FORSCHUNGSERGEBNISSE DER
WIRTSCHAFTSUNIVERSITÄT WIEN

Rainer Quante

Management of Stochastic Demand in Make-to-Stock Manufacturing

PETER LANG

FORSCHUNGSERGEBNISSE DER
WIRTSCHAFTSUNIVERSITÄT WIEN

Rainer Quante

Management of Stochastic Demand in Make-to-Stock Manufacturing

Up to now, demand fulfillment in make-to-stock manufacturing is usually handled by advanced planning systems. Orders are fulfilled on the basis of simple rules or deterministic planning approaches not taking into account demand fluctuations. The consideration of different customer classes as it is often done today requires more sophisticated approaches explicitly considering stochastic influences. This book reviews current literature, presents a framework that addresses revenue management and demand fulfillment at once and introduces new stochastic approaches for demand fulfillment in make-to-stock manufacturing based on the ideas of the revenue management literature.

Rainer Quante, 2000–2005 Study of Business Computing at the University of Paderborn; 2004–2005 Diploma student at a car manufacturer in Ulm; 2005–2008 Doctoral student at the Vienna University of Economics and Business Administration; 2008 to present Supply Chain Controller at an Austrian Airline in Vienna.

Retrodigitization in 2018

www.peterlang.com

Management of Stochastic Demand
in Make-to-Stock Manufacturing

Forschungsergebnisse der Wirtschaftsuniversität Wien

Band 37

PETER LANG
Frankfurt am Main · Berlin · Bern · Bruxelles · New York · Oxford · Wien

Rainer Quante

Management of Stochastic Demand in Make-to-Stock Manufacturing

PETER LANG
Internationaler Verlag der Wissenschaften

Bibliographic Information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data is available in the internet at <http://www.d-nb.de>.

This book is available Open Access thanks to the kind support of ZBW – Leibniz-Informationszentrum Wirtschaft.

Cover design:
Atelier Platen according to a design of the advertising agency Publique.

University logo of the Vienna University of Economics and Business Administration.
Printed with kindly permission of the University.

Sponsored by the Vienna University of Economics an Business Administration.

ISSN 1613-3056
ISBN 978-3-631-59409-4
ISBN 978-3-631-75386-6 (eBook)

www.peterlang.de

Für meine Eltern

Contents

List of Figures

List of Tables

Nomenclature

aATP	Allocated Available-to-Promise
AP	Allocation Planning
APS	Advanced Planning Systems
ATO	Assemble-to-Order
ATP	Available-to-Promise
BOP	Batch Order Processing
CDF	Cummulative Distribution Function
cf.	confer, compare
CODP	Customer Order Decoupling Point
CPU	Central Processing Unit
CTP	Capable-to-Promise
CV	Coefficient of Variation
DLP	Deterministic Linear Programming
DM	Demand Management
EPO	Enterprise Profit Optimization
ERP	Enterprise Resource Planning
FCFS	First-Come-First-Served
GHz	Gigahertz
GLPK	GNU Linear Programming Kit
GOP	Global Order Processing
GSL	GNU Scientific Library
IP	Integrated Pricing
IR	Inventory Rationing
KPI	Key Performance Indicator
LP	Linear Programming

MB	Megabyte
MRP	Material Resource Planning
MTO	Make-to-Order
NB	Negative Binomial Distribution
ODF	Origin-Destination-Fare
PC	Personal Computer
PTP	Profitable-to-Promise
RLP	Randomized Linear Programming
RM	Revenue Management
SIC	Stochastic Inventory Control
SOPA	Single Order Processing After Allocation Planning
SOPA_A	Aggregated SOPA
SOPA_D	Disaggregated SOPA
TRM	Traditional Revenue Management
uATP	Unallocated Available-to-Promise

Chapter 1

Introduction

1.1 Research Topic and Motivation

"The theory of inventory control tells us how much safety stock is necessary for fulfilling 99% of the orders in time, but not how to select the 1%, maybe some tens of orders per day, which are postponed or cancelled." (Fleischmann and Meyr, 2004, p. 14)

Although not explicitly stated, the authors indicate that there is more to demand management than just achieving a high service level. Rejection or postponement of orders are decisions that should be properly considered, since they play a critical role for any enterprise. Hill (2000) concisely sums it up when stating that "the most important orders are the ones that you turn down".

A number of concepts and methods emerged in the past decades addressing this issue by trying to more actively manage demand. One prominent example for successful demand management is the emergence of revenue management, which was first applied in the pricing strategies of airline tickets.

In the late 1970s, deregulation of the American airline market allowed new airlines to enter the market. Specialized only on the most profitable routes, the new airlines were highly successful and gained substantial market shares, so that the established airlines had to react to the increased competition. As most of them operated a large network with manifold destinations, they could not compete in a conventional manner against the highly-specialized new-comers able to offer much lower prices: due to specialization, the new airlines had less infrastructure costs, less maintenance costs and by focusing on the most popular destinations they reached a high seat utilization. In contrast, on many flights of the established airlines, seats were not completely sold—especially on the less popular routes and on weekends.

American Airlines was the first player to react to the new market conditions by an innovative pricing strategy. Instead of a cost-covering pricing of seats, they set the prices for some tickets on the less-utilized flights on the basis of marginal costs. Since marginal costs of an additional passenger are close to zero, American Airlines realized that it is better to sell the seat for a very low price, instead of leaving it empty.

Looking back, the new ticket pricing strategy developed by American Airlines sounds intuitive, but at that time it was an innovative way of thinking.

It was a simple idea that enabled American Airlines to offer competitive and even lower prices than the competing low-cost airlines. The only problem was to identify which seats could be sold for normal prices and which seats to sell for low prices because they would stay empty otherwise. American Airlines tied the availability of low-price tickets to conditions which were fulfilled only by leisure customers usually not willing to pay the normal prices. For example, low-price tickets had to be bought 30 days in advance, preventing business travelers from buying these tickets. Thus, the introduction of specialized tickets designed for a specific customer class enabled American Airlines to skim much more revenues from the total possible market potential.

In the last decades, revenue management (RM) has become a very popular method of managing demand to increase profitability. This is not astonishing given the high revenue increasing potentials of RM. Boyd (1998, p. 29) for instance states that "revenue improvements from implementing a revenue management system can range from 2–8 percent (or more) depending on the carrier". The German airline Lufthansa AG reported an increase in revenues of €715 Mio. in 1997 (Klophaus, 1998, p. 150)—approximately equal to the result of normal operations in this year (Kimms and Klein, 2005, p. 2). As seen in the case of American Airlines, the success of RM essentially relies on identifying and exploiting differences in the customers' willingness to pay.

However, RM is mainly deployed in service industries—as for example airlines, car rentals, or hotels. It has not (yet) proven to be as successful in other domains of application as, e.g., in manufacturing. In those industries, different demand management concepts evolved in the past (for an overview see Fleischmann and Meyr, 2004). Demand management in manufacturing is often handled by a demand fulfillment module of the so-called advanced planning systems (APS). This module takes into account production quantities determined by a mid-term master planning module and short-term production planning. Based on these quantities, the demand fulfillment module decides on the basis of simple rules which customer to fulfill at which time, e.g. rules such as the first-come-first-served (FCFS) principle. As these rules are rather simple and created with a focus on general applicability, the results of demand fulfillment in APS leave space for improvements. Therefore, demand management in manufacturing might learn from the experiences gained in the service industries during the last decades.

Accordingly, practitioners as well as researchers put more and more effort in exploring ways to adapt RM concepts to the specific needs of manufacturing (Harris and Pinder, 1995, Swann, 1999, Arslan et al., 2007, Gupta and Wang, 2007). The core idea is that customer differentiation is beneficial also in a manufacturing environment. Additionally, the building block of RM in the service industries—perishable assets—corresponds to perishable capacity in manufacturing industries. For instance, an empty seat in an airplane can be compared to a machine standing still due to an insufficient number of orders.

However, the majority of scientific research focuses on adapting RM concepts to make-to-order environments because of the mentioned correspondence of perishable assets and perishable capacities in make-to-order manufacturing. In make-to-stock manufacturing environments, this correspondence does not apply as the machines schedules are based on forecasts instead of specific customer orders. The aim of this thesis is thus to analyze the current state-of-the-art in demand management (irrespective of a specific industry), and then relating the ideas found in the literature to make-to-stock manufacturing environments.

Our starting point is the current process of demand fulfillment in APS for make-to-stock manufacturing. In the case of make-to-stock, production planning is done on the basis of demand forecasts: when a customer order arrives, it can be either fulfilled from on-hand inventory or postponed to later arriving supply. The basic question to be answered in this thesis is to decide if it pays off to refuse a low margin customer order in expectancy of future more profitable orders.

The analysis relies on a number of assumptions as summarized in the following:

- Make-to-stock manufacturing environment with scarce capacities
- Deterministic future incoming supply
- Customers with different priorities
- Immediate order confirmation required
- Customers are willing to accept a late delivery under a price discount

In the short-term, it is assumed that the later arriving supply quantities are known and can be promised to arriving customers. Additionally, we assume a setting of scarce capacities, because the case of oversupply in make-to-stock manufacturing reduces to simply accepting and fulfilling all arriving customers orders. A further assumption in this work is that customers can be segmented according to their different willingness to pay, different costs of fulfillment, or different strategic importance. The first case typically applies to airlines when they charge different prices according to the remaining booking time and other factors like remaining capacity. In the second case, the costs of serving a customer order can be used as a differentiator. Note that only those costs which can still be influenced when accepting the order are relevant here. This includes, for example, transportation costs, taxes, and any variable costs of downstream production. The third case, the discrimination according to the strategic importance of customers may go beyond immediate costs and revenues. For example, loyal customers may be extremely important and should be given more favorable terms than occasional customers (see Quante et al. (2009b, Sect. 3.1.5) for a further discussion). In addition, customers are assumed to require an immediate response to their order, but are willing to accept a late delivery under a price discount. Note that these assumptions are equivalent to those in the work of Meyr (2009).

1.2 Organization, Objectives and Contributions

The idea of this thesis originates from the current state of the art of demand fulfillment in make-to-stock manufacturing. where in general APS are used as supporting tools. Therefore, this work starts in Chapter 2 with a description of the current state-of-the-art in demand fulfillment and introduces the required terms and definitions.

In order to search the literature for alternative approaches and concepts suitable for make-to-stock manufacturing, we decided to systematically classify the literature dealing with demand management. The focus was explicitly also beyond manufacturing when reviewing the literature, since we want to search in other disciplines for further ideas. We start introducing a framework for demand management (DM) in Chapter 3 and identify generic model types. In addition, a classification of commercial software solutions is presented in order to get an idea of how these solutions work.

Subsequently, based on the framework of Chapter 3, the general types are aligned to the specific requirements of make-to-stock manufacturing at the beginning of Chapter 4 and shortcomings of the respective model types are identified. A detailed analysis of specific models follows with a focus on manufacturing environments, but without concentrating on make-to-stock systems at this point.

Based on the literature review, Chapter 5 presents new models that reflect important characteristics of order fulfillment in make-to-stock production environments, namely customer heterogeneity, limited short-term supply flexibility, and short-term allocation flexibility. Previous literature has not addressed the interplay between these factors. The presented models are primarily based on the ideas of revenue management. We prove structural properties of the models and derive an optimal demand fulfillment policy. The result links order fulfillment in make-to-stock manufacturing to revenue management concepts. By this, we provide a way to unite the currently distinct concepts.

As these models of Chapter 5 explicitly take into account stochastic demand, we compare the developed models with existing deterministic ones described in Section 4.3. Before we conduct an extensive numerical study assessing the performance of various models in Chapter 7, we introduce the used simulation environment in Chapter 6. We show the superiority of the developed approaches in stochastic environments over the simple FCFS policy and the deterministic models of Section 4.3. Additionally, we identify which key influence factors drive the potential benefits.

This work concludes in Chapter 8 with a discussion of the results and issues for future research.

Chapter 2

Demand Fulfillment in Make-to-Stock Manufacturing

In this chapter, we provide a description of the current state of demand fulfillment in make-to-stock manufacturing. In order to do this, we start with an introduction of basic concepts as the customer order decoupling point. Subsequently, we discuss the structure of advanced planning systems and show the activities involved in an exemplary demand fulfillment process in make-to-stock manufacturing. Additionally, we explain the notion of available-to-promise.

2.1 Make-to-Stock and the Customer Order Decoupling Point

In the work of Fleischmann and Meyr (2004) it is shown that demand fulfillment is strongly related to the position of the customer order decoupling point (CODP). Therefore—before we explain demand fulfillment and advanced planning systems in more detail—we give a short explanation of the CODP. As seen in Fig. 2.1, the CODP divides the supply chain into forecast-driven and order-driven processes (Sharman, 1984, Hoekstra and Romme, 1992).

The CODP holds the inventory that is needed to hedge against forecast errors and replenishment uncertainty. The CODP plays a pivotal role in our analysis since many decisions upstream of the CODP are dependent on the available inventory and on future replenishment orders. For a more detailed discussion of the CODP concept and its impact in different production environments (including make-to-order (MTO), assemble-to-order (ATO) and

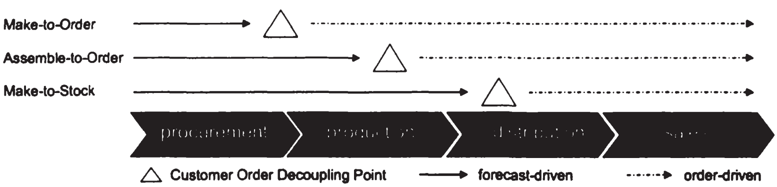

Figure 2.1: Customer Order Decoupling Point

make-to-stock (MTS)) we refer to Fleischmann and Meyr (2004). Note that the supply chain may contain additional production processes downstream of the CODP, but this is not the case for MTS production, since the CODP holds already the final product.

2.2 Structure of Advanced Planning Systems

Tasks of demand management are well-established in make-to-stock manufacturing companies and are usually supported by modern information systems, in particular advanced planning systems. These systems provide support for the entire planning tasks along the supply chain, from long-term strategic decision making to short-term operational decisions. Rohde et al. (2000) classify the supply chain decisions in a two-dimensional matrix (Fig. 2.2). This matrix is vertically structured according to the planning horizon and horizontally structured according to the sequence of planning tasks—from upstream to downstream in the supply chain.

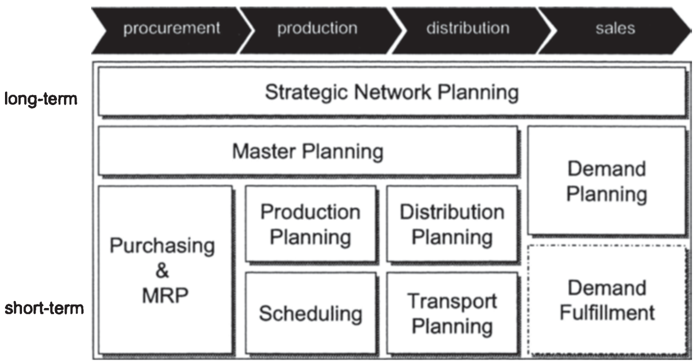

Figure 2.2: Structure of Advanced Planning Systems (Rohde et al., 2000)

The planning tasks in the supply chain extend from procurement of raw materials over production and distribution towards selling of final products. Since all these tasks are interrelated, a consecutive processing will not lead to optimal plans. For instance, selling of products can only be done with information about production and distribution in order to generate reliable due dates. On the other hand, the production task requires reliable demand forecasts in order to decide about lot-sizes and working times. To be able to support all these tasks, APS usually have several modules structured according to planning horizon and forecast accuracy. In order to cope with the interdependencies, the modules are organized in a hierarchical order. Long-term decisions based on low forecast accuracy provide the limits of the lower tasks which are done on the basis of better information. In order to do so, the upper tasks anticipate decision making in lower levels (Schneeweiß, 2003, Sect. 2.1). From top to down, the decisions become more accurate. Fig. 2.2 illustrates that on the top level, the *strategic network planning* is responsible for coordinating the entire supply chain. One the levels below, the tasks become more specialized.

The planning tasks associated with this thesis are located in the lower-right corner and are supported from the so-called *demand fulfillment* module. In order to illustrate how demand fulfillment for make-to-stock manufacturing is supported by advanced planning systems, Fig. 2.3 illustrates the involved supply chain activities and information flow in an exemplary consumer goods MTS supply chain.

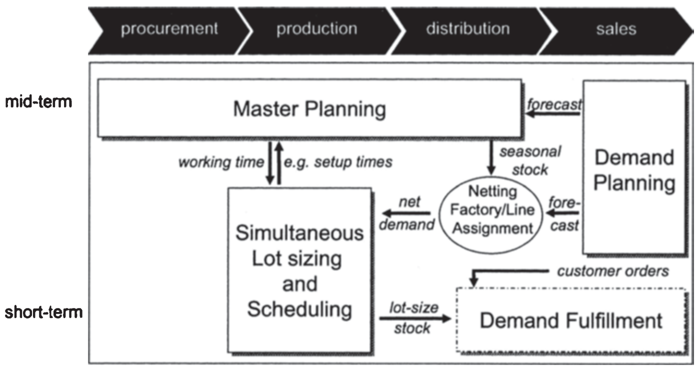

Figure 2.3: Exemplary Structure of an Advanced Planning System for Make-to-Stock Manufacturing (Fleischmann et al., 2008, p. 97)

In the mid-term, the *demand planning* activity is responsible for generating aggregate demand forecasts which are handed over to the *master planning* activity. It receives deterministic demand forecasts and prices as inputs from the demand planning (Kilger and Wagner, 2008) and then determines the best combinations of sales, production and replenishment quantities and the corresponding inventories under given capacity constraints. The planning horizon usually ranges from a few weeks up to several months. Therefore, data accuracy is low and the planning can only be done on the basis of aggregated data, i.e., products and customers are aggregated to groups or classes, respectively. In order to be able to balance supply with demand, the master planning activity basically determines the extent of seasonal stocks and possible changes of working times. A detailed description of the master planning activity can be found in the work of Rohde and Wagner (2008).

The generated master plan is handed over to the short-term activities. As demand forecasts usually become more accurate on the short-term, the planned stocks (from the master plan) and the current inventory are matched with the short-term demand forecasts. The result of this planning step is the *net demand*, which serves as input for the *lot-sizing and scheduling* task. This step additionally considers the working time restrictions from the master plan and computes the lot-sizes of the anticipated production quantities. Note, as this is a make-to-stock environment, all planning tasks up to here are performed on the basis of forecasts.

The last activity—which is the essential one in the scope of this work—is *demand fulfillment*. This activity takes into account the planned production quantities and quotes incoming customer orders according to their desired delivery date. This activity is the only one considered so far that is not based on forecasts but on the actual customer orders. In traditional material resource planning systems (MRP), order quotes are generated on the basis of available inventory. If there is no stock on-hand, the orders are quoted against the production lead-time. Kilger and Meyr (2008, p. 182) give a simple example that illustrates the weakness of such an order promising mechanism: it is not guaranteed that capacity constraints are not violated and a feasible plan is generated. Therefore, modern APS make use of more sophisticated methods. The inner logic of these methods is based on the notion of available-to-promise (ATP). In the following section, we introduce this key functionality of modern APS demand fulfillment solutions.

2.3 Available-to-Promise

2.3.1 Definition

The notion of available-to-promise is strongly related to advanced planning systems, and, with the success of APS solutions in the past years, is facing increasing attention. Nevertheless, it is not a very new concept. Fischer (2001) and Kilger and Meyr (2008) mention the work of Schwendinger (1979) as the earliest reference to available-to-promise. Since there is already a detailed analysis and literature review of ATP (see Fischer, 2001), we will not cite any references before the year 2001 and put an emphasis on the years thereafter.

Kilger and Meyr (2008) define ATP as "... the current and future availability of supply and capacity that can be used to accept new customer orders". A different definition of ATP is provided in the work of Ball et al. (2004), in which ATP is defined as "... a business function [...] directly linking customer orders with enterprise resources". The commercial software vendor SAP defines ATP quantity as the "... quantity available to MRP for new sales orders" and

ATP check as a "...function used to check [...] if a product can be confirmed" (SAP Help Portal, 2008).

Throughout this work, we follow the definition of Kilger and Meyr (2008) when referring to the term ATP and see it as a quantity rather than a functionality. It is important to note that a positive ATP quantity does not mean that there is stock on-hand, because ATP takes also future available quantities into account that are still to be produced. In the work of Fischer (2001), four different functions associated with the notion of ATP are identified:

- Availability check of products and evaluation of alternative solutions
- Order confirmation and due date assignment
- Steps taken in case of temporary inability to deliver
- Due date monitoring and order repromising

In the first step, the ATP quantities are used to check the availability of products. In cases when the product is currently not available, the next time the product will be available can directly be derived from the ATP quantities (as future supplies are considered). The availability check allows for a direct confirmation of customer orders including the determination of due dates. These two functions are essential for this work. The other two functions involve more activities as demand fulfillment and are beyond the scope of this work. For example, in case of temporary inability to deliver, a re-planning of the production (lot-sizing and scheduling) has to be done.

2.3.2 Dimensions of ATP

As seen before, the master planning activity calculates aggregated plans structured according to certain categories, e.g., product or customer groups. As the master plan is generally the basis for calculating ATP quantities, ATP exhibits a similar structure. Kilger and Meyr (2008, p. 184) mention several alternative possibilities to structure ATP quantities, e.g., location, sourcing type, region, market etc. but refer to *product, time* and *customer* as the most important ones. In the following, we give a short description of these three dimensions.

The level of detail of the ATP quantities in the ***product dimension*** corresponds to the location of the CODP. In case of MTO/ATO, products are customer specific and production/assembly is driven by incoming customer orders. The current and future supply and capacity that is taken into account to accept new customer orders is represented by the production capacity after the CODP and the inventory of semi-finished goods stored at the CODP. In the case of MTS manufacturing, the production is entirely based on forecasts. The current and future supply taken into account when the customer order arrives is hence represented by finished goods.

The master plan is generally created for a few weeks up to a few months. Accordingly, the granularity of the ***time dimension*** ranges from days to months, mostly depending on the forecast accuracy. The ATP quantities are therefore also structured according to certain time buckets, from which the customer orders are fulfilled.

The third important dimension is the ***customer dimension***. Kilger and Meyr (2008) distinguish this dimension according to a supply- or demand constraint mode. Since the motivation of this thesis is driven by the problem to decide which customer to fulfill first when capacity is scarce, the supply constraint mode is the important one in this thesis. In this mode, not all orders can be fulfilled and the ATP functionality has to provide means to decide about customer priorities. Thus, APS assign customers to certain classes in order to have a customer hierarchy.

The allocation of the quantities from the master plan to customer hierarchies is usually done in APS on the basis of simple rules. Kilger and Meyr (2008, Sect. 9.4.3) distinguishes three important rules:

- *Rank based*: Allocation of quantities according to predefined ranks. The customer with the highest rank gets the forecasted quantities, the lower ranks the remaining quantities.
- *Per committed*: The quantities are allocated proportionally to the forecasts of the different customers. If customer A has forecasted 100, and B 200, then B gets twice as much as A irrespective of the actual available quantities.
- *Fixed split*: This rule is independent of the demand forecasts. Quantities are allocated according to a predefined fixed ratio, e.g., customer A gets 60%, B gets 40%.

The process of determining the available quantities is called *allocation planning* (AP).

To find a reliable due date for a customer order, it is searched through demand fulfillment alternatives in the mentioned dimensions. This means, e.g., to search in the time dimension, checking for ATP back- or forwards in time, in the product dimension, checking for substitute products, and in the customer dimension, checking for availability in other priority classes (Kilger and Meyr, 2008). This process is called *order promising.* Usually, simple rules are defined as search strategies for the different dimensions (Meyr et al., 2008a, Sect. 18.3.1). Profitability of different fulfillment alternatives is generally not taken into account during this search. However, recent software systems do not only consider *available-to-promise* (ATP) quantities (available inventory) or *capable-to-promise* (CTP) quantities (available capacity), but also follow a *profitable-to-promise* (PTP) logic that enables them to compare customer orders and fulfillment alternatives according to their priority. CTP quantities are defined as "... the remaining capacity of the assembly lines, if this capacity is a potential bottleneck" (Fleischmann and Meyr, 2004). As they are only applicable in MTO and ATO production systems, CTP is not relevant here.

Chapter 3

A Framework for Demand Management

In this chapter, we propose a classification framework for demand management. This framework is subsequently used to identify general types of models and software in relation to the most important key-decision variables in supply chain management. To demonstrate the correspondence of the identified general types of models with current scientific research, we present and discuss important review papers for each of the identified types. In contrast to the subsequent Chapter 4, we will not discuss single research papers in this chapter.

We start this chapter with a definition of demand management as it is understood in this work, following by a detailed description of the framework. Once the framework is introduced, we discuss the identified model types supplemented by references to important review papers.

3.1 Demand Management Defined

Demand management as it is understood in this work is closely related to the previously described demand fulfillment in APS. As we have seen, the notion of demand fulfillment is common in manufacturing, but similar concepts exist in other industries, e.g., the earlier mentioned revenue management in the service industries. Therefore—as we do not solely build on the notion of demand fulfillment—we refer to the concepts developed in this work as *demand management* concepts.

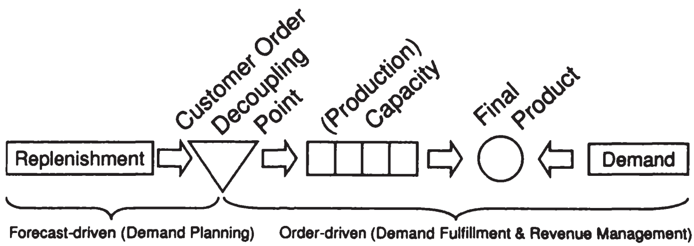

Figure 3.1: Supply Chain Elements (Quante et al., 2009b)

In order to define demand management in the following, we adopt the supply chain framework of Quante et al. (2009b) as shown in Fig. 3.1. Relevant in this work are the processes in the supply chain downstream of the CODP (order-driven processes). The framework shows the analogy between demand fulfillment and revenue management as it extends elements like the CODP (which is important for demand fulfillment) with elements representing revenue management concepts. For instance, pricing decisions can be linked to the item *final product* and capacity allocation decisions are linked to the item *demand*.

Demand management is closely tied to decisions in the depicted supply chain and therefore depends on its specific characteristics. For example, the current inventory at the CODP or the remaining production capacity may influence pricing decisions or promised due dates. By capturing these characteristics, the depicted elements of the supply chain framework provide a systematic basis for identifying demand management models and software.

3.2 General Model Types for Demand Management

3.2.1 Classifying Demand Management Models

This section introduces a framework for generic demand management types developed to classify different literature streams. The structure of the framework follows the work of Quante et al. (2009b). The authors distinguish demand management models according to demand or price decisions, and control over the replenishment quantity (as shown in Fig. 3.2). With these two dimensions, distinct model types corresponding to common research streams are identified. Quante et al. (2009b) note that these two dimensions are the "key decision variables regarding demand and supply". Before reviewing each model type in detail, the two dimensions are described in the following. Models in the first row of Fig. 3.2 take demand as entirely exogenous. They satisfy demand first-come-first-served (FCFS) at a given price. In particular, these models do not consider a segmentation of customers. The middle and bottom row of Fig. 3.2 represent a more active management of demand. The models in the second row influence the demand by adjusting prices. Like in the first row, the customers are not segmented into specific classes and treated all equal just distinguished according to their willingness to pay. The bottom row entails models that explicitly consider heterogeneous customers. In response to a customer request, these models face a trade-off between accepting a current, low-priority customer now versus reserving the resources for high-priority customers expected in the future. The distinction between the middle and bottom row in Fig. 3.2 corresponds with the classification of RM models of Talluri and van Ryzin

Demand / Price consideration	Replenishment consideration: None	Replenishment consideration: Data	Replenishment consideration: Decision variable
Data	-	Order Promising	Stochastic Inventory Control
Price-based	Markdown / Pricing / Auctions	Trade Promotions	Integrated Pricing
Quantity-based	Traditional RM	aATP	Inventory Rationing

Figure 3.2: Types of Demand Management Models (Quante et al., 2009b)

(2004). In order to follow their terminology, models are labeled as *price-based* or *quantity-based.*

The columns in Fig. 3.2 reflect the way models decide about inventory replenishment at the CODP. In the first column, models do not consider replenishments. As for example, in airline revenue management, the seats in an airplane can not be replenished when sold out. In the middle column, inventory replenishments are considered but not actively influenced. The right column entails models that actively decide about the quantities they want to replenish.

In the following, we briefly discuss each of the aforementioned model types. We will refer to available review papers for each of the model types. A detailed discussion of individual articles relevant in this work is following in Chapter 4.

3.2.2 Single-Class Exogenous Demand Models

No models were found fitting in the upper left cell of Fig. 3.2. This is not surprising since models with a given price and no consideration of replenishment or inventory, respectively, have nothing to decide on, neither on the demand nor on the replenishment side of the supply chain.

In the next cell to the right, the so-called ***order promising*** models consider price (i.e. demand), current inventory, and future replenishment quantities as given. This results in information about product availability and delivery times. For each incoming customer order the model decides real-time on the due date. The decision is made in a greedy fashion, based on availability of goods. As an example, the basic order promising model of Sect. 4.3.1 falls in this category. An introduction and overview of this so-called "real-time mode" or "single-order-processing" models is given in Ball et al. (2004), Chen et al. (2001), and Fleischmann and Meyr (2004). Additionally, a broad overview of due date management models with an emphasis on stochastic models is included in the work of Keskinocak and Tayur (2004).

The upper right cell of Fig. 3.2 holds the vast class of ***stochastic inventory control*** (SIC) models, which focus on optimal inventory replenishment. Some of these models primarily address the structure of optimal replenishment policies, as for example the famous (s, S)-policy proven by Scarf (1960). Other models seek to determine optimal control parameters of such policies, such as the optimal ordering time, order quantity and inventory review intervals. Many SIC models build on the classical *newsvendor* model, which seeks to determine the optimal order quantity for a perishable product under stochastic demand. An overview of single-period newsvendor problems is given by Khouja (1999). Silver (1981) provides an overview and typology of many standard inventory problems, such as the ones mentioned above. General up-to-date overviews of inventory models can be found in the textbooks by Silver et al. (1998), Porteus (2002) and Tempelmeier (2006).

3.2.3 Price-Based Demand Models

The model types in the middle row of Fig. 3.2 treat price as a decision, which influences the demand. Pure ***pricing*** models aim to determine an optimal selling price, without considering replenishments. For example, given a price-demand relation, the goal is to find the price which maximizes total revenues. Mild et al. (2006) review factors influencing demand and show how to find optimal prices.

Markdown models determine the right price path for inventory clearance for a given amount of inventory, which cannot be replenished during the planning horizon. Elmaghraby and Keskinocak (2003) classify several dynamic pricing models with and without replenishment decisions, the latter ones including markdown models.

Auctions, as discussed for example by Talluri and van Ryzin (2004, Sect. 6), take a fundamentally different approach to pricing. They provide a price-discovery mechanism and thereby an alternative to posting fixed prices. This approach is particularly valuable if little demand information is available. The aforementioned authors discuss the close connection between auctions and dynamic pricing.

Trade promotion models represent a type of pricing models that consider replenishments as an exogenous input and therefore fit in the second column of Fig. 3.2. Neslin (2002) provides an overview and discusses the reasons for promotions.

Research in ***integrated pricing*** (IP) models dates back to Whitin (1955) who extends the EOQ-formula as well as the classical newsvendor model with price decisions. This field has seen extensive research in the last decades, which is summarized, for example, by Petruzzi and Dada (1999). Recent research focuses on multiple period models, which are discussed in the well-known literature reviews of Chan et al. (2004), Elmaghraby and Keskinocak (2003) and Yano and Gilbert (2003). Few models exist for environments in which replenishment, prices and due dates are set simultaneously. Some models of this type and other models dealing with due date setting can be found, for example, in the previously mentioned review paper by Keskinocak and Tayur (2004).

From an application-oriented perspective it is worthwhile comparing IP and a successive application of pricing & SIC models. While IP models recognize the interdependence between pricing and replenishment and therefore determine decisions simultaneously, they do so at the cost of a more simplified demand and supply representation. Pure pricing models may include sophisticated demand functions, including reference price effects, promotion effects, and competition (Mild et al., 2006). Similarly, SIC models consider factors such as multiple suppliers and quantity-discounts. IP models typically cannot deal with these factors due to tractability (Elmaghraby and Keskinocak, 2003, Sect. 4).

3.2.4 Quantity-Based Demand Models

Models in the bottom row of Fig. 3.2 take prices as exogenous but manage demand by means of rationing strategies. In contrast to the models of the top and middle row, the models distinguish multiple customer classes and prioritize them rather than fulfilling orders in an FCFS manner.

The type ***traditional revenue management*** (TRM) in the first cell of the third row refers to models that are common in airline applications. In these models, given units of a perishable product (e.g., seats on a flight on a specific day) are allocated to customers with different priorities or different willingness to pay. The basic question is whether to accept a given order or to reserve capacity in anticipation of more profitable future orders. McGill and van Ryzin (1999) and Pak and Piersma (2002) provide an overview and a short history of research in traditional revenue management with a focus on airline applications. Boyd and Bilegan (2003) discuss models focusing on e-commerce applications. The recent review by Chiang et al. (2007) includes an overview of RM practices in different industries.

Models of the type *Allocated available-to-promise* (***aATP***) are similar to the order promising type of the top row except for differentiating between multiple customer classes. Scarce resources (inventory on hand, planned stock at the CODP or capacity downstream of the DP) are allocated to these classes according to customer profitability or other priority measures. Within each class, customer requests are usually handled FCFS, just as in traditional order promising. Examples of these type have been discussed before in Sect. 4.3.2. Guerrero and Kern (1988) introduce the general problem of accepting and refusing orders and discuss the requirements and implications of order promising mechanisms. For reviews of the mostly deterministic models of this type the reader is referred to Kilger and Meyr (2008) and Pibernik (2005).

If customer requests do not have to be answered instantaneously, several customer orders can be collected and jointly promised in a *batch*, thereby creating higher degrees of freedom for selecting the most important or profitable orders within a simultaneous optimization process. Overviews of these so-called "batch order promising models" can again be found in the work of Ball et al. (2004), Chen et al. (2001) or Fleischmann and Meyr (2004).

A review of integrated due-date management and job-scheduling models with deterministic orders is provided by Gordon et al. (2002). The article considers batch-models in which due dates are determined according to current capacity and the desired delivery date. Keskinocak and Tayur (2004) give a general overview of due-date setting models.

aATP and TRM models are similar in that they decide about demand fulfillment with respect to different customer classes. The most significant difference concerns the perishability of resources. TRM considers "perishable" products, e.g., empty seats on a specific flight, which are lost after the departure date, whereas the ATP quantities managed in aATP models are generally storable. Another difference concerns the time horizon. TRM models typically consider a fixed day of capacity availability, e.g., the departure date of a flight. In contrast, aATP models consider multiple periods linked through the storability of excess inventory. Furthermore, aATP models usually assume deterministic demand whereas demand in TRM models is stochastic.

The last model type within the grid concerns ***inventory rationing*** (IR) models. Similar to the relationship between aATP and order promising, IR models extend SIC models by distinguishing and prioritizing multiple customer classes. For an early review refer to Kleijn and Dekker (1998). As traditional SIC models, IR models may consider deterministic or stochastic replenishment lead times. A further distinction within this class of models concerns the number of demand classes considered, which may be general or limited to two classes.

IR and aATP models differ in terms of exogenous versus endogenous replenishment. Specifically, IR models consider replenishment decisions with stationary deterministic or stochastic lead times. In contrast, aATP typically considers capacitated, dynamic and deterministic arrivals of push-based production (=replenishment) quantities. To this end, aATP usually assumes deterministic and dynamic demand forecasts whereas IR models assume stochastic demand.

In addition to the model types captured in Fig. 3.2, a few recent research streams aim to combine several types by simultaneously considering multiple attributes. For example, Kocabıyıkoğlu and Popescu (2005) jointly analyze price and allocation decisions with two customer classes. Since most quantity-based models assume exogenous prices, this seems to be a promising direction for future research. Bitran and Caldentey (2003) formulate a general model of this problem and review the current state of research. Another approach is pursued in Ding et al. (2006) in which trade promotion models are combined with inventory rationing models. The authors denote the resulting new problem type by ADP, referring to the allocation of available stock, discounting and prioritization of customers.

3.3 General Software Types for Demand Management

3.3.1 Classifying Demand Management Software

The software market for demand and supply chain solutions has changed in recent years. For many years the focus was on the supply side. The interest is now, however, turning to end-to-end solutions including the demand side. Big supply chain solution providers like Oracle and SAP are investing large amounts in the acquisition of demand-related know-how. For example, in 2005 SAP took over Khimetrics, a leading vendor of markdown, price, and promotion-optimization solutions. Oracle—after taking over one of its largest competitors in supply chain solutions, Peoplesoft, in 2005—simultaneously invested in the demand solutions of Demantra (2006), ProfitLogic (2005) and Retek (2005), all of them leading vendors of retail revenue management software. Another big consolidation occurred in 2006 when JDA Software—a provider of specialized retail solutions—took over Manugistics, a supply chain solution provider focusing on profit optimization in the consumer goods industry.

The scope of our current analysis is restricted to software supporting short-term decision making in DM. These solutions draw data from other software systems, such as Customer Relationship Management systems on the demand side (Buttle, 2004) and Enterprise Resource Planning systems (Stadtler and

Kilger, 2008) on the supply side. Since these systems themselves do not focus on decision making we do not include them in our analysis.

As discussed in the previous section, scientific optimization models are fairly well described in the literature. One can easily identify data, decision variables, restrictions and solution strategies. Moreover, the solution quality is often analyzed in detailed numerical studies. This is different for commercial software solutions. Usually, available information is scarce and reveals little of the underlying technology. Software users can only assess the supported input data, available options, and the resulting output that is automatically calculated. The solution quality can hardly be evaluated objectively and is usually judged by user experience.

Our analysis of software modules reflects this limited availability of objective information. We build our characterization of software types and functionalities primarily on available software reviews and whitepapers. As a starting point we use essentially the same dimensions as for the scientific models. Model data and decisions roughly correspond with software input and output, respectively.

Fig. 3.3 structures software types along the same axes as the model types of Section 3.2. We choose names according to the functionality of commercial software modules on the market. The remainder of this section briefly reviews each of these software types.

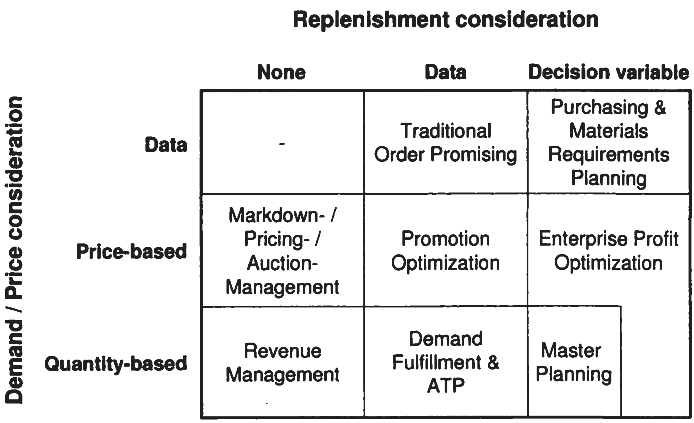

Figure 3.3: Types of Demand Management Software (Quante et al., 2009b)

3.3.2 Single-Class Exogenous Demand Solutions

The mid upper cell of Fig. 3.3 denoted by ***traditional order promising*** contains traditional software modules for short-term order promising under known inventory availability. When a customer order arrives, the software simply determines whether the order can be satisfied out of available inventory. If not, the order is backlogged according to a standard lead time without considering future capacity or additional incoming supply. It is easy to see that this approach can lead to an order peak after the standard lead time and thus to severe capacity problems in the future. Kilger and Meyr (2008) illustrate this situation in a simple example.

Refilling of inventory is usually left to ***purchasing & materials requirements planning*** modules, which are part of enterprise resource planning (ERP) systems. Essentially, these systems support refilling of non-bottleneck material and components from a single vendor. An overview of these classical systems can be found, for example, in the textbook of Vollmann et al. (2005). Since these classical systems provide sufficient solution quality only for very simple settings, specialized inventory modules consider extensions such as capacitated replenishment, stochastic demand, and multiple suppliers (Stadtler, 2008). Such modules usually are part of larger advanced planning and supply chain planning software suites. Additionally, there are specialized vendors of supply-chain wide inventory optimization tools, such as Optiant (Optiant, 2007) with its inventory suite Powerchain and Smartops (Smartops, 2007).

3.3.3 Price-Based Solutions

Markdown management systems are mainly used in retail, for example for end-of-season stock clearance. An example of markdown management systems is B_Line, described by Mantrala and Rao (2001) under the name MARK. The system takes possible prices and corresponding demand probability distributions for each period as inputs and can find both markdown and markup price paths. The output consists of a specific price in each period. Furthermore, MARK is capable of finding a suitable amount of initial inventory by iterating through a discrete set of possible inventory levels. Elmaghraby and Keskinocak (2003, Sect. 3.2) describe the capability of markdown solutions.

Software systems of the type ***pricing management*** are relatively new. This is due to improvements in computing power and increased availability of past sales data. The rise of data warehouses and cheap computing power has recently allowed the use of automated pricing systems for many applications. Pricing management systems are based on complex price-demand functions for which suitable parameters have to be estimated, a process requiring vast amounts of past sales data. For example, to estimate price elasticity, the sales data must include a certain degree of diversity, corresponding with at least a few past price changes. Capacity or inventory restrictions are usually not considered in these types of software (see for example Mild et al., 2006).

The quick expansion of e-commerce applications has boosted the use of ***auction systems***. The large number of different systems merits a review in its own right and exceeds the scope of our analysis. We refer to Kambil and van Heck (2002) for a systematic introduction to this field. Vakali et al. (2001) discuss the characteristics of internet-based auction systems and present a short survey of popular applications.

Similar to the previously described markdown systems, ***promotion optimization*** is also used in retail environments, as described by Elmaghraby and Keskinocak (2003, Sect. 3.2). Very detailed information about the capability of such systems can be found on-line, for example from the vendors mentioned at the beginning of this section.

The term ***enterprise profit optimization*** (EPO) was coined by the software company Manugistics, who claims to be the first vendor offering an integrated pricing and supply solution (Manugistics, 2002). Furthermore, Manugistics software is meant to be able to allocate scarce resources to the most profitable customers, thus simultaneously applying ideas of quantity-based DM. Demand and supply planning is realized in many solutions, but not in an integrated way and not including price decisions. Most APS forecast demand for different price levels and then successively analyze—within the context of mid-term planning—several what-if scenarios and their effects on the total supply chain.

3.3.4 Quantity-Based Solutions

APS software modules that support mid-term, aggregated supply and demand decisions are known as ***master planning*** modules (Meyr et al., 2008b) as previously mentioned in Sect. 2.2. They receive deterministic demand forecasts and prices as inputs from the demand planning module of APS (Kilger and Wagner, 2008) and then determine the best combinations of sales, production and replenishment quantities and the corresponding inventories under given capacity constraints. Quantities can be allocated to different customer classes. In terms of the supply chain framework in Fig. 3.1, master planning modules deal with forecast-driven planning activities (e.g., push-based replenishment of the CODP) and therefore fall outside the scope of our definition of DM. However, we feel that they deserve mention since their resulting allocations serve as the primary input for the short-term, capacity-checked order promising, executed by the Demand Fulfillment and ATP modules of APS. A detailed list of options considered in master planning modules can be found in the work of Rohde and Wagner (2008).

The type in the lower middle cell takes capacity and inventory replenishments into account and corresponds to the ***demand fulfillment & ATP*** modules of APS previously described in Chapter 2. These modules extend the aforementioned traditional order promising and determine due dates for incoming customer orders, which promise to be more reliable than simple standard lead times. In addition, if ATP quantities are allocated to customer priority classes—in the usually implemented aggregated way—order promising differentiates with respect to customer importance, based on customer profitability or strategic impact.

Revenue management software is widely used by airlines, hotel chains, and car rental agencies. RM software systems basically take the given capacity and offered tariffs as input and decide on acceptance or rejection of customer orders. One of the main differences with demand fulfillment & ATP is that RM software focuses on revenues rather than costs. Furthermore, RM systems usually forecast demand in much more detail than demand fulfillment modules, e.g., for each flight, on each day, and for each customer class. These forecasts require a large amount of historical sales data in order to be reliable. Modern revenue management systems can handle many additional industry-specific issues, such as overbooking and connecting flights in the airline context (Talluri and van Ryzin, 2004, Sect. 10.1.3, 11.2).

Chapter 4

Demand Management Models in MTS Manufacturing

In this chapter, we align the identified general types of demand management described in Chapter 3 with the characteristics of manufacturing environments. By this, we can further narrow the domain of demand management models by excluding those ones that appear hardly suitable for MTS production systems. After identification of the most suitable models, we complement the general review from the previous chapter by a more detailed review of specialized papers taking into account the characteristics of manufacturing environments.

4.1 Matching of Model and Software Types to the Requirements of MTS Manufacturing

In this section, we identify alignments and misalignments between models and software within the domain of MTS manufacturing. Specifically, the goal of this section is to identify the most appropriate models and software types for DM decisions in MTS manufacturing. We also seek to highlight remaining research needs and provide the motivation for the models presented in Chapter 5.

We build our discussion around the structure of Figures 3.2 and 3.3. Specifically, we compare the supply flexibility in MTS manufacturing with the way that replenishment decisions are supported by different model and software types. Similarly, the demand flexibility is tied to the supported demand management. In this way, we identify the most appropriate cell(s) in Figures 3.2 and 3.3. This allows us to recognize empty spots and future research needs.

In manufacturing, production processes are the most important—and usually most costly—process steps. The customer order decoupling point as the interface between forecast-driven demand planning and customer-oriented demand fulfillment describes whether a certain production process is operated under demand (un)certainty, what type of stocks (raw material, components, final products) have to be held, where the main bottlenecks (stocks, production capacity) can be expected, and how long customer service times will be.

In MTS environments, all production processes are executed in a forecast-based way. Due to upstream capacity limitations, production planning decides on short-term replenishment of CODP inventories in a push-based, "vendor-driven" manner. Thus, models including replenishment decisions (third column of Fig. 3.2) can only support the mid-term, forecast-based demand planning, but not the short-term demand fulfillment. In order to make use of the (uncertain) information on future CODP inventory replenishments, as implied by the production plans, demand fulfillment models of the second column appear the most appropriate.

Because of the MTS market conditions and contracting practice, pricing decisions typically have to be taken on a mid-term basis. For example, the Demand Planning module of an APS forecasts several price-demand scenarios, including different alternatives for price discounts or promotions. These scenarios are passed on to a master planning module, which checks each of them with respect to supply chain constraints, selects the most profitable one, and generates directives for the (forecast-based) short-term production planning. Thus, short-term pricing flexibility is rather limited, which rules out the models on the second row of Fig. 3.2. Price-based approaches appear mainly applicable on the mid-term planning level, e.g. to determine demand forecasts in conjunction with optimal prices.

The remaining models are order promising and aATP models as the most applicable ones for DM in MTS manufacturing. Both of them consider the current level of CODP inventory. Order promising in an MTS environment searches through the ATP quantities in an FCFS manner to be able to fulfill a customer order. Newer approaches process several customer orders in a batch and allocate ATP to the most profitable customer orders. aATP models overcome the disadvantage of batch order promising, namely not providing a real-time order promise and forcing the customer to wait. However, they are dependent on forecast-based information on CODP inventories that is provided by the master and production plans, and on the possibility of customer segmentation. The ATP search rules that are used to consume the allocated ATP quantities of the different customer classes follow similar ideas as traditional RM methods. However, since products in MTS manufacturing are durable and can be stored, they also have to be able to "search over time", i.e. to take future inventory replenishments into account.

The results of this section can be summarized as follows: Replenishment decisions are not relevant for short-term demand fulfillment in MTS manufacturing, which rules out the third column in Fig. 3.2. The second row is also not appropriate for short-term applications due to the low pricing flexibility. Traditional RM methods do not consider storability of inventory. The

Replenishment consideration

	None	Data	Decision variable
Quantity-based	Traditional RM *+ storability* *+ several replenishments*	aATP *+ stochastics*	Inventory Rationing *+ deterministic supplies* *+ no ordering*

Figure 4.1: Quantity-Based Demand Management Models

matching of the generalized model and software types of Chap. 3 to the specific characteristics of MTS manufacturing leads to the following research questions:

How can demand management models be adapted to the specific needs of make-to-stock manufacturing? Unlike in service environments, order promising in manufacturing industries is a multi-period problem, i.e. production earlier or deliveries later than the customers' requested date are possible. Therefore, traditional RM techniques typically cannot be applied as is in MTS manufacturing industries. They have to be adapted to deal with the holding and (future) replenishment of CODP inventory and costs. The allocation and ATP consumption rules currently used in APS serve this purpose, but are very basic. Furthermore, stochastic influences from the demand side are usually ignored. Inventory rationing models cannot be applied as is to MTS manufacturing due to their short-term replenishment decisions. Excluding replenishment decisions from IR models and using the "remaining" rationing policies for demand fulfillment seems to be a way of adapting such models. Fig. 4.1 summarizes the match of models and software to the requirements of MTS manufacturing. Subsequently, such new models are shown in Chap. 5.

Since it is not clear how different models perform under realistic conditions, a comprehensive performance analysis considering the specific characteristics of MTS manufacturing will be made in Chap. 7. The analysis should help to assess the different models and rules taken from commercial software solutions and should determine the main influence factors on the performance.

4.2 Quantity-Based Demand Management in Manufacturing

We will discuss the literature of the three identified types in more detail, as basically review papers have been mentioned before. Additionally, we take into account the characteristics of MTS manufacturing when reviewing the literature.

Table 4.1: Overview of Publications on TRM in Manufacturing (Method: C=Conceptual, M=Simulation/analytical models, R=Review, S=Case study)

Publication	CODP	Classes	Method
Harris and Pinder (1995)	ATO	2	C,M
Kalyan (2002)	MTO/ATO/MTS	–	C
Kuhn and Defregger (2004)	MTO	2	S
Rehkopf and Spengler (2004)	MTO	several	C,M
Barut and Sridharan (2005)	MTO	several	M
Kimms and Klein (2005)	MTO	–	C
Jalora (2006)	MTO	several	M
Kolisch and Zatta (2006)	MTO	–	C
Rehkopf (2006)	MTO	–	C,M,R
Kumar and Frederick (2007)	MTO/MTS	3	S
Spengler et al. (2007)	MTO	several	M
Chiang et al. (2007)	–	–	R
Specht and Gruß (2007)	ATO	–	C,S

4.2.1 Traditional Revenue Management

Overviews of *TRM* and its applications in different industries (including manufacturing) are provided by Kimms and Klein (2005) and Chiang et al. (2007). These overviews reveal that the literature on quantity-based RM in manufacturing is basically limited to MTO and ATO production environments. This can be explained by the correspondence of *production capacities* and *perishable assets* which are the building blocks of RM in the service industries (Weatherford and Bodily, 1992). Usually, TRM papers for manufacturing determine optimal protection levels for production capacities (located downstream of the CODP).

Table 4.1 shows important papers on TRM in manufacturing classified according to the CODP, the number of customer classes considered and the methodology applied in the paper. We distinguish between four methodologies: *conceptual* (C), *simulation/analytical models* (M), *review* (R), and *case studies* (S). Conceptual papers are understood as those describing and developing new concepts and ideas without defining a specific model. This is done in the next methodology of defining simulation- or analytical models. Review papers focus on providing a comprehensive state-of-the-art of the literature within a special field of research. They are often combined with conceptual work, e.g., the development of a classification framework. Finally, case studies focus on a detailed description of real-world applications.

A seminal paper for the adoption of RM concepts in manufacturing is the one from Harris and Pinder (1995). The authors describe the analogies between ATO production and the service industries and proposes a TRM model for a repair facility and a sports apparel manufacturer. Unlike many papers in TRM, the author describes not only how to determine protection levels for a two-class problem, but also how to determine the optimal selling prices.

Similar to the previous paper, Kalyan (2002) establishes a relationship between airline revenue management and its counterparts in manufacturing. The author introduces the notion of MAV (minimal acceptable value) and shows how MAV can be used in many different applications. MAV is understood as the value of a resource and is known in revenue management as the bid-price, a threshold for accepting those orders that generate higher revenues than the bid-price. The basic idea of the paper is that all assets or resources are associated with a certain value (MAV). If this value is known, it can be compared to any incoming order. If the price of the order is higher than the MAV, the order is accepted, and rejected otherwise. The paper gives many examples and shows how MAV can be calculated in manufacturing applications.

Kuhn and Defregger (2004) discuss a case of a paper production company with a fixed capacity on the short-term. Production is triggered by incoming orders (MTO) and customers are divided into two classes differentiated in resulting revenues (€ 50,000 vs. € 10,000) and required lead times (2 days vs. 3 days). If a customer order is accepted, the paper machine is busy for two days. The problem is to decide which of the low-margin customers to reject in anticipation of later arriving high-margin customers. The authors propose a Markov decision process solved by means of a linear program. A numerical study shows a profit increase of 5.4% compared to a simple FCFS policy.

The work of Rehkopf and Spengler (2004) is a further paper identifying the characteristics between the service industries and MTO manufacturing according to the application of revenue management. The authors focus on the iron and steel industry and show how a network revenue management problem (Talluri and van Ryzin, 2004) can be solved with a linear program in order to allocate capacity. This conceptual work is deepened in the later released paper of Spengler et al. (2007) which is discussed further below.

The work of Barut and Sridharan (2005) is selected as a representative of an entire line of research in scheduling and accepting orders in an MTO manufacturing setting. Rather than to focus just on the acceptance decisions, the authors—in contrast to the previously mentioned papers—take into account the scheduling and exact due date determination of incoming orders. They propose a heuristic procedure (called DCAP) that is able to process orders in distinct lots and handles more than two customer classes. The results indicate that DCAP performs better than an FCFS policy. Jalora (2006) presents a similar model as shown in Barut and Sridharan (2005) but extends their work by showing a dynamic order acceptance and scheduling policy evolving over time when the exact demand realizations are known.

Kolisch and Zatta (2006) present a short introduction into revenue management for the process industry. Their work includes an empirical study over 124 companies of the process industry. The large majority of companies confirmed that revenue management is a suitable method to cope with overcapacities and increasing competition. 80% of the surveyed companies currently apply methods of revenue management, the majority of them (74%) for capacity management and 15% price management.

Kumar and Frederick (2007) show the case of Andersen, an American window manufacturer. The case describes a simplified decision problem Andersen is facing. The company sells its products over three distinct distribution channels. In two of the three channels, Andersen produces windows upon order placement and generates only low or medium margins due to long lead times. In the third channel, the customer buys the window directly in a warehouse and pays a surplus due to the short lead times. Andersen is producing on-stock in this third channel. The aim is to find an optimal amount of inventory stored in the warehouse. A closed-form solution is presented in this study resembling the famous newsvendor problem.

Spengler et al. (2007) propose a model for the iron and steel industry in an MTO setting. The authors define a stochastic dynamic program based on a vector of CTP quantities. The optimal policy is to accept an order if its current contribution margin exceeds the expected marginal profits. As the problem size is too large to be solved directly with dynamic programming, the authors approximate the expected marginal profits by a multi-dimensional knapsack formulation. In numerical tests the proposed method results in increased contribution margins of 5.3% compared to a simple FCFS rule. This work as well as the previous mentioned work of Rehkopf and Spengler (2004) is included in the dissertation of Rehkopf (2006).

Specht and Gruß (2007) describe the successful application of a revenue management system in the Ford Motor Company. The described system consists of three distinct tools: (1) a package optimizer to create bundles of promising car accessories, (2) a price optimizer to generate incentive schemes, and (3) a decision support tool for car dealers that provides information about what cars to order. Additionally to this case study, Specht and Gruß (2007) discuss the requirements of future revenue management systems in the automotive industry. The authors mention that such a system should generate buying alternatives based on the current capacity utilization in the supply chain and acceptable lead times.

Table 4.2: Overview of Publications on aATP in Manufacturing (Method: C=Conceptual, M=Simulation/analytical models, R=Review, S=Case study)

Publication	CODP	Classes	Method
Chen et al. (2001)	MTS	–	M
Fischer (2001)	MTS	several	C,M,R
Chen et al. (2002)	MTS	–	M
Gordon et al. (2002)	MTO	–	R
Ball et al. (2004)	–	–	C,R
Fleischmann et al. (2004)	–	–	C,R
Keskinocak and Tayur (2004)	MTO	–	R
Pibernik (2005)	–	–	C,R
Lee (2006)	MTO/ATO	4 to 8	M,S
Pibernik (2006)	MTS	several	C,M
Pibernik and Yadav (2009)	MTS	2	M
Kilger and Meyr (2008)	–	–	C,R
Meyr (2009)	MTS	several	C,M

4.2.2 Allocated Available-to-Promise

Models of the type aATP allocate capacity/inventory to customers according to their profitability. These models are usually found in manufacturing environments and are often directly applicable in APS to be used in real-world applications. In this section, we complement the review of aATP models by discussing important examples of this type of models.

The most relevant paper in this thesis is the one by Meyr (2009). The model presented here follows a two-step approach: In a first step, quotas are assigned to customer classes (Allocation Planning) which are consumed later on in real-time by the incoming orders (ATP consumption). Due to the importance of this paper, we give a detailed description of the presented model and research approach in the subsequent Section 4.3.

In contrast to the two-step approach of Meyr (2009), batch order promising models use batch intervals to be able to select the most profitable orders. Examples of batch models can be found in the work by Chen et al. (2001), Chen et al. (2002), and Fischer (2001). As shown in these articles, the general drawback of batch models—the long response time to customer requests—is usually handled by first answering the request with an approximative due date and then refining this date in a subsequent step.

Lee (2006) describes a simulation tool developed for IBM's hardware business in order to evaluate the performance of the entire demand management activities—from calculating ATP quantities over scheduling of orders up to due date promising. The author claims that this is one of only a few works considering the entire process, including ATP generation. The considered dimensions of ATP are the type of product, demand classes, supply classes and time buckets. ATP quantities are calculated according to information about component availability, current inventory, work-in-process, the master plan, supplier commitments and production capacity. When a customer order arrives, the model searches through the ATP dimensions according to predefined search rules. The aim of the study is to show the change in key performance indicators when shifting from an MTO to a configure-to-order (CTO) business. The study reveals that customer service increases dramatically in the CTO setting.

Pibernik (2006) discusses order fulfillment strategies in case of stock out situations. This work is one of the few papers taking deterministic future replenishments into account. The author proposes a model to allocate a set of orders to ATP quantities when the costs of short- and long-term consequences of delaying or rejecting an order can be quantified. Similar to the work of Meyr (2009), the author describes a way to pre-allocate ATP quantities to customer classes in order to increase profitability. The author also mentions the *rank-based* allocation strategy as earlier mentioned in Section 2.3.2. It starts allocating all available quantities to the highest customer class. Then, it iterates in descending order through the lower classes allocating the remaining quantities.

Together with the previous described work, the paper of Pibernik and Yadav (2009) is the second paper found considering customer classes and deterministic future replenishments . The authors propose a model that selects among two customer classes under a service-level constraint. The described model in this paper is closely linked to our setting, but reveals some distinctive features. The most prominent distinction is the service-level objective as we maximize the expected profit. Furthermore, the model assumes stochastic due dates, two customer classes and lost sales where we allow more than two classes and both, lost sales and backlogging.

4.2.3 Inventory Rationing

There are two broad research streams within the type of IR models. Ha (1997a) and De Véricourt et al. (2002) propose models with multiple demand classes and stochastic replenishment times, thus assuming limited production (=replenishment) capacities. In contrast, Melchiors et al. (2000) and Arslan et al. (2007) model deterministic replenishment lead times and unbounded replenishment quantities. All of these models take decisions on ordering and rationing

Table 4.3: Overview of Publications on IR in Manufacturing (Method: C=Conceptual, M=Simulation/analytical models, R=Review, S=Case study)

Publication	CODP	Classes	Method
Ha (1997b)	MTS	2	M
Ha (1997a)	MTS	several	M
Kleijn and Dekker (1998)	–	–	R
Melchiors et al. (2000)	MTS	2	M
De Véricourt et al. (2002)	MTS	several	M
Arslan et al. (2007)	MTS	several	M
Gupta and Wang (2007)	MTO/MTS	2	M
Möllering and Thonemann (2007)	MTS	2	M
Defregger and Kuhn (2007)	MT0	several	M
Teunter and Klein Haneveld (2008)	MTS	2	M,R

levels, which are typically expressed in policies like (s, S, R) where s is the reorder point, S the order-up-to level, and R the protection level between customer classes. When inventory falls below s, it is filled up to S. In a two-class setting, demand from the low margin class is fulfilled as long as the inventory is above the protection level R.

Although some of the following work is also linked to revenue management, we classify all models deciding about replenishments as IR models. Defregger and Kuhn (2007) propose a model in an MTO setting using a Markov decision process. The authors consider a decision maker who has to decide whether to accept an order—taking into account the order specific profit and the orders maximum lead time—and whether to raise the inventory level by additional production. In contrast to the conventional approaches in MTO settings, the authors include a finished goods inventory in their model in order to be able to fulfill high-margin orders with a short lead-time. However, it is still an MTO setting since production is triggered by incoming customer orders.

Gupta and Wang (2007) study a setting with two customer classes, distinguished in transactional and more valuable contractual orders. The manufacturer can choose which transactional orders to fulfill but has to meet the demand of contractual orders at a short notice. The authors study two scenarios, one in which the contractual orders are produced-to-stock and one with make-to-order production. In the MTS scenario, the manufacturers decision is twofold: (1) how many units to accept from the transactional orders (contractual have to be fulfilled) and (2) how much to produce to raise finished goods inventory.

Möllering and Thonemann (2007) focus on analyzing optimal backorder clearing in critical level policies. In contrast to models assuming lost sales, backorder models have to decide about which backorders to fulfill under different inventory levels. In the more realistic setting of order sizes greater than one, the authors prove an optimal way of backorder clearing.

The work of Teunter and Klein Haneveld (2008) includes an up-to-date literature review and classification of inventory rationing models. Additionally, the authors propose a dynamic rationing model with two customer classes taking into account the remaining time until the next customer order arrives. The authors mention the easy implementation of their approach as rationing levels are provided in charts and lookup-tables.

Table 4.3 gives an overview about models on IR. Although the mentioned papers in the table are not exhaustive, it represents an overview about current research in this field. The papers in the table suggest that there is a further distinction in research between papers considering two customer classes and papers considering more than two classes. Ha (1997b) explains this when he states "as the number of customer classes increases, the optimal policy will be difficult to compute because of the curse of dimensionality and will be even more difficult to implement". For practical applications, many of these models are tailored to a specific problem and usually not applicable in general settings considered by APS.

4.3 A Selected Deterministic Allocation and Order Promising Model

In this section we discuss the work of Meyr (2009) in which an LP model for demand fulfillment is presented which is the building block for the later developed stochastic models of Chapter 5. The author distinguishes between two different types of models: models with customer segmentation and models without customer segmentation. The models support the demand fulfillment functionality of APS as described in Sect. 2.2. They receive the ATP information from a production planning module and promise due dates to incoming customer orders as shown in Fig. 4.2. In the following, we provide a mathematical formulation of these models. In addition, we propose a classification of search rules for ATP consumption which consume allocated ATP quantities. Due to an aggregated or disaggregated consumption of aATP quantities, we analyze how far these rules are applicable in stochastic environments.

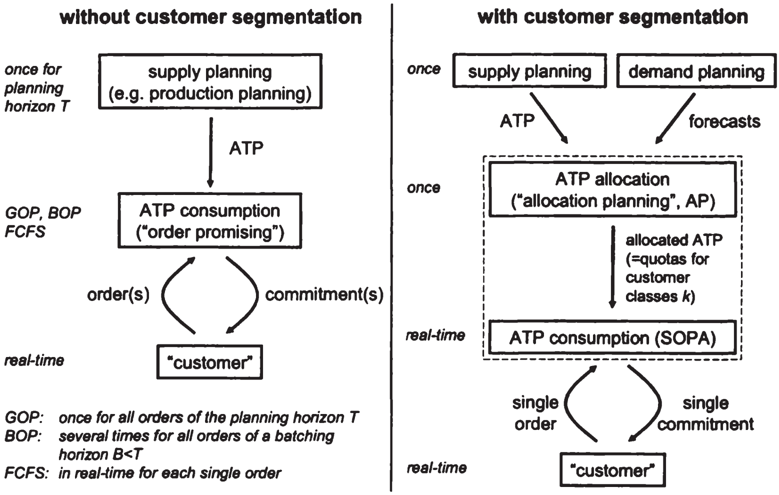

Figure 4.2: Modeling Environment (Meyr, 2009)

4.3.1 Models Without Customer Segmentation

The following basic order promising model without customer segmentation assigns customer orders to ATP quantities according to profitability. In order to model the successive arrival of customer orders the proposed LP formulation has several iteration steps. Meyr (2009) defines a set I^s of currently known orders in iteration s and distinguishes accordingly between three different cases: (1) if the model is executed after each order, i.e., I^s contains only the new order, the model resembles an FCFS policy. One can imagine that this model does not lead to an optimal profit as customers with a low level of profitability might buy units better kept for later arriving high margin customers. (2) If the model is executed after a certain *reasonable* amount of time, i.e., I^s contains all orders arrived in that certain time period, the approach is called *batch order processing* (BOP) referring to a batching horizon B in which the orders are collected to be processed in the next run. (3) If the batching horizon equals the total planning horizon, the model is not executed until all orders are known. This approach results in an optimal solution and is therefore called *global order processing* (GOP). The following equations show the *basic order promising model*:

$$\max \sum_{i \in I^s} \sum_{t=1}^{T+1} p_{it} o_{it}^s \tag{4.1}$$

subject to

$$\sum_{t=1}^{T+1} o_{it}^s = q_i \; \forall i \in I^s \tag{4.2}$$

$$\sum_{i \in I^s} o_{it}^s \leq ATP_t^s \; \forall t = 1, \dots, T. \tag{4.3}$$

The objective function 4.1 maximizes the profit over the planning horizon when satisfying the orders of set I^s. The profit p_{it} includes the revenues p_i generated by the customer i minus optional backlogging costs b_i or holding costs h, i.e., $p_{it} = p_i - h(d_i - t)$ if the required delivery date d_i of the order is after the fulfillment date t and $p_{it} = p_i - b_i(t - d_i)$ in case of a late delivery (and $p_{it} = p_i$ otherwise).

The model includes two constraints. 4.2 guarantees that the total quantity q_i of each order is fulfilled (either from *real* supply or from the infinite supply of period $T + 1$). 4.3 guarantees that, in each period, not more than the total supply ATP_t is assigned to all orders.

In the cases of BOP or FCFS where several iterations are executed, the ATP quantity has to be updated after each iteration in order to cope for already assigned demand quantities. This step corresponds to the previously described *Netting* activity (see Sect. 2.2). Formally, the ATP quantities are updated after iteration s with the following formula: $ATP_t^{s+1} = ATP_t^s - \sum_{i \in I^s} o_{it}^s \; \forall t = 1, \dots, T$. Table 4.4 summarizes the used symbols and notation.

Table 4.4: Notation of the Basic Order Promising Model

Indices:	
$s = 1, \dots, S$	Iterations
$t = 1, \dots, T+1$	Periods of the planning horizon (T+1: dummy period)
$i = 1, \dots, I$	Customer orders
I^s	Set of orders that are promised in iteration s
Decision variables:	
$o_{it}^s \geq 0$	Part of order i which is served by ATP of period t and promised during iteration s
Data:	
ATP_t^s	Not yet assigned supply that becomes available in period t and can still be promised to customers during iteration s
q_i	Requested order quantity of order i
p_{it}	Per unit profit of order i if satisfied by ATP of period t
p_i	Per unit revenue of order i
d_i	Requested delivery date of order i
b_i	Per unit and period backlogging costs of order i
h	Per unit and time holding costs

For the sake of comparability with the ongoing development of stochastic models, we slightly reformulate the basic order promising model. In the current formulation, holding and backlogging costs are part of the profit term p_{it}. This leads to the situation that the costs are only charged when a customer order is fulfilled. In case of backlogging, this seems intuitive, but not for holding costs. Therefore, we introduce a fictitious order ϕ with a delivery date $d_\phi = T$ where all remaining quantities of ATP are assigned to. The profit term $p_{\phi t}$ only includes the holding cost term, i.e., $p_{\phi t} = -h(T-t)$. The *adapted order promising model* can be stated as

$$\max \sum_{i \in I^s \cup \phi} \sum_{t=1}^{T+1} p_{it} o_{it}^s \tag{4.4}$$

subject to

$$\sum_{t=1}^{T+1} o_{it}^s = q_i \ \forall i \in I^s \tag{4.5}$$

$$\sum_{i \in I^s \cup \phi} o_{it}^s = ATP_t^s \ \forall t = 1, \dots, T. \tag{4.6}$$

Table 4.5: Notation of the Network Flow Model

Indices:	
I_t^s	Set of orders i with a desired delivery date in period t
Decision variables:	
$x_t^s \geq 0$	Downstream flow of inventory in period t and iteration s
$y_t^s \geq 0$	Upstream flow of inventory in period t and iteration s
$z_i^s \geq 0$	Lost sales quantity of order i and iteration s
Data:	
b	Per unit and period backlogging costs (independent of order i)

Note that the original and the adapted formulation result in equal profits in case of high scarcity (for which the model was intended). Even if large LP models can be solved with modern LP solvers, calculation time increases rapidly when solving the GOP model. Especially in the rolling horizon simulations of Chapter 7, the problem size gets very large. We propose a network flow formulation of the above adapted order promising model that can be solved for larger instances within a reasonable time. However, this model comes with the disadvantage that the backlogging costs are not order dependent and, therefore, have to be changed to b. Additionally, this model is not able to show from which ATP quantity a certain order is fulfilled.

For the mathematical model formulation, we define a subset $I_t^s \subseteq I^s$ with $I_t^s := \{i \in I^s | d_i = t\}$ as the set all orders $i \in I^s$ with a desired delivery date in period t. The *network flow order promising model* is henceforth given as

$$\max \sum_{i \in I^s} p_i(q_i - z_i^s) - \sum_{t=1}^{T+1} (hx_t^s + by_t^s) \tag{4.7}$$

subject to

$$x_{t+1}^s - y_{t+1}^s - x_t^s + y_t^s = ATP_t^s - \sum_{i \in I_t^s} (q_i - z_i^s) \; \forall t = 1, \ldots, T \tag{4.8}$$

$$z_i^s \leq q_i \; \forall i = 1, \ldots, I^s \tag{4.9}$$

$$x_1^s = 0, y_1^s = 0, y_{T+1}^s = 0. \tag{4.10}$$

The symbols and notation of the network flow problem are shown in Table 4.5.

The objective function 4.7 maximizes the profit resulting from a flow of ATP quantities between the periods towards incoming orders. If the ATP quantities flow downwards (i.e., to later periods) holding costs are charged. If the flow goes to the opposite direction, backlogging costs are charged. The inventory balance constraint 4.8 controls the flow between the periods. Incoming quantities are consisting of the current period's ATP quantities (ATP_t^s), the downstream flow from the previous period (x_t^s) and the upstream flow from the next period (y_{t+1}^s). Outgoing quantities are consisting of the downstream flow to the next period (x_{t+1}^s), the upstream flow to the previous period (y_t^s) and the flow to the orders arriving in this period (minus lost sales) ($\sum_{i \in I_t^s} (q_i - z_i^s)$). Fig. 4.3 illustrates the flow of quantities. The further constraints in 4.10 set the initial quantities of the first and last period. x_1^s can also be used to set the initial inventory.

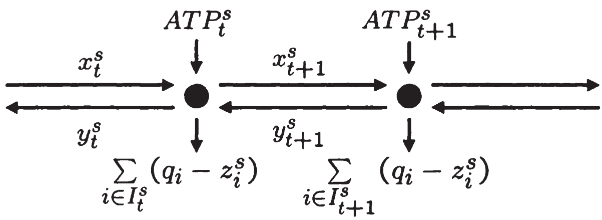

Figure 4.3: Illustration of the Inventory Balance Constraint

4.3.2 Models With Customer Segmentation

The models with customer segmentation try to overcome the disadvantages of the previous models. In case of FCFS, it is not guaranteed that a reasonable solution is reached. BOP models overcome this drawback but result in longer waiting times for the customer (as due dates are not determined until the end of the batch interval is reached). The general structure of models with customer segmentation is shown on the right side in Fig. 4.2. These models usually follow a two-step approach: in a first step (*allocation planning*), quotas are assigned to a given number of customer classes based on ATP quantities and demand forecasts. In a second step (*ATP consumption*), the quotas are consumed in real-time by incoming customer orders.

In the following, we present an approach by Meyr (2009) based on an LP model. Instead of a rule-based approach to determine the quotas (as described in Section 2.3.2), the author proposes a model allocating ATP quantities according to the profitability of customer classes. The determined quotas in the AP step are subsequently called *aATP*. Meyr (2009) uses the acronym *SOPA* for his approach, referring to *s*ingle *o*rder *p*rocessing after *a*llocation planning.

Table 4.6: Notation of the Allocation Planning Model

Decision variables:	
$z_{kt\tau} \geq 0$	Part of demand of priority class k in period τ which is satisfied by ATP of period t
$f_t \geq 0$	Still unallocated part of ATP of period t
Data:	
$d_{k\tau}^{min}$	Lower bound on sales to priority class k in period t
$d_{k\tau}^{max}$	Upper bound on sales to priority class k in period t
$p_{kt\tau}$	Per unit profit if ATP of period t satisfies demand of priority class k in period τ

With the symbols and notation of Table 4.6 we can formulate the *allocation planning model* as

$$\max \sum_{k=1}^{K} \sum_{t=1}^{T+1} \sum_{\tau=1}^{T} p_{kt\tau} z_{kt\tau} \tag{4.11}$$

subject to

$$d_{k\tau}^{min} \leq \sum_{t=1}^{T+1} z_{kt\tau} \leq d_{k\tau}^{max} \;\forall k, \tau = 1, \ldots, T \tag{4.12}$$

$$\sum_{k=1}^{K} \sum_{\tau=1}^{T} z_{kt\tau} + f_t = ATP_t^1 \;\forall t = 1, \ldots, T. \tag{4.13}$$

The profit resulting from allocating ATP quantities to customer classes is maximized in the objective function 4.11. Differently to p_{it} in the basic order promising model, the profit term $p_{kt\tau}$ represents the profit of a customer class and not an individual order. Meyr (2009) suggests to calculate $p_{kt\tau}$ as the average profit of all orders assigned to customer class k. Constraint 4.12 assures that not more than the maximum expected demand is allocated to a customer class and period (or not less than the minimum, respectively). In this constraint, the demand forecasts generated by the demand planning activity are considered in the model. Constraint 4.13 restricts the allocated quantities to the available ATP.

In the following, we show how the quotas or aATP quantities, respectively, can be extracted from the optimal solution of the AP model, denoted by "*". Meyr (2009) proposes two different granularities of aATP quantities. Firstly, a fine granularity in which the aATP quantities are structured according to the customer class k, the availability time of ATP t and the assignment of ATP to a time slot τ. The aATP in this granularity can directly be taken from the optimal values of the decision variables $z_{kt\tau}$ as shown in 4.14.

$$aATP^1_{kt\tau} := z^*_{kt\tau} \;\forall k, t, \tau. \tag{4.14}$$

The second granularity represents an aggregated version of the aATP quantities in which the time slots τ are aggregated to one time period. Equation 4.15 shows how the aggregated aATP quantities are extracted from the optimal values of the decision variables $z_{kt\tau}$. This version especially makes sense when forecast accuracy is low and the allocation of ATP to certain time slots results in poor profits. The proposed aggregation provides a way to compensate forecast errors. However, this aggregation allows customers to consume ATP quantities far away from their desired delivery date.

$$aATP^1_{kt} := \sum_{\tau=1}^{T} z^*_{kt\tau} \;\forall k, t. \tag{4.15}$$

Additionally to the aATP quantities, the AP model might also result in unallocated ATP, e.g., when the maximum expected demand is low in comparison to available ATP. These quantities are henceforth called uATP as shown in 4.16.

$$uATP^1_t := f^*_t \;\forall t. \tag{4.16}$$

The second step of the SOPA approach processes the incoming customer orders in real-time according to the aATP quantities. Let Ξ_i be the set of accessible priority classes from order i. Additionally, let $\hat{\imath}(s)$ denote the current (and only) order processed in the iteration step s. Accordingly, we define the *order processing model* as

$$\max \sum_{k \in \Xi_{\hat{\imath}(s)}} \sum_{t=1}^{T+1} p_{\hat{\imath}(s),t} \overline{o}^s_{kt} + \sum_{t=1}^{T} p_{\hat{\imath}(s),t} x^s_t \tag{4.17}$$

Table 4.7: Notation of the SOPA Model

Indices:	
Ξ_i	Set of priority classes which can be consumed by order i
Decision variables:	
$\overline{o}^s_{kt} \geq 0$	Part of allocated ATP of period t for priority class k which is in iteration s assigned to order $\hat{i}(s)$ showing a desired delivery date $d_{\hat{i}(s)}$
$x^s_t \geq 0$	Part of unallocated ATP of period t which is in iteration s assigned to order $\hat{i}(s)$ showing a desired delivery date $d_{\hat{i}(s)}$
Data:	
$aATP^s_{kt\tau}$	ATP that becomes available in period t and has been allocated to orders in priority class k with a requested delivery date in period τ
$uATP^s_t$	ATP that becomes available in period t but has not yet been allocated to any priority class or planned delivery date
$\hat{i}(s)$	The (single) order considered in iteration s

subject to

$$\sum_{k \in \Xi_{\hat{i}(s)}} \sum_{t=1}^{T+1} \overline{o}^s_{kt} + \sum_{t=1}^{T} x^s_t = q_{\hat{i}(s)} \tag{4.18}$$

$$\overline{o}^s_{kt} \leq aATP^s_{ktd_{\hat{i}(s)}} \quad \forall k \in \Xi_{\hat{i}(s)}, t = 1, \ldots, T \tag{4.19}$$

$$x^s_t \leq uATP^s_t \quad \forall t = 1, \ldots, T. \tag{4.20}$$

Table 4.7 summarizes the symbols and notation of the above model. 4.17 maximizes the profit of the quantities assigned to order $\hat{i}(s)$. 4.18 guarantees that the assigned quantities to order $\hat{i}(s)$ are equal to the demand (if not enough supply is available, the order is fulfilled from the fictitious supply of period $T+1$). 4.19 and 4.20 restrict the consumed quantities of order $\hat{i}(s)$ to the available allocated and unallocated quantities. The SOPA model takes as input a given number of customer classes and corresponding assignments of orders to classes. The assignment of orders to classes has to be done before the actual order arrives, otherwise the SOPA approach will not work. As this seems to be confusing on the first sight, Meyr (2009) argues that not a single order has to be assigned to a class, but the associated customer handing in this order. Long-term relationships between vendors and their customers usually make it possible for such an assignment to be done before the actual order processing starts. All orders i of a customer assigned to class k are then also part of this class. The set Ξ_i of accessible classes of order i includes at least the class of the customer posting the order i, but might be extended when a customer is allowed to consume quantities of lower classes. Meyr (2009) analyzes which approaches perform best.

As the above order processing model solves an LP for each single order that arrives, it seems reasonable for practical implementations to mimic the SOPA search process through the aATP quantities by using simple rules. A description of such rules can be found in the following Section 4.3.3.

A remaining task is still to determine an appropriate number of customer classes k. As the number of classes gets larger, a finer determination of quotas is possible. However, the demand forecast $d_{k\tau}^{max}$ becomes less accurate with increasing k. As this is a medium-term planning task and is usually not part of the short-term demand fulfillment, we will not focus on this trade-off but nevertheless provide some insights to this problem in the numerical studies of Chapter 7.

The benefits of SOPA result from the immediate response to customer requests and the ability to increase profitability by selecting among customer classes. The major drawback of this model is the negligence of demand variability as it assumes demand as deterministic.

4.3.3 Search Rules for ATP Consumption

This section builds on the model described in Sect. 4.3.2, but instead of using an LP for the ATP consumption step, simple rules are proposed. The rules take as input the $aATP$ quantities generated in the Allocation Planning step. The presented rules emulate the order processing step for the disaggregated $aATP$ quantities (cf. Equation 4.14) and for the aggregated quantities (cf. Equation 4.15). Additionally, we show a modified rule adapted to the specifics of stochastic environments.

Disaggregated Single Order Processing

The search rule presented in this section takes as input the disaggregated quantities $aATP_{kt\tau}^1$. When not otherwise stated, we use the notation of Section 4.3, hence $\hat{i}(s)$ denotes the currently processed order. The order is allowed to consume $aATP$ quantities from its own and lower priority classes reserved for the requested delivery date $d_{\hat{i}(s)}$. The search rule and the order of execution is depicted in Fig. 4.4. The rule starts searching in the orders own priority class

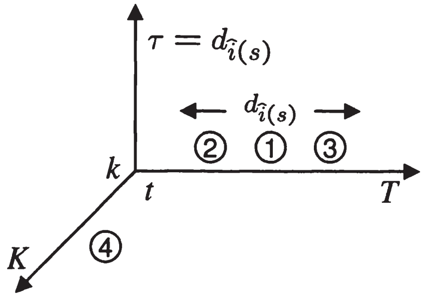

Figure 4.4: Rule-Based Order Processing (Disaggregated)

in quantities reserved for the requested delivery date ($\tau = d_{\hat{i}(s)}$) and moves through the availability times starting at the delivery date ($t = d_{\hat{i}(s)}$) (cf. ①), then before ($t < d_{\hat{i}(s)}$) (cf. ②) and then after the delivery date ($t > d_{\hat{i}(s)}$) (cf. ③). Then it searches through all availability times in lower classes (cf. ④) at the delivery date. If not all requested quantities can be fulfilled, it is searched through the unallocated quantities available at the delivery date of the order, then through quantities available prior to the delivery date and then through the quantities available after the delivery date.

Aggregated Single Order Processing

In contrast to the disaggregated $aATP$, the quantities reserved for a specific period τ are aggregated and can be consumed by all incoming orders. The corresponding aggregated search rule hence merely searches in the dimensions t (availability time of ATP) and k (customer class). Again, consumption of quantities from lower classes is allowed. The search procedure is depicted in Fig. 4.5. In order to prevent high backlogging costs, the search rule starts searching through those quantities available at the requested delivery date of the order ($t = d_{\hat{i}(s)}$) (cf. ①). If there are no quantities available in the orders own priority class, then the search moves to lower classes (cf. ②). If the order is still not completely fulfilled, the search goes to ATP quantities (of the orders own priority class) available before the delivery date ($t < d_{\hat{i}(s)}$) (cf. ③). The search is repeated for $aATP$ quantities reserved for lower classes and available before the delivery date. If there are still quantities requested, it is searched through the $aATP$ quantities available after the delivery date ($t > d_{\hat{i}(s)}$) (cf. ④), first in the orders own priority class and then in lower

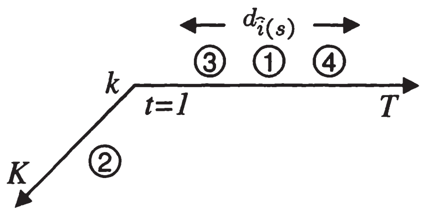

Figure 4.5: Rule-Based Order Processing (Aggregated)

classes. If not all requested quantities can be fulfilled, it is searched through the unallocated quantities as described in the previous section.

A Classification of Search Rules

We have seen two different search rules that mimic the behavior of the order processing model of Sect. 4.3.2. In order to classify different search rules and their performance in stochastic environments, we propose a classification of search rules according to the following set of attributes. All search rules are based on the disaggregated $aATP$ quantities of Equation 4.14.

First, we distinguish search rules according to the attribute *nested aATP*. If the search is nested, consumption of quantities from lower classes is allowed.

The second and third attributes concern the aggregation according to the requested delivery date of the order. As in the case of the search rule described 4.3.3, an order is only allowed to consume quantities reserved for the order's own delivery date. However, as backlogging of orders is usually much more expensive than an early availability of goods (and associated holding costs), we distinguish the aggregation according to quantities reserved for periods after the delivery date of the current order and quantities reserved for periods before the delivery date. These attributes are called *aggregate future aATP* or *aggregate past aATP*, respectively. For instance, in case *aggregate future aATP* is not, but *aggregate past aATP* is allowed, an order is allowed to consume all quantities reserved for periods at and prior to its delivery date. In this way, the expensive backlogging of orders is limited to a minimum. If both attributes are set to true, then the order is allowed to consume both quantities, prior to and after the requested delivery date.

Since the allocation planning step might also result in unallocated ATP quantities ($uATP$), search rules have to define how these quantities are consumed. Since $uATP$ is allocated to neither delivery dates nor customer classes, the only way to differentiate it is the availability time. The attribute *future uATP available* states whether $uATP$ available after the delivery date of the order can be consumed or not.

Table 4.8: Classification of Search Rules

Attribute	SOPA_D	SOPA_A
nested *aATP*	✓	✓
aggregate future *aATP*	–	✓
aggregate past *aATP*	–	✓
future *uATP* available	✓	✓

According to these attributes, we classify two different search rules. First, we denote the search rule of Sect. 4.3.3 as SOPA_D, referring to a disaggregated single order processing after allocation planning. As access to lower classes is allowed, *nested aATP* is set to true. Accordingly, *aggregate future/past aATP* both are set to false, as only the quantities reserved for the delivery date can be consumed. Also, *future uATP available* is set to true as consumption of all unallocated quantities is allowed.

Second, we denote an aggregated version of SOPA as SOPA_A. This rule resembles the search rule described in Sect. 4.3.3, but is slightly different. In contrast to Sect. 4.3.3, SOPA_A takes as input the disaggregated $aATP$ quantities (cf. 4.14), but is allowed to consume all quantities reserved for periods prior and after the requested delivery date. The search proceeds as follows: First, SOPA_A searches through the disaggregated quantities like the rule SOPA_D. Then it moves to quantities reserved for periods prior to the delivery date ($\tau < d_{\hat{i}(s)}$). Afterwards, it searches through quantities reserved for periods after the delivery date. An overview of the discussed search rules is given in Table 4.8.

4.4 Summary

The discussion of selected papers revealed that only very little research exists for the specific setting described in this work. When it comes to stochastic approaches, nothing was found so far. The literature on TRM features many papers with a manufacturing focus, but nearly all models are developed for an MTO setting. There was only one paper considering stochastic demand and a similar setting as ours (Pibernik and Yadav, 2009), but it differs in terms of focusing on a service level objective. The manifold literature on inventory rationing shows that there is much research in this field, but usually the approaches are adapted to very specific settings and not generally applicable. This is—on the other hand—the advantage of aATP models, as they are developed to be used in APS. However, these models are usually kept simple in order to preserve general applicability. Thus, these models still offer many possibilities for further improvements.

Chapter 5

New Demand Management Approaches

In this chapter, we address the identified shortcomings of the TRM, aATP and IR models and show new approaches designed specifically for the characteristics of MTS manufacturing. First, we show an approach based on revenue management ideas and show the optimal demand fulfillment policy. However, for large problem sizes this approach soon reaches its limits since it requires extensive calculation time. In addition, we propose an approximative model based on an adapted version of the SOPA approach which uses ideas from the network revenue management literature, called randomized linear programming (RLP). Due to the high applicability of LP models, randomized linear programming combines a conventional LP with stochastic demand information.

5.1 A Revenue Management Approach

5.1.1 Model formulation

Recall our setting of a make-to-stock manufacturing system facing demand from multiple customer classes. Customer classes differ in the price per unit that they pay. Scheduled inventory replenishments are known. Given this information, the manufacturer decides for each order whether to satisfy it from stock, backorder it at a penalty cost, or reject it. The objective is to maximize profit over a finite planning horizon, taking into account sales revenues, inventory holding costs, and backorder penalties.

In order to achieve this goal the manufacturer has to make a trade-off for each order whether to satisfy it or whether to save the supply for potentially more profitable future orders. We make the following assumptions to model this situation.

Assumption 5.1. *Orders from a given customer class follow a compound Poisson process. The order processes of different classes are mutually independent and they are independent of the available supply.*

Let λ denote the expected number of orders of the Poisson process in a given period, then the probability that exactly k orders arrive can be calculated with

$$f(k, \lambda) = \frac{\lambda^k e^{-\lambda}}{k!}. \tag{5.1}$$

The Poisson assumption is common in many RM models, specifically in so-called dynamic demand models (see Talluri and van Ryzin (2004)). As Lautenbacher and Stidham (1999), we allow non-unit order sizes, which appears appropriate in a manufacturing environment. In our analysis, we discretize the planning horizon in such a way that the probability of receiving multiple orders within a single period is negligible. Let T denote the length of the planning horizon and t the period index. Moreover, let $c = 1, \dots C$ identify the different customer classes.

Assumption 5.2. *Inventory replenishments are exogenous and known.*

This assumption reflects the APS planning hierarchy. Inventory replenishments are determined in mid-term and short-term production planning and then serve as input for order promising decisions. Let x_i be the available supply arriving at the beginning of period i, $i = 1, \dots, T$ and let $\bar{x} = (x_1, \dots, x_T)$ be the vector of all of these replenishments. Note that at time t, x_i corresponds with inventory on-hand if $i \leq t$ and with a future scheduled replenishment otherwise. In APS terminology, $\bar{x}$ denotes the ATP quantities.

Assumption 5.3. *Order due dates are equal to order arrival times, but orders can be backlogged at a price discount.*

This assumption reflects the MTS context. Customers expect immediate delivery, in principle. Late deliveries are only acceptable at a price discount. Let r_c denote the unit revenue from satisfying an order of class c from stock. Delaying an order gives rise to unit backorder costs b per period. Analogously, holding costs h are incurred for all units of inventory on hand at the end of a period. Note that unit backorder and holding costs are independent of time and customer class.

Assumption 5.4. *Partial order fulfillment is allowed.*

This assumption includes splitting an order for partial delivery in different periods. This is a technical assumption, which we need for tractability. We discuss its impact and potential relaxations later on. Let d denote the order quantity, u_i the amount of supply arriving in period i used to satisfy a given customer order, and let $\bar{u} = (u_1, \dots, u_T)$. Note that for an order arriving in period t, u_i corresponds with delivery from stock if $i \leq t$ and with backlogging otherwise.

Table 5.1 summarizes the above notation. We can now formulate our problem as a stochastic dynamic program with state variable $\bar{x}$ and decision variable $\bar{u}$. In principle, one can drop all entries from these vectors, for which $x_i = 0$. I.e., the dimension of the state space corresponds with the number of scheduled replenishments. However, for ease of notation we use $\bar{x}$ and $\bar{u}$ as defined above.

The profit $\hat{P}_t(\bar{x}, d, c, \bar{u})$ earned in period t depends on the available supply, order size, customer class, and fulfillment decision as follows

$$\hat{P}_t(\bar{x}, d, c, \bar{u}) = r_c \sum_{i=1}^{T} u_i - b \sum_{i=1}^{T} u_i (i-t)(1-\delta_{it}) - h \sum_{i=1}^{T} (x_i - u_i)\delta_{it}, \qquad (5.2)$$

where δ_{it} is defined as 1 if $i \leq t$ and 0 otherwise, and $\bar{u}$ has to satisfy $u_i \leq x_i$ for all i and $\sum_i u_i \leq d$.

The first term in Equation 5.2 calculates the revenues received from satisfying the current order of class c. The second term computes backlogging costs that occur when using supply that arrives later than the customer order, i.e. when $\delta_{it} = 0$. These costs are computed for the total length of the delay $(i-t)$. The third term represents holding costs that are charged for the on-hand inventory at the end of period t. Note that unlike the backlogging costs, which are charged for the total customer waiting time, holding costs only cover the current period t. To simplify subsequent calculations, we define $P_t(i, c)$ as the incremental profit per unit of supply i used to satisfy one unit of an order of class c in period t. Collecting the terms in Equation 5.2 that depend on u_i yields

$$P_t(i, c) = r_c - b(i-t)(1-\delta_{it}) + h\delta_{it} \qquad (5.3)$$

and

$$\hat{P}_t(\bar{x}, d, c, \bar{u}) = \sum_{i=1}^{T} P_t(i, c) u_i - h \sum_{i=1}^{T} x_i \delta_{it}. \qquad (5.4)$$

In addition to the current period's profit, we also have to take into account the impact of a fulfillment decision $\bar{u}$ on future profits. The state transition is given by $\bar{x} \rightarrow \bar{x} - \bar{u} = (x_1 - u_1, \ldots, x_T - u_T)$. Letting $V_t(\bar{x})$ denote the maximum

Table 5.1: Notation of the Revenue Management Approach

Indices:	
$t = 1, \dots, T$	Periods of the planning horizon
$i = 1, \dots, T$	Period of inventory replenishment
State variables:	
$\bar{x} = (x_1, \dots, x_T)$	Vector of available supply quantities
Decision variables:	
$\bar{u} = (u_1, \dots, u_T)$	Vector of supply quantities used to fulfill a given order
Random variables:	
c	Customer class
d	Order quantity
$F(c, d)$	Joint CDF of customer class c and order quantity d
Data:	
r_c	Unit revenue from customer class c
b	Unit backorder costs per period
h	Unit holding cost per period

expected profit-to-go from period t to the end of the planning horizon T for a given supply vector $\bar{x}$ we then obtain the following Bellman recursion

$$V_t(\bar{x}) = E_{d,c}\left[\max_{\substack{\bar{u} \\ 0 \le u_i \le x_i, \sum_{i=1}^{T} u_i \le d}} \left\{\sum_{i=1}^{T}(u_i P_t(i,c) - h x_i \delta_{it}) + V_{t+1}(\bar{x} - \bar{u})\right\}\right] \quad (5.5)$$

with the boundary condition $V_{T+1}(\bar{x}) = 0$.

5.1.2 Structural properties and optimal policy

We now analyze structural properties of the value function of the dynamic program defined in the previous section. This will then allow us to characterize the optimal fulfillment policy. All proofs of this section are given in the appendix. We start by defining marginal profits:

Definition 5.1. $\Delta_i V_t(\bar{x}) := V_t(\bar{x}) - V_t(\bar{x} - \bar{e}_i)$ *for* $x_i \ge 1$,

where $\bar{e}_i$ denotes the i-th unit vector. Definition 5.1 concerns the expected marginal value of a unit of supply arriving in period i or, equivalently, the opportunity costs of selling this unit. Using this definition, we can rewrite the Bellman recursion of Equation 5.5 as follows.

$$\begin{aligned}
V_t(\bar{x}) &= E_{d,c}\left[\max_{\substack{\bar{u} \\ 0\le u_i\le x_i,\sum_{i=1}^{T}u_i\le d}}\left\{\sum_{i=1}^{T}(u_iP_t(i,c)-hx_i\delta_{it})+V_{t+1}(\bar{x}-\bar{u})\right\}\right]\\
&= E_{d,c}\left[\max_{\substack{\bar{u} \\ 0\le u_i\le x_i,\sum_{i=1}^{T}u_i\le d}}\left\{\sum_{i=1}^{T}(u_iP_t(i,c)-hx_i\delta_{it})\right.\right.\\
&\quad\left.\left.+V_{t+1}(\bar{x})-\sum_{i=1}^{T}\sum_{z=1}^{u_i}\Delta_iV_{t+1}\left(\bar{x}-\bar{e}_i(z-1)-\sum_{j=1}^{i-1}(\bar{e}_ju_j)\right)\right\}\right]\\
&= V_{t+1}(\bar{x})-h\sum_{i=1}^{t}x_i+E_{d,c}\left[\max_{\substack{\bar{u} \\ 0\le u_i\le x_i,\sum_{i=1}^{T}u_i\le d}}\left\{\sum_{i=1}^{T}\left(\sum_{z=1}^{u_i}(P_t(i,c)\right.\right.\right.\\
&\quad\left.\left.\left.-\Delta_iV_{t+1}(\bar{x}-\bar{e}_i(z-1)-\sum_{j=1}^{i-1}(\bar{e}_ju_j))\right)\right\}\right]. \qquad (5.6)
\end{aligned}$$

Note that this formulation decomposes $\bar{u}$ into single-unit steps. In this way, the maximization in Equation 5.6 reflects the trade-off between the profit of selling a unit of supply now and the corresponding opportunity cost. Also note that a similar decomposition is well-known for the classical single-leg airline yield management problem (see Talluri and van Ryzin (2004), page 59). What is different in Equation 5.6 is the summation over i, which introduces an additional dimension into the problem.

We now identify properties of the value function that help us evaluate the above maximization expression. The first step is to compare the marginal values of supplies arriving in different periods.

Proposition 5.1. *For all $m < n$ and for all $\bar{x}$ with $x_m, x_n \geq 1$ the value function satisfies:*
a) $\Delta_m V_t(\bar{x}) - \Delta_n V_t(\bar{x}) \leq b(\max(n,t) - \max(m,t)) = b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t))$
b) $V_t(\bar{x} + \bar{e}_m) - V_t(\bar{x} + \bar{e}_n) \leq b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t))$.

Proposition 5.1 states that the difference between the marginal value of one unit of a supply arriving in period m and one unit arriving in period n is bounded by the difference in backordering costs of using each of these supplies in period t. This relationship implies the following important monotonicity property, regarding the alternative fulfillment options.

Proposition 5.2. *For all $m < n$ and for all $\bar{x}$ with $x_m, x_n \geq 1$ it holds that* $P_t(m,c) - \Delta_m V_{t+1}(\bar{x}) \geq P_t(n,c) - \Delta_n V_{t+1}(\bar{x}), \forall c.$

The terms on the left-hand-side of the inequality can be interpreted as the net benefit of the current revenues from selling a unit of supply arriving in period m minus the opportunity cost of not having that unit available in the future. Proposition 5.2 states that this net benefit is decreasing in the arrival time of the supply. Therefore, an order should always be either fulfilled using the earliest available supply or not at all (if the left-hand-side becomes negative). The next important property concerns the concavity of the value function along certain axes.

Proposition 5.3. *Let $\bar{x}$ be such that $\sum_i x_i \geq 2$. Furthermore, let $m = \min\{i|x_i > 0\}$ and let $n = m$ if $x_m > 1$ and $n = \min\{i|i > m, x_i > 0\}$ otherwise. Then*
$P_t(m,c) - \Delta_m V_{t+1}(\bar{x}) \geq P_t(n,c) - \Delta_n V_{t+1}(\bar{x} - \bar{e}_m), \forall c.$

Proposition 5.3 implies in particular that the value function V_t is concave in the quantity of the earliest available supply. This is intuitive since one would expect available supply to have decreasing marginal benefits.

The above properties allow us to characterize the optimal fulfillment policy. The optimal policy turns out to be a generalization of the well-known booking-limit policies in traditional revenue management (Talluri and van Ryzin (2004)). We summarize this result in the following theorem.

Theorem 5.1. *Define the following set of critical levels:*
For $i = 1, \ldots, T$ let $\bar{y}_i = \bar{x} - \sum_{k=1}^{i} \bar{e}_k x_k$,
and let $L_t(c,i,\bar{y}_i) = \max\{k|P_t(i,c) < \Delta_i V_{t+1}(\bar{y}_i + k\bar{e}_i)\}$.
Then the following fulfillment decision is optimal in period t, given an order quantity d from customer class c:
Start with $i = 1$;
set $u_i = \max\left(\min\left(x_i - L_t(c,i,\bar{y}_i), d - \sum_{k=1}^{i-1} u_k\right), 0\right)$;
if $u_i < x_i$ set $u_k = 0$ for all $k > i$ and stop, otherwise repeat for i+1.

Intuitively, this policy successively consumes units of supply, in the order of their arrival, until the immediate marginal profit drops below the opportunity costs. This very much resembles a traditional booking-limits policy. The values $L_t(c,i,.)$ set nested protection levels that bar some amount of supply i from consumption by classes c and higher. Note that we have separate protection levels for each class and supply arrival. Also note that the protection levels of supply i depend on the available quantities of subsequent arrivals $\bar{y}_i$. However, $L_t(c,i,.)$ is independent of x_i (and of all earlier arrivals x_j for $j < i$) and therefore indeed acts as a protection level. The amount of supply i exceeding

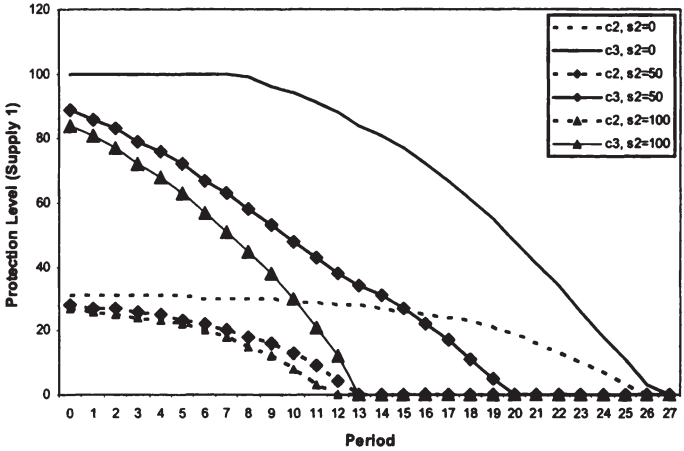

Figure 5.1: Non-Increasing Protection Levels

$L_t(c, i, .)$ is available for consumption by customer class c at time t. It is worth pointing out that even the most valuable customer class, i.e. $c = 1$, may face non-trivial booking limits for future supplies, i.e. for $i > t$: While this class can always consume supply on hand it is not necessarily optimal to backlog demand from this class, due to the incurred backorder penalties. We illustrate the various protection levels $L_t(c, i, .)$ graphically in an example in the next section.

The following proposition shows that, as in the classical case, protection levels decrease in time. This is intuitive since a shorter remaining planning horizon implies less selling opportunities and therefore available supply is of less value. It is worth mentioning, however, that this result is only true for stationary demand. Unlike in traditional revenue management models, the holding cost term in our model may destroy the monotonicity of the protection levels if the demand distribution changes across periods.

Proposition 5.4. *The protection levels $L_t(c, i, \bar{y}_i)$ defined in Theorem 5.1 are non-increasing in t.*

5.1.3 A Numerical Example

Fig. 5.1 illustrates the protection levels of supply one for different customer classes and different levels of the second supply. For example, the line "$c2, s2 = 0$" shows the protection levels for class two and no remaining quantities of the second supply. In the first period, 31 units of supply one are protected from consumption of class two (and, accordingly, class three). As the value of supply

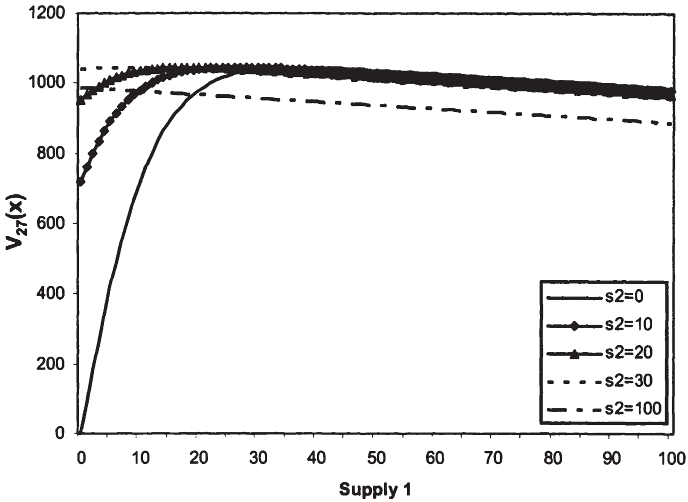

Figure 5.2: Concave Value Function

decreases as times goes on, the protection levels decrease. At the end of the planning horizon (period 27), no supplies have to be protected anymore. As seen in Proposition 5.4, protection levels are non-increasing in time.

Proposition 5.3 states that the value function is concave. This behavior is illustrated in Fig. 5.2 which shows the values of $V_t(x)$ in period 27 for different remaining quantities of the second supply. In case of $s2 = 0$ and small amounts of the first supply, each additional unit of supply one contributes much to the expected profit which explains the huge slope of the line from $s1 = 0$ to 20. After the maximum is reached the slope of the line becomes negative as holding costs are expected that reduce the expected profit. In the extreme case of $s2 = 100$ each additional unit of supply never contributes in a positive way to the profit and the line steadily decreases.

5.2 Approximations Based on Linear Programming

The idea behind the models presented in this section originates from the work on network revenue management, as described in Talluri and van Ryzin (2004). In network revenue management, resources are bundled together to form the products that customers buy. As an example, airlines offer so-called *origin-destination-fares* (ODF), wich often consist of several single flight legs. The complexity of network RM results from the fact that single flight legs might be part of more than one ODF. If two customers want to buy the same flight leg, for example the first customer from Vienna to Chicago over Amsterdam and the second customer from Vienna to Amsterdam, it has to be decided which customer brings more profit.

Many factors contributed to the development of these approaches. First, optimal solutions of realistically-sized networks can not easily be computed as the curse of dimensionality prevents fast computations. Second, as airlines usually have an enormous amount of demand information available, approximating the optimal solution via deterministic approaches would simply mean a waste of resources. Third, methods of linear programming are known to be efficient and are easily to be implemented. As an introduction to the randomized linear programming approach of Section 5.2.2, we first introduce a deterministic linear programming model (DLP) which also serves as an approximation method to the optimal solution but does not take into account demand variability.

5.2.1 Deterministic Linear Programming

The following deterministic linear programming model resembles the allocation planning model described in Sect. 4.3.2, only differing in the demand constraint (5.8). Instead of restricting the allocated quantities to the demand forecast, the *deterministic linear program* uses the expected demand $E\left[d_{k\tau}\right]$ and is formulated as shown in the following.

$$\max \sum_{k=1}^{K} \sum_{t=1}^{T+1} \sum_{\tau=1}^{T} p_{kt\tau} z_{kt\tau} \tag{5.7}$$

subject to

$$d_{k\tau}^{min} \leq \sum_{t=1}^{T+1} z_{kt\tau} \leq E\left[d_{k\tau}\right] \; \forall k, \tau = 1, \ldots, T \tag{5.8}$$

$$\sum_{k=1}^{K} \sum_{\tau=1}^{T} z_{kt\tau} + f_t = ATP_t^1 \; \forall t = 1, \ldots, T. \tag{5.9}$$

Despite the fact that the DLP is very efficient to solve and is easily applicable in practical settings, it is not the final answer for stochastic problems as it has an important disadvantage: the DLP neglects demand variations and simply considers the expected demand. All information included in demand distributions are not taken into account.

To illustrate the weaknesses of the DLP approach, consider a case with extreme supply shortage, two customer classes and a large difference between the profits of class one and class two. The DLP will certainly allocate the highest possible number of units to class one—as this is the most profitable class—which equals the expected demand of this class. The rest will be allocated to the less profitable class two. In the case that the demand of class one exceeds the expected demand, the DLP results in poor profits since not all class one customers are satisfied. The profit loss in not satisfying class one customers is substantial as class two generates only very low profits compared to class one.

This example resembles the situation in the famous newsvendor problem. Here, one has to decide how many newspapers to order on the day before the actual demand is realized. It is assumed that the distribution of the demand is known when the decision about the quantity is made. An intuitive solution to this problem is to order exactly the expected demand as the probability of resulting in too high or too low inventory is minimized. However, this intuitive solution does not always yield the highest profit, only in case the costs of having too much or too less are equal. In a situation in which it is very important to satisfy all customers because costs for loss of goodwill are high, the optimum order quantity will be higher than the expected demand. In the following numerical example, we show that this argumentation is also valid for a simple allocation planning and order promising problem.

Example 5.1. Consider a two-customer class problem with an expected demand per period of class one and two of $E[D] = 5$. Available supply is 10 units per period. For the sake of simplicity, all units that are not consumed are lost at the end of a period. We consider three scenarios of different profitabilities: (1) class one yields a profit of $p_1 = 5$, class two of $p_2 = 1$, (2) $p_1 = 2$ and $p_2 = 1$ and (3) $p_1 = 1$ and $p_2 = 1$. Furthermore, we illustrate the effects of different allocation strategies, from the one extreme making all units available to the first class to the other extreme case of allocating everything to the second class. We do not consider nesting, i.e. class one is not allowed to consume from class two. To calculate profitability of each scenario and allocation strategy, we simulated a Poisson distributed demand stream over 10,000 periods.

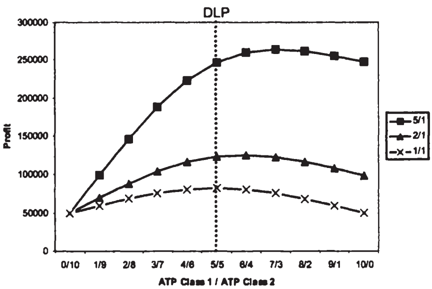

Figure 5.3: DLP and the Optimal Solution

Fig. 5.3 illustrates the effects of the different allocation strategies. In case of equal profitabilities, the best allocation strategy is the one from the DLP approach, which means allocating the expected demand quantities. However, if class one is more profitable as class two, it is beneficial to make more units available to class one. In these cases, the DLP strategy does not perform well as the optimum moves away from the expected demand quantities. These results are in compliance with the previously discussed newsvendor problem.

5.2.2 Randomized Linear Programming

We have seen that the optimal solution may not be reached by simply inserting the expected demand in the AP step. The RLP approach combines the easy to use LP formulation of the DLP, and the available stochastic demand information. The idea is to repetitively solve the DLP, not with the expected, but with a random demand drawn from the known stochastic demand distribution. Let $D^i_{k\tau}$ denote a random variable following a known stochastic distribution of the demand quantity. Then we can formulate the *randomized linear programming problem* as

$$H_t(D^i_{k\tau}) = \max \sum_{k=1}^{K} \sum_{t=1}^{T+1} \sum_{\tau=1}^{T} p_{kt\tau} z^i_{kt\tau} \tag{5.10}$$

subject to

$$d^{min}_{k\tau} \le \sum_{t=1}^{T+1} z^i_{kt\tau} \le D^i_{k\tau} \ \forall k, \tau = 1, \ldots, T \tag{5.11}$$

$$\sum_{k=1}^{K} \sum_{\tau=1}^{T} z^i_{kt\tau} + f_t = ATP^1_t \ \forall t = 1, \ldots, T. \tag{5.12}$$

Note that the resulting optimal solution $H^*_t(D^i_{k\tau})$ is a random variable. The optimal solution of the stochastic problem can be approximated by calculating the expected value $E[H^*_t(D^i_{k\tau})]$. We estimate the allocated quantities $z_{kt\tau}$ by using the weighted average of the resulting quantities over N simulation runs, as shown in 5.13.

$$z_{kt\tau} = \frac{\sum_{i=1}^{N} H^*_t(D^i_{k\tau}) \times z^i_{kt\tau}}{\sum_{i=1}^{N} H^*_t(D^i_{k\tau})} \tag{5.13}$$

In the subsequent numerical analysis of the RLP approach, we choose $N = 30$ which showed robust results in some preliminary tests. This is in line with the chosen parameters of Talluri and van Ryzin (1999, p. 213), although he mentions that this behavior might also be problem dependent.

Numerical studies of RLP applications in the airline industry does not show promising results. It is often the case that the DLP approach dominates the RLP solution. De Boer et al. (2002) analyzed this result in more detail and found an answer to this problem. The authors state that "... the observed domination of the deterministic model over probabilistic techniques is a fortunate by-product of ignoring the uncertainty related to demand. This phenomenon is based on nesting".

As the RLP approach is an alternative version for the allocation planning step, we still have to decide about how to consume the resulting aATP quantities. In Section 4.3.3 we discussed different consumption rules and mentioned that SOPA_D might be fragile to stochastic demand streams. In contrast, SOPA_A aggregates the aATP quantities in order to compensate forecast errors. Thus, we choose SOPA_D as the consumption rule after running the RLP because RLP already considers stochastic demand.

Chapter 6

Simulation Environment

In order to prepare the numerical studies of Chapter 7, we start by introducing the developed simulation environment. We divide this chapter into two parts. First, we discuss the technical settings and issues regarding the implementation, and second, we introduce the parameter set required to run simulations.

6.1 Technical Settings and Implementation Issues

6.1.1 Test Environment

All models and algorithms are part of a single simulation environment implemented in C++ and compiled under Microsoft Visual Studio 2005. The LP models (AP and GOP) are solved using the open source GNU Linear Programming Kit (GLPK) version 4.28. The GLPK library is accessed over a C++ interface in order to run a large number of simulations consecutively. All simulations are executed on a standard PC with a 2.4GHz Intel Core 2 CPU and 512MB memory.

Uniform random numbers are generated using the "Mersenne Twister" algorithm provided by the open source GNU Scientific Library (GSL) version 1.9. In extensive simulations, it was realized that the standard C++ random number generator does not yield satisfying results, so we turned to the Mersenne Twister algorithm described by Matsumoto and Nishimura (1998). To generate random numbers, we use the *inverse transformation method*: when X is a random variable with a cumulative distribution function F, and U is a Uniform distributed random number between 0 and 1, then X can be computed with $X = F^{-1}(U)$.

6.1.2 Implementation Issues

As the implementation usually bears some complications, we illustrate the RM approach of Section 5.1 in pseudo code. In particular, we show how to compute the value function of Equation 5.5 which is required to compute the critical levels $L_t(c, i, \bar{y}_i)$. The value function $V_t(\bar{x})$ has to be computed for all realizations of $\bar{x}$ in all periods t. As $V_t(\bar{x})$ can only be computed when $V_{t+1}(\bar{x})$ is known, the algorithm iterates from the last period T to the first period.

This step is illustrated in Algorithm 1. The algorithm *recursion* as shown in Line 2 contains the iterations through $\bar{x}$ and is explained below.

Algorithm 1: Stochastic Dynamic Program

---

input : The $ATP(t)$ quantities
input : The number of periods T
input : A zero-initialized capacity vector $\bar{x}$
input : A zero-initialized vector $v_t(\bar{x})$ for the value function
output: The vector $v_t(\bar{x})$ filled with the expected profit-to-go for each realization of $\bar{x}$

1 **for** $t \leftarrow T$ **to** 1 **do**
2   `recursion`$(0, t, \bar{x}, v_t(\bar{x}))$;
3 **end**

Algorithm 2 shows a recursive iteration through all realizations of $\bar{x}$. A recursive algorithm is necessary because the number of periods T—and accordingly, possible supply arrivals—might change from one run of the algorithm to another one, and for each period a loop through the supply quantities arriving in this and all other periods is required.

For each specific realization of $\bar{x}$, the expected value $E_{d,c}$ has to be computed. This is done by iterating through all possible demand realizations d and c. As the theoretical limit of d goes to infinity, we have to find an upper limit for d, denoted by D. If the limit D is chosen too high, the computation time is above practical limits. If the limit is chosen too low, the results are not satisfying. In all simulations in this work we chose a limit D such that 99.9% of all possible values d are in the range between 0 and D. D can be computed easily with the *inverse transformation method*: $D = F^{-1}(0.999)$.

In the following, we go through the steps of the algorithm and explain the data variables. In Line 1, the algorithm iterates through the supply quantities arriving in the period *atp_arrival_time*. In Line 2, the current supply quantity of ATP arriving in *atp_arrival_time* is stored in the vector $\bar{x}$. In the next step, the stopping criteria is checked: when the last period is reached, there are no further supplies to be considered. Prior to the last period, a new recursion is started (Line 4) with the next arriving supply. At the time the stopping criteria is met and the "else" sector is reached (Line 5), the vector $\bar{x}$ contains the supply information of all arriving supplies. In the next steps, the expected value is calculated by iterating through all possible demand quantities (Line 6) and customer classes (Line 7).

In Line 8, the policy of Theorem 5.1 is evaluated in the Function *acceptance Rule*. An order is accepted stepwise as long as the profit is larger than the opportunity costs. Note that in this step, any other policy can also be used,

Algorithm 2: recursion(atp_arrival_time, period, $\bar{x}$, $v_t(\bar{x})$)

```
   input : The atp_arrival_time of the current considered ATP
           quantity
   input : The current considered period t
   input : The vector of the current considered capacities x̄
   input : A zero-initialized vector v_t(x̄) for the value function
   output: The vector v_t(x̄) filled with the expected profit-to-go for each
           realization of x̄ in the current period t

 1 for cap ← 0 to ATP_atp_arrival_time do
 2     x_atp_arrival_time = cap;
 3     if atp_arrival_time < T then
 4         recursion(atp_arrival_time + 1, t, x̄, v_t(x̄));
 5     else
 6         for d ← 0 to D do
 7             for c ← 1 to C do
 8                 profit ←acceptanceRule(x̄, d, c);
 9                 expprofit ← expprofit+Prob(d, c) ×profit;
10             end
11         end
12         v_t(x̄) ← expprofit;
13     end
14 end
```

e.g. FCFS or the SOPA approach. Therefore, we do not describe this step in more detail. The function *acceptanceRule* returns the profit associated with the fulfillment decisions of the current order represented by the demand quantity d and class c. For each demand realization (possible values of d and c), the probability $Prob(d, c)$ is multiplied to the profit resulting from the acceptance decision (Line 9). After the iterations through all demand realizations, the profit-to-go for the current considered realization of vector $\bar{x}$ is stored in $v_t(\bar{x})$ (Line 12).

6.2 Simulation Issues

6.2.1 Data Generation

Kimms and Müller-Bungart (2007) present a review on papers dealing with demand data generation with a focus on different assumptions on demand. The authors state that "... assuming demand data that follows a non-homogeneous

Poisson process is more or less standard nowadays". However, the Poisson assumption requires that demand and variance are equal, which might not be true in many applications. For instance, Lawless (1987) states that count data often display extra variation beyond the scope of Poisson distributions. In such cases, the negative Binomial distribution (NB) exhibits certain advantages, as NBs have a long tradition in the marketing and operations literature. It is often mentioned that NBs fit best to observed customer demand (e.g. Ehrenberg, 1959, Agrawal and Smith, 1998).

We distinguish two settings of demand data streams. First, we present a general setting representing demand fulfillment decisions of one year executed on a rolling horizon. Second, we consider a setting which is consistent with the assumptions of the RM approach of Section 5.1. For instance, in the RM setting, only one order is allowed to arrive per period. Table 6.1 displays the options considered in the two settings and shows the used symbols and underlying assumptions.

Table 6.1: Data Stream Options

Description	**Distribution**		**Symbol**
	General setting	RM setting	
Demand data stream			
#Orders per period	$Poisson(\mu)$	*Deterministic*	n_t
Order quantity	*Deterministic*	$1 + NB(\mu - 1, \sigma^2)$	d_i
Revenues per order	$U^c(a, b)$	$U^d \in (r_1, \ldots, r_K)$	r_i
Requested due date	*Deterministic*	*Deterministic*	γ_i
No-arrival probability	–	*Deterministic*	p_0
Supply data stream			
Inter-arrival time	*Deterministic*	*Deterministic*	ϕ
Supply quantity	*Deterministic*	*Deterministic*	s_t

In the general setting, the number of orders n_t arriving in period t is modeled as a Poisson process. The demand quantity per order d_i is set to a fixed value. The revenues r_i of order i are uniformly distributed in the range a to b. As we focus on make-to-stock environments, we assume in the simulation runs that an immediate order confirmation is required (i.e., $\gamma_i = 0$).

Supply data is generated in both settings without any stochastic influences. This is due to our assumption that on the short-term, supply can be considered as given. The supply quantities s arrive in specific intervals, defined by the inter-arrival time ϕ. We model the supply arrival process s_t with $s_t = s$ if $t \bmod \phi = 0$ and $s_t = 0$ otherwise.

The RM setting is different to the general setting in order to account for the different demand assumptions. First of all, the number of orders n_t is deterministic and set to 1 (cf. Assumption 5.1). However, in order to be able to capture the effects of no-arrival probabilities, we introduce p_0 denoting the probability that no customer order arrives in a period. Then, the arrival process can be modeled as such that one customer class arrives per period with an arrival probability of $(\frac{1-p_0}{\#classes})$ for each class (all classes have the same arrival probability). As the customer demand is expected to be discrete and positive, we chose an NB to model the demand quantity. Additionally, it allows us to analyze the effects of large variations in demand. The demand quantity is modeled to be strictly positive with $d_i \sim 1 + NB(\mu - 1, \sigma^2)$. Note that in contrast to the general setting, the demand quantity is always positive (i.e. > 0).

In contrast to the general setting, the revenues of all orders within the same class are equal. The assignment of orders to classes is stochastic and uniformly distributed. Due to Assumption 5.3, orders require an immediate fulfillment, but are willing to accept later deliveries with a price discount. Therefore, $\gamma_i = 0$ holds in all simulations.

To study the impact of supply shortage on the performance, we define *shortage* as a ratio of the total supply throughout the simulation horizon and the total expected demand, more formally stated as

$$sr = 1 - \frac{\sum_{t=1}^{T} s_t}{(1-p_0) \times E[n_t] \times E[d_i] \times T}. \tag{6.1}$$

6.2.2 Simulation Options

The simulation runs of Chapter 7 can be distinguished according to the data stream options of the previous section and simulation options as shown in Table 6.2. We describe the simulation options in the following.

Table 6.2: Simulation Options

Description	**Symbol**
Number of customer classes	K
Simulation horizon	T
Planning window	W
Replanning frequency	F
Backlogging costs	b
Holding costs	h
Forecast error	e

The *number of customer classes* K is considered as a given input and can be changed from one simulation run to the next. By this, we can assess the trade-off between finer customer differentiation and increasing forecast accuracy. In this regard, the assignment of customers to classes is done based on a simple procedure which results in a well-balanced amount of customers in each class. Refer to the work of Meyr (2008) for a more detailed analysis of different clustering methods. The procedure works as follows: first, determine the average number of customers per class by dividing the number of customers by the number of classes. Second, sort all customers according to their profits. Then, assign the most valuable customer to the first class. If the first class contains the average number of customers, move to the next class. Repeat this process until all customers are assigned.

The next three options are used to define a rolling horizon. The *simulation horizon* T covers all periods of a simulation run, e.g. one year. The *planning window* W covers the periods in which the short-term demand management decisions are simulated. In practical settings, the planning window covers a few weeks to month. The *replanning frequency* F determines the amount of time that lies between two consecutive planning windows. Usually, the replanning frequency is shorter than the planning window in order to have overlapping time periods. Figure 6.1 illustrates the three options.

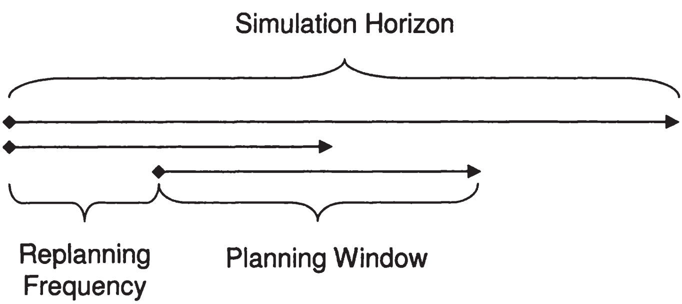

Figure 6.1: Simulation Horizon, Window and Frequency

Backlogging and *holding costs* (b and h, respectively) are not dependent on a specific order and, therefore, are part of the simulation options. We did not implement class- or order-dependent backlogging and holding cost because neither the RM approach (cf. Sect. 5.1) nor the network version of GOP (cf. Sect. 4.3.1) support it.

The last option determines how demand forecasts are generated during a simulation run. We distinguish between forecasts based on the mean demand, and forecasts that are generated according to a predefined *forecast error* e. In the first case, we assume that the mean demand is known (e.g. by exponential smoothing) and is used as a forecast for the demand in SOPA and the RM approach. In case of RM, we additionally assume that the demand variance is known as well. If the mean demand is chosen for generating demand forecasts, we denote this by M. The resulting forecast errors in this case are hence equal to the standard deviation of the demand stream. For example, if the demand is Poisson distributed with $\lambda = 10$, then the standard deviation is 3.33 which is also the mean deviation from the forecasts.

In the second case, we generate forecasts according to the predefined error and the true realization of the demand, denoted by d^*_{kt}. Let d^{max}_{kt} denote the demand forecast for class k in period t and ϵ_{kt} the forecast error, then the forecast error is distributed as

$$\epsilon_{kt} \sim Norm(0, \frac{\sigma_e}{\sqrt{K}}), \tag{6.2}$$

and the demand forecast can be calculated with $d^{max}_{kt} = \max\{0, d^*_{kt} + \epsilon_{kt}\}$. Note that negative demand is changed to 0. σ_e is chosen as a percentage of the mean demand during the planning window. Let m denote the mean demand in the planning window over all classes, then σ_e can be calculated with $\sigma_e = e \times m$.

6.2 shows that the variability of the forecast error within a specific class decreases with an increasing number of classes. However, the variability of the forecast error over all classes $\sum_{k=1}^{K} \epsilon_{kt}$ increases with an increasing number of classes. This behavior resembles common forecasting methods used in practical settings.

For illustration of this rather counterintuitive behavior, consider a mean demand of 1,000 units per period with a standard deviation of 500. If the demand is equally segmented into 1,000 individual classes, a mean of one unit per period per class can be expected. According to Formula 6.2, the standard deviation of the demand per class changes to $\frac{500}{\sqrt{1{,}000}} = 15.81$. This small example shows that the standard deviation of the demand of one class has to be smaller than the standard deviation of the total demand.

6.2.3 Output and Key Performance Indicators

Since the discussed approaches of Chapter 5 focus on profit maximization, we choose the expected profit as the key performance indicator (KPI) in order to compare the different approaches. Expected profits of the RM approach described in Section 5.1 can be directly calculated by solving the value function 5.5. An implementation of the value function is discussed in Algorithm 2. If the assumptions of the RM approach hold, expected profits of SOPA and FCFS can be calculated by means of the value function, as well, just replacing the

RM acceptance rule (Line 8 in Algorithm 2) with the SOPA or FCFS rules. Equation 6.3 shows the value function adapted to FCFS or SOPA.

$$V_t(\bar{x}) = E_{d,c}\left[\underset{0\leq u_i\leq x_i,\sum_{i=1}^{T}u_i\leq d}{\underset{\bar{u}}{\text{FCFS / SOPA}}}\left\{\sum_{i=1}^{T}(u_iP_t(i,c) - hx_i\delta_{it}) + V_{t+1}(\bar{x}-\bar{u})\right\}\right]. \tag{6.3}$$

However, in the simulation runs of Chapter 7 we do not calculate expected profits by means of the value function, but rather run a "reasonable" large number of simulations with different random variates (drawn from the discussed distributions). Subsequently, we calculate the average profit over all simulation runs as an approximation of the real expected profit. The reasons for this are: (1) in some of the scenarios, the demand assumptions of the RM approach do not hold. (2) The GOP solution cannot be calculated by means of the value function, because the Bellman principle of optimality does not hold in this case. The principle states that the optimal decision $\bar{u}$ in period t only depends on the current state $\bar{x}$ and the expected profits in period $t+1$. However, in case of GOP, an optimal decision in period t depends on the actual demand realizations of all periods. (3) In large scenarios, computing the value function is practically not possible due to the curse of dimensionality. In these cases, the approximate expected profit is sufficient to identify trends in the numerical results.

A problem remains how to choose a reasonable number of simulation runs, denoted by Ω in the following. If Ω is chosen too large, the simulation time increases too much. If chosen too low, the approximation of the expected profits is insufficient and the results might be due to random influences. In scenarios with a large variance in the demand, Ω must be chosen high in order to prevent random outcomes. In the simulation runs of Chapter 7, we have manually chosen Ω in the way that a larger Ω would not contribute very much to the trends seen in the different figures.

In order to capture how much of the performance can be attributed to the different approaches directly and not to differences in the simulation settings, we predominantly show profit deviations to GOP in the numerical results, instead of showing absolute profits. These profit deviations, or relative profits respectively, are calculated according to the formula $\sum_{i=1}^{\Omega}\frac{GOP_i^*-S_i^*}{GOP_i^*}$, when S_i^* represents the optimal profit of a certain approach S in simulation run i and GOP_i^* the optimal profit of GOP in simulation run i.

We complement the analysis of expected profits with an analysis of service rates. We distinguish between four different service rates: (1) order quantities fulfilled *in time*, order quantities fulfilled too *early*, order quantities fulfilled too late (*backlogging*), and rejected order quantities (*lost sales*). We choose quantity-oriented service measures (Tempelmeier, 2006) adapted to our setting with simultaneous lost sales and backlogging. Furthermore, since order quantities can be split, we did not focus on orders as a whole but calculated the service measures according to single units. For instance, the amount of *backlogging* covers all units (not orders) with a delivery date in the planning window that are backlogged.

Chapter 7

Numerical Analysis

We have seen in Chapter 2 that demand fulfillment with customer segmentation is able to substantially increase profitability. However, in the literature these approaches have been analyzed only in deterministic settings neglecting stochastic demand. Further approaches have been developed in Chapter 5 to explicitly account for stochastic demand. Our aim in this chapter is to assess both types of demand management approaches under realistic conditions, i.e. to analyze the performance of the described approaches under stochastic influences. Due to technical reasons, we split the analysis in two parts: first, we analyze the different SOPA approaches in a large scenario imitating the demand management process of a company over one year. Second, we show a smaller scenario following the demand assumptions of the RM approach in Section 5.1 to show the relative performance of the approaches. The numerical results underlying the respective figures displayed in the following chapter are shown in Appendix B. Note that the exact shapes of the figures depend on the chosen input data. However, the movements of the curves under parameter changes provide interesting insights independent of the actual scenario settings.

7.1 SOPA in Stochastic Environments

We start with a detailed analysis of SOPA under stochastic demand compared with the optimal solution GOP and a simple FCFS rule. In a preliminary study, it turned out that FCFS made use of the possibility to backlog, without considering the resulting costs. Thus, for reasons of fairness, we assume that future arriving supplies are not taken into account by FCFS. Hence, orders are only fulfilled if inventory is on-hand and therefore extensive backlogging is prevented. For our analysis, we define a base case in the first step and vary the parameters in the subsequent sections. The notation and parameters follow the description of the simulation environment in Chapter 6.

Example 7.1. Consider a problem with uniform distributed revenues $r_i \sim U(90, 110)$. The customer arrival process follows a Poisson distribution with $n_t \sim P(10)$ and a fixed order quantity of $d_i = 12$ units per order. Immediate order fulfillment is assumed (i.e. $\gamma_i = 0$). The supply inter-arrival time is $\phi = 14$ periods starting in period 1, each supply with a quantity of $s = 1,000$

units. The simulation horizon is $T = 365$ periods, the planning window $W = 28$ periods and the replanning frequency is $F = 7$ periods. Per unit and period backlogging costs are $b = 10$ and holding costs $h = 1$. If not otherwise stated, we consider three customer classes ($K = 3$) and assume that the mean demand per class is known ($e = M$).

In Example 7.1, we show a demand management process executed over one year based on a rolling horizon. The parameters are chosen to mimic a realistic demand management process for a single product. In case of backlogging costs, it is not easy to find realistic assumptions. Anderson et al. (2006) show a method for determining stockout costs, which they apply in an empirical study in a retail environment. Depending on whether short- or long-term effects are included, they report stockout costs of 8.76$ and 22.69$, respectively, for a product selling for 51.06$. This corresponds to costs of 17.15% and 44.43% of the product's price. According to these findings, we decide to give roughly 10% discount ($b = 10$ and $E[r_i] = 100$) in each period a product is not immediately available (backlogging). The shortage ratio in this example is $sr = 1 - \frac{27 \times 1{,}000}{10 \times 12 \times 365} = 0.3836$ (cf. Equation 6.1), i.e., it can be expected that 38.36% of all orders are not satisfied.

7.1.1 Base Case Analysis

We start by applying FCFS, GOP, and the two versions of SOPA to the base case. The base case simulated average profit of SOPA_D equals 2,543,725, and the average profit of SOPA_A equals 1,774,840. Table 7.1 compares this value to the average profits of the two benchmark approaches GOP and FCFS. We observe that SOPA_D outperforms both FCFS and SOPA_A. However, the difference to FCFS is marginal. The profits of SOPA_A indicate that the aggregation of ATP quantities leads to poor results. We will analyze the reasons for the poor results of SOPA_A below. The second row in the table shows the relative performance of the approaches in relation to the optimal solution GOP. Note that the relative profits can not be obtained from the absolute values shown in the table, because relative profits are calculated as an average over all single simulation runs.

Table 7.1: Base Case Average Profits for Different Approaches

	GOP	FCFS	SOPA_D	SOPA_A
Absolute	2,617,006	2,515,394	2,543,725	1,774,840
Relative		3.88%	2.80%	32.18%

Figure 7.1 helps to explain the observed performance differences. The figure specifies the service levels achieved by each approach. Specifically, it shows for each policy the average fraction of orders lost and backordered of each customer class. Due to the large size of the scenario, the solution of GOP was calculated with the network flow LP (cf. 4.7 described in Section 4.3.1) which is not able to calculate class-specific service rates.

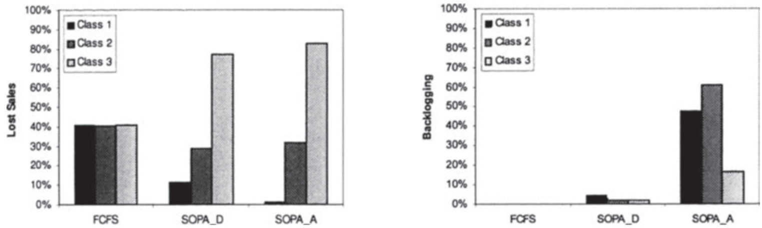

(a) Lost Sales (b) Backlogged Quantities

Figure 7.1: Base Case Service Rates I

Figure 7.1 shows that FCFS rejects customers independent of their class and results in equal amounts of lost sales (Figure 7.1 (a)). The two versions of SOPA mainly differ in terms of how they treat the most profitable class. The aggregated SOPA rejects nearly no quantities of this customer class, but this behavior does not explain the large differences in profits.

Figure 7.1 (b) displays the amount of backlogged quantities. As said before, in FCFS backlogged quantities are zero by definition because backlogging is not allowed. Similarly, SOPA_D very rarely makes use of backlogging, basically in the highest and most profitable class. In contrast, SOPA_A allows consumption of quantities reserved for all periods and therefore results in extensive backlogging. Due to the high backlogging costs, SOPA_A performs poor in terms of profitability.

In Figure 7.2, the service rates for too early and timely deliveries are displayed. As expected, SOPA_D has high rates for the on-time delivery in all customer classes. Regarding early deliveries, due to the holding costs service rates differ considerably between the highest and the lowest customer class. SOPA_A, however, produces high service rates regarding on-time deliveries in the highest class, but—due to nesting—low rates for the other two classes. Considering early deliveries, SOPA_A has low rates, especially for the two lowest classes. This effect goes in hand with the high amount of backlogging, which leaves no space for early deliveries. The results of FCFS are intuitive, since backlogging is not allowed and therefore early deliveries have the highest

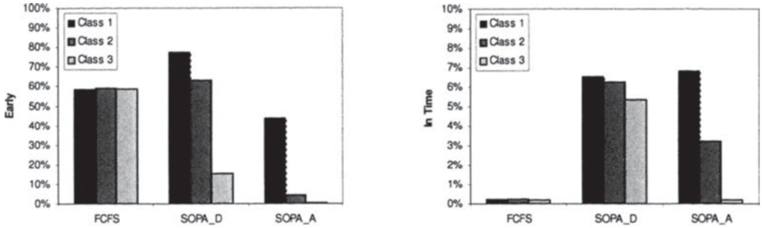

(a) Early Delivery (b) On-Time Delivery

Figure 7.2: Base Case Service Rates II

probability to occur if supply quantities are available. In the following, we change parameters of the base case and analyze the resulting effects. For the sake of comparability, we include the results of the base case in the following figures. The base case is always denoted by *.

7.1.2 Impact of Customer Classes

In this section, we analyze the behavior of the profits when the number of customer classes increases. It is expected that with an increasing number of classes, the SOPA results improve due to a finer allocation, but on the other hand are negatively affected by decreasing forecast accuracy. Figure 7.3 shows the results of FCFS, GOP, SOPA_A, and SOPA_D. Note that * indicates the base case.

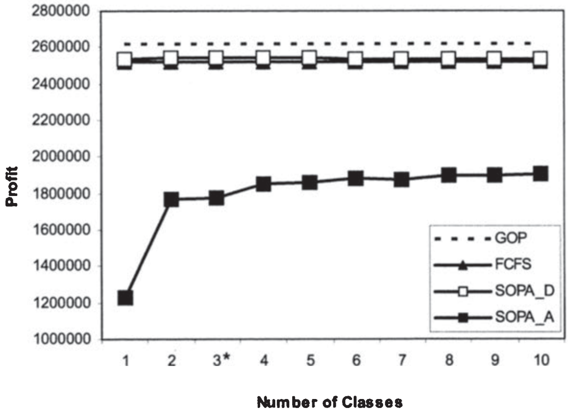

Figure 7.3: Average Profits with Varying Number of Classes

First, note that FCFS and GOP do not consider customer segmentation and, hence, are not influenced by the increasing number of classes. The figure shows that the aggregated version of SOPA is far away from the profits of FCFS, GOP and SOPA_D, which are located close together. Nevertheless, SOPA_A benefits from an increasing number of classes, especially regarding the step from one to two classes. The expected trade-off between a finer segmentation and an increasing forecast error can be seen in the case of SOPA_D where the profits first slightly increase up to two or three customer classes and then decrease again. This effect is not visible from the graph, but can better be seen in the data Table B.3 in the Appendix. Despite the small dimensions of the phenomenon, it has proven robust to parameter changes.

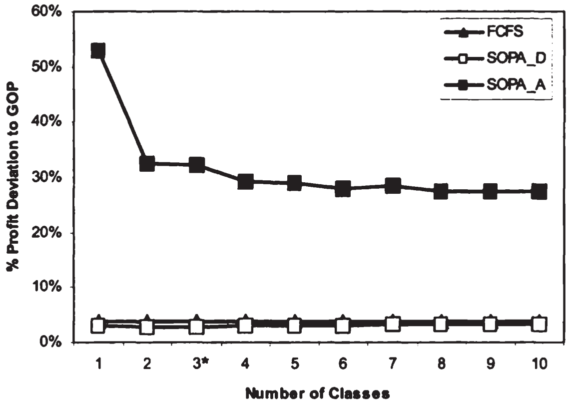

Figure 7.4: Average Relative Profits with Varying Number of Classes

Figure 7.4 displays the relative performance of the approaches compared to the optimal solution (GOP) with the same scenario as used in Figure 7.3. We see that SOPA_D and FCFS are roughly 3-4% worse than the optimal solution, whereas SOPA_A produces poor results. These are due to the high amount of backlogging which is not prevented by SOPA_A. Since FCFS seems to be highly competitive in this scenario, we will analyze further settings in the subsequent chapters.

7.1.3 Impact of Customer Heterogeneity

We have seen in the base case that FCFS is competitive to the SOPA approaches. In the following, we analyze the effect of an increased customer heterogeneity regarding profitability. We expect that higher variation in customer profitability positively influences the performance of approaches with customer segmentation. We vary the heterogeneity by changing the revenues r_i. In the base case, we had a 10% deviation from the mean, upwards and downwards ($r_i \sim U(90, 110)$). Additionally, we now consider cases with 30%, 50%, and 70% deviation.

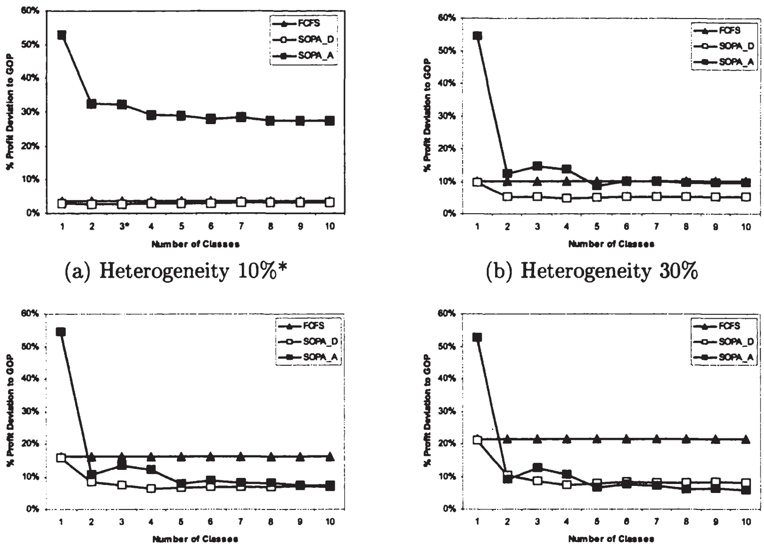

(a) Heterogeneity 10%*

(b) Heterogeneity 30%

(c) Heterogeneity 50%

(d) Heterogeneity 70%

Figure 7.5: Variation of Customer Heterogeneity

Figure 7.5 displays the three scenarios with varied customer heterogeneity in addition to the base case with 10% deviation (denoted by *). Again, the figure shows the deviation from the optimal solution. As expected, FCFS is not able to manage increased heterogeneity. The results get worse from (a) (with approximately 3% profit deviation to GOP) to (d) (approximately 21% profit deviation to GOP). In contrast, both SOPA approaches benefit a lot—especially around two to five customer classes. The jagged appearance of SOPA_A is due to the stochastic input data and shows that it is less robust than SOPA_D.

The trade-off between forecast accuracy and a finer segmentation can better be seen now in the curve of SOPA_D. In the case of four customer classes and high heterogeneity, SOPA_D gets closest to GOP. Interestingly, SOPA_A approaches the profits of SOPA_D in case of very high heterogeneities.

7.1.4 Impact of Forecast Errors

In this section, we analyze the influence of different forecast errors and their effect on the expected profit. In the base case, (again denoted by * in Figure 7.6), we assume that the mean demand in a customer class is known and can be used as a forecast. In addition to the base case, we analyze a setting with no forecast errors 7.6 (a), i.e., deterministic forecasts, and settings with 50% and 100% forecast errors. Note that we consciously do not want to simulate a forecast procedure, but rather choose to directly calculate forecast errors. In this respect, our analysis is independent of the chosen forecast procedure (and its resulting forecast error), which allows us to analyze a much larger spectrum of different forecast errors. The calculation of forecast errors is explained in Equation 6.2.

The setting M in the upper right corner results in forecast errors which are determined by the standard deviation of the demand data stream. In our Example 7.1, we have a Poisson distributed arrival process of customers with $\lambda = 10$ per period with a fixed quantity of 12 units each. Hence, on average 120 units are ordered per period with a standard deviation of $37.95 = \sqrt{10} \times 12$. I.e., if we use the mean demand of 120 units as forecast, we result in errors of 37.95 units on average, approximately 32%. Therefore, in Figure 7.6 we arrange the setting M in between the scenario with 0% forecast errors and 50% forecast errors.

Figure 7.6 shows the behavior of SOPA with different forecast errors. Since FCFS does not require demand forecasts, its profits are not affected by increasing forecast errors and are, hence, equal throughout all figures. The curve of SOPA_D in Part 7.6 (a) approaches the GOP solution with an increasing number of customer classes. This is intuitive since when the number of classes equals the number of orders, the AP step optimally allocates ATP quantities to each order. In case of only one class, low-margin customers might consume quantities from the single class that are better kept for later high-margin customers. Therefore, SOPA_D cannot yield optimal profits in case of few customer classes.

With increasing forecast errors, SOPA_D becomes worse and is soon outperformed by FCFS. In the cases of 50% and 100% error, SOPA_D does not benefit from a larger number of classes since the imprecision of the forecasts of each class increases, which overcompensates the finer customer segmentation.

The behavior of SOPA_A is rather different to SOPA_D. Due to the aggregation of quantities, SOPA_A even benefits from the forecast errors and the performance improves with an increasing number of classes. Nevertheless, the results of SOPA_A stay in the range between 10% and 20% deviation from the optimum and at around 10% deviation from FCFS. In the the case of very high

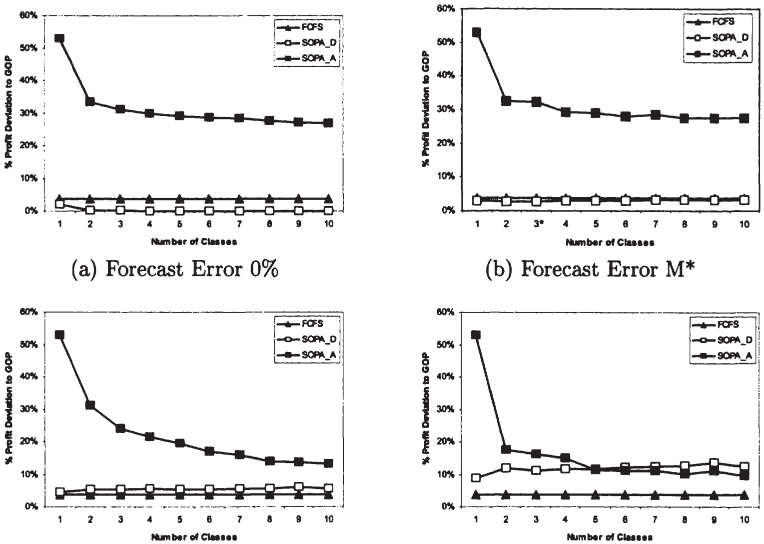

(a) Forecast Error 0%

(b) Forecast Error M*

(c) Forecast Error 50%

(d) Forecast Error 100%

Figure 7.6: Variation of Forecast Errors

forecast errors, SOPA_D is in a similar range as SOPA_A. In this respect, we conclude that in case of high errors, the allocation of ATP to customer classes in the AP step is rather a random procedure. The consumption rules based on the allocated ATP quantities reveal a random outcome a long way from an optimal solution.

7.1.5 Impact of Backlogging Costs

We have seen already in the previous sections that profits are considerably influenced by the amount of backlogged quantities. In the following, we analyze the effects of different backlogging costs. Figure 7.7 displays the behavior of FCFS, SOPA_A, and SOPA_D when backlogging costs are increased, starting from zero up to ten. Note that the last point represents the base case.

Some interesting effects are visible: First, GOP is only slightly influenced by increasing backlogging costs confirming that the optimal allocation of ATP to orders works fine. SOPA_D behaves similarly to GOP which indicates that the allocation of ATP to specific due dates works well regardless of the backlogging costs. FCFS is not influenced at all since we forbid backlogging

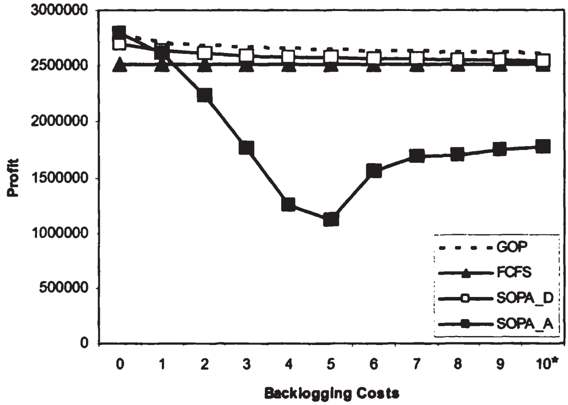

Figure 7.7: Variation of Backlogging Costs

in this case. In the case of low backlogging costs, FCFS is outperformed by GOP and SOPA_D, but both approaches approximate the curve of FCFS with further increasing backlogging costs.

SOPA_A exhibits interesting behavior. In case of no backlogging costs, the results of SOPA_A are equal to the results of GOP. This result can be explained by the fact that the aggregation compensates all forecast errors, whereas late fulfillment of orders does not influence the results due to the zero backlogging costs. Early availability of ATP actually results in holding costs, but due to the assumed scarce capacities this does not play a role in these scenarios as ATP certainly never becomes available before the actual due date. With increasing backlogging costs the profit of SOPA_A decreases drastically. However, this behavior is reversed at backlogging costs of five. We explain this behavior by the interplay of errors and backlogging costs: in the range of low backlogging costs, wrong allocations and forecast errors do not play a pivotal role as they are compensated by the aggregation of ATP quantities. With medium backlogging costs, backlogged customer orders decrease the profit. In case of high backlogging costs, the AP step allocates ATP more carefully to the different classes which results in higher profits.

7.2 Analysis of the Revenue Management Approach

In the following, we numerically analyze the revenue management approach of Section 5.1. As noted earlier, we choose a scenario following the demand and supply assumptions of Section 5.1.1. For reasons of practicability, the scenario is much smaller than in the previous analysis. This is due to the large computation time of the stochastic dynamic program. The following Example 7.2 shows the base case scenario.

Example 7.2. Consider a three-customer class problem with revenues $p_1 = 100$, $p_2 = 90$, and $p_3 = 80$. Let $h = 1$ and $b = 10$. The simulation and planning horizon is $T = 28$ periods with two receipts of supply of 100 units each, the first in the start period $t = 1$ and the second in period $t = 15$. We assume exactly one customer arrival per period ($p_0 = 0$) with the demand quantity following a negative Binomial distribution with mean $\mu = 12$ and a standard deviation of $\sigma = 8$.

In the next subsection, we evaluate our RM procedure for this base case and compare it to the FCFS, SOPA_D, and GOP benchmarks. In the subsequent sections, we investigate the impact of several key model parameters on the relative performance of RM, thereby identifying conditions that are particularly conductive to the use of revenue management in MTS demand management. Specifically, we address the impact of demand variability (Section 7.2.2), customer heterogeneity (Section 7.2.3), and supply shortage (Section 7.2.4).

7.2.1 Base Case Analysis

The base case simulated average profit of the RM policy equals 17,635.90. Table 7.2 compares this value to the average profits of the three benchmark policies. We observe that RM outperforms both FCFS and SOPA_D by about 2%. In terms of relative profits to GOP, FCFS looses more than 3% to GOP, SOPA_D roughly 3%, and RM around 1%.

Table 7.2: Base Case Average Profits for Different Approaches

	GOP	FCFS	SOPA_D	RM
Absolute	17,843.03	17,247.27	17,327.95	17,635.90
Relative		3.33%	2.91%	1.16%

Figure 7.8 helps to explain the observed performance differences by displaying the service levels graphically. Again, we measure customer service in terms of lost sales (Figure 7.8 (a)) and backlogging of orders (Figure 7.8 (b)).

A number of differences between the policies stand out. First of all, FCFS does not differentiate between customer classes, which is reflected in roughly uniform service levels across classes. All other policies clearly prioritize high-value orders. While FCFS achieves the lowest total number of lost sales, it loses relatively many high-value orders. Comparing RM and SOPA_D, we see that SOPA_D rejects even more orders of classes two and three. However, the resulting decrease in lost sales of class one—relative to RM—is insufficient to compensate the lost revenues. It appears that in the face of demand uncertainty, the prioritization of the SOPA_D rule is slightly too aggressive. Another difference between RM and SOPA_D concerns the backordering behavior. SOPA_D backorders many more orders from lower customer classes. Under demand uncertainty, it appears preferable to reserve more future supply for potential future orders, thereby avoiding backorder penalties.

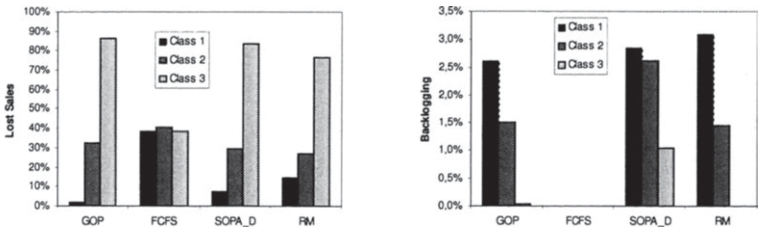

(a) Lost Sales (b) Backlogged Quantities

Figure 7.8: Customer Service Levels in the Base Case

7.2.2 Impact of Demand Variability

We now investigate how the above results depend on a number of key parameters. We start by addressing demand variability. To this end, we vary the standard deviation of the order size $\sigma$from 4 to 16. In addition, we consider the case of a constant order size of $\mu = 12$ units. This range corresponds with a CV of the order size between 0 and 1.33. The scenario $CV = 0.67$ / $\sigma = 8$ corresponds with the base case analyzed in detail in the previous section (again denoted by *). Note that in addition to the order size, we also have uncertainty in the order arrivals. Therefore, even the scenario with a constant order size is not entirely deterministic.

Figure 7.9 shows the expected profits for different CVs of GOP, RM, FCFS, and SOPA using the setting of example 7.2. We observe the following: (1) The negative effect of increasing variance is maximal if using the purely deterministic SOPA approach whereas (2) FCFS seems unaffected from the increase at least in case of low variability. (3) In the case of deterministic demand quantities ($CV = 0$), SOPA is competitive to RM. (4) Overall, the RM approach performs best, but the consideration of stochastic demand does not prevent its profit from decreasing with increasing demand variability.

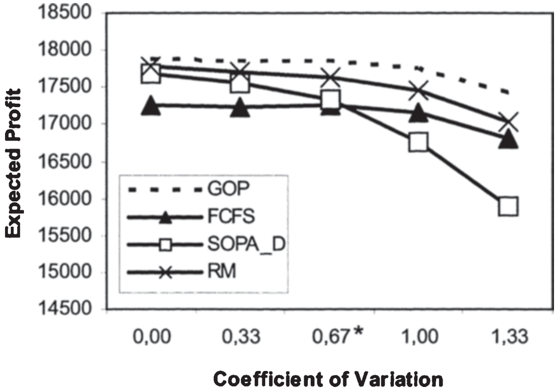

Figure 7.9: Average Profits for Different Levels of Demand Variability

To summarize, we conclude that the effectiveness of customer differentiation in demand fulfillment decreases with increasing demand variability and that forecasting errors should be taken into account in the fulfillment decision.

Note that some of the variations in the average profits can be explained by changes in the maximum attainable ex-post profit GOP. To eliminate this effect, we consider again the relative deviation of FCFS, SOPA_D, and RM with respect to GOP in the subsequent analyses.

7.2.3 Impact of Customer Heterogeneity

In Figure 7.10 we show how increasing customer heterogeneity changes the relative performance of FCFS, RM and SOPA_D under increasing demand variability. In the low heterogeneity case (7.10 (a)), RM and SOPA_D hardly outperform FCFS. In particular, SOPA_D is only competitive in cases with low variability. In the cases with medium and high customer heterogeneity (Figures 7.10 (b) and 7.10 (c)), SOPA_D and RM in general perform better than FCFS. The relative performance of RM and SOPA_D is only weakly influenced by customer heterogeneity under low demand variability (less than 0.33), but we find a strong influence under high demand variability. As expected, the difference between SOPA_D and RM rises with increasing values of heterogeneity and demand variability. Figure 7.10 (c) reveals that the deterministic SOPA_D is not able to benefit from increasing heterogeneity when the variability is high. This is in line with the previous findings. Another interesting result is that the relative performance of FCFS, though not influenced by increasing variability, is strongly affected by increasing heterogeneity. The latter is intuitive since FCFS treats all customers equally and does not care for the larger profit potential. Summarizing the results we can say that RM outperforms FCFS and SOPA_D, especially under high variability. In case of low variability, SOPA_D seems to be competitive to RM. FCFS as the simplest rule is only appropriate if heterogeneity is low. The point in which SOPA_D and FCFS approximately perform equally moves towards the upper right corner with increasing heterogeneity (from 7.10 (a) to 7.10 (c)).

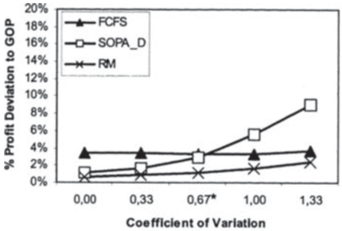

(a) Heterogeneity 10%

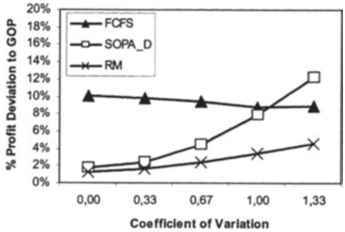

(b) Heterogeneity 20%

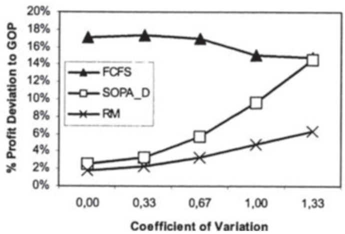

(c) Heterogeneity 30%

Figure 7.10: Varied Customer Heterogeneity for Different Levels of Demand Variability

7.2.4 Impact of Supply Shortage

Figure 7.11 displays the relative performance under different no-arrival probabilities (and—ceteris paribus—different shortage rates). Three different cases have been picked to analyze shortage: We start from a setting where FCFS and SOPA_D have an approximate equal profit ($CV = 0.67$). The influence of different shortage rates in this situation is depicted in Figure 7.11 (a). Contrasting to this, Figure 7.11 (b) shows the setting with lower demand variability ($CV = 0.33$) and 7.11 (c) shows the case with higher customer heterogeneity.

The left-hand side of each figure shows the case of supply shortage (shortage rates > 0) whereas the right side depicts the case of excess supply (shortage rates < 0). We can see in all cases that FCFS—in contrast to RM and SOPA_D—declines with decreasing shortages, regardless of whether we assume under- or oversupply. This is intuitive since the benefit of selecting among the most profitable orders is continuously decreasing with decreasing shortages. Regarding RM and SOPA_D, they also approximate the GOP value in the case of excess shortage. In the extreme case of 138% excess supply, the benefit of rationing vanishes since all orders can be satisfied so that the differences between the different approaches vanish. In the other extreme (large shortages), SOPA_D and RM benefit from their rationing strategy. While the performance of FCFS diminishes with increasing shortage, SOPA_D and RM reach higher profits when the shortage rate shifts from 21% to 41%. In Figures 7.11 (b) and 7.11 (c), both approaches outperform FCFS. Comparing the shapes of the SOPA_D and RM curves, it can be seen that in all scenarios SOPA_D is more dependent on the shortage rate, as it can only compete with FCFS in cases of very high or very low shortage. For the RM approach, we can see that it performs better than both, SOPA_D and FCFS in all cases, even in those with oversupply.

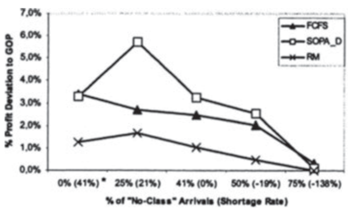

(a) $CV = 0.67$ and Low Heterogeneity

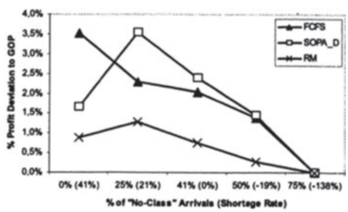

(b) $CV = 0.33$ and Low Heterogeneity

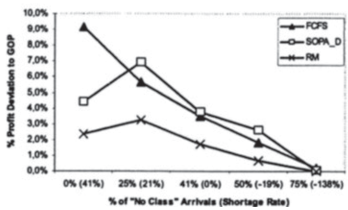

(c) $CV = 0.67$ and Medium Heterogeneity

Figure 7.11: Impact of Shortage Rate on Average Profit Deviation from GOP for Different Scenarios

7.3 Analysis of Randomized Linear Programming

The RLP approach presented in Section 5.2.2 aims at combining the general applicability of SOPA with consideration of stochastic demand. Therefore, now we focus on analyzing the effects of demand variability on the performance of RLP. We start with the base case as described in Example 7.1 and run the RLP approach with 30 iterations. In each iteration, the RLP approach draws random variates from the known demand distribution and uses them as forecasts in the allocation planning step. The second step (order promising) is done in the same way as in the SOPA_D approach.

In order to capture the effects of high demand variability, we complement the base case scenario (Poisson arrival process with a mean and variance of 10) with scenarios that follow a negative Binomial distributed process with the same mean ($\mu = 10$) but higher variances ($\sigma^2 = 64$ and 144). Furthermore, we consider a scenario with a fixed number of customers per day. As in the base case, we assume in all settings that the mean demand is known. In case of the RLP approach, we additionally assume that the variance of the demand distribution is known as well.

Given these parameters, we can calculate the resulting forecast errors. In Section 7.1.4, we result in a forecast error of 32% in the base case scenario if we assume that the mean demand is known. Accordingly, the forecast error with variance 64 is $80\% = \frac{\sqrt{64} \times 12}{120} \times 100$ and $120\% = \frac{\sqrt{144} \times 12}{120} \times 100$ with variance 144. In case of a fixed number of customers per day, we result in no forecast errors since the known mean demand is equal to the real demand. All these calculations are based on the setting of Example 7.1 in which we assume on average 10 customers per day with a fixed order quantity of 12 units (= 120 units total demand per period).

Analog to Section 7.1.3 we consider different customer heterogeneities of 10%, 30%, 50%, and 70%. The results can be seen in Figure 7.12. From the upper left hand side to the lower right hand side the variability of the demand increases. Correspondingly, the forecast errors increase as well and, hence, the AP step results in less reliable allocations. Figure 7.12 (a) shows the setting in which we have a fixed number of customers per period and, therefore, no variability. SOPA_D and RLP result in the same profits because RLP draws random variates from a distribution with no variance and therefore uses always the mean.

Like in the analysis before, FCFS is unaffected by demand variability and its curve does not seem to move from Part 7.12 (a) to 7.12 (d). This behavior is intuitive since FCFS fulfills incoming demand irrespective on fluctuations.

The profits of SOPA_D are affected negatively by increasing variability, but benefit from increasing heterogeneity. In contrast, SOPA_A seems to bene-

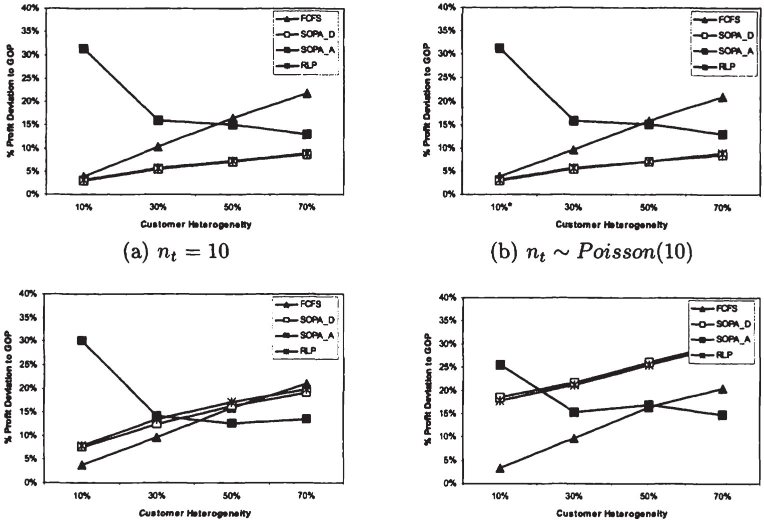

(a) $n_t = 10$ (b) $n_t \sim Poisson(10)$

(c) $n_t \sim NB(10, 64)$ (d) $n_t \sim NB(10, 144)$

Figure 7.12: RLP with Different Forecast Errors and Varied Customer Heterogeneity

fit from increasing variability in case of low heterogeneity and high demand variability. The profits of SOPA_A move from nearly 33% deviation from the optimal solution GOP with no variability to nearly 26% with 144 demand variability, while in case of high heterogeneities (70%), there is a slight deterioration from 12% to 15% deviation from GOP.

In general the RLP approach performs equally well as SOPA_D. A reason for the strong resemblance between RLP and SOPA_D is the same way of aATP consumption. Astonishingly, the more complex RLP approach brings no real benefit in contrast to SOPA_D. The consideration of the variance in the RLP approach seemingly is no suitable way to cope with stochastic demand. RLP even results in slightly less profits in case of $\sigma^2 = 64$ demand variance. However, in the last scenario with the highest considered variability ($\sigma^2 = 144$), RLP performs slightly better than SOPA_D, but on a low level which is even far away from the FCFS solution.

The results are in line with the previously mentioned findings of De Boer et al. (2002). The authors claim that the stochastic nature of the demand is already compensated by the "nested" protection levels and cannot further be improved by the RLP approach. In order to check this argument, we run additional simulations that forbid the nested consumption of ATP quantities. The setting of Figure 7.13 is in all terms equal to the one of Figure 7.12, but nested consumption is not allowed in the RLP approach and the two versions of SOPA.

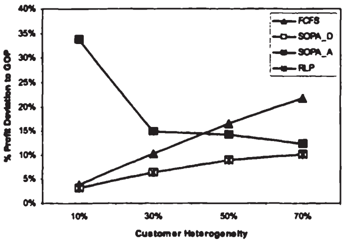

(a) $n_t = 10$

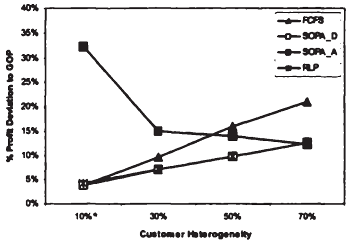

(b) $n_t \sim Poisson(10)$

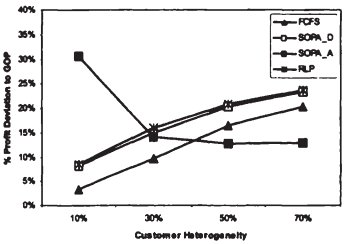

(c) $n_t \sim NB(10, 64)$

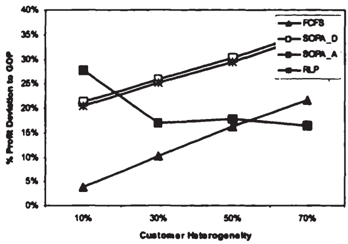

(d) $n_t \sim NB(10, 144)$

Figure 7.13: RLP and SOPA with Non-Nested Order Consumption

Note that in Figure 7.13, the FCFS line is equal to the one of the previous figure but the lines of RLP and SOPA changed. In general it can be said that the prohibition of nesting leads to slightly lower profits than in the previous scenarios with nesting. However, the decrease in profits lies around 2% and is even hard to see in the figure. Refer to the data tables B.12 and B.13 for a better illustration of this issue. Interestingly, the prohibition of nesting affects all four considered scenarios in the same way, even the constant demand scenario 7.13 (a). In this respect, the claim of De Boer et al. (2002) that nesting compensates for high demand variability cannot be supported in our settings as it brings profit improvements irrespective on the level of variability.

7.4 Summary

We have seen in this chapter that the performance of demand management approaches in stochastic environments is essentially driven by the degree of customer heterogeneity and the quality of demand forecasts. The amount of supply shortage and the degree of stochasticity are further important influencing factors. In general it can be said that SOPA is beneficial in comparison with FCFS when customer heterogeneity is high and forecasts are good. However, the use of SOPA has to be carefully considered that the effort to segment customers and to create demand forecasts pays off.

The RLP approach performs similar to the analyzed SOPA_D approach. Therefore, a use of the more complex RLP instead of SOPA_D cannot be recommended, as the gathering of information about demand distribution requires much more effort.

The RM approach dominates the benchmark approaches throughout the different scenarios. However, RM requires much more effort and is dependent on certain assumptions on the demand. The use of RM can be recommended when the effort of gathering reliable demand data and spending computing time stays in a reasonable proportion to the expected profit improvements.

Chapter 8

Conclusion

In this work, we discussed the problem of how to effectively manage stochastic demand in make-to-stock manufacturing. Specifically, we considered the situation of a manufacturer who decides on the quantities he is willing to sell to different customer classes. The order acceptance decisions take into account on-hand inventory as well as already planned production quantities scheduled to arrive in the future. For each order, the manufacturer has to decide—based on its profitability—whether to accept the order, to reject it, or to backlog it against a price discount. The problem is motivated by the demand fulfillment task in advanced planning systems. A key characteristic of the problem setting is that production orders cannot be changed in the short term. This is in line with the hierarchical planning approach of most advanced planning systems and reflects the reality of many manufacturers.

We presented a literature classification and overview of research in demand management. It turned out that the majority of models considering stochastic demand focuses on make-to-order environments. We adopted ideas from the classified literature, especially from traditional revenue management approaches, and transferred them to make-to-stock manufacturing. To our knowledge, this work is the first to apply revenue management in this context.

We developed two different approaches considering stochastic demand. First, we model the make-to-stock demand fulfillment problem as a stochastic dynamic program. We proved that the optimal policy in this model has a simple, intuitive structure, which can be interpreted as an extension of the well-known booking-limit policies in classical revenue-management problems. By explicitly capturing demand uncertainty, our model differs from the rule-based deterministic models commonly underlying the demand fulfillment modules of advanced planning systems. Second, we combined conventional LP-based models with stochastic demand information by repetitively solving the LP with different random variates. This idea stems from the randomized linear programming of network revenue management problems.

We tested the models numerically and compared them against a first-come-first-served rule and against a deterministic optimization approach. Our results show that explicitly accounting for demand uncertainty significantly improves the performance of demand fulfillment. The results also show that customer differentiation can yield a substantial profit increase, in particular if differences in profitability are large across orders and if supply is scarce. In conclusion, our results highlight substantial opportunities for improving the current practice of order fulfillment in make-to-stock manufacturing.

Our models make a first step in this direction. However, many challenges remain. For the RM approach, the biggest limitation is that it is not easily scalable. For large-scale applications, the full evaluation of the value function is computationally intractable. The RLP approach, however, is applicable even in large-scale applications, but seems to bring no further benefits towards deterministic linear programming approaches. In order to gain further insight into the relative performance of different methods, they should also be compared based on empirical data, in addition to theoretical demand distributions. In addition, further analysis of the mentioned rules in Section 2.3.2 is necessary. As they are commonly applied in APS, FCFS alone is not a fair benchmark for RLP and the RM approach.

Another relevant extension to our model would be to include different customer due-dates. It is not immediately clear which effect this will have on the structure of the optimal fulfillment policy. Another direction for future research will be to include short-term price incentives, which complement the order acceptance decisions addressed here.

Appendix A

Proofs of the Structural Properties of the RM approach[1]

A.1 Proof of Proposition 5.1

We show part a) by induction. For $t = T+1$ the inequality holds since $V_{T+1} \equiv 0$ and since the right-hand side is non-negative for $n > m$.

Now assume that inequality a) holds for $t+1$. We show that it also holds for t. We can rewrite

$$\begin{aligned} &\Delta_m V_t(\bar{x}) - \Delta_n V_t(\bar{x}) \leq b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t)) \\ \Leftrightarrow\; &V_t(\bar{x} - \bar{e}_n) - V_t(\bar{x} - \bar{e}_m) \leq b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t)). \end{aligned}$$

We show that this inequality holds for any given values of c and d, which implies that it also holds in expectation. To this end, let $W_t^{\bar{u}}(\bar{x}, c, d)$ denote the maximum expected profit-to-go when starting in period t with a supply vector $\bar{x}$, receiving demand d from customer class c and taking the fulfillment decision $\bar{u}$. Furthermore, let $\bar{u}_t^*(\bar{x}, d, c)$ be an optimal decision in period t under the same conditions.

Using this notation, we have to show that

$$\begin{aligned} &W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_n,d,c)}(\bar{x} - \bar{e}_n, c, d) - W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_m,d,c)}(\bar{x} - \bar{e}_m, c, d) \\ &\leq b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t)). \end{aligned}$$

We show that there is a feasible decision $\bar{u}_t(\bar{x}-\bar{e}_m, d, c)$ for which this inequality holds. This suffices since $W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_m,d,c)} \geq W_t^{\bar{u}_t(\bar{x}-\bar{e}_m,d,c)}$.

[1]The proofs in this chapter are part of the paper of Quante et al. (2009a).

By definition, we have

$$\begin{aligned}W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_n,d,c)}(\bar{x}-\bar{e}_n,c,d) &= \sum_{i=1}^{T}(u_i^* P_t(i,c) - hx_i\delta_{it}) \\ &\quad + h\delta_{nt} + V_{t+1}((\bar{x}-\bar{e}_n)-\bar{u}^*) \\ &= \cdots + u_m^* P_t(m,c) - hx_m\delta_{mt} + \cdots + u_n^* P_t(n,c) \\ &\quad - h(x_n-1)\delta_{nt} + \cdots + V_{t+1}(\bar{x}-\bar{e}_n-\bar{u}^*), \quad \text{(A.1)}\end{aligned}$$

where we have omitted the arguments of u_i^* for notational convenience. We now construct an appropriate feasible decision for state $\bar{x}-\bar{e}_m$. We distinguish two cases.

Case (i): $u_m^* > 0$

In this case, the decision $\bar{u}_t^*(\bar{x}-\bar{e}_n,d,c)-\bar{e}_m+\bar{e}_n$ is feasible in state $\bar{x}-\bar{e}_m$ and we get

$$\begin{aligned}W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_n,d,c)-\bar{e}_m+\bar{e}_n}(\bar{x}-\bar{e}_m,c,d) &= \sum_{i=1}^{T}(u_i^* P_t(i,c) - hx_i\delta_{it}) \\ &\quad + P_t(n,c) - P_t(m,c) \\ &\quad + h\delta_{mt} + V_{t+1}((\bar{x}-\bar{e}_m)-(\bar{u}^*-\bar{e}_m+\bar{e}_n)) \\ &= \cdots + (u_m^*-1)P_t(m,c) - h(x_m-1)\delta_{mt} \\ &\quad + \cdots + (u_n^*+1)P_t(n,c) - hx_n\delta_{nt} + \ldots \\ &\quad + V_{t+1}(\bar{x}-\bar{e}_n-\bar{u}^*). \quad \text{(A.2)}\end{aligned}$$

Taking the difference between Equations A.1 and A.2, the profits-to-go after period $t+1$ vanish and we are left with the difference in current profits, which equals

$$P_t(m,c) - P_t(n,c) + h\delta_{nt} - h\delta_{mt} = b(n-m+\delta_{nt}(t-n)+\delta_{mt}(m-t)),$$

by definition of P_t.

Case (ii): $u_m^* = 0$

In this case, the decision $\bar{u}_t^*(\bar{x}-\bar{e}_n,d,c)$ is feasible in state $\bar{x}-\bar{e}_m$ and we get

$$\begin{aligned} W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_n,d,c)}(\bar{x}-\bar{e}_m,c,d) = &\sum_{i=1}^{T}(u_i^* P_t(i,c) - hx_i\delta_{it}) + h\delta_{mt} \\ &+ V_{t+1}((\bar{x}-\bar{e}_m)-\bar{u}^*) \\ = &\cdots + u_m^* P_t(m,c) - h(x_m-1)\delta_{mt} + \dots \\ &+ u_n^* P_t(n,c) - hx_n\delta_{nt} + \cdots + V_{t+1}(\bar{x}-\bar{e}_m-\bar{u}^*). \end{aligned} \tag{A.3}$$

Calculating the difference in current profits between Equations A.1 and A.3 yields

$$\begin{aligned} &-hx_m\delta_{mt} - h(x_n-1)\delta_{nt} - (-h(x_m-1)\delta_{mt} - hx_n\delta_{nt}) \\ &= h(\delta_{nt}-\delta_{mt}) \leq 0, \end{aligned}$$

where the inequality follows from $n > m$. For the difference in expected future profits we obtain

$$\begin{aligned} &V_{t+1}(\bar{x}-\bar{e}_n-\bar{u}^*) - V_{t+1}(\bar{x}-\bar{e}_m-\bar{u}^*) \\ &\leq b(n-m+\delta_{nt+1}(t+1-n)+\delta_{mt+1}(m-t-1)) \\ &\leq b(n-m+\delta_{nt}(t-n)+\delta_{mt}(m-t)), \end{aligned} \tag{A.4}$$

where the first inequality follows from the induction assumption and the second inequality follows since A.4 is decreasing in t for $m < n$. This completes the proof of Part a).

Part b) follows immediately from Part a) by replacing $\bar{x}$ with $\bar{x}+\bar{e}_m+\bar{e}_n$ and using the definition of $\Delta_i V_t(\bar{x})$. □

A.2 Proof of proposition 5.2

We have the following equivalences

$$\begin{aligned} &P_t(m,c) - \Delta_m V_{t+1}(\bar{x}) \geq P_t(n,c) - \Delta_n V_{t+1}(\bar{x}) \\ &\Leftrightarrow P_t(m,c) - P_t(n,c) \geq \Delta_m V_{t+1}(\bar{x}) - \Delta_n V_{t+1}(\bar{x}) \\ &\overset{5.3}{\Leftrightarrow} b(n-m+\delta_{nt}(t-n)+\delta_{mt}(m-t)) + h(\delta_{mt}-\delta_{nt}) \geq \\ &\Delta_m V_{t+1}(\bar{x}) - \Delta_n V_{t+1}(\bar{x}). \end{aligned}$$

From Part a) of Proposition 5.1 we know that $\Delta_m V_{t+1}(\bar{x}) - \Delta_n V_{t+1}(\bar{x}) \leq b(n - m + \delta_{nt+1}(t + 1 - n) + \delta_{mt+1}(m - t - 1))$. As in Case (ii) of the proof of Proposition 5.1, this implies the desired result since for $n > m$ we have $\delta_{nt+1}(t+1-n)+\delta_{mt+1}(m-t-1) \leq \delta_{nt}(t-n)+\delta_{mt}(m-t)$ and $\delta_{mt}-\delta_{nt} \geq 0$. □

A.3 Proof of Proposition 5.3 and Theorem 5.1

We show both properties jointly by induction. For $t = T$, Proposition 5.3 holds since $V_{T+1}(.) \equiv 0$ and $P_T(m, c) = P_T(n, c)$ for $m, n \leq T$.

Now assume that Proposition 5.3 holds for Period t. We first show that Theorem 5.1 then holds for Period t and subsequently that Proposition 5.3 holds for Period $t - 1$.

Equation 5.6 shows that one can decompose the fulfillment decision into unit steps and that selling a given unit is beneficial if immediate profits outweigh the opportunity cost of losing this unit. Proposition 5.2 shows that there is an optimal policy $\bar{u}^*$ for which $u_i^* > 0$ implies $x_j - u_j^* = 0$ for all $j < i$. Any optimal policy that does not satisfy this property can be modified by swapping one unit of supply j against one unit of supply i. Proposition 5.2 implies that this modification does not decrease the objective function value. Therefore, one can obtain an optimal solution through a line search, starting with the earliest available supply, i.e. the smallest i, for which $x_i > 0$. The induction assumption of Proposition 5.3 shows that the objective function is concave along this search line. Therefore, one can stop the search as soon as immediate profits drop below the opportunity costs. This proves that the procedure defined in Theorem 5.1 yields an optimal policy.

We now show Proposition 5.3 for Period $t-1$. We use the same notation as in the proof of Proposition 5.1. Let $\bar{u}^1 := \bar{u}_t^*(\bar{x}, c, d)$, $\bar{u}^2 := \bar{u}_t^*(\bar{x} - \bar{e}_m, c, d)$, and $\bar{u}^3 := \bar{u}_t^*(\bar{x} - \bar{e}_m - \bar{e}_n, c, d)$ denote the optimal decisions in states $\bar{x}$, $\bar{x} - \bar{e}_m$, and $\bar{x} - \bar{e}_m - \bar{e}_n$, respectively, for a given customer class c and demand quantity d. Furthermore, let

$$A := W_t^{\bar{u}^1}(\bar{x}, c, d) - W_t^{\bar{u}^2}(\bar{x} - \bar{e}_m, c, d) \quad \text{and}$$
$$B := W_t^{\bar{u}^2}(\bar{x} - \bar{e}_m, c, d) - W_t^{\bar{u}^3}(\bar{x} - \bar{e}_m - \bar{e}_n, c, d).$$

We rewrite Proposition 5.3 for Period $t - 1$ as $\Delta_m V_t(\bar{x}) - \Delta_n V_t(\bar{x} - \bar{e}_m) \leq P_{t-1}(m, c) - P_{t-1}(n, c)$ and show that

$$A - B \leq P_t(m, c) - P_t(n, c) - h(\delta_{mt} - \delta_{nt})$$
$$\leq P_{t-1}(m, c) - P_{t-1}(n, c),$$

for any values of c and d, which implies that these inequalities also hold in expectation.

The second inequality follows directly from the definition of $P_t(.,.)$ for $m < n$. For the first inequality, we distinguish three cases, based on the value of $\bar{u}^1$.

Case (i): $\bar{u}^1 \equiv 0$

Theorem 5.1 implies $\bar{u}_2 = \bar{u}_3 \equiv 0$. From the definition of W_t we get

$$\begin{aligned} A - B &= \Delta_m V_{t+1}(\bar{x}) - h\delta_{mt} - \Delta_n V_{t+1}(x - \bar{e}_m) + h\delta_{nt} \\ &\leq P_t(m,c) - P_t(n,c) - h(\delta_{mt} - \delta_{nt}), \end{aligned}$$

where the inequality follows from the induction assumption.

Case (ii): $0 < \sum_{i=1}^{T} u_i^1 < d$

Theorem 5.1 implies $\bar{u}^2 = \bar{u}^1 - \bar{e}_m$. There are two possibilities for $\bar{u}^3$. If $\bar{u}^2 \neq 0$ then $\bar{u}^3 = \bar{u}^2 - \bar{e}_n$, otherwise $\bar{u}^3 \equiv 0$. In the first case, we get

$$A - B = P_t(m,c) - P_t(n,c) - h(\delta_{mt} - \delta_{nt}).$$

In the second case, $B = \Delta_n V_{t+1}(\bar{x} - \bar{e}_m) - h\delta_{nt}$ and Theorem 5.1 implies that this value is smaller than $P_t(n,c) - h\delta_{nt}$ since it is optimal not to sell another unit of supply n. Thus, $A - B$ satisfies the desired inequality in either case.

Case (iii): $\sum_{i=1}^{T} u_i^1 = d$

Theorem 5.1 implies that either $\bar{u}^2 = \bar{u}^1 - \bar{e}_m$ or $\bar{u}^2 = \bar{u}^1 - \bar{e}_m + \bar{e}_k$ for some $k \geq n$. The first alternative leads to the same calculations as in Case (ii) above. The second alternative leaves two options for $\bar{u}^3$, namely either $\bar{u}^3 = \bar{u}^2 - \bar{e}_n + \bar{e}_l$ for some $l \geq k$ or $\bar{u}^3 = \bar{u}^2 - \bar{e}_n$. In the first case, we get

$$\begin{aligned} A - B = {} & P_t(m,c) - P_t(k,c) + \Delta_k V_{t+1}(\bar{x} - \bar{u}^1) \\ & - P_t(n,c) - P_t(l,c) + \Delta_l V_{t+1}(\bar{x} - \bar{u}^1 - \bar{e}_k) - h(\delta_{mt} - \delta_{nt}) \\ \leq {} & P_t(m,c) - P_t(n,c) - h(\delta_{mt} - \delta_{nt}), \end{aligned}$$

where the inequality follows from the induction assumption. The other case regarding $\bar{u}_3$ yields $B = P_t(n,c) - h\delta_{nt}$ and therefore

$$\begin{aligned} A - B &= P_t(m,c) - P_t(k,c) + \Delta_k V_{t+1}(\bar{x} - \bar{u}_1) - P_t(n,c) - h(\delta_{mt} - \delta_{nt}) \\ &\leq P_t(m,c) - P_t(n,c) - h(\delta_{mt} - \delta_{nt}), \end{aligned}$$

where the inequality follows from the fact that it is optimal in state $\bar{x} - \bar{u}_1$ to sell an additional unit of supply k. □

A.4 Proof of Proposition 5.4

We show by induction that $\Delta_i V_{t+1}(\bar{x}) > 0$ implies $\Delta_i V_t(\bar{x}) \geq \Delta_i V_{t+1}(\bar{x})$. Since $P_t(i,c)$ is increasing in t this assures non-increasing protection levels for any i and c.

For $t = T$ the condition is empty since $\Delta_i V_{T+1}(\bar{x}) = 0$ for all i, c, and $\bar{x}$. In other words, all protection levels vanish at the end of the planning horizon.

Assume now that Proposition 5.4 holds for Period $t+1$. We show that it also holds for Period t. To this end, assume that $\Delta_i V_{t+1}(\bar{x}) > 0$. Using the Bellman recursion of V_t we have

$$\Delta_i V_t(\bar{x}) = -h\delta_{it} + E_{c,d}\left[\max\left\{P_t(i,c)\delta_{d>0}\ ;\ \Delta_i V_{t+1}(\bar{x})\right\}\right], \qquad \text{(A.5)}$$

where $\delta_{d>0}$ equals unity if $d > 0$ and zero otherwise. For $t < i$ the holding-cost term vanishes and we immediately get $\Delta_i V_t(\bar{x}) \geq \Delta_i V_{t+1}(\bar{x})$. For $t \geq i$ we rewrite (A.5) for $t+1$

$$\Delta_i V_{t+1}(\bar{x}) = -h\delta_{it+1} + E_{c,d}\left[\max\left\{P_{t+1}(i,c)\delta_{d>0}\ ;\ \Delta_i V_{t+2}(\bar{x})\right\}\right], \qquad \text{(A.6)}$$

and compare the individual terms. We have $-h\delta_{it} = -h\delta_{it+1}$ and $P_t(i,c) = P_{t+1}(i,c)$ since $t \geq i$. In addition, the induction assumption implies that $\Delta_i V_{t+1}(\bar{x}) \geq \Delta_i V_{t+2}(\bar{x})$ if the maximum in (A.6) is attained by the last term. Therefore, $\Delta_i V_t(\bar{x}) \geq \Delta_i V_{t+1}(\bar{x})$, which completes the proof. □

Appendix B

Data Tables

Table B.1: Data of Fig. 7.1

(a) Lost Sales

Cl.	FCFS	SOPA_D	SOPA_A
1	40.89%	11.76%	1.56%
2	40.52%	28.73%	31.83%
3	40.97%	77.30%	82.98%

(b) Backlogged Quantities

Cl.	FCFS	SOPA_D	SOPA_A
1	0.00%	4.28%	47.59%
2	0.00%	2.09%	60.59%
3	0.00%	1.96%	16.37%

Table B.2: Data of Fig. 7.2

(a) Early Delivery

Cl.	FCFS	SOPA_D	SOPA_A
1	58.84%	77.43%	43.99%
2	59.21%	62.95%	4.36%
3	58.82%	15.40%	0.45%

(b) On Time Delivery

Cl.	FCFS	SOPA_D	SOPA_A
1	0.27%	6.53%	6.86%
2	0.27%	6.23%	3.22%
3	0.20%	5.35%	0.20%

Table B.3: Data of Fig. 7.3

Cl.	GOP	FCFS	SOPA_D	SOPA_A
1	2,617,006	2,515,394	2,535,015	1,231,004
2	2,617,006	2,515,394	2,544,101	1,767,756
3	2,617,006	2,515,394	2,543,725	1,774,840
4	2,617,006	2,515,394	2,540,218	1,849,463
5	2,617,006	2,515,394	2,539,859	1,860,748
6	2,617,006	2,515,394	2,535,119	1,882,588
7	2,617,006	2,515,394	2,533,737	1,870,950
8	2,617,006	2,515,394	2,531,967	1,894,752
9	2,617,006	2,515,394	2,531,083	1,894,894
10	2,617,006	2,515,394	2,533,138	1,900,758

Table B.4: Data of Fig. 7.4

Cl.	FCFS	SOPA_D	SOPA_A
1	3.88%	3.13%	52.96%
2	3.88%	2.79%	32.45%
3	3.88%	2.80%	32.18%
4	3.88%	2.93%	29.33%
5	3.88%	2.95%	28.90%
6	3.88%	3.13%	28.06%
7	3.88%	3.18%	28.51%
8	3.88%	3.25%	27.60%
9	3.88%	3.28%	27.59%
10	3.88%	3.21%	27.37%

Table B.5: Data of Fig. 7.5

(a) Heterogeneity 10%

Cl.	FCFS	SOPA_D	SOPA_A
1	3.88%	3.13%	52.96%
2	3.88%	2.79%	32.45%
3	3.88%	2.80%	32.18%
4	3.88%	2.93%	29.33%
5	3.88%	2.95%	28.90%
6	3.88%	3.13%	28.06%
7	3.88%	3.18%	28.51%
8	3.88%	3.25%	27.60%
9	3.88%	3.28%	27.59%
10	3.88%	3.21%	27.37%

(b) Heterogeneity 30%

Cl.	FCFS	SOPA_D	SOPA_A
1	10.24%	9.67%	54.44%
2	10.24%	5.42%	12.52%
3	10.24%	5.38%	14.69%
4	10.24%	4.91%	13.61%
5	10.24%	5.01%	8.50%
6	10.24%	5.27%	10.09%
7	10.24%	5.28%	10.11%
8	10.24%	5.32%	9.52%
9	10.24%	5.36%	9.70%
10	10.24%	5.27%	9.55%

(c) Heterogeneity 50%

Cl.	FCFS	SOPA_D	SOPA_A
1	16.29%	15.63%	54.40%
2	16.29%	8.28%	10.51%
3	16.29%	7.23%	13.49%
4	16.29%	6.33%	12.16%
5	16.29%	6.50%	7.80%
6	16.29%	6.94%	8.78%
7	16.29%	6.76%	8.07%
8	16.29%	6.94%	8.05%
9	16.29%	7.02%	7.37%
10	16.29%	6.90%	7.36%

(d) Heterogeneity 70%

Cl.	FCFS	SOPA_D	SOPA_A
1	21.61%	21.08%	52.63%
2	21.61%	10.41%	9.04%
3	21.61%	8.63%	12.76%
4	21.61%	7.47%	10.53%
5	21.61%	7.81%	6.60%
6	21.61%	8.40%	7.66%
7	21.61%	8.02%	7.14%
8	21.61%	8.05%	6.07%
9	21.61%	8.31%	6.37%
10	21.61%	8.20%	5.94%

Table B.6: Data of Fig. 7.6

(a) Forecast Error 0%

Cl.	FCFS	SOPA_D	SOPA_A
1	3.88%	1.94%	52.96%
2	3.88%	0.36%	33.43%
3	3.88%	0.15%	31.04%
4	3.88%	0.10%	30.00%
5	3.88%	0.06%	29.13%
6	3.88%	0.04%	28.66%
7	3.88%	0.03%	28.33%
8	3.88%	0.02%	27.67%
9	3.88%	0.02%	27.19%
10	3.88%	0.02%	26.83%

(b) Forecast Error M

Cl.	FCFS	SOPA_D	SOPA_A
1	3.88%	3.13%	52.96%
2	3.88%	2.79%	32.45%
3	3.88%	2.80%	32.18%
4	3.88%	2.93%	29.33%
5	3.88%	2.95%	28.90%
6	3.88%	3.13%	28.06%
7	3.88%	3.18%	28.51%
8	3.88%	3.25%	27.60%
9	3.88%	3.28%	27.59%
10	3.88%	3.21%	27.37%

(c) Forecast Error 50%

Cl.	FCFS	SOPA_D	SOPA_A
1	3.88%	4.54%	52.96%
2	3.88%	5.42%	31.21%
3	3.88%	5.41%	23.94%
4	3.88%	5.61%	21.51%
5	3.88%	5.40%	19.60%
6	3.88%	5.24%	16.87%
7	3.88%	5.63%	15.84%
8	3.88%	5.54%	13.90%
9	3.88%	6.01%	13.69%
10	3.88%	5.54%	13.13%

(d) Forecast Error 100%

Cl.	FCFS	SOPA_D	SOPA_A
1	3.88%	8.88%	52.96%
2	3.88%	11.79%	17.47%
3	3.88%	11.07%	16.21%
4	3.88%	11.63%	14.84%
5	3.88%	11.74%	11.49%
6	3.88%	12.10%	11.02%
7	3.88%	12.39%	11.14%
8	3.88%	12.58%	10.19%
9	3.88%	13.59%	11.16%
10	3.88%	12.49%	9.55%

Table B.7: Data of Fig. 7.7

Backlogging Costs	GOP	FCFS	SOPA_D	SOPA_A
0	2,799,229	2,515,394	2,701,440	2,788,465
1	2,713,034	2,515,394	2,638,449	2,606,074
2	2,683,542	2,515,394	2,608,421	2,231,926
3	2,666,971	2,515,394	2,590,617	1,760,303
4	2,655,540	2,515,394	2,579,133	1,253,369
5	2,646,701	2,515,394	2,569,319	1,112,489
6	2,639,198	2,515,394	2,563,497	1,556,233
7	2,632,968	2,515,394	2,558,761	1,682,605
8	2,627,311	2,515,394	2,554,074	1,702,957
9	2,621,976	2,515,394	2,548,068	1,745,037
10	2,617,006	2,515,394	2,543,725	1,774,840

Table B.8: Data of Fig. 7.8

(a) Lost Sales

Cl.	GOP	FCFS	SOPA_D	RM
1	1.82%	38.65%	7.82%	14.62%
2	32.73%	40.40%	29.45%	26.91%
3	86.41%	38.53%	83.62%	76.89%

(b) Backlogged Quantities

Cl.	GOP	FCFS	SOPA_D	RM
1	2.62%	0.00%	2.84%	3.10%
2	1.51%	0.00%	2.62%	1.45%
3	0.03%	0.00%	1.04%	0.00%

Table B.9: Data of Fig. 7.9

CV	GOP	FCFS	SOPA_D	RM
0.00	17,890.36	17,265.44	17,682.46	17,768.62
0.33	17,851.67	17,229.71	17,553.97	17,694.21
0.67	17,843.03	17,247.27	17,327.95	17,635.90
1.00	17,748.89	17,164.66	16,767.71	17,449.12
1.33	17,440.88	16,803.89	15,880.23	17,031.49

Table B.10: Data of Fig. 7.10

(a) Heterogeneity 10%

CV	FCFS	SOPA_D	RM
0.00	3.49%	1.17%	0.68%
0.33	3.48%	1.68%	0.88%
0.67	3.33%	2.91%	1.16%
1.00	3.28%	5.58%	1.69%
1.33	3.65%	9.10%	2.37%

(b) Heterogeneity 20%

CV	FCFS	SOPA_D	RM
0.00	10.04%	1.72%	1.25%
0.33	9.77%	2.44%	1.63%
0.67	9.42%	4.50%	2.45%
1.00	8.84%	7.95%	3.47%
1.33	8.89%	12.21%	4.63%

(c) Heterogeneity 30%

CV	FCFS	SOPA_D	RM
0.00	17.07%	2.50%	1.73%
0.33	17.28%	3.32%	2.26%
0.67	16.95%	5.71%	3.29%
1.00	15.10%	9.58%	4.86%
1.33	14.82%	14.57%	6.31%

Table B.11: Data of Fig. 7.11

(a) $CV = 0.67$ and Low Heterogeneity

Shortage Rate	FCFS	SOPA_D	RM
0% (41%)	3.38%	3.28%	1.24%
25% (21%)	2.70%	5.73%	1.64%
41% (0%)	2.47%	3.23%	1.02%
50% (-19%)	2.03%	2.56%	0.49%
75% (-138%)	0.31%	0.10%	0.00%

(b) $CV = 0.33$ and Low Heterogeneity

Shortage Rate	FCFS	SOPA_D	RM
0% (41%)	3.52%	1.66%	0.89%
25% (21%)	2.29%	3.53%	1.28%
41% (0%)	2.03%	2.39%	0.76%
50% (-19%)	1.38%	1.44%	0.28%
75% (-138%)	0.01%	0.00%	0.00%

(c) $CV = 0.67$ and Medium Heterogeneity

Shortage Rate	FCFS	SOPA_D	RM
0% (41%)	9.18%	4.41%	2.36%
25% (21%)	5.64%	6.93%	3.25%
75% (-138%)	3.50%	3.75%	1.71%
41% (0%)	1.83%	2.63%	0.66%
50% (-19%)	0.23%	0.07%	0.00%

Table B.12: Data of Fig. 7.12

(a) $n_t = 10$

Heterogeneity	FCFS	SOPA_D	SOPA_A	RLP
10%	3.81%	1.65%	32.60%	1.65%
30%	10.28%	3.34%	14.81%	3.34%
50%	16.47%	4.37%	13.79%	4.37%
70%	21.78%	5.11%	12.45%	5.11%

(b) $n_t \sim Poisson(10)$

Heterogeneity	FCFS	SOPA_D	SOPA_A	RLP
10%	3.83%	2.98%	31.24%	3.19%
30%	9.65%	5.50%	15.86%	5.64%
50%	15.90%	7.05%	15.07%	7.20%
70%	20.97%	8.62%	12.92%	8.81%

(c) $n_t \sim NB(10, 64)$

Heterogeneity	FCFS	SOPA_D	SOPA_A	RLP
10%	3.40%	7.45%	29.98%	7.87%
30%	9.77%	12.42%	14.15%	13.49%
50%	16.46%	16.26%	12.63%	17.04%
70%	20.37%	18.98%	13.45%	19.65%

(d) $n_t \sim NB(10, 144)$

Heterogeneity	FCFS	SOPA_D	SOPA_A	RLP
10%	3.83%	18.59%	25.56%	17.75%
30%	10.33%	21.65%	15.41%	21.23%
50%	16.36%	26.08%	16.87%	25.56%
70%	21.63%	30.00%	14.78%	29.50%

Table B.13: Data of Fig. 7.13

(a) $n_t = 10$

Heterogeneity	FCFS	SOPA_D	SOPA_A	RLP
10%	3.81%	3.17%	33.69%	3.17%
30%	10.28%	6.45%	14.92%	6.45%
50%	16.47%	8.94%	14.15%	8.94%
70%	21.78%	10.10%	12.34%	10.10%

(b) $n_t \sim Poisson(10)$

Heterogeneity	FCFS	SOPA_D	SOPA_A	RLP
10%	3.83%	3.93%	32.23%	4.12%
30%	9.65%	7.02%	14.97%	6.99%
50%	15.90%	9.75%	14.01%	9.70%
70%	20.97%	12.66%	12.32%	12.59%

(c) $n_t \sim NB(10, 64)$

Heterogeneity	FCFS	SOPA_D	SOPA_A	RLP
10%	3.40%	8.18%	30.56%	8.65%
30%	9.77%	14.96%	14.05%	16.01%
50%	16.46%	20.26%	12.81%	20.81%
70%	20.37%	23.29%	13.01%	23.65%

(d) $n_t \sim NB(10, 144)$

Heterogeneity	FCFS	SOPA_D	SOPA_A	RLP
10%	3.83%	21.37%	27.81%	20.57%
30%	10.33%	25.92%	17.04%	25.21%
50%	16.36%	30.28%	17.90%	29.45%
70%	21.63%	35.50%	16.51%	34.63%

Bibliography

Agrawal, Narendra, Stephen A. Smith. 1998. Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Research Logistics* **43**(6) 839–861.

Anderson, Eric T., Gavan J. Fitzsimons, Duncan Simester. 2006. Measuring and mitigating the costs of stockouts. *Management Science* **52**(11) 1751–1763.

Arslan, Hasan, Stephen C. Graves, Thomas Roemer. 2007. A single-product inventory model for multiple demand classes. *Management Science* **53**(9) 1486–1500.

Ball, Michael O., Chien-Yu Chen, Zhen-Ying Zhao. 2004. Available-to-promise. David Simchi-Levi, S. David Wu, Zuo-Jun Shen, eds., *Handbook of Quantitative Supply Chain Analysis – Modeling in the E-Business Era*, chap. 11. Kluwer Academic Publishers, 447–483.

Barut, Mehmet, V. Sridharan. 2005. Revenue management in order-driven production systems. *Decision Sciences* **36**(2) 287–316.

Bitran, Gabriel, René Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing & Service Operations Management* **5**(3) 203–229.

Boyd, Andrew. 1998. Airline alliance revenue management. *OR/MS Today* **25**(5) 28–31.

Boyd, E. Andrew, Ioana C. Bilegan. 2003. Revenue management and e-commerce. *Management Science* **49**(10) 1363–1386.

Buttle, Francis. 2004. *Customer relationship management*. 1st ed. Butterworth Heinemann, Oxford.

Chan, L. M. A., Z. J. Max Shen, David Simchi-Levi, Julie L. Swann. 2004. Coordination of pricing and inventory decisions: A survey and classification. David Simchi-Levi, S. David Wu, Zuo-Jun Shen, eds., *Handbook of Quantitative Supply Chain Analysis – Modeling in the E-Business Era*, chap. 9. Kluwer Academic Publishers, 335–392.

Chen, Chien-Yu, Zhen-Ying Zhao, Michael O. Ball. 2001. Quantity and due date quoting available to promise. *Information Systems Frontiers* **3**(4) 477–488.

Chen, Chien-Yu, Zhenying Zhao, Michael O. Ball. 2002. A model for batch advanced available-to-promise. *Production and Operations Management* **11**(4) 424–440.

Chiang, Wen-Chyuan, Jason C. H. Chen, Xiaojing Xu. 2007. An overview of research on revenue management: current issues and future research. *International Journal of Revenue Management* **1**(1) 97–128.

De Boer, Sanne V., Richard Freling, Nanda Piersma. 2002. Mathematical programming for network revenue management revisited. *European Journal of Operational Research* **137**(1) 72–92.

De Véricourt, Francis, Fikri Karaesmen, Yves Dallery. 2002. Optimal stock allocation for a capacitated supply system. *Management Science* **48**(11) 1486–1501.

Defregger, Florian, Heinrich Kuhn. 2007. Revenue management for a make-to-order company with limited inventory capacity. *OR Spectrum* **29**(1) 137–156.

Ding, Qing, Panos Kouvelis, Joseph M. Milner. 2006. Dynamic pricing through discounts for optimizing multi-class fulfillment. *Operations Research* **54**(1) 169–183.

Ehrenberg, A. S. C. 1959. The pattern of consumer purchases. *Applied Statistics* **8**(1) 26–41.

Elmaghraby, Wedad, Pinar Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science* **49**(10) 1287–1309.

Fischer, Markus E. 2001. *"Available-to-Promise": Aufgaben und Verfahren im Rahmen des Supply Chain Management*. Theorie und Forschung Band 704, Wirtschaftswissenschaften Band 63, Roderer, Regensburg.

Fleischmann, Bernhard, Herbert Meyr. 2004. Customer orientation in advanced planning systems. Harald Dyckhoff, Richard Lackes, Joachim Reese, eds., *Supply Chain Management and Reverse Logistics*, chap. 3. Springer, 297–321.

Fleischmann, Bernhard, Herbert Meyr, Michael Wagner. 2008. Advanced planning. Hartmut Stadtler, Christoph Kilger, eds., *Supply Chain Management and Advanced Planning*, 4th ed. Springer, Berlin, Heidelberg, 81–106.

Fleischmann, Moritz, Joseph Hall, David Pyke. 2004. Smart pricing: Linking pricing decisions with operational insights. *Erim report series research in management* .

Gordon, Valery, Jean-Marie Proth, Chengbin Chu. 2002. A survey of the state-of-the-art of common due date assignment and scheduling research. *European Journal of Operational Research* **139**(1) 1–25.

Guerrero, Hector H., Gary M. Kern. 1988. How to more effectively accept and refuse orders. *Production and Inventory Management Journal* **29**(4) 59–63.

Gupta, Diwakar, Lei Wang. 2007. Capacity management for contract manufacturing. *Operations Research* **55**(2) 367–377.

Ha, Albert Y. 1997a. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. *Management Science* **43**(8) 1093–1103.

Ha, Albert Y. 1997b. Stock-rationing policy for a make-to-stock production system with two priority classes and backordering. *Naval Research Logistics* **44**(5) 457–472.

Harris, Frederick H., Jonathan P. Pinder. 1995. A revenue management approach to demand management and order booking in assemble-to-order manufacturing. *Journal of Operations Management* **13**(4) 299–309.

Hill, Terry. 2000. *Manufacturing Strategy*. 3rd ed. McGraw-Hill, Boston, Mass.

Hoekstra, Sjoerd, Jac E. Romme. 1992. *Integral Logistic Structures: Developing Customer-Oriented Goods Flow*. McGraw-Hill, London.

Jalora, Anshu. 2006. Order acceptance and scheduling at a make-to-order system using revenue management. Ph.D. thesis, Texas A&M University.

Kalyan, Vibhu. 2002. Dynamic customer value management: Asset values under demand uncertainty using airline yield management and related techniques. *Information Systems Frontiers* **4**(1) 101–119.

Kambil, Ajit, Eric van Heck. 2002. *Making Markets: How Firms Can Design and Profit from Online Auctions and Exchanges*. Harvard Business School Press, Boston.

Keskinocak, Pinar, Sridhar Tayur. 2004. Due-date management policies. David Simchi-Levi, S. David Wu, Zuo-Jun Shen, eds., *Handbook of Quantitative Supply Chain Analysis – Modeling in the E-Business Era*, chap. 12. Kluwer Academic Publishers, 485–547.

Khouja, Moutaz. 1999. The single-period (news-vendor) problem: Literature review and suggestions for future research. *Omega* **27**(5) 537–553.

Kilger, Christoph, Herbert Meyr. 2008. Demand fulfilment and ATP. Hartmut Stadtler, Christoph Kilger, eds., *Supply Chain Management and Advanced Planning*, 4th ed. Springer, Berlin, Heidelberg, 181–198.

Kilger, Christoph, Michael Wagner. 2008. Demand planning. Hartmut Stadtler, Christoph Kilger, eds., *Supply Chain Management and Advanced Planning*, 4th ed. Springer, Berlin, Heidelberg, 133–160.

Kimms, Alf, Robert Klein. 2005. Revenue Management im Branchenvergleich. *Zeitschrift für Betriebswirtschaft* **1** 1–30.

Kimms, Alf, Michael Müller-Bungart. 2007. Simulation of stochastic demand data streams for network revenue management problems. *OR Spectrum* **29** 5–20.

Kleijn, Marcel J., Rommert Dekker. 1998. An overview of inventory systems with several demand classes. Econometric institute report 9838/a, RSM Erasmus University.

Klophaus, R. 1998. Revenue Management: Wie die Airline Ertragswachstum schafft. *Absatzwirtschaft* **Sondernummer Oktober** 146–155.

Kocabıyıkoğlu, Ayşe, Ioana Popescu. 2005. Joint pricing and revenue management with general stochastic demand. Working paper, INSEAD.

Kolisch, Rainer, Danilo Zatta. 2006. Revenue-Management in der Sachgüterproduktion. *Marketing Journal* **12** 38–41.

Kuhn, Heinrich, Florian Defregger. 2004. Revenue Management in der Sachgüterproduktion. *Wirtschaftswissenschaftliches Studium* **5** 319–324.

Kumar, Sameer, Jared L. Frederick. 2007. Revenue management for a home construction products manufacturer. *Journal of Revenue & Pricing Management* **5**(4) 256–270.

Lautenbacher, Conrad J., Shaler Jr. Stidham. 1999. The underlying markov decision process in the single-leg airline yield-management problem. *Transportation Science* **33**(2) 136–146.

Lawless, Jerald F. 1987. Negative binomial and mixed poisson regression. *The Canadian Journal of Statistics* **15**(3) 209–225.

Lee, Young M. 2006. Simulating impact of available-to-promise generation on supply chain performance. *WSC '06: Proceedings of the 38th conference on Winter simulation*. Winter Simulation Conference, Monterey, California, 621–626.

Mantrala, Murali K., Surya Rao. 2001. A decision-support system that helps retailers decide order quantities and markdowns for fashion goods. *Interfaces* **31**(3) 146–165.

Manugistics. 2002. Enterprise profit optimization. URL `http://www.manugistics.com/documents/epo_whitepaper.pdf`.

Matsumoto, Makoto, Takuji Nishimura. 1998. Mersenne twister: A 623-dimensionally equidistributed uniform pseudorandom number generator. *ACM Transactions on Modeling and Computer Simulation* **8**(1) 3–30.

McGill, Jeffrey I, Garrett van Ryzin. 1999. Revenue management: Research overview and prospects. *Transportation Science* **33**(2) 233–256.

Melchiors, P., R. Dekker, M.J. Kleijn. 2000. Inventory rationing in an (s, Q) inventory model with lost sales and two demand classes. *Journal of the Operational Research Society* **51**(1) 111–122.

Meyr, Herbert. 2008. Clustering methods for rationing limited resources. Lars Mönch, Giselher Pankratz, eds., *Intelligente Systeme zur Entscheidungsunterstützung*. SCS Publishing House e.V., 19–32.

Meyr, Herbert. 2009. Customer segmentation, allocation planning and order promising in make-to-stock production. *OR Spectrum* **31**(1) 229–256.

Meyr, Herbert, Heidrun Rosič, Christian Seipl, Michael Wagner, Ulrich Wetterauer. 2008a. Architecture of selected APS. Hartmut Stadtler, Christoph Kilger, eds., *Supply Chain Management and Advanced Planning*, 4th ed. Springer, Berlin, Heidelberg, 349–366.

Meyr, Herbert, Michael Wagner, Jens Rohde. 2008b. Structure of Advanced Planning Systems. Hartmut Stadtler, Christoph Kilger, eds., *Supply Chain Management and Advanced Planning*, 4th ed. Springer, Berlin, Heidelberg, 109–116.

Mild, Andreas, Martin Natter, Thomas Reutterer, Alfred Taudes, Jürgen Wöckl. 2006. Retail revenue management. Peter Schnedlitz, Renate Buber, Thomas Reutterer, Arnold Schuh, Christoph Teller, eds., *Innovationen in Marketing und Handel*. Linde, Wien, 124–143.

Möllering, Karin T., Ulrich W. Thonemann. 2007. An optimal critical level policy for inventory systems with two demand classes. Tech. rep., University of Cologne.

Neslin, Scott A. 2002. Sales promotion. Barton A. Weitz, Robin Wensley, eds., *Handbook of Marketing*, chap. 13. Sage Publications, London, 310–338.

Optiant. 2007. URL http://www.optiant.com.

Pak, Kevin, Nanda Piersma. 2002. Airline revenue management: An overview of OR techniques 1982-2001. Econometric Institute Report 256, Erasmus University Rotterdam.

Petruzzi, Nicholas C., Maqbool Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Operations Research* **47**(2) 183–194.

Pibernik, Richard. 2005. Advanced available-to-promise: Classification, selected methods and requirements for operations and inventory management. *International Journal of Production Economics* **93–94** 239–252.

Pibernik, Richard. 2006. Managing stock-outs effectively with order fulfilment systems. *Journal of Manufacturing Technology Management* **17**(6) 721–736.

Pibernik, Richard, Prashant Yadav. 2009. Inventory reservation and real-time order promising in a make-to-stock system. *OR Spectrum* **31**(1) 281–307.

Porteus, Evan L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press.

Quante, Rainer, Moritz Fleischmann, Herbert Meyr. 2009a. A stochastic dynamic programming approach to revenue management in a make-to-stock production system. Working paper, Vienna University of Economics and Business Administration.

Quante, Rainer, Herbert Meyr, Moritz Fleischmann. 2009b. Revenue management and demand fulfillment: Matching applications, models, and software. *OR Spectrum* **31**(1) 31–62.

Rehkopf, Stefan. 2006. *Revenue Management-Konzepte zur Auftragsannahme bei kundenindividueller Produktion*. 1st ed. DUV.

Rehkopf, Stefan, Thomas Spengler. 2004. Revenue management in a make-to-order environment. H. Fleuren, D. den Hertog, Kort P., eds., *Operations Research Proceedings 2004*. Springer, 470–478.

Rohde, Jens, Herbert Meyr, Michael Wagner. 2000. Die Supply Chain Planning Matrix. *PPS Management* **5** 10–15.

Rohde, Jens, Michael Wagner. 2008. Master planning. Hartmut Stadtler, Christoph Kilger, eds., *Supply Chain Management and Advanced Planning*, 4th ed. Springer, Berlin, Heidelberg, 161–180.

SAP Help Portal. 2008. URL http://help.sap.com.

Scarf, Herbert. 1960. The optimality of (s,S) policies for the dynamic inventory problem. K. J. Arrow, S. Karlin, P. Suppes, eds., *Mathematical Methods in the Social Sciences*. Stanford University Press, Stanford, California, 196–202.

Schneeweiß, Christoph. 2003. *Distributed decision making*. 2nd ed. Springer.

Schwendinger, j. 1979. Master production scheduling's available-to-promise. *APICS Conference Proceedings* 316–330.

Sharman, Graham. 1984. The rediscovery of logistics. *Harvard Business Review* **62**(5) 71–79.

Silver, Edward A. 1981. Operations research in inventory management: A review and critique. *Operations Research* **29**(4) 628–645.

Silver, Edward A., David F. Pyke, Rein Peterson, eds. 1998. *Inventory Management and Production Planning and Scheduling*. 3rd ed. Wiley.

Smartops. 2007. URL `http://www.smartops.com`.

Specht, Dieter, Christian M. F. Gruß. 2007. Revenue Management Ũ eine Strategie für die Produktion in der Automobilindustrie? Dieter Specht, ed., *Strategische Bedeutung der Produktion: Tagungsband der Herbsttagung 2006 der Wissenschaftlichen Kommission Produktionswirtschaft im VHB*, 1st ed. DUV, Wiesbaden, 61–72.

Spengler, Thomas, Stefan Rehkopf, Thomas Volling. 2007. Revenue management in make-to-order manufacturing – an application to the iron and steel industry. *OR Spectrum* **29**(1) 157–171.

Stadtler, Hartmut. 2008. Purchasing and material requirements planning. Hartmut Stadtler, Christoph Kilger, eds., *Supply Chain Management and Advanced Planning*, 4th ed. Springer, Berlin, Heidelberg, 217–230.

Stadtler, Hartmut, Christoph Kilger, eds. 2008. *Supply Chain Management and Advanced Planning*. 4th ed. Springer, Berlin, Heidelberg.

Swann, Julie L. 1999. Flexible pricing policies: Introduction and a survey of implementation in various industries. Contract report, General Motors Corporation.

Talluri, Kalyan T., Garrett J. van Ryzin. 1999. A randomized linear programming method for computing network bid prices. *Transportation Science* **33**(2) 207–216.

Talluri, Kalyan T., Garrett J. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers.

Tempelmeier, Horst. 2006. *Inventory Management in Supply Networks. Problems, Models, Solutions*. 1st ed. Norderstedt: Books on Demand.

Teunter, Ruud H., Willem K. Klein Haneveld. 2008. Dynamic inventory rationing strategies for inventory systems with two demand classes, poisson demand and backordering. *European Journal of Operational Research* **190**(1) 156–178.

Vakali, Athena, Lefteris Angelis, Dimitra Pournara. 2001. Internet based auctions: a survey on models and applications. *ACM SIGecom Exchanges* **2**(1) 6–15.

Vollmann, Thomas E., William L. Berry, D. C. Whybark. 2005. *Manufacturing Planning and Control for Supply Chain Management*. 5th ed. Mcgraw-Hill.

Weatherford, Lawrence R., Samuel E. Bodily. 1992. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. *Operations Research* **40**(5) 831–844.

Whitin, T. M. 1955. Inventory control and price theory. *Management Science* **2**(1) 61–68.

Yano, Candace Arai, Stephen M. Gilbert. 2003. Coordinated pricing and production/procurement decisions: A review. Amiya K. Chakravarty, Jehoshua Eliashberg, eds., *Managing Business Interfaces: Marketing, Engineering and Manufacturing Perspectives*, chap. 3. International Series in Quantitative Marketing, Kluwer Academic Publishers, Norwell, Massachusetts, 65–103.

Forschungsergebnisse der Wirtschaftsuniversität Wien

Herausgeber: Wirtschaftsuniversität Wien –
vertreten durch a.o. Univ. Prof. Dr. Barbara Sporn

Band 1 Stefan Felder: Frequenzallokation in der Telekommunikation. Ökonomische Analyse der Vergabe von Frequenzen unter besonderer Berücksichtigung der UMTS-Auktionen. 2004.

Band 2 Thomas Haller: Marketing im liberalisierten Strommarkt. Kommunikation und Produktplanung im Privatkundenmarkt. 2005.

Band 3 Alexander Stremitzer: Agency Theory: Methodology, Analysis. A Structured Approach to Writing Contracts. 2005.

Band 4 Günther Sedlacek: Analyse der Studiendauer und des Studienabbruch-Risikos. Unter Verwendung der statistischen Methoden der Ereignisanalyse. 2004.

Band 5 Monika Knassmüller: Unternehmensleitbilder im Vergleich. Sinn- und Bedeutungsrahmen deutschsprachiger Unternehmensleitbilder – Versuch einer empirischen (Re-)Konstruktion. 2005.

Band 6 Matthias Fink: Erfolgsfaktor Selbstverpflichtung bei vertrauensbasierten Kooperationen. Mit einem empirischen Befund. 2005.

Band 7 Michael Gerhard Kraft: Ökonomie zwischen Wissenschaft und Ethik. Eine dogmenhistorische Untersuchung von Léon M.E. Walras bis Milton Friedman. 2005.

Band 8 Ingrid Zechmeister: Mental Health Care Financing in the Process of Change. Challenges and Approaches for Austria. 2005.

Band 9 Sarah Meisenberger: Strukturierte Organisationen und Wissen. 2005.

Band 10 Anne-Katrin Neyer: Multinational teams in the European Commission and the European Parliament. 2005.

Band 11 Birgit Trukeschitz: Im Dienst Sozialer Dienste. Ökonomische Analyse der Beschäftigung in sozialen Dienstleistungseinrichtungen des Nonprofit Sektors. 2006

Band 12 Marcus Kölling: Interkulturelles Wissensmanagement. Deutschland Ost und West. 2006.

Band 13 Ulrich Berger: The Economics of Two-way Interconnection. 2006.

Band 14 Susanne Guth: Interoperability of DRM Systems. Exchanging and Processing XML-based Rights Expressions. 2006.

Band 15 Bernhard Klement: Ökonomische Kriterien und Anreizmechanismen für eine effiziente Förderung von industrieller Forschung und Innovation. Mit einer empirischen Quantifizierung der Hebeleffekte von F&E-Förderinstrumenten in Österreich. 2006.

Band 16 Markus Imgrund: Wege aus der Insolvenz. Eine Analyse der Fortführung und Sanierung insolventer Klein- und Mittelbetriebe unter besonderer Berücksichtigung des Konfigurationsansatzes. 2007.

Band 17 Nicolas Knotzer: Product Recommendations in E-Commerce Retailing Applications. 2008.

Band 18 Astrid Dickinger: Perceived Quality of Mobile Services. A Segment-Specific Analysis. 2007.

Band 19 Nadine Wiedermann-Ondrej: Hybride Finanzierungsinstrumente in der nationalen und internationalen Besteuerung der USA. 2008.

Band 20 Helmut Sorger: Entscheidungsorientiertes Risikomanagement in der Industrieunternehmung. 2008.

Band 21 Martin Rietsch: Messung und Analyse des ökonomischen Wechselkursrisikos aus Unternehmenssicht: Ein stochastischer Simulationsansatz. 2008.

www.peterlang.de