008). Automatic time series forecasting: The forecast package for R. *Journal of Statistical Software*, 26(3):1–22.

Hyndman, R. J. and Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International journal of forecasting*, 22(4):679–688.

Hyndman, R. J., Koehler, A. B., Ord, J. K., and Snyder, R. D. (2008). *Forecasting with exponential smoothing: The state space approach*. Springer, New York.

Hyndman, R. J., Koehler, A. B., Snyder, R. D., and Grose, S. (2002). A state space framework for automatic forecasting using exponential smoothing methods. *International Journal of Forecasting*, 18(3):439–454.

Kadar, C. and Pletikosa, I. (2018). Mining large-scale human mobility data for long-term crime prediction. *EPJ Data Science*, 7(1).

Kaneko, Y. and Yada, K. (2016). A Deep Learning Approach for the Prediction of Retail Store Sales. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*, pages 531–537.

Karlin, S. (1960a). Dynamic inventory policy with varying stochastic demands. *Management Science*, 6(3):231–258.

Karlin, S. (1960b). Optimal Policy for Dynamic Inventory Process with Stochastic Demands Subject to Seasonal Variations. *Journal of the Society for Industrial and Applied Mathematics*, 8(4):611–629.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *Advances in Neural Information Processing Systems 30*, pages 3146–3154. Curran Associates, Inc.

Kingma, D. P. and Ba, J. L. (2015). Adam: A method for stochastic optimization. In *International Conference on Learning Representations 2015*, San Diego.

Kleywegt, A. J., Shapiro, A., and Homem-de Mello, T. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2):479–502.

Koenker, R. (2005). Quantile regression. Cambridge University Press, New York.

Kök, A. G. and Fisher, M. L. (2007). Demand estimation and assortment optimization under substitution : Methodology and application. *Operations Research*, 55(6):1001–1021.

Kök, A. G., Fisher, M. L., and Vaidyanathan, R. (2015). Assortment Planning: Review of Literature and Industry Practice. In Agrawal, N. and Smith, S. A., editors, *Retail Supply Chain Management*, volume 122, pages 99–153.

Kourentzes, N., Barrow, D. K., and Crone, S. F. (2014). Neural network ensemble operators for time series forecasting. *Expert Systems with Applications*, 41(9):4235–4244.

Kunnumkal, S. and Topaloglu, H. (2008). Using stochastic approximation methods to compute optimal base-stock levels in inventory control problems. *Operations Research*, 56(3):646–664.

Lau, H.-S. and Lau, A. H. L. (1996). Estimating the demand distributions of single-period items having frequent stockouts. *European Journal of Operational Research*, 92(2):254–265.

Levi, R., Perakis, G., and Uichanco, J. (2015). The data-driven newsvendor problem: New bounds and insights. *Operations Research*, 63(6):1294–1306.

Levi, R., Roundy, R. O., and Shmoys, D. B. (2007). Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research*, 32(4):821–839.

Liyanage, L. H. and Shanthikumar, J. G. (2005). A practical inventory control policy using operational statistics. *Operations Research Letters*, 33(4):341–348.

Morton, T. E. and Pentico, D. W. (1995). The finite horizon nonstationary stochastic inventory problem: near-myopic bounds, heuristics, testing. *Management Science*, 41(2):334–343.

Musalem, A., Olivares, M., Bradlow, E. T., Terwiesch, C., and Corsten, D. (2010). Structural Estimation of the Effect of Out-of-Stocks. *Management Science*, 56(7):1180–1197.

Netessine, S. and Rudi, N. (2003). Centralized and Competitive Inventory Models with Demand Substitution. *Operations Research*, 51(2):329–335.

Oroojlooyjadid, A., Snyder, L. V., and Takác, M. (2018). Applying Deep Learning to the Newsvendor Problem.

Parlar, M. and Goyal, S. K. (1984). Optimal Ordering Decisions for two Substitutable Products with Stochastic Demand. *OPSEARCH*, 21(1):1–15.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Perakis, G. and Roels, G. (2008). Regret in the newsvendor model with partial information. *Operations Research*, 56(1):188–203.

Prak, D. and Teunter, R. (2019). A general method for addressing forecasting uncertainty in inventory models. *International Journal of Forecasting*, 35(1):224–238.

Prak, D., Teunter, R., and Syntetos, A. (2017). On the calculation of safety stocks when demand is forecasted. *European Journal of Operational Research*, 256(2):454–461.

Qin, Y., Wang, R., Vakharia, A. J., Chen, Y., and Seref, M. M. H. (2011). The newsvendor problem: Review and directions for future research. *European Journal of Operational Research*, 213(2):361–374.

R Core Team (2017). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Reinsel, D., Gantz, J., and Rydning, J. (2018). The Digitization of the World - From Edge to Core. Technical report, International Data Corporation (IDC), Framingham.

REWE (2020). Konzernlagebericht für das Geschäftsjahr 2019. Technical report, Rewe-Zentalfinanz EG, Köln.

Sachs, A. L. (2015). Data-Driven Order Policies with Censored Demand and Substitution in Retailing. In *Retail Analytics*, pages 57–78.

Sachs, A. L. and Minner, S. (2014). The data-driven newsvendor with censored demand observations. *International Journal of Production Economics*, 149:28–36.

Scarf, H. (1959). The optimality of (s, S) policies in the dynamic inventory problem. In Arrow, K. J., Karlin, S., and Suppes, P., editors, *Mathematical Methods in the Social Science*. Stanford University Press, Stanford.

Scarf, H. E. (1958). A min-max solution of an inventory problem. In Arrow, K. J., Karlin, S., and Scarf, H. E., editors, *Studies in the Mathematical Theory of Inventory and Production*, pages 201–209. Stanford University Press, Stanford.

Schlapp, J. and Fleischmann, M. (2018). Technical Note — Multiproduct Inventory Management Under Customer Substitution and Capacity Restrictions. *Operations Research*, 66(3):740–747.

Shapiro, A. (2003). Monte carlo sampling methods. In Ruszczynski, A. and Shapiro, A., editors, *Handbooks in Operations Research and Management Science*, volume 10, pages 353–425. Elsevier Science B.V., Boston, USA.

Silver, E. (1978). Inventory control under a probabilistic time-varying, demand pattern. *AIIE Transactions*, 10(4):371–379.

Silver, E. A., Pyke, D. F., and Thomas, D. J. (2017). *Inventory and production management in supply chains*. Taylor and Francis, New York.

Takeuchi, I., Le, Q. V., Sears, T. D., and Smola, A. J. (2006). Nonparametric quantile estimation. *Journal of Machine Learning Research*, 7:1231–1264.

Taylor, J. W. (2000). A quantile regression approach to estimating the distribution of multiperiod returns. *The Journal of Forecasting*, 19:299–311.

Taylor, J. W. (2007). Forecasting daily supermarket sales using exponentially weighted quantile regression. *European Journal of Operational Research*, 178(1):154–167.

Thomassey, S. and Fiordaliso, A. (2006). A hybrid sales forecasting system based on clustering and decision trees. *Decision Support Systems*, 42(1):408–421.

Topaloglu, H. (2013). Joint stocking and product offer decisions under the multinomial logit model. *Production and Operations Management*, 22(5):1182–1199.

Trapero, J. R., Cardós, M., and Kourentzes, N. (2019). Empirical safety stock estimation based on kernel and GARCH models. *Omega*, 84:199–211.

van Ryzin, G. J. and Mahajan, S. (1999). On the Relationship Between Inventory Costs and Variety Benefits in Retail Assortments. *Management Science*, 45(11):1496–1509.

Van Woensel, T., Van Donselaar, K., Broekmeulen, R., and Fransoo, J. (2007). Consumer responses to shelf out-of-stocks of perishable products. *International Journal of Physical Distribution & Logistics Management*, 37(9):704–718.

Vulcano, G., van Ryzin, G., and Chaar, W. (2010). Choice-Based Revenue Management: An Empirical Study of Estimation and Optimization. *Manufacturing & Service Operations Management*, 12(3):371–392.

Vulcano, G., van Ryzin, G., and Ratliff, R. (2012). Estimating primary demand for substitutable products from sales transaction data. *Operations Research*, 60(2):313–334.

Wan, M., Huang, Y., Zhao, L., Deng, T., and Fransoo, J. C. (2018). Demand estimation under multi-store multi-product substitution in high density traditional retail. *European Journal of Operational Research*, 266(1):99–111.

Zhang, G., Patuwo, B. E., and Hu, M. Y. (1998). Forecasting with artificial neural networks:: The state of the art. *International Journal of Forecasting*, 14(1):35–62.

Zhang, J., Xie, W., and Sarin, S. C. (2018). Multi-Product Newsvendor Problem with Customer-driven Demand Substitution : A Stochastic Integer Program Perspective. Working Paper.

Zhang, Y. and Gao, J. (2017). Assessing the performance of deep learning algorithms for newsvendor problem. In *ICONIP 2017*, volume 10634 LNCS, pages 912–921.

Zipkin, P. (1989). Critical Number Policies for Inventory Models with Periodic Data. *Management Science*, 35(1):71–80.

Curriculum Vitae

Sebastian Müller

Professional Experience

since 07/2020	Expert Forecasting, Optimization, and Simulation DB Analytics, Deutsche Bahn AG, Frankfurt, Germany
11/2015 - 08/2020	Research Assistant Chair of Logistics and Supply Chain Management, University of Mannheim, Mannheim, Germany

Education

02/2016 - 11/2020	Doctoral Studies in Business Administration (Dr. rer. pol.) University of Mannheim, Mannheim, Germany
10/2013 - 09/2015	Studies in Industrial Engineering and Management (M.Sc.) Karlsruhe Institute of Technology, Karlsruhe, Germany
09/2011 - 12/2011	Studies in Business Development Engineering (Erasmus Scholarship) Ecole Centrale de Lyon, Lyon, France
10/2009 - 03/2013	Studies in Industrial Engineering and Mechanical Engineering (B.Sc.) Technical University of Darmstadt, Darmstadt, Germany
08/1999 - 07/2008	Abitur Tilemannschule, Limburg, Germany